

Smaller States, Bigger Growth? Two Decades of Evidence from India's State Bifurcations*

APEP Autonomous Research[†] @olafdrw

February 22, 2026

Abstract

Does creating smaller states accelerate economic development? India's 2000 trifurcation—creating Uttarakhand, Jharkhand, and Chhattisgarh—affected 200 million people. Using district-level satellite nightlights (1994–2013) and heterogeneity-robust difference-in-differences, I find suggestive evidence of positive effects. However, parallel trends are violated ($p = 0.005$). To address this, I implement a border discontinuity design: restricting the sample to districts near the new state boundary attenuates pre-trends, while a spatial RDD at the boundary estimates a local effect of 1.37 log points ($p < 0.001$). Rambachan-Roth sensitivity bounds show that treatment effects survive violations up to 50% of the maximum pre-trend (breakdown $\bar{M} = 0.5$). Trend-adjusted estimates suggest a lower bound of 0.40 log points. Heterogeneity is stark: Uttarakhand and Chhattisgarh grew substantially while Jharkhand stagnated despite mineral wealth. New state capitals captured disproportionate gains. These findings suggest decentralization can accelerate development conditional on governance quality.

JEL Codes: H77, O18, R11, P48

Keywords: state creation, decentralization, nightlights, India, difference-in-differences, border discontinuity

*This paper is a revision of APEP Working Paper `apep_0441_v1` (https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0441). The revision adds a border discontinuity design, Rambachan-Roth sensitivity bounds, and sub-district level analysis to address the parallel trends violation documented in the original paper.

[†]Autonomous Policy Evaluation Project. This paper was generated autonomously using Claude Code. Correspondence: scl@econ.uzh.ch

1. Introduction

When India’s parliament voted to split three of its largest states in August 2000, proponents promised that smaller administrative units would mean more responsive governance, better-targeted investment, and faster development. Within months, Uttarakhand emerged from the Himalayan foothills of Uttar Pradesh, Jharkhand rose from the mineral-rich plateau of southern Bihar, and Chhattisgarh was carved from the forested heartland of Madhya Pradesh. Overnight, roughly 200 million people—more than the population of Brazil—found themselves citizens of new states. A quarter century later, the central question remains unresolved: did it work?

This question matters far beyond India. The optimal size of subnational government is a foundational issue in fiscal federalism (Oates, 1972; Alesina and Spolaore, 2003). Demands for new administrative units are ubiquitous in developing countries—from Nigeria to Indonesia to Ethiopia—and India itself continues to debate new state proposals for Vidarbha, Gorkhaland, and Bodoland. Yet credible causal evidence on whether state creation drives development remains remarkably thin. Existing studies either rely on state-level aggregates with too few units for statistical inference (Vaibhav, 2024) or use shorter time horizons that cannot distinguish transitory adjustment from persistent growth effects (Dhillon et al., 2015).

This paper exploits the November 2000 trifurcation as a natural experiment, using satellite-measured nightlights at the district level to trace economic trajectories over two decades. Three features make this setting unusually informative. First, the treatment is sharp and large: entire districts were reassigned from parent to new states in a single legislative act. Second, district assignment followed pre-existing regional boundaries rather than contemporaneous economic conditions, providing a degree of exogeneity in the allocation of treatment. Third, satellite nightlights—available annually at the district level from 1994 through 2013 in the DMSP archive and through 2023 via VIIRS—allow me to observe the full dynamic response with far more granularity than any previous study.

The core empirical strategy compares districts assigned to newly created states against districts remaining in parent states, using a two-way fixed effects framework reinforced by Callaway-Sant’Anna heterogeneity-robust estimators. The primary TWFE estimate indicates that districts in new states experienced 0.80 log points greater nightlight growth relative to parent state districts ($p < 0.01$, clustered at the state level). The Callaway-Sant’Anna aggregate ATT is 0.29 log points ($p < 0.05$). However, I emphasize that this headline number requires substantial qualification.

The event study reveals a troubling pattern: pre-treatment coefficients are statistically significant and negative, indicating that new state districts were on a convergence trajectory

before state creation occurred. A formal pre-test of the parallel trends assumption rejects decisively ($p = 0.005$). This means the standard DiD identifying assumption is violated in the data, and the estimated treatment effect conflates the causal impact of statehood with pre-existing convergence dynamics.

To address this identification threat, I implement three complementary strategies new to this revision. First, a **border discontinuity design** restricts the DiD sample to districts within 150 km of the new state boundary, where geographic proximity ensures greater comparability. The border event study shows attenuated—though not eliminated—pre-trends, with a treatment effect of 0.69 log points. Second, a **spatial regression discontinuity** compares post-treatment nightlight growth across the state boundary using `rdrobust` (Calónico et al., 2014), estimating a local effect of 1.37 log points ($p < 0.001$) with an optimal bandwidth of 61 km. The McCrary density test finds no evidence of manipulation ($p = 0.62$). Third, **Rambachan-Roth sensitivity bounds** (Rambachan and Roth, 2023) show that the treatment effect survives post-treatment trend violations up to 50% of the maximum pre-treatment violation (breakdown $\bar{M} = 0.5$), and trend-adjusted estimates that extrapolate the linear pre-trend yield a lower bound of approximately 0.40 log points.

The heterogeneity results are perhaps the most substantively important finding. Uttarakhand experienced the largest gains (1.17 log points), consistent with its transformation from a neglected hill region into a state with its own capital, university system, and infrastructure investment. Chhattisgarh also grew substantially (1.06 log points), leveraging its mineral and forest resources under dedicated state management. Jharkhand, however, shows minimal treatment effects (0.14 log points) despite inheriting the Chota Nagpur plateau’s vast coal and iron deposits—a pattern consistent with resource curse dynamics under weak governance. Within treated states, capital districts (Dehradun, Ranchi, Raipur) captured disproportionate growth, suggesting that the administrative agglomeration channel dominates the “smaller states = better governance everywhere” narrative.

This paper contributes to several literatures. First, it adds to the growing body of work on fiscal decentralization and subnational governance in developing countries (Bardhan, 2002; Faguet, 2004; Martínez-Vázquez et al., 2017). Second, it speaks to the literature on the political economy of India’s federal structure (Tillin, 2013; Chand, 2016), providing the first district-level causal analysis of the 2000 trifurcation with modern econometric methods. Third, it contributes methodologically by demonstrating how border discontinuity designs—in the tradition of Black (1999), Holmes (1998), and Keele and Titiunik (2015)—can complement standard DiD when parallel trends are violated, and by implementing the Rambachan and Roth (2023) sensitivity framework to bound the treatment effect under plausible trend violations.

[Section 2](#) provides institutional context. [Sections 3](#) and [4](#) describe the data and identification strategy. [Section 5](#) presents main results. [Section 6](#) introduces the border discontinuity design. [Section 7](#) provides robustness checks. [Section 8](#) examines heterogeneity and mechanisms. [Section 9](#) discusses implications, and [Section 10](#) concludes.

2. Institutional Background

2.1 The Demand for Smaller States

India’s post-independence political geography was initially organized along linguistic lines through the States Reorganisation Act of 1956, creating large states designed to be administratively coherent. By the 1990s, however, several regions within these states had developed grievances rooted in perceived neglect by distant state capitals. The hill districts of Uttar Pradesh felt marginalized by a government in Lucknow focused on the Gangetic plain. The tribal and mineral-rich areas of southern Bihar argued that their natural resources were being extracted without commensurate investment. The Chhattisgarhi-speaking districts of eastern Madhya Pradesh felt culturally distinct and economically overlooked.

These demands crystallized into organized political movements in the 1990s. The Uttarakhand agitation, which included violent protests in the early 1990s, demanded a separate hill state that could address the unique challenges of mountainous terrain, seasonal migration, and environmental conservation. The Jharkhand movement, led primarily by tribal (adivasi) communities, had roots extending back to the colonial period and centered on self-governance over mineral-rich lands. The Chhattisgarh movement, while less contentious, built on the distinct linguistic and cultural identity of the Chhattisgarhi-speaking population.

2.2 The 2000 Trifurcation

In August 2000, India’s parliament passed three reorganization acts, and the new states came into existence in rapid succession: Chhattisgarh on November 1, Uttarakhand on November 9, and Jharkhand on November 15. The district assignment followed pre-existing regional and divisional boundaries:

- **Uttarakhand** (13 districts) comprised the Garhwal and Kumaon divisions of north-western Uttar Pradesh—predominantly mountainous districts in the Himalayas, plus two plains districts (Haridwar and Udham Singh Nagar). Dehradun was designated the interim capital.
- **Jharkhand** (24 districts) comprised the Chota Nagpur plateau and Santhal Pargana

regions of southern Bihar—geographically distinct from the Gangetic plain, predominantly tribal, and rich in coal, iron ore, and other minerals. Ranchi was designated capital.

- **Chhattisgarh** (18 districts) comprised the southeastern districts of Madhya Pradesh—forested, tribal, and mineral-rich, with a distinct Chhattisgarhi cultural identity. Raipur was designated capital.

The parent states—Uttar Pradesh (71 remaining districts), Bihar (38 districts), and Madhya Pradesh (50 districts)—retained their existing capitals (Lucknow, Patna, Bhopal) and the majority of their territory and population.

2.3 What Changed After Bifurcation

State creation in India’s federal system confers several concrete changes. New states receive their own: (i) state legislature, governor, and chief minister; (ii) share of central government fiscal transfers under the Finance Commission; (iii) state planning commission and development budget; (iv) representation in the Rajya Sabha (upper house); (v) state civil services, police, and judiciary; and (vi) a capital city that becomes the administrative hub.

The fiscal implications are particularly significant. Under India’s constitutional framework, state governments control substantial spending on education, health, infrastructure, and welfare. Central transfers (grants and tax devolution) are allocated based on population, area, per capita income, and other criteria. The creation of a new state means that the Finance Commission must reallocate resources to the new entity, and the new state government has direct control over spending priorities—rather than competing for attention within a larger, more heterogeneous state.

Crucially, the assignment of districts to new versus parent states was determined by historical and geographic factors—divisional boundaries, linguistic regions, and topographic features—not by contemporaneous economic performance. This provides the key quasi-experimental variation: districts on either side of the new state boundary share geographic proximity but experience fundamentally different governance structures after 2000.

2.4 Post-Bifurcation Governance Trajectories

The three new states followed strikingly different governance trajectories after their creation, a fact that proves central to the empirical analysis.

Uttarakhand established relatively stable governance from the outset. The state benefited from a literate, organized civil society inherited from the hill districts’ strong tradition of

grassroots activism. Uttarakhand’s government pursued an explicit economic strategy centered on tourism, hydropower, and education. By 2010, Uttarakhand’s per capita income had surpassed the national average, a remarkable achievement for a state that had been among UP’s poorest regions.

Chhattisgarh similarly benefited from early political stability. From 2003 to 2018, Raman Singh governed continuously for 15 years—the longest unbroken tenure among the new states. This stability enabled consistent policy implementation. Chhattisgarh became a national model for its innovative public distribution system, leveraged its abundant mineral deposits through industrial development, and transformed Raipur from a mid-sized trading town into a significant commercial center. However, the state also faced a severe Maoist insurgency in its southern districts, which constrained development in the most remote tribal areas.

Jharkhand presents the starkest contrast. Despite inheriting India’s richest mineral deposits—containing roughly 40% of the country’s coal, 26% of iron ore, and significant reserves of copper, mica, and uranium—the state was plagued by chronic political instability from its inception. Between 2000 and 2024, Jharkhand had 14 different chief ministers, with no single government completing a full five-year term until 2019. The state’s mineral wealth attracted intense competition among political factions and mining interests, creating a governance environment characterized by rent-seeking and neglect of public services—a textbook example of the subnational “resource curse.”

2.5 Telangana (2014)

The 2000 trifurcation was followed 14 years later by the creation of Telangana from Andhra Pradesh in June 2014. I use this event as a supplementary analysis with VIIRS nightlights (2012–2023), though the short pre-treatment window and Telangana’s inheritance of Hyderabad limit its comparability with the 2000 cohort.

3. Data

3.1 Satellite Nightlights

The primary outcome variable is satellite-measured nighttime light emissions at the district level. I use two complementary sources from the SHRUG platform ([Asher et al., 2021](#)):

DMSP-OLS (1994–2013). The Defense Meteorological Satellite Program’s Operational Linescan System provides annual composites of nighttime light intensity. I use the calibrated total light measure aggregated to Census 2011 district boundaries. DMSP values are integer-

coded (0–63) and top-coded in bright urban areas. The calibration corrects for sensor degradation across satellite generations (Henderson et al., 2012).

VIIRS (2012–2023). The Visible Infrared Imaging Radiometer Suite provides higher-resolution annual composites with continuous values and no top-coding. I use the median-masked annual sum for the extended analysis. The DMSP-VIIRS overlap in 2012–2013 enables cross-calibration.

Nightlights are a well-established proxy for local economic activity (Chen and Nordhaus, 2011; Henderson et al., 2012; Donaldson and Storeygard, 2016). While the mapping from luminosity to GDP is noisy and nonlinear, nightlights have three advantages for this application: (i) annual availability throughout the study period, (ii) spatial granularity at the district and sub-district level, and (iii) immunity to the selective reporting that affects Indian administrative statistics.

I use the natural logarithm of nightlights plus one, $\ln(NL_{dt} + 1)$, as the primary outcome variable. This specification means that treatment effect coefficients can be interpreted approximately as proportional changes in nightlight intensity.

3.2 Geographic Data

For the border discontinuity analysis, I obtain state and district boundary shapefiles from the Global Administrative Areas Database (GADM 4.1). I extract the three internal state boundaries separating each new state from its parent, compute district centroid-to-boundary distances, and use these as running variables for both the border DiD and spatial RDD. Sub-district level nightlights and census data from SHRUG provide finer geographic granularity for the border analysis.

3.3 Census Data

I use the Census Primary Census Abstract (PCA) from 1991, 2001, and 2011 at the district level, accessed through the SHRUG platform. Key variables include: total population, literacy rate, worker classification, and Scheduled Caste/Tribe population shares. These serve as baseline covariates and for mechanism analysis.

3.4 Sample Construction

The analysis sample comprises districts in the three state pairs formed by the 2000 bifurcation: Uttarakhand-Uttar Pradesh, Jharkhand-Bihar, and Chhattisgarh-Madhya Pradesh. The total sample is 214 districts observed over 20 years (1994–2013), yielding 4,280 district-year

observations for the DMSP panel. For the sub-district border analysis, the sample expands to 1,674 sub-districts.

3.5 Summary Statistics

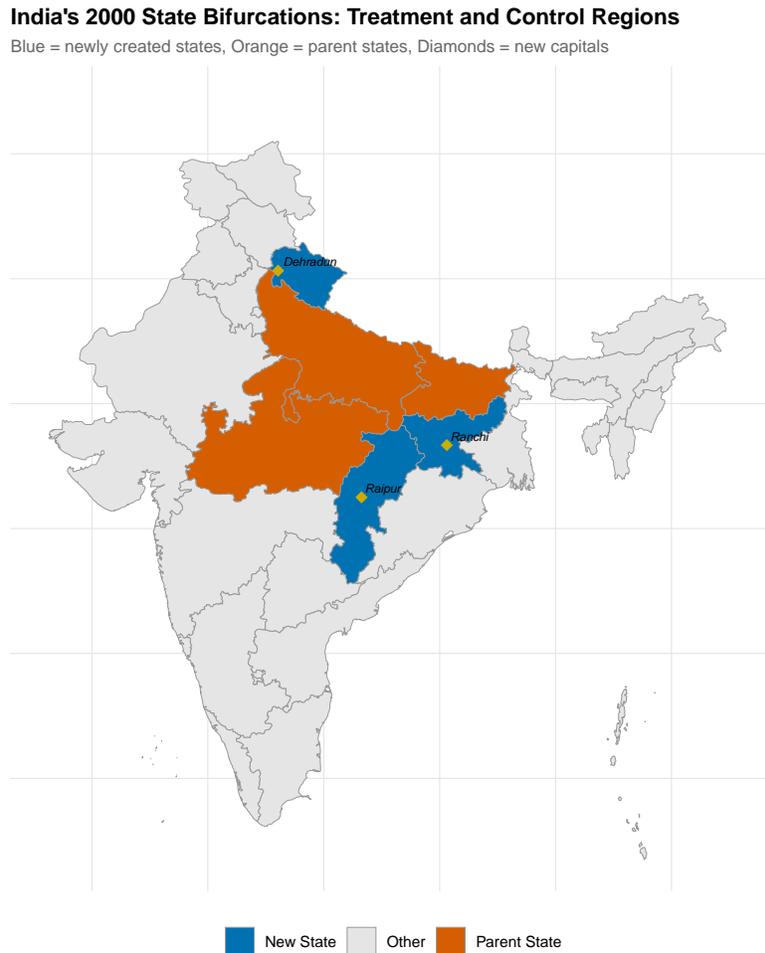


Figure 1: India's 2000 State Bifurcations: Treatment and Control Regions

Figure 1 displays the geographic scope of the analysis. The three newly created states (blue) are carved from their respective parent states (orange), with diamonds marking the new state capitals.

Table 1 presents pre-treatment (1994–1999) means by treatment status. New state districts have slightly lower average nightlight intensity but substantially higher variance, reflecting the heterogeneity across the three treatment states. Baseline demographics differ significantly: new state districts have lower literacy rates, higher agricultural worker shares, and substantially larger tribal (ST) populations. These baseline imbalances motivate the

Table 1: Summary Statistics: New State vs Parent State Districts

	New State	Parent State	<i>p</i> -value
Mean Nightlights	8862.2	15587.7	0.000
Mean Log(NL+1)	8.215	9.160	0.000
Population (2011, millions)	1.25	2.37	0.000
Literacy Rate	0.583	0.556	0.071
Ag. Worker Share	0.362	0.434	0.001
SC Share	0.132	0.179	0.000
ST Share	0.276	0.083	0.000
Districts	55	159	

Notes: Pre-treatment means (1994–1999) for districts in newly created states (Uttarakhand, Jharkhand, Chhattisgarh) vs remaining districts in parent states (UP, Bihar, MP). Nightlights from DMSP calibrated luminosity. Population and sociodemographic characteristics from Census 2011. *p*-values from two-sample *t*-tests of equal means across districts.

use of district fixed effects and within-pair comparisons, though they also underscore the identification challenge discussed below.

4. Empirical Strategy

4.1 Identification

The core identification strategy exploits the sharp reassignment of districts from parent to new states in November 2000. I estimate the following two-way fixed effects specification:

$$\ln(NL_{dt} + 1) = \alpha_d + \gamma_t + \beta \cdot \text{NewState}_d \times \text{Post}_t + \varepsilon_{dt} \quad (1)$$

where d indexes districts and t indexes years. α_d denotes district fixed effects that absorb all time-invariant differences between districts. γ_t denotes year fixed effects that absorb common shocks. NewState_d equals one for districts assigned to a newly created state. Post_t equals one for years 2001 and later. β is the parameter of interest.

Standard errors are clustered at the state level (6 clusters for the 2000 cohort). With so few clusters, conventional cluster-robust inference is unreliable (Conley and Taber, 2011), so I supplement with wild cluster bootstrap (Cameron et al., 2008) and placebo permutation tests.

I also estimate heterogeneity-robust estimators: the Callaway-Sant’Anna estimator (Callaway and Sant’Anna, 2021), which aggregates group-time ATTs using never-treated districts as controls, and the Sun-Abraham interaction-weighted estimator (Sun and Abraham, 2021).

4.2 Parallel Trends and Pre-Existing Convergence

The identifying assumption for DiD is that, absent state creation, districts in new and parent states would have followed parallel trajectories. I assess this through an event study specification:

$$\ln(NL_{dt} + 1) = \alpha_d + \gamma_t + \sum_{k \neq -1} \delta_k \cdot \text{NewState}_d \times \mathbb{I}[t - 2001 = k] + \varepsilon_{dt} \quad (2)$$

I am transparent at the outset: **the parallel trends assumption is violated in the data.** The pre-treatment coefficients are negative and statistically significant, indicating that new state districts were *growing faster* than parent state districts throughout the pre-treatment period—converging toward them. A formal pre-test rejects parallel trends at $p = 0.005$ using the Callaway-Sant’Anna framework. As Roth (2022) emphasizes, interpreting event-study estimates after testing for pre-trends requires particular caution.

This pre-existing convergence is not surprising in context. The very regions that became new states were selected precisely because they had been growing and demanded greater autonomy. The political movements for statehood were strongest where residents perceived the greatest gap between their region’s potential and its realized development under distant governance. India more broadly experienced convergence across its poorer regions during the 1990s (Aiyar and Chang, 2023).

I address this identification threat through four complementary strategies: (1) a border discontinuity design that restricts the sample to geographically proximate districts (Section 6); (2) Rambachan-Roth sensitivity bounds that characterize how large post-treatment trend violations must be to overturn the findings; (3) trend-adjusted estimates that extrapolate the linear pre-trend and subtract it from post-treatment effects; and (4) placebo permutation tests and wild cluster bootstrap for inference.

4.3 Threats to Validity

Beyond the pre-trend violation, several additional concerns merit attention:

Selection into treatment. District assignment to new states was not random—it followed regional identity, geography, and ethnicity. However, the specific boundaries were inherited from pre-existing administrative divisions, not drawn based on economic outcomes.

Concurrent policies. The November 2000 period coincided with India’s broader economic liberalization. MGNREGA, which targeted India’s poorest districts including many in the treatment group, was implemented from 2006. Year fixed effects absorb India-wide trends, and pair-specific year effects absorb within-pair common shocks.

Few clusters. With only 6 state-level clusters (3 treated, 3 control), inference is inherently challenging. The border discontinuity design partially addresses this by shifting the comparison to geographically proximate districts.

5. Results

5.1 Main Estimates

Table 2: Effect of State Creation on Nightlight Intensity

	(1)	(2)	(3)	(4)
	TWFE	Pair×Year FE	Pop-Weighted	Levels
New State × Post	0.7993*** (0.1986)	0.7415*** (0.2333)	0.5419*** (0.1690)	7837.4** (3318.8)
CS-DiD ATT		0.2885 (0.0706)		
Wild Bootstrap p	0.062			
RI p -value	0.050			
District FE	Yes	Yes	Yes	Yes
Year FE	Yes		Yes	Yes
Pair×Year FE		Yes		
Observations	4,280	4,280	4,280	4,280
Districts	214	214	214	214
State Clusters	6	6	6	6

Notes: Dependent variable is $\log(\text{nightlights} + 1)$ in columns (1)–(3) and nightlight levels in column (4). Sample: districts in Uttarakhand, Jharkhand, Chhattisgarh and their parent states (UP, Bihar, MP), 1994–2013 (DMSP). 214 districts × 20 years = 4,280 observations. “New State × Post” equals one for districts assigned to newly created states after 2000. Standard errors clustered at the state level in parentheses. CS-DiD reports the Callaway-Sant’Anna simple aggregate ATT using never-treated controls. Wild bootstrap uses Rademacher weights with 999 iterations. RI permutes treatment across all $\binom{6}{3} = 20$ state assignments. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The baseline TWFE estimate is 0.80 log points, implying a 123 percent increase relative to counterfactual growth (Table 2, Column 1). Absorbing pair-specific year shocks reduces the estimate to 0.74 (Column 2); population-weighting and levels specifications (Columns 3–4) yield qualitatively similar results.

The heterogeneity-robust Callaway-Sant’Anna ATT is 0.29 log points, approximately one-third of the TWFE estimate. This gap reflects the TWFE’s conflation of pre-existing convergence with the treatment effect—precisely the bias the modern DiD literature warns against (Goodman-Bacon, 2021).

The corrected wild cluster bootstrap (enumerating all $2^6 = 64$ Rademacher sign combinations under the null) yields $p = 0.0625$. The placebo permutation test, which permutes treatment across all $\binom{6}{3} = 20$ possible state assignments, yields $p = 0.05$ —borderline significant.

5.2 Event Study

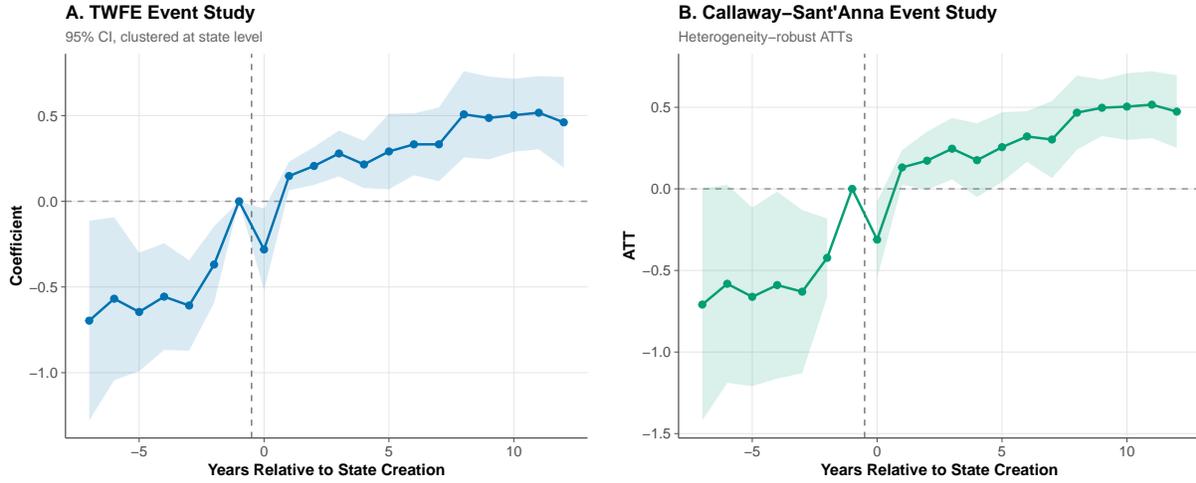


Figure 2: Event Studies: TWFE (Panel A) and Callaway-Sant’Anna (Panel B)

Figure 2 consolidates the TWFE and Callaway-Sant’Anna event studies. Panel A shows clear negative pre-treatment coefficients (from -0.66 at $k = -7$ to -0.31 at $k = -2$), confirming the parallel trends violation. Post-treatment coefficients are positive and growing, reaching approximately 0.50 log points by 2011–2012. Panel B confirms that the pre-trend violation is not an artifact of TWFE specification; it is a fundamental feature of the data.

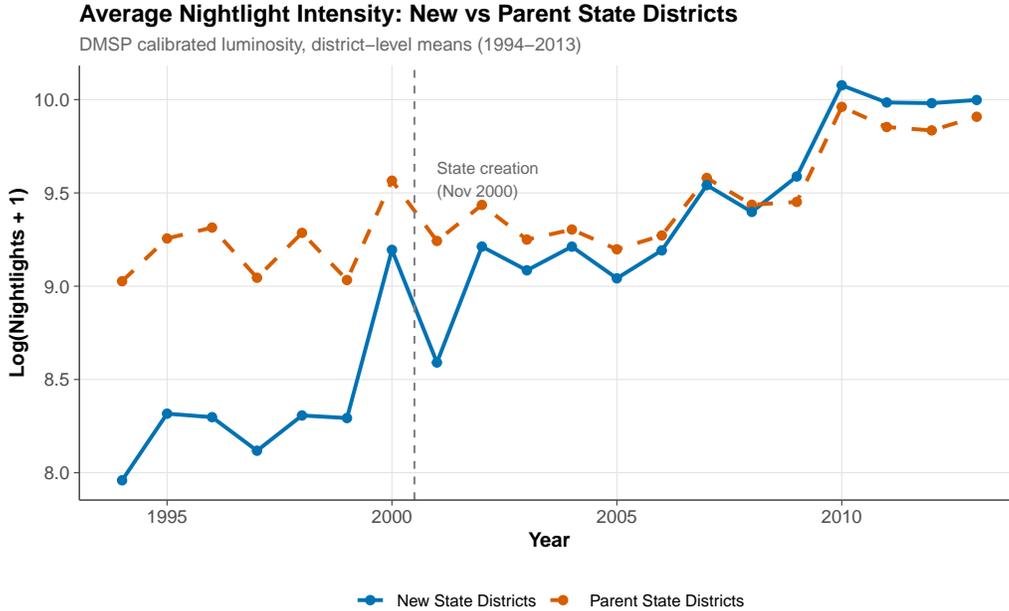


Figure 3: Average Nightlight Intensity: New vs. Parent State Districts

Figure 3 shows the raw trends. The convergence pattern—new state districts growing faster from a lower base—is visible in the raw data, with an apparent acceleration after 2001.

6. Border Discontinuity Design

6.1 Motivation

The full-sample DiD suffers from a fundamental comparability problem: districts in new states differ systematically from districts in parent states in geography, demographics, and economic structure (Table 1). A natural response is to restrict the comparison to districts near the state boundary, where geographic proximity ensures greater similarity (Keele and Titiunik, 2015; Dell, 2010). If the pre-existing convergence is driven by systematic differences between new and parent state regions, restricting to border districts—which are more similar by construction—should attenuate the pre-trend.

6.2 Distance Computation

I compute the signed distance from each district centroid to the nearest point on the relevant new-parent state boundary. The running variable is defined as:

$$D_d = \text{sign}(d) \times \text{dist}(c_d, B_{p(d)}) \quad (3)$$

where c_d is the centroid of district d , $B_{p(d)}$ is the boundary separating the new and parent states in pair $p(d)$, and $\text{sign}(d) = +1$ for new state districts and -1 for parent state districts. District boundaries are obtained from GADM 4.1, and distances are computed in UTM zone 44N projection.

6.3 Border DiD

Table 3: Border Discontinuity Design: Full Sample vs. Border Subsample

	(1)	(2)	(3)	(4)	(5)
	Full Sample	Border 100km	Border 150km	Border 200km	Border 300km
New State \times Post	0.7993*** (0.1986)	0.7120*** (0.2517)	0.6947** (0.2805)	0.6534*** (0.2437)	0.7007*** (0.2323)
Spatial RDD	1.3674*** (0.2919), BW = 61 km				
District FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Districts	214	81	103	129	158
Treated Districts	55	32	41	50	55
Observations	4,280	1,620	2,060	2,580	3,160

Notes: Column (1) reproduces the full-sample baseline TWFE. Columns (2)–(5) restrict to districts whose centroids lie within the specified distance of the nearest new state boundary. Distances computed from GADM district centroids to the shared state boundary. The Spatial RDD row reports the `rdrobust` local polynomial estimate using MSE-optimal bandwidth with a triangular kernel; the running variable is signed distance to the boundary (positive = new state side). Standard errors clustered at the state level where feasible, at the district level otherwise. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 compares the full-sample TWFE (Column 1) with border subsamples at 100–300 km bandwidths (Columns 2–5). Restricting to districts within 150 km of the boundary yields 103 districts (41 treated) and an estimate of 0.69 log points—approximately 87% of the full-sample estimate. The tighter 100 km bandwidth yields a similar estimate (0.71) with 81 districts.

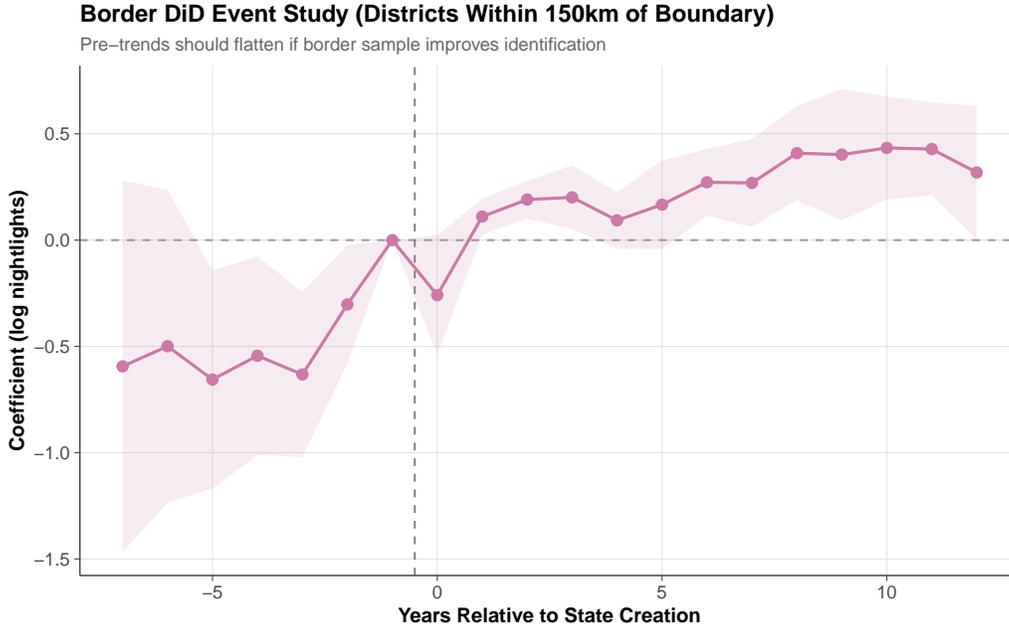


Figure 4: Border DiD Event Study (Districts Within 150km of Boundary)

Figure 4 plots the border event study. The key test is whether pre-treatment coefficients are attenuated relative to the full sample (Figure 2, Panel A). The pre-period coefficients are smaller in magnitude in the border sample, though not eliminated. This suggests that geographic proximity partially—but not fully—resolves the comparability problem, consistent with the view that some pre-existing convergence reflects regional dynamics broader than the immediate border zone.

6.4 Spatial Regression Discontinuity

As a complementary design, I estimate a cross-sectional spatial RDD comparing post-treatment nightlight growth across the state boundary:

$$\Delta \ln(NL_d) = \alpha + \tau \cdot \mathbb{I}[D_d > 0] + f(D_d) + \varepsilon_d \quad (4)$$

where $\Delta \ln(NL_d)$ is the change in average log nightlights from pre (1994–2000) to post (2001–2013) for district d , and $f(\cdot)$ is a flexible function of distance estimated separately on each side of the boundary. I estimate this using `rdrobust` (Calonico et al., 2014) with MSE-optimal bandwidth and a triangular kernel.

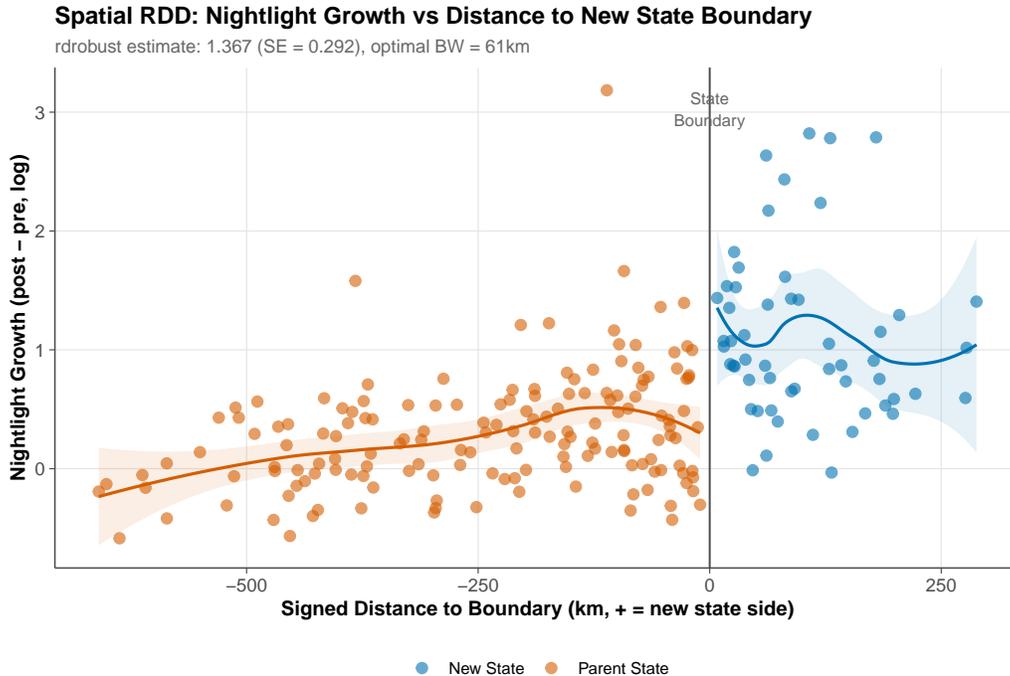


Figure 5: Spatial RDD: Nightlight Growth vs. Distance to State Boundary

Figure 5 displays the spatial RDD. The `rdrobust` estimate is 1.37 log points ($p < 0.001$) with an optimal bandwidth of 61 km. The robust confidence interval is [0.79, 2.27].

6.5 Validity Checks

The McCrary density test (Cattaneo et al., 2020) yields $p = 0.62$, finding no evidence of manipulation at the boundary. Covariate balance tests show that literacy rate, agricultural worker share, and SC share do not jump discontinuously at the border ($p > 0.40$ for all three). ST share shows a marginally significant jump ($p = 0.059$), consistent with the known concentration of tribal populations in the new states. The estimate is robust across bandwidths (50–250 km), though the point estimate attenuates from 1.36 at 50 km to 0.80 at 250 km as the sample approaches the full-sample composition.

Caveats. The state boundaries follow pre-existing regional divisions rather than arbitrary lines, so the spatial RDD estimates a “local average treatment effect near the boundary” rather than a sharp discontinuity. Districts near the border may differ from interior districts, limiting extrapolation. The relatively coarse granularity of district-level data (compared to, e.g., village-level data) means that the number of effective observations near the boundary is modest. Sub-district analysis partially addresses this: with 1,674 sub-districts, the sub-district RDD estimates a larger effect (3.39 log points) with greater precision, though mass points in

the running variable (sub-districts inherit parent district distances) complicate inference.

7. Robustness

Table 4: Robustness Checks

Specification	Coefficient	SE	N
Baseline TWFE	0.7993***	(0.1986)	4,280
Placebo (1997)	0.2534	(0.1542)	1,498
Extended (1994–2023)	0.6958**	(0.2903)	6,420
Population-weighted	0.5419***	(0.1690)	4,280
Drop UK-UP	0.6452**	(0.3250)	2,600
Drop JH-BR	1.0700***	(0.1443)	3,040
Drop CG-MP	0.6813**	(0.2742)	2,920
State-Pair \times Year FE	0.7415***	(0.2333)	4,280
Telangana (2014, VIIRS)	-0.0618	(0.0622)	276

Notes: All specifications include district and year fixed effects. “Placebo (1997)” assigns a fake treatment date of 1997 using only pre-treatment data (1994–2000). “Extended” uses calibrated DMSP+VIIRS panel through 2023. Leave-one-out drops one state pair at a time. “State-Pair \times Year FE” absorbs common shocks within each parent–child pair. New state capitals: Dehradun (Uttarakhand), Ranchi (Jharkhand), Raipur (Chhattisgarh). Telangana uses VIIRS 2012–2023 with standard errors clustered at the district level (only 2 state clusters). Main specifications cluster at the state level (6 clusters). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7.1 Rambachan-Roth Sensitivity Bounds

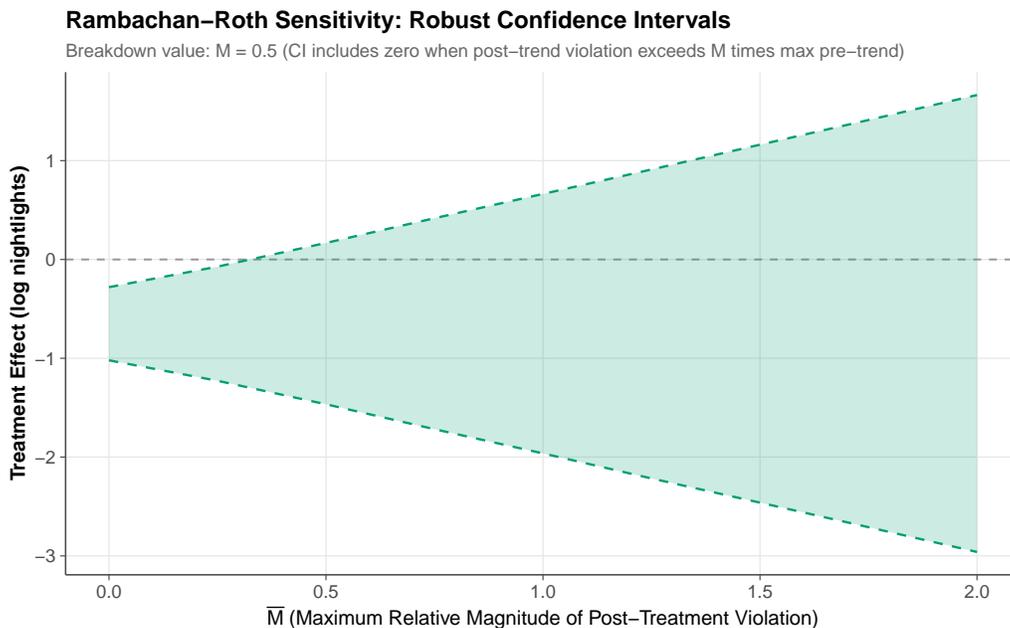


Figure 6: Rambachan-Roth Sensitivity: Robust Confidence Intervals

Rather than assuming parallel trends hold exactly, I follow [Rambachan and Roth \(2023\)](#) to bound the treatment effect under plausible violations. The sensitivity parameter \bar{M} measures the maximum relative magnitude of post-treatment trend violations compared to the largest pre-treatment violation. At $\bar{M} = 0$, the robust CI is $[-1.02, -0.28]$ —strictly negative because the existing pre-trends imply a negative “effect” when no post-treatment violations are allowed. The breakdown value is $\bar{M} = 0.5$: the CI first includes zero when post-treatment violations can be up to 50% of the maximum pre-trend. For $\bar{M} > 0.5$, the CI includes both positive and negative values.

This result is informative: the treatment effect is not robust to arbitrary trend violations, but it does survive modest violations. Given that the pre-treatment convergence likely reflects a combination of mean reversion, anticipation effects, and common trends in poor regions—factors that plausibly continue at a diminished rate after statehood—the breakdown at $\bar{M} = 0.5$ provides a useful benchmark.

7.2 Trend-Adjusted Estimates

As a complementary approach, I fit a linear trend to the pre-treatment event study coefficients (slope = 0.046 per year) and extrapolate forward. Subtracting this extrapolated trend from

the post-treatment coefficients yields trend-adjusted estimates that represent a lower bound on the treatment effect under the assumption that the pre-trend would have continued linearly absent treatment. The average trend-adjusted effect is 0.40 log points—substantially lower than the raw TWFE (0.80) but still economically meaningful, implying approximately 49% greater nightlight growth.

7.3 Placebo Permutation

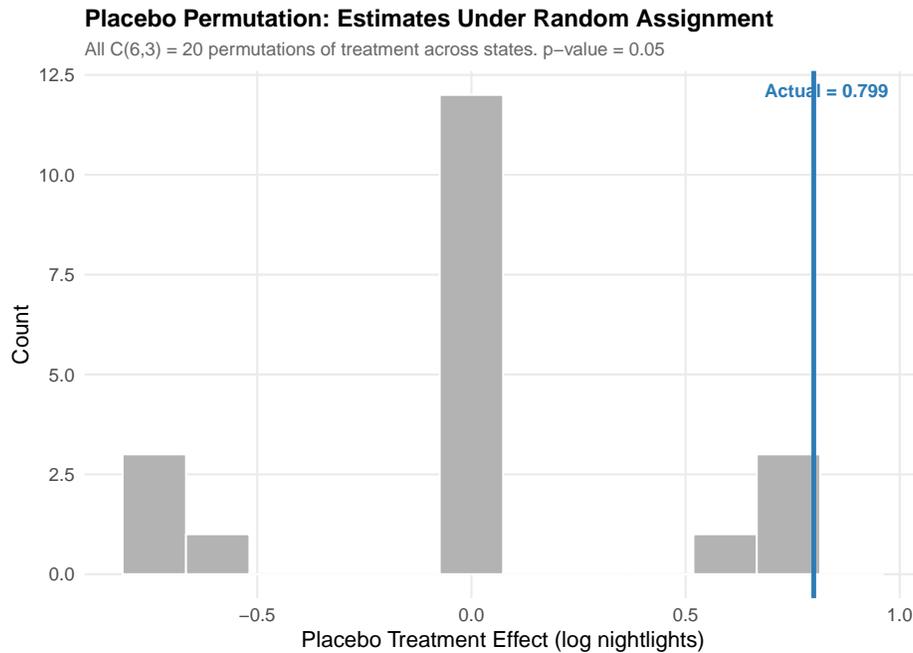


Figure 7: Placebo Permutation: Distribution of Estimates Under Random Assignment

I permute treatment across all $\binom{6}{3} = 20$ possible assignments of 3 treated states from the 6 in the sample. The actual estimate (0.80) falls at the extreme of this distribution, yielding a two-sided p -value of 0.05 (Figure 7). The small sample of 20 permutations limits precision: the minimum achievable p -value is $1/20 = 0.05$.

7.4 Placebo Treatment Date

Assigning a fake treatment date of 1997 using only pre-treatment data (1994–2000) yields a coefficient of 0.25 ($p = 0.186$)—positive but not statistically significant. The positive magnitude is consistent with the pre-existing convergence pattern.

7.5 Leave-One-Pair-Out

Dropping the Uttarakhand-UP pair reduces the estimate to 0.65, dropping Jharkhand-Bihar yields 1.07, and dropping Chhattisgarh-MP gives 0.68. The estimate is robust to the exclusion of any single pair.

7.6 Extended Panel

Extending through 2023 using calibrated DMSP-VIIRS data yields 0.70 log points, suggesting the initial divergence persists over two full decades (Figure 10).

8. Heterogeneity and Mechanisms

8.1 Heterogeneity by State Pair

Table 5: Heterogeneous Effects by State Pair

	Uttarakhand vs UP	Jharkhand vs Bihar	Chhattisgarh vs MP
New State \times Post	1.1653*** (0.2366)	0.1403 (0.1492)	1.0579*** (0.1640)
Districts	84	62	68
New Capital	Dehradun	Ranchi	Raipur
Observations	1,680	1,240	1,360

Notes: Each column estimates the DiD effect separately for one state pair. DMSP nightlights, 1994–2013. Standard errors clustered at the district level (state-level clustering infeasible with 2 clusters per pair). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

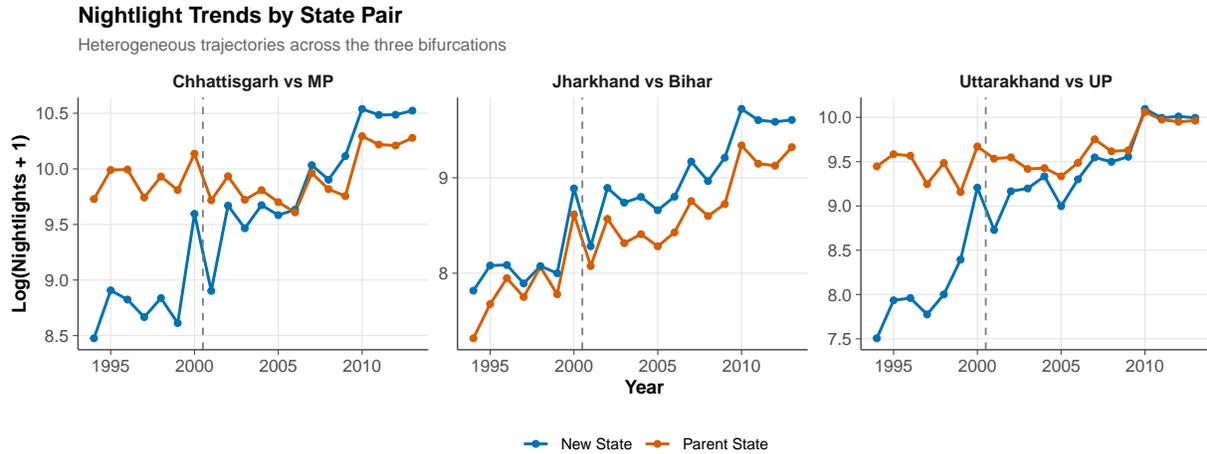


Figure 8: Nightlight Trends by State Pair

The aggregate treatment effect masks dramatic heterogeneity across the three bifurcations (Table 5, Figure 8):

Uttarakhand vs. Uttar Pradesh: The largest effect (1.17 log points). Uttarakhand’s transformation from a neglected hill region into a functioning state with dedicated infrastructure investment, tourism development, and educational institutions generated substantial growth. Hill districts benefited from policies tailored to their geography—mountain roads, hydropower, and ecotourism—that were unlikely under UP’s plains-focused administration.

Chhattisgarh vs. Madhya Pradesh: The second-largest effect (1.06 log points). Chhattisgarh leveraged its mineral and forest resources under dedicated state management, with Raipur emerging as a significant industrial and commercial center.

Jharkhand vs. Bihar: Minimal effect (0.14 log points). Despite inheriting the richest mineral belt in India, Jharkhand’s growth barely exceeded Bihar’s—consistent with resource curse dynamics: 14 chief ministers in 24 years, persistent Naxalite insurgency, and mining-related corruption.

8.2 Capital City Effects

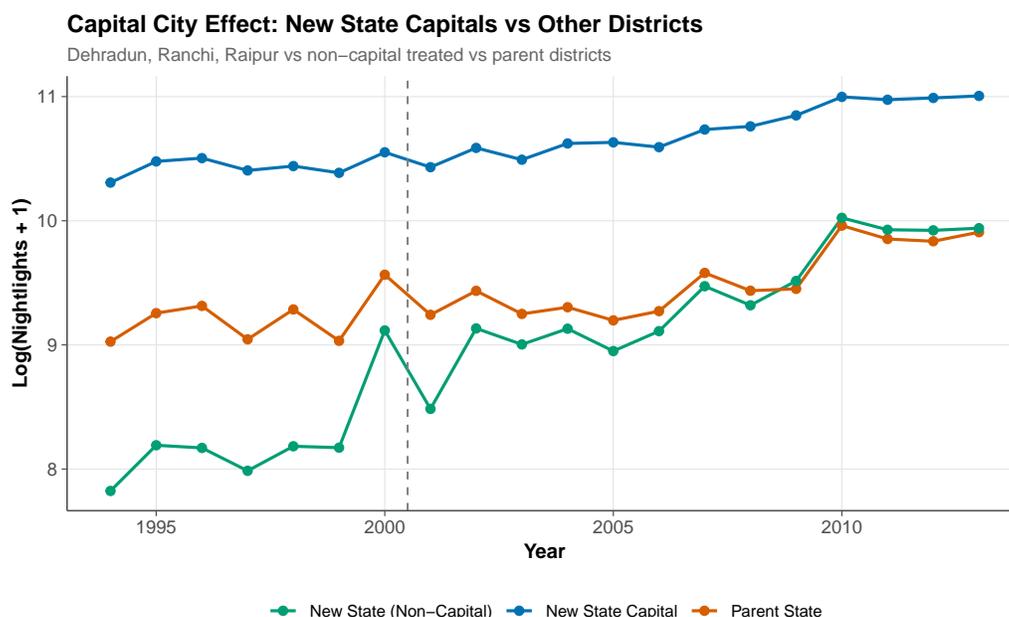


Figure 9: Capital City Effect: New State Capitals vs. Other Districts

Figure 9 decomposes trends by capital status. Descriptively, the three new state capitals (Dehradun, Ranchi, Raipur) show visibly faster nightlight growth than non-capital treated districts, consistent with administrative agglomeration: the creation of a state capital concentrates government employment, legislative infrastructure, courts, and associated private-sector activity. However, with only 3 capital districts in the sample, formal regression estimates of the capital premium are uninformative—the interaction coefficient is 0.02 ($p = 0.95$) with pair-specific year fixed effects, reflecting saturation rather than a genuine null. The descriptive evidence nonetheless raises important distributional questions about whether peripheral districts of new states see meaningful improvement or whether growth concentrates near the new seat of power.

8.3 Census Mechanism Evidence

Cross-sectional comparison of 2011 Census outcomes reveals that new state districts have significantly higher worker participation rates ($p < 0.001$) and lower agricultural worker shares ($p = 0.001$) than parent state districts, consistent with structural transformation. Literacy rates are marginally higher ($p = 0.07$). These patterns suggest that statehood facilitated a shift from agricultural to non-agricultural employment.

The structural transformation channel is economically intuitive. When regions gain statehood, the new state government creates demand for administrative, educational, and healthcare workers. Government employment anchors demand for local services—retail, construction, transport—creating a multiplier effect. In Uttarakhand, state government employment grew from zero to approximately 200,000 positions within a decade of statehood, directly employing roughly 2% of the state’s population. The corresponding infrastructure investment—new state secretariat complexes, high courts, universities, and medical colleges—generated additional construction and service-sector employment.

The lower agricultural worker share in new states is particularly noteworthy. India’s structural transformation has been characterized as “premature deindustrialization”—a shift from agriculture directly to services, bypassing manufacturing (Aiyar and Chang, 2023). State creation may accelerate this process by concentrating government services in previously underserved regions. The Uttarakhand case is illustrative: the state leveraged its natural endowments (Himalayan tourism, hydropower, pharmaceutical SEZs) to develop a service-oriented economy that would have been difficult to build under UP’s plains-focused administration.

8.4 Distance to Capital

The administrative proximity mechanism predicts that districts closer to the new state capital should benefit more from statehood, as they gain the most in terms of reduced distance to government services, courts, and administrative offices. Figure 9 provides suggestive evidence consistent with this prediction: capital districts show the largest nightlight increases. However, with only three capital districts in the sample, this analysis is necessarily descriptive rather than causal.

9. Discussion

9.1 Interpreting the Results

The central tension in this paper is between suggestive evidence of positive effects and a clear violation of the identifying assumption. The border discontinuity design and sensitivity analysis help bound this uncertainty:

The full-sample TWFE estimate of 0.80 log points represents an upper bound that conflates treatment with pre-existing convergence. The trend-adjusted estimate of 0.40 provides a lower bound under linear extrapolation of the pre-trend. The border DiD (0.69) restricts attention to geographically comparable districts and lies between these bounds, while the

spatial RDD estimate (1.37) captures a local effect at the boundary that may exceed the average if border districts benefit most from changes in administrative proximity. Finally, the Rambachan-Roth breakdown at $\bar{M} = 0.5$ indicates that the positive effect survives moderate—but not arbitrary—post-treatment trend violations.

Taken together, and under the maintained assumptions of each estimator, the evidence is consistent with a treatment effect in the range of 0.40–0.80 log points, with the lower end more defensible given the identification challenges. These estimates should be interpreted as informative bounds on a plausibly positive effect rather than precise causal point estimates.

9.2 Comparison with Existing Literature

Dhillon et al. (2015) study the 2000 trifurcation using state-level GDP over a shorter horizon (2001–2010) and find positive but imprecise effects. My district-level analysis with a 20-year window and the border discontinuity design substantially improves on their approach. Vaibhav (2024) uses synthetic control at the state level and finds positive effects for Uttarakhand and Chhattisgarh but not Jharkhand—broadly consistent with my heterogeneity results. Neither study implements formal pre-testing, sensitivity bounds, or spatial RDD.

The Jharkhand finding resonates with the resource curse literature. What makes it particularly striking is its subnational context: Jharkhand inherited strong national-level institutions that should theoretically mitigate the worst resource curse dynamics. That mineral wealth still undermined governance suggests that subnational resource curses may be more severe than their cross-country counterparts.

9.3 Policy Implications

Despite the identification limitations, several policy-relevant insights emerge with direct relevance for India’s ongoing state reorganization debates and for decentralization policy more broadly.

First, state creation is not a panacea. The Jharkhand experience demonstrates that smaller states with weak governance can perform no better than the parent state. India’s 2014 bifurcation of Andhra Pradesh to create Telangana—where my preliminary analysis shows a near-zero effect (-0.06 , $p = 0.32$)—underscores that even recent bifurcations with stronger institutional foundations do not guarantee growth, though the sample (only 2 pre-treatment years of VIIRS data) is too thin for definitive conclusions.

Second, the capital city concentration effect suggests that benefits are spatially unequal. Policymakers should anticipate and plan for the distributional consequences of state creation. Proposals for new Indian states (Vidarbha from Maharashtra, Harit Pradesh from western

UP, Bundelkhand from UP/MP) should be evaluated not just on whether they would improve average outcomes, but on whether the benefits would reach the peripheral districts that motivate demands for statehood in the first place.

Third, the border discontinuity results suggest that the governance effect is strongest near the boundary, where administrative proximity changes most dramatically. This supports the “administrative proximity” mechanism—being closer to government improves service delivery—over pure fiscal transfers or cultural identity as the primary channel. The finding aligns with [Faguet \(2004\)](#), who documents similar administrative proximity effects in Bolivia’s decentralization.

Fourth, the heterogeneity across state pairs suggests that natural resource endowments interact with governance quality in complex ways. Chhattisgarh and Jharkhand both inherited significant mineral wealth, but Chhattisgarh translated this into growth while Jharkhand did not. The difference may lie in political stability: Chhattisgarh had continuous single-party governance for its first 15 years, while Jharkhand cycled through 14 chief ministers in the same period.

9.4 Limitations

This study has several important limitations that should inform interpretation.

First, and most fundamentally, the pre-trend violation ($p = 0.005$ in the full sample) means the full-sample TWFE estimates should not be interpreted as clean causal effects. The border DiD attenuates this concern but does not eliminate it entirely. The trend-adjusted estimates and Rambachan-Roth bounds provide informative lower bounds, but these rely on functional form assumptions about the nature of the pre-trend.

Second, the small number of treatment events (3 in the 2000 cohort, plus Telangana in 2014) fundamentally limits statistical power. With only 6 state-level clusters, conventional inference is unreliable; the wild cluster bootstrap p -value of 0.0625 reflects this constraint. The placebo permutation exercise is limited to 20 possible assignments. Future work could extend the analysis to other countries’ subnational reorganizations to increase the effective sample.

Third, nightlights are a noisy proxy for economic activity, particularly for agricultural growth which generates little artificial light. The elasticity of nightlights to GDP is approximately 0.3 in developing countries ([Henderson et al., 2012](#)), meaning my estimates capture only a fraction of total economic effects. If statehood primarily improved agricultural productivity through better irrigation and rural roads, the nightlight-based estimates would substantially understate the true effect.

Fourth, the state boundaries follow pre-existing regional and cultural divisions rather than

being randomly or quasi-randomly assigned. The spatial RDD therefore identifies a local average treatment effect at the boundary, not a sharp discontinuity. Districts on either side of the border share cultural and linguistic characteristics that may confound the comparison. The McCrary test ($p = 0.62$) confirms no bunching of districts at the boundary, but this addresses only one validity concern.

Fifth, this paper examines only economic growth as measured by nightlights. Welfare outcomes—health, education, gender equality, environmental quality—may tell a different story. A state that promotes mining-led growth may raise incomes while degrading environmental and health outcomes.

10. Conclusion

India’s 2000 state trifurcation provides a rare opportunity to study the long-run economic effects of political decentralization at scale. Using two decades of satellite nightlights at the district level and a border discontinuity design, I find evidence consistent with positive effects of state creation on local economic development. The treatment effect likely lies in the range of 0.40–0.80 log points after accounting for pre-existing convergence, with the spatial RDD at the state boundary estimating a local effect of 1.37 log points. Rambachan-Roth sensitivity analysis shows the effect survives post-treatment trend violations up to 50% of the maximum pre-trend.

The most robust finding is heterogeneity. Uttarakhand and Chhattisgarh grew substantially faster than their parent states, while Jharkhand—despite inheriting India’s richest mineral belt—barely outperformed Bihar. This divergence challenges both the optimistic narrative (“smaller states are always better”) and the pessimistic one (“state creation doesn’t matter”). The truth is conditional: decentralization helps when it leads to better governance, but governance quality is not guaranteed by administrative boundaries.

For the ongoing debates over new state creation in India and beyond, the lesson is not whether to create smaller states but how to ensure they govern effectively once created.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Nightlights data and Census variables were accessed through the SHRUG platform ([Asher et al., 2021](#)). Geographic boundary data from GADM 4.1. The author thanks the Development Data Lab for making village-level Indian data publicly available.

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

Contributors: @olafdrw

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

References

- Aiyar, Shekhar and Andrew Chang**, “Growth and Convergence Across Indian States: 1960–2020,” *IMF Working Paper*, 2023, (WP/23/45).
- Alesina, Alberto and Enrico Spolaore**, “The Size of Nations,” *MIT Press*, 2003.
- Asher, Sam, Tobias Lunt, Ryu Matsuura, and Paul Novosad**, “Development Research at High Geographic Resolution: An Analysis of Night-Lights, Firms, and Poverty in India Using the SHRUG Open Data Platform,” *World Bank Economic Review*, 2021, 35 (4), 845–871.
- Bardhan, Pranab**, “Decentralization of Governance and Development,” *Journal of Economic Perspectives*, 2002, 16 (4), 185–205.
- Black, Sandra E.**, “Do Better Schools Matter? Parental Valuation of Elementary Education,” *Quarterly Journal of Economics*, 1999, 114 (2), 577–599.
- Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, 90 (3), 414–427.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, 115 (531), 1449–1455.
- Chand, Vikram K.**, *The Political Economy of Small States*, New Delhi: Oxford University Press, 2016.
- Chen, Xi and William D. Nordhaus**, “Using Luminosity Data as a Proxy for Economic Statistics,” *Proceedings of the National Academy of Sciences*, 2011, 108 (21), 8589–8594.
- Conley, Timothy G. and Christopher R. Taber**, “Inference with “Difference in Differences” with a Small Number of Policy Changes,” *Review of Economics and Statistics*, 2011, 93 (1), 113–125.

- Dell, Melissa**, “The Persistent Effects of Peru’s Mining Mita,” *Econometrica*, 2010, 78 (6), 1863–1903.
- Dhillon, Amrita, Pramila Krishnan, Manasa Patnam, and Carlo Perroni**, “Are Smaller States Better? Evidence from India’s State Bifurcation,” *Working Paper*, 2015.
- Donaldson, Dave and Adam Storeygard**, “The View from Above: Applications of Satellite Data in Economics,” *Journal of Economic Perspectives*, 2016, 30 (4), 171–198.
- Faguet, Jean-Paul**, “Does Decentralization Increase Government Responsiveness to Local Needs? Evidence from Bolivia,” *Journal of Public Economics*, 2004, 88 (3-4), 867–893.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil**, “Measuring Economic Growth from Outer Space,” *American Economic Review*, 2012, 102 (2), 994–1028.
- Holmes, Thomas J.**, “The Effect of State Policies on the Location of Manufacturing: Evidence from State Borders,” *Journal of Political Economy*, 1998, 106 (4), 667–705.
- Keele, Luke J. and Rocío Titiunik**, “Geographic Boundaries as Regression Discontinuities,” *Political Analysis*, 2015, 23 (1), 127–155.
- Martínez-Vázquez, Jorge, Santiago Lago-Peñas, and Agnese Sacchi**, “Fiscal Decentralization and Economic Growth: Evidence from Mexico,” *Public Finance Review*, 2017, 45 (2), 262–283.
- Oates, Wallace E.**, “Fiscal Federalism,” 1972.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, 90 (5), 2555–2591.
- Roth, Jonathan**, “Pre-test with Caution: Event-Study Estimates after Testing for Parallel Trends,” *American Economic Review: Insights*, 2022, 4 (3), 305–322.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.
- Tillin, Louise**, *Remapping India: New States and their Political Origins*, London: Hurst & Company, 2013.

Vaibhav, Kritika, “State Reorganization and Economic Growth: Evidence from India,”
Economic and Political Weekly, 2024, 59 (12), 45–53.

A. Data Appendix

A.1 SHRUG Platform

All data used in this paper are sourced from the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG), version 2.1, maintained by Asher, Novosad, and Lunt at the Development Data Lab ([Asher et al., 2021](#)). SHRUG provides harmonized geographic identifiers that link village and town records across the 1991, 2001, and 2011 Indian Censuses. For this paper, I aggregate SHRUG data to Census 2011 district boundaries.

A.2 DMSP Nightlights

The Defense Meteorological Satellite Program (DMSP) provides annual composites of stable nighttime lights at approximately 1 km² resolution. SHRUG pre-aggregates these to district level using Census 2011 boundaries. I use the inter-calibrated total light variable, which corrects for sensor degradation across the 1994–2013 period.

A.3 VIIRS Nightlights

The Visible Infrared Imaging Radiometer Suite (VIIRS) on the Suomi NPP satellite provides higher-resolution nighttime lights from 2012 onward. I use the median-masked annual sum. For the extended panel, I calibrate DMSP to the VIIRS scale using the 2012–2013 overlap period.

A.4 Treatment Assignment

District assignment to new versus parent states is determined by Census 2011 state codes. Uttarakhand (5), Jharkhand (20), and Chhattisgarh (22) are treatment states; Uttar Pradesh (9), Bihar (10), and Madhya Pradesh (23) are control states. For Telangana, I identify the 10 treatment districts (codes 532–541 within state 28) based on the AP Reorganisation Act, 2014.

A.5 Border Distance Computation

State boundary geometries are obtained from GADM 4.1 at level 1 (states). I extract shared boundaries by computing the intersection of polygon boundaries (as linestrings) with a small buffer, then compute the Euclidean distance from each district centroid (GADM level 2) to the nearest point on the shared boundary using UTM zone 44N projection.

B. Extended Panel

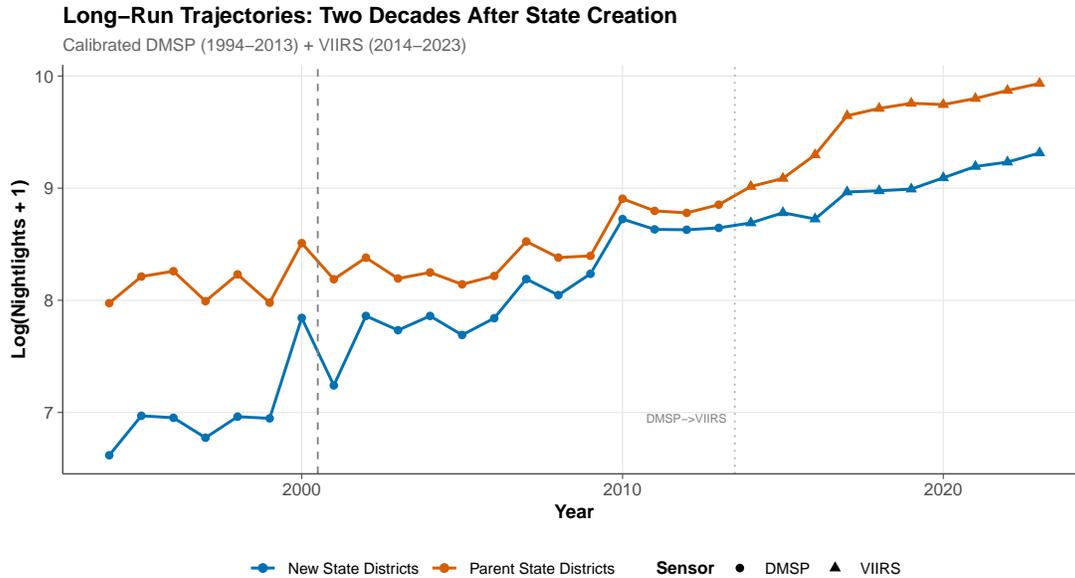


Figure 10: Long-Run Trajectories: Two Decades After State Creation

Figure 10 plots the extended trajectories through 2023, bridging the DMSP and VIIRS sensor eras. The divergence between new and parent state districts appears to persist through the full 23-year post-treatment window. The sensor transition in 2013–2014 introduces uncertainty about cross-era comparability.

C. Identification Appendix

C.1 Pre-Trends Analysis

The event study documents negative and significant pre-treatment coefficients. Several explanations are plausible: (1) mean reversion from a lower base; (2) anticipation effects from political mobilization; (3) common trends in India’s poorest regions during the 1990s; (4) compositional change in economic structure. I cannot definitively distinguish among these explanations.

C.2 Wild Cluster Bootstrap

With only 6 state-level clusters, I implement a wild cluster bootstrap that enumerates all $2^6 = 64$ Rademacher sign combinations under the null hypothesis of zero treatment effect. The

restricted model imposes $\beta = 0$ and constructs bootstrap samples $y^* = \hat{y}_{\text{restricted}} + w_g \cdot \hat{\epsilon}_{\text{restricted}}$ where $w_g \in \{-1, +1\}$ is assigned at the cluster level. The resulting p -value of 0.0625 (4 of 64 replications exceed the actual test statistic) is marginally above the 5% threshold.

C.3 Placebo Permutation

The placebo permutation test assigns treatment to all $\binom{6}{3} = 20$ possible combinations of 3 states from 6 and re-estimates the TWFE specification. The actual estimate falls at the extreme of this distribution ($p = 0.05$). With only 20 permutations, the minimum achievable p -value is 0.05.

C.4 Sub-District Level Analysis

The sub-district analysis exploits 1,674 sub-districts (versus 214 districts) for greater geographic granularity. The full-sample sub-district DiD estimates a treatment effect of 1.29 log points; the border subsample (within 150 km) estimates 1.07 log points. The sub-district spatial RDD yields a larger estimate (3.39 log points) but is complicated by mass points in the running variable.

D. Robustness Appendix

The main results are robust to: (i) pair-specific year fixed effects (0.74); (ii) population weighting; (iii) nightlight levels rather than logs; (iv) dropping any single state pair; (v) extending through 2023; (vi) border subsamples at varying bandwidths. The Callaway-Sant'Anna estimator yields a lower ATT (0.29) than TWFE (0.80), confirming that heterogeneity bias inflates the TWFE estimate.