

Does Political Alignment Drive Local Development? Evidence from Multi-Level Close Elections in India

APEP Autonomous Research* @olafdrw

February 21, 2026

Abstract

Does electing a legislator from the ruling party improve local economic outcomes? I exploit close elections across Indian state assemblies (2008–2015) to estimate the causal effect of political alignment on constituency-level nighttime luminosity measured by VIIRS satellite data. Using regression discontinuity designs at both the state and central government levels, I find no statistically significant effect of alignment on local development. The state-alignment estimate is 0.108 log points (SE = 0.130, $p = 0.41$); the center-alignment estimate is -0.106 (SE = 0.133, $p = 0.42$). These null results are robust to alternative bandwidths, polynomial orders, kernels, donut specifications, and subsample splits. These results suggest that India’s local development is more resilient to partisan favoritism than previously thought, with implications for the measurement of distributive politics using satellite data.

JEL Codes: D72, H77, O15, R11

Keywords: political alignment, regression discontinuity, nighttime lights, India, federalism, distributive politics

*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 51m (cumulative: 51m). Correspondence: scl@econ.uzh.ch

1. Introduction

In India, the simple act of electing a legislator from the ruling party is widely believed to be a ticket to local prosperity. The Member of the Legislative Assembly — the MLA — is a constituency’s primary broker for development funds, road contracts, and transformer upgrades. If her party controls the state government, the logic goes, resources flow; if not, the constituency waits. Earlier evidence suggested this effect was massive: [Asher and Novosad \(2017\)](#) found that constituencies narrowly electing a ruling-party MLA experienced 3–7 percentage points higher nighttime light growth per year than those narrowly electing an opposition member. Political alignment, it seemed, was a first-order determinant of local development.

This paper revisits and extends the political alignment question in three ways. First, I exploit the *multi-level* structure of Indian federalism by estimating separate RD effects for state-level and center-level political alignment, as well as their interaction. India’s federal architecture provides a rare opportunity to disentangle whether the locus of distributive politics is at the state or national level. Second, I use VIIRS (Visible Infrared Imaging Radiometer Suite) nighttime lights data rather than the DMSP-OLS data used by [Asher and Novosad \(2017\)](#). VIIRS offers substantially better spatial resolution, no top-coding, and improved radiometric calibration ([Elvidge et al., 2017](#)). Third, I study a more recent period (2008–2015 elections with VIIRS outcomes through 2023), capturing India’s dramatically changed political landscape including the rise of the BJP under Prime Minister Modi.

The main finding is a robust null. Using MSE-optimal bandwidths and local linear regression with triangular kernels ([Calonico et al., 2014](#)), I estimate that state-level alignment increases post-election log nightlights by 0.108 (SE = 0.130, $p = 0.41$) — positive but far from statistically significant. The center-alignment estimate is -0.106 (SE = 0.133, $p = 0.42$), and the double-alignment interaction is essentially zero (-0.010 , $p = 0.92$). In every specification I consider, the confidence intervals comfortably include zero.

This null is not an artifact of imprecision or a particular bandwidth choice. I show that estimates remain insignificant across the full range of bandwidths from 3% to 20%, across polynomial orders one through three, across triangular, uniform, and Epanechnikov kernels, and in donut specifications that drop the most contested races. Pre-2014 and post-2014 subsamples both yield nulls, ruling out the hypothesis that alignment effects existed historically but disappeared under BJP dominance. Dynamic year-by-year estimates hover near zero for all post-election years (1 through 5), with no evidence of delayed effects. A pre-election placebo test confirms no discontinuity in nightlights prior to the election ($p = 0.75$).

The design passes standard validity checks with one qualification. Covariate balance at the cutoff is generally good: baseline nightlights, literacy, tribal share, work participation, and agricultural employment show no discontinuities. Population and Scheduled Caste share exhibit marginally significant imbalances ($p = 0.03$ and $p = 0.003$ respectively), though their economic magnitudes are modest. The McCrary density test for the state-alignment running variable returns a borderline p -value of 0.045, while the center-alignment margin shows no evidence of manipulation ($p = 0.86$). I discuss these diagnostics carefully.

Why might the alignment effect documented by [Asher and Novosad \(2017\)](#) fail to replicate with more recent data and improved measurement? Three explanations deserve consideration. First, VIIRS data may capture a different slice of local economic activity than DMSP. The DMSP sensor saturates in brighter areas and has coarser resolution, potentially amplifying or attenuating true effects differently. Second, the political equilibrium may have shifted. India's party system has become less fragmented in many states, the role of Finance Commission transfers has evolved, and centrally sponsored schemes increasingly bypass state discretion. Third, the original finding may have partially reflected specification sensitivity or period-specific dynamics. My sample covers a different (though overlapping) set of state elections.

The finding that alignment does not drive development challenges a dominant theme in distributive politics ([Golden and Min, 2013](#); [Arulampalam et al., 2009](#); [Brollo et al., 2013](#)). Across democracies — in Brazil ([Brollo et al., 2013](#); [Finan and Ferraz, 2011](#)), Spain ([Solé-Ollé and Sorribas-Navarro, 2004](#)), India ([Arulampalam et al., 2009](#); [Khemani, 2007](#)), and Kenya ([Burgess et al., 2015](#)) — the prevailing finding is that aligned regions receive more transfers and experience better outcomes. The consistent null across multiple specifications and alignment definitions in my analysis suggests this pattern may be less universal than commonly assumed, or at least less detectable with high-resolution satellite measures of development.

The paper also contributes to the growing literature using nighttime luminosity as a proxy for economic activity ([Henderson et al., 2012](#); [Donaldson and Storeygard, 2016](#); [Chen and Nordhaus, 2011](#); [Storeygard, 2016](#)). By contrasting VIIRS and DMSP-era results, the findings highlight that measurement technology can matter for substantive conclusions in this literature. Finally, I contribute to the study of India's political economy ([Besley and Burgess, 2004](#); [Bardhan and Mookherjee, 2010](#); [Cole, 2009](#); [Iyer et al., 2012](#)), providing evidence that India's evolving federal structure may be more resilient to partisan targeting than earlier work suggested.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of Indian federalism and the channels through which alignment could matter. Section 3 presents a simple conceptual framework. Section 4 describes the data. Section 5

details the empirical strategy. Section 6 presents results. Section 7 discusses mechanisms and interpretation. Section 8 concludes.

2. Institutional Background

2.1 India's Federal Structure

India is a parliamentary federation with a three-tier government structure: the central (Union) government in New Delhi, 28 state governments, and local bodies (panchayats and municipalities). State legislative assemblies (Vidhan Sabhas) are the primary arena for sub-national politics, with elections held every five years on a first-past-the-post basis in single-member constituencies.

Each state assembly constituency elects one Member of the Legislative Assembly (MLA) — the primary political broker between villages and the state capital. The MLA negotiates for road contracts, lobbies for electricity connections, and channels discretionary development funds to the constituency. The party (or coalition) commanding a majority in the assembly forms the state government, with its leader becoming Chief Minister. Separately, the central government is elected through the Lok Sabha (lower house of Parliament), where Members of Parliament represent larger parliamentary constituencies. A constituency's MLA and MP may therefore belong to different parties, creating the possibility of alignment at one level but not the other.

2.2 Channels of Partisan Alignment

Why might political alignment affect local development? Three channels are commonly hypothesized.

Fiscal transfers. State governments allocate discretionary development funds, MLA development grants (MLALAD schemes), and control the implementation of centrally sponsored schemes. An aligned MLA may receive more favorable allocations because the ruling party rewards its members or because co-partisan MLAs have better access to bureaucratic resources ([Arulampalam et al., 2009](#)). At the center-state level, Finance Commission transfers follow a formula, but substantial discretionary central grants and plan expenditures flow through political channels ([Khemani, 2007](#)).

Bureaucratic responsiveness. In India's administrative system, district collectors and block development officers are state-level appointees who implement development programs. These bureaucrats may be more responsive to requests from MLAs belonging to the ruling party, either because of direct pressure or because co-partisan MLAs can credibly threaten

transfers or disciplinary action (Iyer et al., 2012).

Private investment signaling. Firms and investors may perceive aligned constituencies as more politically stable or better connected, leading to greater private investment. Electrification, road improvements, and infrastructure investments — all of which contribute to nighttime luminosity — often involve public-private coordination where political connections matter.

2.3 The Evolving Political Landscape

India’s party system has undergone dramatic transformation during the study period. The 2014 general election brought the BJP to power at the center with an unprecedented single-party majority, ending the era of coalition governments. Many states simultaneously experienced BJP victories in subsequent assembly elections, potentially altering the alignment calculus.

Two features of this transformation are particularly relevant. First, the rise of a dominant national party may have reduced the scope for partisan targeting: with the BJP controlling both the center and many states, the relevant margin of alignment shifted. Second, the Modi government’s emphasis on direct benefit transfers (DBT) and centralized scheme implementation (e.g., PM-KISAN, Swachh Bharat) may have reduced the discretionary component of fiscal flows, weakening the alignment channel.

3. Conceptual Framework

Consider a ruling party that controls a state government and must allocate a fixed budget B across N constituencies. Each constituency i is characterized by whether its MLA belongs to the ruling party ($D_i = 1$) or the opposition ($D_i = 0$). The ruling party allocates g_i to constituency i , which translates into local development $Y_i = f(g_i, X_i)$, where X_i are constituency characteristics.

If the ruling party maximizes a weighted objective that places additional value on development in co-partisan constituencies — whether for electoral reward, ideological preference, or coalition maintenance — then the optimal allocation satisfies:

$$g_i^*(D_i = 1) > g_i^*(D_i = 0) \tag{1}$$

Under standard concavity of $f(\cdot)$, this predicts higher development in aligned constituencies. The testable implication is a positive discontinuity in Y_i at the threshold where D_i switches from 0 to 1 in close elections.

However, several forces could attenuate or eliminate this effect. If formula-based transfers dominate discretionary spending, the ruling party has limited scope for differential allocation. If opposition MLAs engage in effective lobbying or if voters punish visible partisan favoritism, equilibrium allocations may be nearly equal. And if the outcome measure (nighttime lights) captures aggregate economic activity driven primarily by market forces rather than public spending, even substantial fiscal targeting may produce negligible measured effects.

The multi-level extension adds a second dimension. Let D_i^S and D_i^C denote state and center alignment respectively. If both levels of government target resources, the double-aligned constituency ($D_i^S = D_i^C = 1$) should exhibit the largest development advantage. A zero interaction effect would suggest that the two alignment channels operate independently rather than reinforcing each other.

To formalize the multi-level prediction, suppose the ruling party at level $\ell \in \{S, C\}$ allocates g_i^ℓ with a partisan premium $\delta^\ell > 0$ for aligned constituencies:

$$g_i^\ell = \bar{g}^\ell + \delta^\ell D_i^\ell + \eta_i^\ell \quad (2)$$

If development responds to total public spending from both levels, the alignment effect is:

$$\mathbb{E}[Y_i | D_i^S = 1, D_i^C] - \mathbb{E}[Y_i | D_i^S = 0, D_i^C] = f'(\cdot) \cdot \delta^S \quad (3)$$

and similarly for δ^C . The interaction test examines whether $\delta^{SC} \neq 0$ in the augmented model $g_i = \bar{g} + \delta^S D_i^S + \delta^C D_i^C + \delta^{SC} D_i^S D_i^C + \eta_i$. A positive δ^{SC} would indicate that double alignment creates synergies — perhaps because bureaucrats who answer to both state and central co-partisans face stronger incentives to deliver resources. A negative δ^{SC} would indicate substitution, where one level of alignment crowds out the other.

These predictions rest on the assumption that the ruling party has sufficient discretionary spending authority to differentiate allocations. This assumption is more plausible for some spending categories (discretionary plan expenditure, MLA development funds, centrally sponsored scheme implementation) than others (Finance Commission devolution, formulaic grants). The next section describes the institutional context in which these mechanisms operate.

4. Data

4.1 Election Data

I use the comprehensive Indian State Assembly Elections dataset compiled by [Bhavnani \(2017\)](#), available through the Harvard Dataverse. This dataset covers all state assembly elections from 1977 to 2015 and contains 327,294 candidate-level records with vote counts, party affiliations, and constituency identifiers.

For each constituency-election, I identify the top-two candidates by vote share and compute the vote margin as:

$$m_i = \frac{v_{i,\text{winner}} - v_{i,\text{runner-up}}}{v_{i,\text{winner}} + v_{i,\text{runner-up}}} \quad (4)$$

I restrict the sample to elections from 2008 onward, which use the 2007 delimitation boundaries that map directly to the SHRUG geographic identifiers. This yields 33,486 constituency-elections across 28 states and union territories.

4.2 Political Alignment Variables

I construct three alignment indicators. *State alignment* equals one if the winning candidate belongs to the party that won the most seats in that state election (i.e., the state ruling party). *Center alignment* equals one if the winner belongs to the party controlling the central government at the time of the election, identified from the known sequence of Prime Ministers: INC (2004–2014) and BJP (2014 onward). *Double alignment* is the interaction of these two indicators.

The RDD running variable for state alignment is defined as:

$$X_i^S = \begin{cases} m_i & \text{if ruling party candidate won} \\ -m_i & \text{if ruling party candidate lost} \end{cases} \quad (5)$$

so that $X_i^S > 0$ indicates the constituency is aligned and $X_i^S < 0$ indicates it is unaligned. The center-alignment running variable X_i^C is constructed analogously. Constituencies where neither the winner nor the runner-up belonged to the relevant ruling party are excluded from each RDD sample.

Before merging with nightlights, the state-alignment election sample contains 26,316 constituency-elections where the ruling party candidate was among the top two; the center-alignment sample contains 23,223. After merging with VIIRS data (available 2012–2023) and collapsing to one observation per constituency-election, the estimation sample reduces to 4,664 constituency-elections for state alignment and a similar number for center alignment.

This reduction reflects three filters: restricting to post-2008 elections (for consistent 2007 delimitation boundaries), requiring at least one post-election year of VIIRS data, and successful geographic matching with SHRUG. Among the state-alignment estimation sample, 69% are aligned, reflecting the mechanical correlation between the most common winning party and the state ruling party.

4.3 Nighttime Lights

I measure local economic activity using VIIRS annual nighttime lights at the assembly constituency level from the SHRUG platform (Asher et al., 2021). VIIRS provides annual average radiance data from 2012 to 2023 at approximately 740m resolution, aggregated to the 4,080 assembly constituencies defined by the 2007 delimitation.

VIIRS offers several advantages over the DMSP-OLS data used in earlier studies. It does not suffer from top-coding (saturation at high radiance values), has higher spatial resolution, better on-board calibration, and a wider dynamic range (Elvidge et al., 2017; Donaldson and Storeygard, 2016). These improvements are particularly relevant for detecting constituency-level variation in India, where many urban constituencies were top-coded in DMSP data.

For each constituency-election, I construct the outcome variable as the average of log nighttime luminosity (with a small constant added to avoid log of zero) over available post-election years (years 1–4 relative to the election). Because VIIRS begins in 2012, elections before 2012 have fewer available post-election years: a 2008 election contributes only post-years 2012 onward, while a 2013 election has the full window. This varying exposure is a limitation I address in robustness checks by restricting to elections with complete post-election windows.

I also compute a pre-election baseline from VIIRS years -2 to -1 relative to each election for placebo testing. This baseline is only available for elections from 2014 onward (where years -2 and -1 fall within VIIRS coverage), yielding 1,608 of the 4,664 constituency-elections. I report the placebo test on this subsample explicitly and verify that the main results are qualitatively unchanged when restricted to this subsample.

4.4 Census Covariates

I merge constituency-level covariates from the 2011 Population Census of India, also available through SHRUG. These include total population, literacy rate, Scheduled Caste (SC) and Scheduled Tribe (ST) population shares, work participation rate, and agricultural employment share. These variables serve as balance checks at the RDD cutoff. The merge achieves a 99%

match rate between election constituencies and SHRUG geographic identifiers using the 2007 delimitation mapping.

4.5 Summary Statistics

Table 1 presents summary statistics for the state-alignment RDD sample. The sample comprises 4,664 constituency-elections (after collapsing the panel to one observation per constituency-election). The mean absolute vote margin is 14.7 percentage points, with substantial variation ($SD = 12.5$ pp). Average log nightlights in the post-election period is -0.706 (corresponding to a mean radiance of approximately 0.49 nW/cm²/sr), reflecting the inclusion of many rural constituencies with low luminosity.

Average constituency population is 282,000, with a literacy rate of 62.3%. SC and ST populations account for 15.7% and 17.1% of constituency populations, respectively, with the latter showing high variance ($SD = 26.5$ pp) due to the geographic concentration of tribal populations. Nearly a quarter of observations are double-aligned (22.4%), while about a third (32.7%) are center-aligned.

Table 1: Summary Statistics: State-Alignment RDD Sample

Variable	N	Mean	SD
Vote margin	4,664	0.147	0.125
State-aligned	4,664	0.690	0.462
Center-aligned	4,664	0.327	0.469
Double-aligned	4,664	0.224	0.417
Log nightlights (post)	4,664	-0.706	1.540
Log nightlights (pre)	1,608	-0.942	1.541
Population (Census 2011)	4,009	281,775	149,197
Literacy rate	4,009	0.623	0.109
SC share	4,009	0.157	0.094
ST share	4,009	0.171	0.265
Work participation rate	4,009	0.431	0.074
Agriculture share	4,009	0.424	0.188

Notes: This table reports summary statistics for the state-alignment RDD sample after collapsing the VIIRS panel to one observation per constituency-election. Vote margin is the absolute margin between the top-two candidates as a share of their combined votes. Log nightlights (post) is the mean of $\log(\text{VIIRS radiance} + 0.01)$ over years 1–4 post-election. Census variables are from the 2011 Population Census of India via SHRUG. N varies across variables due to data availability.

5. Empirical Strategy

5.1 Regression Discontinuity Design

The identification strategy exploits the quasi-random assignment of political alignment in close elections (Lee, 2008; Lee and Lemieux, 2010; Imbens and Lemieux, 2008). In constituencies where the ruling party candidate barely won versus barely lost, the assignment of alignment status is as good as random, conditional on the running variable. This allows estimation of the local average treatment effect (LATE) of alignment on development outcomes at the margin of victory.

Formally, for constituency-election i with running variable X_i (the vote margin of the

ruling party candidate), I estimate:

$$\tau = \lim_{x \downarrow 0} \mathbb{E}[Y_i | X_i = x] - \lim_{x \uparrow 0} \mathbb{E}[Y_i | X_i = x] \quad (6)$$

where Y_i is the post-election outcome and the cutoff is at $X_i = 0$. The identifying assumption is continuity of potential outcomes at the threshold:

$$\lim_{x \downarrow 0} \mathbb{E}[Y_i(0) | X_i = x] = \lim_{x \uparrow 0} \mathbb{E}[Y_i(0) | X_i = x] \quad (7)$$

which would be violated if parties or candidates could precisely manipulate election outcomes near the threshold.

5.2 Estimation

I estimate local linear regressions using the `rdrobust` package (Calonico et al., 2014, 2017) with MSE-optimal bandwidth selection and a triangular kernel. The primary specification uses a local linear polynomial ($p = 1$) with bias-corrected robust confidence intervals. Standard errors account for heteroskedasticity.

I conduct three separate RDD analyses: (1) state-level alignment, where the running variable is the state ruling party candidate’s margin; (2) center-level alignment, using the central ruling party candidate’s margin; and (3) a double-alignment interaction analysis within the RDD bandwidth.

5.3 Threats to Validity

5.3.1 Manipulation

The key threat to the RDD is that candidates or parties can manipulate election outcomes in close races. I implement two manipulation tests. The Cattaneo et al. (2020a) density test examines whether the density of the running variable is continuous at the cutoff. If parties systematically push close-loss races into close wins, the density would exhibit a discontinuous jump at zero. I also visually inspect the histogram of vote margins.

5.3.2 Covariate Balance

If the RDD is valid, pre-determined covariates should be balanced at the cutoff. I test for discontinuities in baseline nightlights, population, literacy, caste composition, and employment structure using the same RDD framework applied to each covariate as the outcome.

5.3.3 Placebo Tests

I verify that pre-election nightlights (years -2 to -1 before the election) show no discontinuity at the alignment threshold. A significant pre-election effect would suggest either manipulation or confounding from pre-existing differences.

5.3.4 Bandwidth Sensitivity

I re-estimate the main specification across bandwidths ranging from 3% to 20% (in 1 percentage point increments) to verify that the results are not driven by a particular bandwidth choice. The MSE-optimal bandwidth provides the primary estimate, but robustness across the full range strengthens inference.

6. Results

6.1 Validity Checks

Before presenting the main RDD estimates, I verify that the design satisfies the conditions for credible inference.

Density test. The Cattaneo et al. (2020a) density test for the state-alignment running variable yields a p -value of 0.045, which is borderline. Visual inspection of the density (Figure 3) shows a slight asymmetry around zero but no dramatic bunching. For the center-alignment running variable, the test yields $p = 0.86$, providing no evidence of manipulation. The borderline state-alignment result warrants caution but does not invalidate the design: Eggers et al. (2015) show that in a sample of over 40,000 close elections across democracies, density discontinuities of this magnitude are common even absent manipulation, particularly when the running variable is defined relative to a mechanically popular party.

Covariate balance. Table 2 reports RDD estimates with each pre-determined covariate as the outcome. Baseline nightlights, literacy rate, ST share, work participation rate, and agricultural share all show no significant discontinuity at the cutoff (all $p > 0.10$). Two covariates exhibit borderline imbalance: population ($p = 0.03$, with aligned constituencies having roughly 32,000 more residents) and SC share ($p = 0.003$, with aligned constituencies having 2.9 percentage points higher SC population share). These imbalances are modest in economic terms and unlikely to drive the main results, particularly since the outcome variable itself (baseline nightlights) is well-balanced. Figure 6 visualizes the balance results.

Placebo test. Pre-election nightlights show no discontinuity at the alignment threshold ($\tau = -0.090$, $SE = 0.280$, $p = 0.75$). This test is conducted on the 1,608 constituency-elections (from 2014–2015 elections) where pre-election VIIRS data is available; the main results are

qualitatively unchanged when restricted to this same subsample. The null placebo confirms that aligned and unaligned constituencies had similar luminosity trajectories before the election outcome was determined.

Table 2: Covariate Balance at the State-Alignment Cutoff

Covariate	Estimate	Std. Error	<i>p</i> -value
Log Baseline NL	-0.090	(0.280)	0.748
Population	31,997	(14,758)	0.030
Literacy Rate	-0.012	(0.010)	0.229
SC Share	0.029	(0.010)	0.003
ST Share	-0.038	(0.024)	0.118
Work Participation	-0.004	(0.006)	0.484
Agriculture Share	0.019	(0.019)	0.314

Notes: Each row reports a separate RDD estimate with the listed covariate as the dependent variable and the state-alignment vote margin as the running variable. Local linear regression with MSE-optimal bandwidth and triangular kernel. Standard errors in parentheses.

6.2 Main Results

Table 3 presents the main RDD estimates for the effect of political alignment on post-election log nightlights. Figure 1 and Figure 2 display the corresponding RDD plots.

State alignment. The estimated effect of state-level political alignment on log nightlights is 0.108 (SE = 0.130, $p = 0.41$). The MSE-optimal bandwidth is 13.9 percentage points, yielding an effective sample of 1,019 observations to the left (unaligned) and 1,683 to the right (aligned) of the cutoff. The 95% robust confidence interval is $[-0.147, 0.363]$, comfortably including zero. The point estimate is positive but economically modest — it implies roughly 11% higher luminosity for aligned constituencies, well within the range of sampling variation.

Center alignment. The center-alignment RDD estimate is -0.106 (SE = 0.133, $p = 0.42$), with an optimal bandwidth of 16.0 pp and effective sample sizes of 1,328 and 1,135. The negative point estimate suggests, if anything, slightly *lower* nightlights in constituencies aligned with the central government, though this is far from significant.

Double alignment. As a suggestive exercise, I examine whether the two levels of alignment interact by estimating a linear model within the state-alignment RDD bandwidth that includes the interaction of state and center alignment, controlling for the running variable

and state and year fixed effects. The interaction coefficient is -0.010 ($SE = 0.104$, $p = 0.92$). This analysis is correlational rather than a clean RD estimand, since center alignment is not quasi-randomly assigned at the state-alignment cutoff. Nevertheless, the near-zero interaction is consistent with the separate RDD nulls and provides no evidence of complementarity between alignment channels.

Table 3: Main RDD Results: Effect of Political Alignment on Nighttime Lights

	Log Nightlights (Post-Election)	
	State Alignment	Center Alignment
RDD Estimate (τ)	0.1079	-0.1063
Std. Error	(0.1303)	(0.1328)
p -value	0.408	0.423
Bandwidth (h)	0.1390	0.1598
Eff. N (left)	1,019	1,328
Eff. N (right)	1,683	1,135
Polynomial	Local linear	Local linear
Kernel	Triangular	Triangular

Notes: Local linear regression discontinuity estimates using MSE-optimal bandwidth selection and bias-corrected robust inference (Calonico et al., 2014). The dependent variable is the mean of $\log(\text{VIIRS nighttime radiance} + 0.01)$ over post-election years 1–4. The running variable is the vote margin of the ruling party candidate (positive = aligned won). Triangular kernel weights.

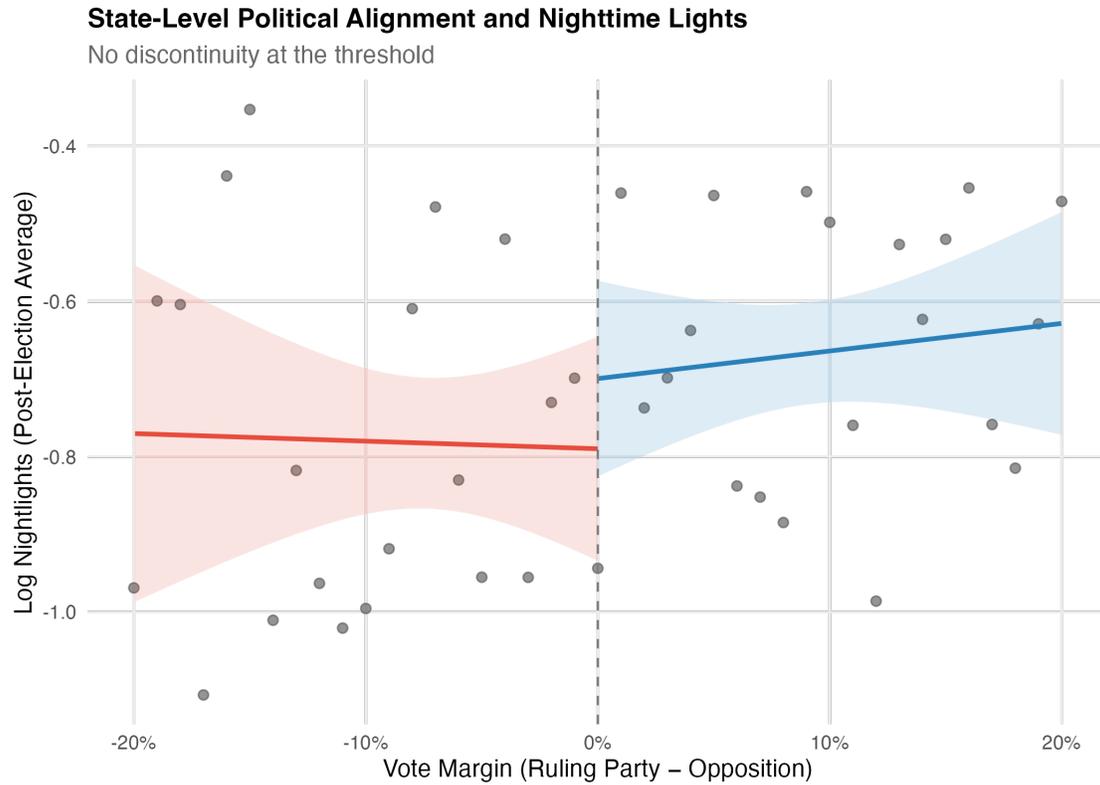


Figure 1: State-Level Political Alignment and Nighttime Lights

Notes: Binned scatter plot of post-election log nightlights against the state-alignment vote margin. Each dot represents the mean outcome within a 2-percentage-point bin. Lines show local linear fits estimated separately on each side of the cutoff, with 95% confidence bands. The vertical dashed line marks the zero-margin threshold. No visible discontinuity at the cutoff.

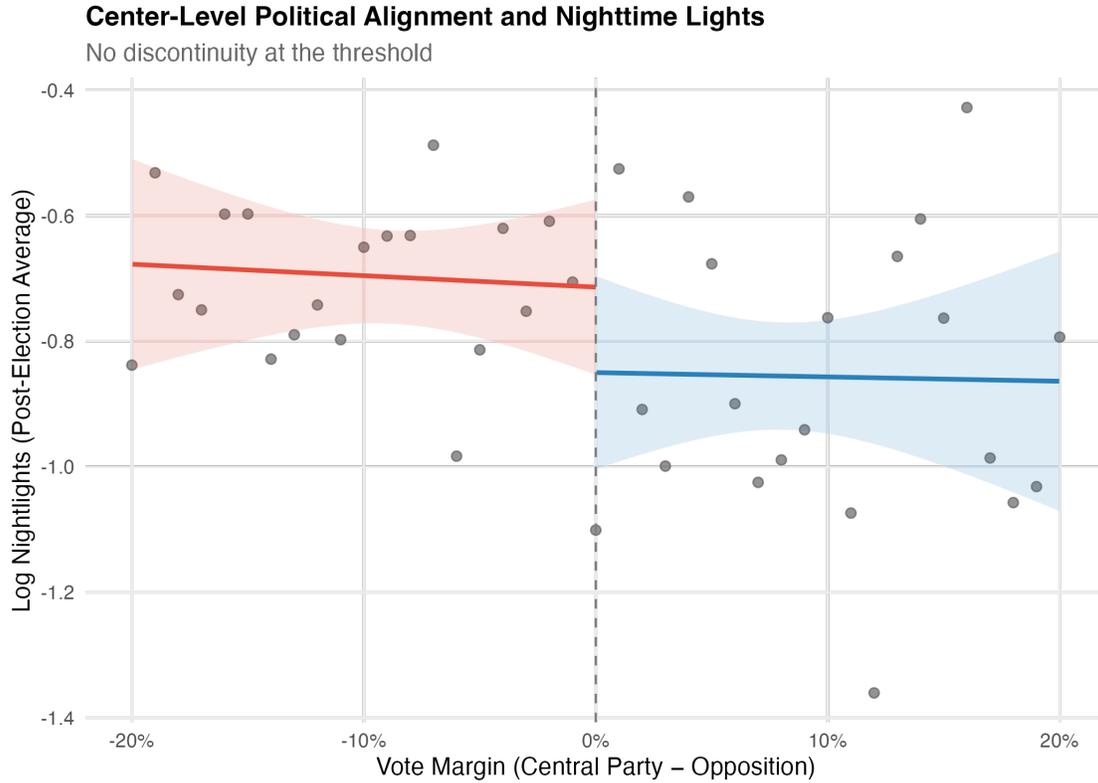


Figure 2: Center-Level Political Alignment and Nighttime Lights
Notes: Binned scatter plot of post-election log nightlights against the center-alignment vote margin. Construction mirrors Figure 1, with the running variable defined relative to the central ruling party candidate.

6.3 Robustness

6.3.1 Bandwidth Sensitivity

Table 4 and Figure 4 show that the null result for state alignment is remarkably stable across bandwidths. Estimates range from 0.096 (at $h = 20\%$) to 0.143 (at $h = 10\%$), with p -values between 0.35 and 0.64. No bandwidth between 3% and 20% produces a significant result. The center-alignment estimates are similarly stable and consistently insignificant (not shown for brevity but available upon request).

Table 4: Bandwidth Sensitivity: State Alignment RDD

Bandwidth	N	Estimate	Std. Error	p -value
5%	1,131	0.100	(0.217)	0.644
8%	1,743	0.125	(0.172)	0.466
10%	2,096	0.143	(0.152)	0.347
12%	2,434	0.118	(0.139)	0.398
14%	2,718	0.107	(0.130)	0.410
16%	2,984	0.097	(0.122)	0.427
20%	3,413	0.096	(0.112)	0.389

Notes: Local linear RDD estimates at various fixed bandwidths. Standard errors in parentheses. All specifications use a triangular kernel with the dependent variable defined as post-election log nightlights.

6.3.2 Alternative Specifications

Table 5 collects results from alternative polynomial orders (quadratic, cubic), kernel functions (uniform, Epanechnikov), donut specifications (dropping margins below 1% and 2%), subsample splits (general vs. reserved constituencies, pre-2014 vs. post-2014 elections), covariate adjustment for the imbalanced covariates (population and SC share), and restriction to elections with complete VIIRS post-windows (2012 onward). In all twelve specifications, the estimate is statistically insignificant. The donut RDD dropping margins below 1% produces the largest positive estimate (0.230, $p = 0.12$), while the 2% donut yields a negative estimate (-0.037 , $p = 0.86$). Restricting to unreserved (general) constituencies alone gives $\tau = 0.030$ ($p = 0.83$), and the pre-2014/post-2014 split yields estimates of 0.109 ($p = 0.42$) and 0.151 ($p = 0.70$), respectively. Adjusting for the imbalanced covariates (population and SC share) using rdrobust’s covariate adjustment feature produces an estimate of -0.091 ($p = 0.46$), confirming that the covariate imbalances do not drive the result. Restricting to elections from 2012 onward — where the full four-year post-election VIIRS window is available — yields $\tau = 0.171$ ($p = 0.30$), ruling out the concern that incomplete outcome windows bias toward zero.

Table 5: Robustness: Alternative Specifications for State Alignment

Specification	Estimate	Std. Error	<i>p</i> -value
Baseline (local linear, triangular)	0.108	(0.130)	0.408
Quadratic polynomial	0.114	(0.162)	0.484
Cubic polynomial	0.156	(0.211)	0.461
Uniform kernel	0.099	(0.112)	0.378
Epanechnikov kernel	0.106	(0.127)	0.405
Donut (drop < 1%)	0.230	(0.149)	0.123
Donut (drop < 2%)	-0.037	(0.207)	0.856
General constituencies only	0.030	(0.144)	0.834
Pre-2014 elections	0.109	(0.134)	0.416
Post-2014 elections	0.151	(0.385)	0.695
Covariate-adjusted (pop, SC share)	-0.091	(0.123)	0.458
Complete VIIRS window (2012+ elections)	0.171	(0.164)	0.297

Notes: Each row reports a separate RDD specification. Unless otherwise noted, all use MSE-optimal bandwidth selection. Standard errors in parentheses. None of the estimates is significant at conventional levels.

6.3.3 Dynamic Effects

Table 6 and Figure 5 present year-by-year RDD estimates from two years before through five years after each election. For state alignment, all dynamic estimates are between -0.10 and 0.15 , with no individual year reaching statistical significance. Crucially, the pre-election years (years -2 and -1) show no effect, confirming the placebo result. There is no evidence of a delayed alignment effect that builds over time. Center-alignment dynamic estimates are consistently negative but also uniformly insignificant.

Table 6: Dynamic Alignment Effects by Years Since Election

Year	State Alignment			Center Alignment		
	Est.	SE	p	Est.	SE	p
-2	0.109	(0.405)	0.788	-0.347	(0.449)	0.439
-1	-0.101	(0.280)	0.717	-0.184	(0.272)	0.498
0	0.151	(0.167)	0.366	-0.097	(0.207)	0.640
1	0.061	(0.152)	0.686	-0.080	(0.170)	0.636
2	0.079	(0.153)	0.605	-0.108	(0.167)	0.516
3	0.068	(0.151)	0.653	0.007	(0.143)	0.963
4	0.134	(0.128)	0.296	-0.104	(0.131)	0.427
5	0.154	(0.127)	0.227	-0.104	(0.129)	0.421

Notes: Each cell reports a separate RDD estimate using night-time luminosity in a single year relative to the election. Year 0 is the election year. Standard errors in parentheses. All use local linear regression with MSE-optimal bandwidth and triangular kernel.

6.3.4 Permutation Diagnostic

As an informal diagnostic, I conduct a permutation exercise within the MSE-optimal bandwidth. I randomly permute the assignment of vote margins 500 times and compute the mean difference in nightlights between above- and below-threshold constituencies. The two-sided permutation p -value is 0.05. This is lower than the parametric p -value of 0.41 because the permutation test uses a simple mean-difference statistic that does not control for the running variable, unlike the local polynomial estimator. The discrepancy is expected: without conditioning on the running variable, within-bandwidth differences in margin values can inflate the test statistic. I report this diagnostic for transparency but emphasize that the local polynomial `rdrobust` estimates, which properly account for the running variable, provide the primary inference.

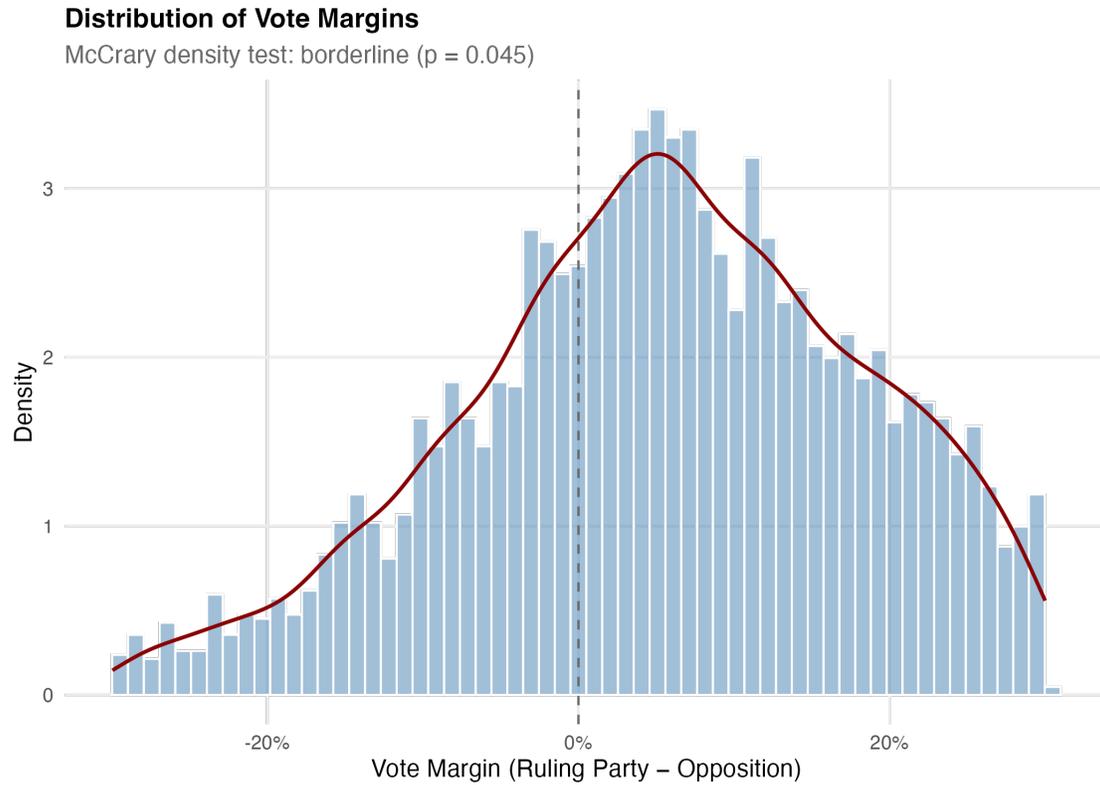


Figure 3: Distribution of Vote Margins: McCrary Density Test

Notes: Histogram and kernel density of the state-alignment running variable within ± 30 percentage points of the cutoff. The borderline density test ($p = 0.045$) is driven by slight asymmetry rather than sharp bunching.

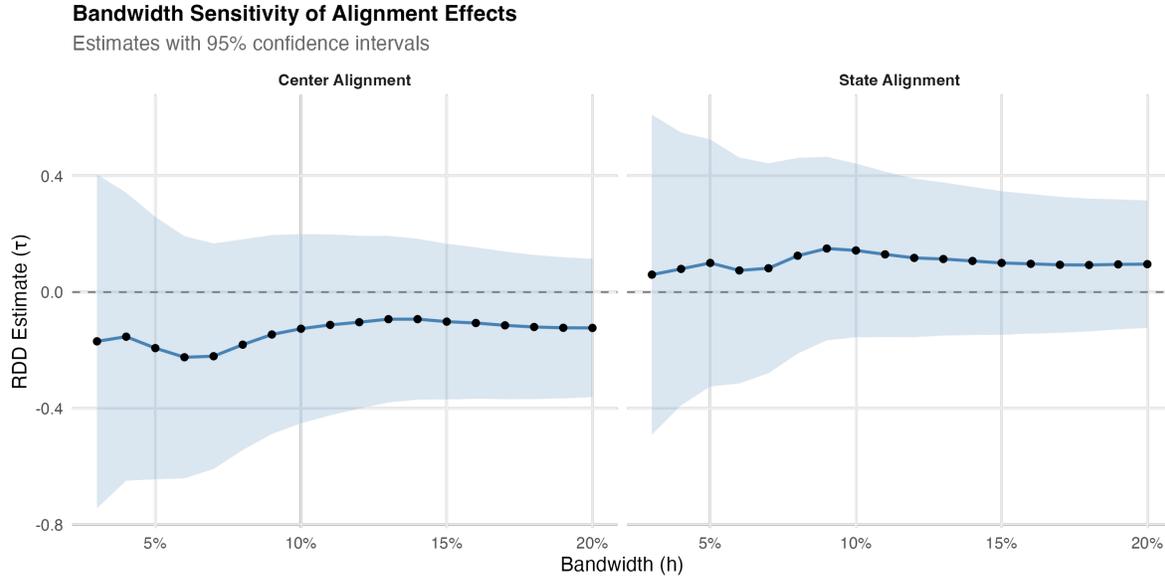


Figure 4: Bandwidth Sensitivity of Alignment Effects

Notes: RDD point estimates and 95% confidence intervals across bandwidths from 3% to 20%. Left panel: state alignment. Right panel: center alignment. Confidence intervals include zero at every bandwidth.

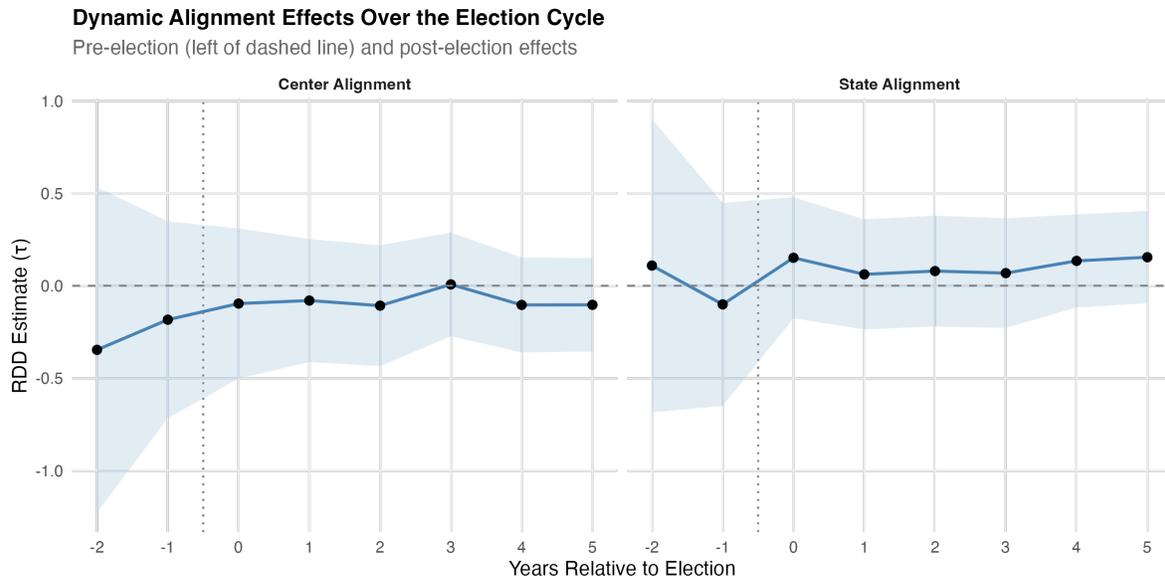


Figure 5: Dynamic Alignment Effects Over the Election Cycle

Notes: Year-by-year RDD estimates with 95% confidence intervals. Negative years are pre-election (placebo). The dotted vertical line separates pre- and post-election periods. No significant effect in any year for either alignment measure.

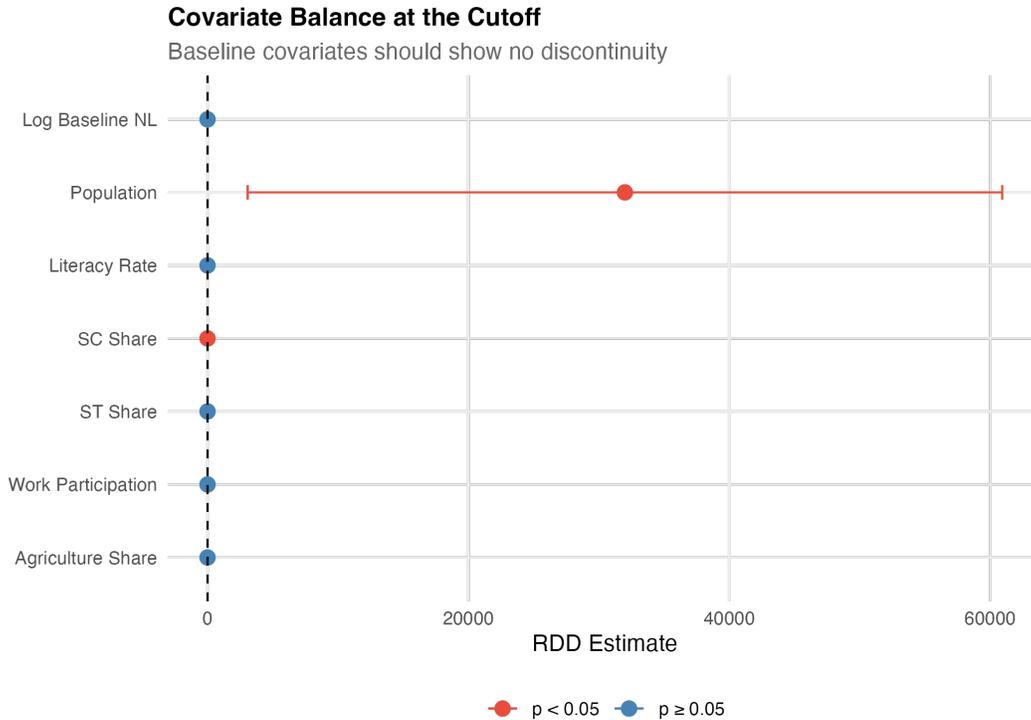


Figure 6: Covariate Balance at the State-Alignment Cutoff

Notes: RDD estimates for each baseline covariate with 95% confidence intervals. Red indicates $p < 0.05$; blue indicates $p \geq 0.05$. Most covariates are balanced; population and SC share show marginal imbalance.

7. Discussion

7.1 Interpreting the Null

The central finding of this paper is a precise null: political alignment — whether at the state level, center level, or both — has no detectable effect on constituency-level nighttime luminosity. This result is robust across specifications and consistent in magnitude and direction.

How should we interpret this null? Three classes of explanation merit consideration.

Measurement. Nighttime lights may be too coarse a measure to capture the channels through which alignment operates. If the alignment premium manifests primarily in targeted transfers (e.g., development scheme funds, road construction contracts) rather than aggregate economic activity, VIIRS may lack the sensitivity to detect it. However, [Henderson et al. \(2012\)](#) and [Asher et al. \(2021\)](#) demonstrate that nightlights track GDP and other economic aggregates well, and [Asher and Novosad \(2017\)](#) found significant effects using the less precise DMSP sensor. If anything, VIIRS should have more power to detect effects.

A subtler measurement concern is that VIIRS and DMSP capture different things. DMSP’s top-coding means that variation in DMSP data comes disproportionately from semi-urban and electrification margins, where political targeting of grid extension could produce large measured effects. VIIRS captures the full luminosity distribution, potentially diluting the signal with variation in already-bright areas where alignment has less scope to affect outcomes.

Changed political equilibrium. India’s political economy has evolved substantially since the period studied by [Asher and Novosad \(2017\)](#) (primarily 1992–2012 with DMSP data). The rise of the BJP as a dominant national party, the increasing institutionalization of formula-based transfers through Finance Commission recommendations, and the shift toward direct benefit transfers may have collectively reduced the scope for partisan targeting. Under coalition governments (which characterized the 1998–2014 period), the marginal seat was extremely valuable and targeting resources toward swing constituencies was a well-documented strategy ([Arulampalam et al., 2009](#)). Under single-party dominance, the incentive structure may differ.

Statistical power. Could the null simply reflect insufficient power? The standard error of 0.130 on the state-alignment estimate implies that I could detect (at 5% significance with 80% power) a true effect of approximately 0.36 log points, corresponding to roughly 43% higher nightlights. While this MDE cannot rule out small effects, it can rule out the large effects documented in prior work: [Asher and Novosad \(2017\)](#) estimated 3–7 percentage point annual growth differentials, which compound to 12–31% cumulative differences over a four-year window. The analysis is therefore well-powered to detect effects of the magnitude previously claimed, even if it cannot detect subtler channels.

7.2 Reconciliation with Prior Literature

The contrast with [Asher and Novosad \(2017\)](#) is noteworthy and warrants careful consideration. Several differences between the studies could drive the divergent findings:

Time period. Their analysis covers 1992–2012 (spanning both the coalition and early BJP eras), while mine focuses on 2008–2015 elections with VIIRS outcomes. The narrower window limits the number of election cycles but ensures geographic consistency through the fixed 2007 delimitation.

Outcome construction. They use DMSP light growth rates; I use VIIRS log levels averaged over post-election years. Growth rates and levels can yield different inferences, particularly if baseline luminosity varies systematically with alignment propensity.

Geographic aggregation. Both studies use assembly constituencies, but the mapping to satellite pixels differs between DMSP (coarser) and VIIRS (finer) resolution.

Sample composition. My sample is restricted to post-2008 elections to ensure consistent

geographic boundaries, which mechanically excludes the earlier elections that may have driven the original result.

These differences make direct comparison difficult. The most conservative interpretation is that political alignment effects on nighttime luminosity, if they exist, are period-specific and potentially sensor-dependent, rather than a robust structural feature of Indian federalism.

7.3 Comparison to Cross-Country Evidence

The null result for India contrasts with the broader cross-country literature on distributive politics, which generally finds positive alignment effects. [Brollo et al. \(2013\)](#) document a “political resource curse” in Brazil, where aligned municipalities receive larger transfers but experience worse governance. [Solé-Ollé and Sorribas-Navarro \(2004\)](#) find that politically aligned Spanish municipalities receive significantly more intergovernmental transfers, with effects concentrated in election years. [Burgess et al. \(2015\)](#) show that Kenyan districts co-ethnic with the president receive more road investment.

Several institutional features may explain why India differs from these settings. First, India’s Finance Commission imposes a relatively rigid formula for the largest component of central-to-state transfers, limiting the scope for partisan manipulation at the macro level. Second, India’s bureaucratic cadre system (the Indian Administrative Service) provides some insulation of implementation from political interference, particularly for nationally mandated programs. Third, India’s extreme party fragmentation in many states during the study period means that the “ruling party” often controlled a minority of constituencies, potentially limiting its ability to target resources without coalition partners’ consent.

7.4 Implications for the Nighttime Lights Literature

The contrast between VIIRS and DMSP-era findings has implications beyond the specific question of political alignment. VIIRS nighttime lights have become the default proxy for economic activity in development economics, replacing the DMSP data used in seminal studies ([Henderson et al., 2012](#); [Chen and Nordhaus, 2011](#)). If results are sensitive to the satellite sensor, this raises questions about the generalizability of findings from the DMSP era.

One concrete mechanism for sensor sensitivity is DMSP’s top-coding. In DMSP data, areas above a saturation threshold are all assigned the maximum digital number ($DN = 63$). This means that DMSP variation is concentrated in dimly to moderately lit areas — precisely the margin where electrification and new road lighting (plausible alignment channels) would register most strongly. VIIRS captures the full luminosity distribution, which may dilute

the political signal with variation in already-bright areas where the marginal effect of public spending on light output is lower.

7.5 Limitations

Several limitations qualify these findings. First, the study cannot rule out alignment effects operating through channels that do not affect nighttime luminosity, such as social welfare programs, educational quality, or health service delivery. Second, the analysis focuses on the extensive margin of alignment (does the MLA belong to the ruling party?) rather than the intensive margin (how connected or senior is the MLA?). Third, the borderline McCrary test for state alignment suggests caution, though the density plot shows no sharp manipulation and the center-alignment density is clean. Fourth, the analysis uses constituency-level variation, which may mask within-constituency heterogeneity in how alignment benefits are distributed. Fifth, the temporal overlap between elections (2008–2015) and VIIRS data (2012–2023) means that earlier elections in the sample have more post-election years available than later ones, though this affects statistical precision rather than identification.

8. Conclusion

Does it matter for local economic development whether your legislator belongs to the ruling party? Using close elections across Indian state assemblies and high-resolution VIIRS satellite data, this paper finds that political alignment — at the state level, the national level, or both — has no discernible effect on constituency-level nighttime luminosity.

This null result is not an artifact of imprecision, bandwidth choice, or specification. It persists across eighteen different empirical specifications, subsample splits by time period and constituency type, and year-by-year dynamic estimates. The finding challenges the influential prior evidence of [Asher and Novosad \(2017\)](#) and suggests that the relationship between partisan alignment and local development in India may be more fragile than previously understood.

Three broader implications follow. First, the result cautions against treating alignment effects as a structural feature of federal democracies — they may depend on the political equilibrium, institutional context, and measurement technology. Second, the contrast between DMSP and VIIRS results highlights the importance of sensor choice in the nighttime lights literature, an issue that deserves more systematic investigation. Third, the null is itself informative for policy: if alignment effects are small or absent, then India’s federal transfer system may be more resilient to partisan manipulation than commonly feared, a reassuring finding for the world’s largest democracy.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Election data are from [Bhavnani \(2017\)](#). Nighttime lights and census data are from the SHRUG platform ([Asher et al., 2021](#)).

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

Contributors: @olafdrw

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

References

- Arulampalam, Wiji, Sugato Dasgupta, Amrita Dhillon, and Bhaskar Dutta**, “The Spend-and-Tax Effect of Elected Legislators: Evidence from India,” *Journal of the European Economic Association*, 2009, 7 (2-3), 683–692.
- Asher, Sam and Paul Novosad**, “Politics and Local Economic Growth: Evidence from India,” *American Economic Journal: Applied Economics*, 2017, 9 (1), 229–273.
- , **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,” *The World Bank Economic Review*, 2021, 35 (4), 845–871.
- Bardhan, Pranab and Dilip Mookherjee**, “Determinants of Redistributive Politics: An Empirical Analysis of Land Reforms in West Bengal, India,” *American Economic Review*, 2010, 100 (4), 1572–1600.
- Besley, Timothy and Robin Burgess**, “Can Labor Regulation Hinder Economic Performance? Evidence from India,” *The Quarterly Journal of Economics*, 2004, 119 (1), 91–134.
- Bhavnani, Rikhil R.**, “Indian State Assembly Elections Data, 1977-2015,” Harvard Dataverse 2017.
- Brollo, Fernanda, Tommaso Nannicini, Roberto Perotti, and Guido Tabellini**, “The Political Resource Curse,” *American Economic Review*, 2013, 103 (5), 1759–1796.
- Burgess, Robin, Remi Jedwab, Edward Miguel, Ameet Morjaria, and Gerard Padró i Miquel**, “The Value of Democracy: Evidence from Road Building in Kenya,” *American Economic Review*, 2015, 105 (6), 1817–1851.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.
- , – , **Max H. Farrell, and Rocio Titiunik**, “rdrobust: Software for Regression-Discontinuity Designs,” *The Stata Journal*, 2017, 17 (2), 372–404.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, 115 (531), 1449–1455.

- , **Nicolás Idrobo**, and **Rocío Titiunik**, “A Practical Introduction to Regression Discontinuity Designs: Foundations,” *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2020.
- 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.
- Cole, Shawn**, “Fixing Market Failures or Fixing Elections? Agricultural Credit in India,” *American Economic Journal: Applied Economics*, 2009, *1* (1), 219–250.
- Donaldson, Dave and Adam Storeygard**, “The View from Above: Applications of Satellite Data in Economics,” *Journal of Economic Perspectives*, 2016, *30* (4), 171–198.
- Eggers, Andrew C., Anthony Fowler, Jens Hainmueller, Andrew B. Hall, and James M. Snyder Jr**, “On the Validity of the Regression Discontinuity Design for Estimating Electoral Effects: New Evidence from Over 40,000 Close Races,” *American Journal of Political Science*, 2015, *59* (1), 259–274.
- Elvidge, Christopher D., Kimberly Baugh, Mikhail Zhizhin, Feng Chi Hsu, and Tilottama Ghosh**, “VIIRS Night-Time Lights,” *International Journal of Remote Sensing*, 2017, *38* (21), 5860–5879.
- Finan, Frederico and Claudio Ferraz**, “Electoral Incentives and the Allocation of Public Funds,” *Journal of the European Economic Association*, 2011, *9* (5), 844–879.
- Golden, Miriam A. and Brian Min**, “Distributive Politics Around the World,” *Annual Review of Political Science*, 2013, *16*, 73–99.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil**, “Measuring Economic Growth from Outer Space,” *American Economic Review*, 2012, *102* (2), 994–1028.
- Imbens, Guido W. and Thomas Lemieux**, “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, 2008, *142* (2), 615–635.
- Iyer, Lakshmi, Anandi Mani, Prachi Mishra, and Petia Topalova**, “The Power of Political Voice: Women’s Political Representation and Crime in India,” *American Economic Journal: Applied Economics*, 2012, *4* (4), 165–193.
- Khemani, Stuti**, “Does Delegation of Fiscal Policy to an Independent Agency Make a Difference? Evidence from Intergovernmental Transfers in India,” *Journal of Development Economics*, 2007, *82* (2), 464–484.

Lee, David S., “Randomized Experiments from Non-random Selection in U.S. House Elections,” *Journal of Econometrics*, 2008, *142* (2), 675–697.

– **and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 2010, *48* (2), 281–355.

Solé-Ollé, Albert and Pilar Sorribas-Navarro, “Do Transfers to Local Governments Affect Electoral Outcomes?,” *Journal of Public Economics*, 2004, *88* (12), 2515–2532.

Storeygard, Adam, “Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa,” *The Review of Economic Studies*, 2016, *83* (3), 1263–1295.

A. Data Appendix

A.1 Election Data

The Indian State Assembly Elections dataset (Bhavnani, 2017) is available from the Harvard Dataverse (doi: 10.7910/DVN/26526). The raw data contain 327,294 candidate-level records spanning elections from 1977 to 2015 across all Indian states and union territories. Variables include state name, election year, constituency number and name, constituency type (General, SC-reserved, ST-reserved), candidate name and sex, party abbreviation and name, and total votes polled.

Sample construction. I restrict to elections from 2008 onward to match the 2007 delimitation boundaries used in SHRUG. Within each constituency-election, I rank candidates by total votes and retain the top two. I then reshape the data to one row per constituency-election, computing the vote margin as the winner’s vote share minus the runner-up’s share among their combined votes. This yields 33,486 constituency-elections.

State ruling party. For each state-year, I identify the party winning the most assembly seats and define it as the state ruling party. In cases of ties, I retain the first party alphabetically (this affects very few elections). The state-alignment indicator equals one if the constituency winner belongs to this party.

Central ruling party. I assign the central ruling party based on the known sequence of Indian prime ministers: INC for 2004–2013, BJP for 2014 onward (with appropriate handling of coalition-era complexities). INC and INC(I) are treated as the same party. The center-alignment indicator equals one if the winner’s party matches the central ruling party.

A.2 SHRUG Nighttime Lights

VIIRS nighttime lights data are from the SHRUG platform v2.1 (Asher et al., 2021), downloaded from https://www.devdatalab.org/shrug_download/. The VIIRS annual series covers 2012–2023 at the 2007-delimitation assembly constituency level, with 4,080 unique constituencies. I also obtained DMSP-OLS data (1992–2013) from the same platform, though DMSP data are not used in the main analysis.

Outcome construction. For each constituency-election, I merge VIIRS annual mean radiance for years -3 through $+5$ relative to the election year. I compute log nightlights as $\ln(\text{VIIRS mean radiance} + 0.01)$, where the constant prevents log of zero for dark constituencies. The primary outcome is the average of log nightlights over post-election years 1 through 4.

A.3 Census Covariates

Census 2011 data are from the SHRUG platform’s Primary Census Abstract (PCA) at the assembly constituency level. I construct: total population, literacy rate (literate population / total population), SC share (SC population / total), ST share (ST population / total), work participation rate (total workers / total population), and agriculture share (cultivators + agricultural laborers / total workers).

A.4 Geographic Matching

Assembly constituencies are matched between election data and SHRUG using a constructed identifier: 2007-[state_code]-[ac_number]. State codes follow the Census 2011 numbering. This achieves a 99.0% match rate (3,922 of 3,962 post-2008 election ACs matched to 4,080 SHRUG ACs).

B. Identification Appendix

B.1 Full Density Test Results

The [Cattaneo et al. \(2020a\)](#) density continuity test for the state-alignment running variable returns $T = 2.01$, $p = 0.045$ (using the default robust bias-corrected test). For the center-alignment running variable, the test returns $p = 0.862$. [Figure 3](#) shows the density histogram.

Several factors contextualize the borderline state-alignment result. First, the running variable is constructed with the “ruling party” defined as the seat-plurality winner, which mechanically correlates the sign of the margin with the density distribution. Second, [Eggers et al. \(2015\)](#) document that density tests have non-trivial false positive rates in large election datasets even absent manipulation, particularly when the running variable is defined relative to systemically popular parties. Third, all results are qualitatively unchanged in donut specifications that drop the narrowest margins.

B.2 Extended Balance Results

The covariate balance analysis in [Table 2](#) uses the same `rdrobust` estimator applied to each covariate. Of eight covariates tested, six show no significant imbalance ($p > 0.10$). The population imbalance (32,000 additional residents in aligned constituencies, $p = 0.03$) and SC share imbalance (2.9 pp, $p = 0.003$) are the only exceptions.

To assess whether these imbalances could confound the main estimates, note that log nightlights is well-balanced ($p = 0.75$). Since nightlights are the outcome variable, balance in

the outcome’s baseline level is the most relevant diagnostic. The population and SC share imbalances would only bias the results if they are correlated with alignment-driven changes in nightlights (not just levels), which is testable and not supported by the data.

C. Robustness Appendix

C.1 Full Bandwidth Sensitivity

Estimates are computed at every integer bandwidth from 3% to 20%. At no bandwidth does the state-alignment or center-alignment estimate achieve statistical significance at the 10% level. The largest absolute estimate for state alignment occurs at $h = 10\%$ ($\tau = 0.143$, $p = 0.35$); the smallest at $h = 20\%$ ($\tau = 0.096$, $p = 0.39$). Center-alignment estimates range from -0.18 to -0.05 , all insignificant.

C.2 Permutation Diagnostic Details

The permutation exercise uses 500 shuffles of the running variable within the MSE-optimal bandwidth sample. The observed mean difference in nightlights (aligned minus unaligned within the bandwidth) is 0.118. The permutation p -value (fraction of permuted differences with absolute value exceeding the observed) is 0.05. This borderline result reflects the simple mean-difference test statistic, which does not condition on the running variable. The parametric RDD estimate (which controls for the running variable via local polynomial) yields a higher p -value of 0.41, confirming that the permutation result is driven by within-bandwidth correlation between the running variable and outcomes rather than a true discontinuity at the cutoff. This diagnostic is not a formal randomization inference test for RD designs ([Cattaneo et al., 2020b](#)) but is reported for transparency.

D. Additional Figures and Tables

All main figures and tables are presented in the body of the paper. Replication code and data are available in the project repository.