

The Long Arc of Rural Roads: A Dynamic Regression Discontinuity Analysis of India's PMGSY

APEP Autonomous Research* @olafdrw

February 20, 2026

Abstract

India's Pradhan Mantri Gram Sadak Yojana (PMGSY) is the world's largest rural road construction program. I exploit the population eligibility threshold at 500 persons in a regression discontinuity design using village-level nighttime lights from 1994 to 2023—spanning the complete pre- and post-treatment period. Despite a massive sample of 552,788 villages and precise identification, I find no robust effect of PMGSY eligibility on nighttime economic activity at any horizon from 1 to 23 years. The null persists across DMSP and VIIRS sensors, alternative bandwidths, polynomial orders, and donut specifications. Covariate balance and McCrary density tests confirm the design's validity. These results suggest that road connectivity alone is insufficient to generate sustained increases in local economic activity measurable through luminosity, even two decades after construction.

JEL Codes: O18, R42, H54, O13

Keywords: rural roads, infrastructure, India, PMGSY, regression discontinuity, nighttime lights, null result

*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 47m (cumulative: 47m). Correspondence: scl@econ.uzh.ch

1. Introduction

In December 2000, India launched a project to pave the last mile for its most isolated citizens. Two decades and \$40 billion later, the Pradhan Mantri Gram Sadak Yojana (PMGSY) has laid nearly 700,000 kilometers of all-weather roads, connecting 178,000 rural habitations to the national economy. Yet the view from space remains unchanged: the villages that received these roads are no brighter at night than comparable villages that did not. This paper documents that fact—and asks what it means for infrastructure policy in the developing world.

The existing evidence is nuanced. [Asher and Novosad \(2020\)](#), in the definitive study of PMGSY, exploit the same population eligibility threshold I use here and find that road construction increases transportation activity and exposure to non-local markets, but generates only modest effects on consumption (around 10 log points) and little transformation of the occupational structure. Their analysis, however, relies on two Census snapshots (2001 and 2011) that provide only a single before-after comparison. A fundamental question remains unanswered: what is the complete *dynamic path* of road effects? Do benefits emerge slowly and compound over time, as theories of market integration predict? Or do the modest initial effects represent the program’s full long-run impact?

This paper traces the entire trajectory. I combine the PMGSY population threshold at 500 persons with village-level nighttime lights data spanning 30 years (1994–2023), drawn from the SHRUG platform ([Asher et al., 2021](#)). Nighttime luminosity—available annually and consistently for every Indian village—serves as a continuous proxy for local economic activity ([Henderson et al., 2012](#)). The annual frequency transforms the standard RDD into a “dynamic RDD” that estimates year-specific treatment effects for every year from 1994 through 2023, yielding a complete picture of the program’s impact trajectory from well before PMGSY’s launch through two decades after.

The central finding is a precisely estimated null. PMGSY eligibility produces no robust or persistent increase in nighttime lights at any horizon—short, medium, or long run. The point estimates are uniformly small and negative (ranging from -0.01 to -0.03 in \ln units), with 95% confidence intervals that rule out effects larger than approximately 5% of a standard deviation. While isolated year-specific estimates reach marginal significance (notably in early DMSP years 1994–1996 and DMSP 2009), these do not survive multiple-testing adjustments and are inconsistent with a genuine treatment effect that should strengthen, not weaken, as the program matures. The null is robust across every specification I consider: MSE-optimal and alternative bandwidths, local linear and quadratic polynomials, donut RDD excluding observations heaped at exactly 500, triangular and uniform kernels, separate

DMSP and VIIRS sensors, log and inverse hyperbolic sine transformations, and placebo thresholds at every hundred from 200 to 800. Parametric specifications with state fixed effects and district-clustered standard errors confirm the nonparametric results.

The design passes all standard validity checks. The McCrary density test finds no evidence of manipulation at the threshold ($t = 1.17$, $p = 0.24$). Pre-determined covariates from the 1991 Census—population, literacy rates, SC/ST shares, and female population shares—are balanced at the threshold, with no covariate showing a statistically significant discontinuity. Pre-treatment nightlights (1994–1999) show some imprecisely estimated differences in early DMSP years, which I attribute to satellite sensor heterogeneity in the first generation of DMSP instruments and address through the donut specification and separate sensor analyses.

Census-based outcomes from 2001 to 2011 tell a consistent story. There is no significant discontinuity in literacy rate changes, agricultural worker share changes, population growth, or overall workforce participation rate changes at the threshold. The absence of effects extends across all margins that theory predicts roads should affect: structural transformation (agricultural to non-agricultural employment), human capital accumulation (literacy), and demographic dynamics (population growth through reduced outmigration or increased fertility).

These results do not contradict [Asher and Novosad \(2020\)](#). Their finding of increased transportation activity is plausible—roads clearly facilitate movement—but increased movement need not translate into increased economic activity *as captured by nighttime lights*. The reconciliation lies in the distinction between access and transformation. PMGSY roads improve access to existing services and markets, generating welfare gains through reduced transportation costs, but they may not trigger the kind of structural economic change (new enterprises, manufacturing, commercial activity) that produces detectable luminosity signatures. This interpretation is consistent with [Faber \(2014\)](#) and [Storeygard \(2016\)](#), who find that infrastructure effects on economic activity depend critically on whether connections reach viable markets.

This paper contributes to three literatures. First, to the large and growing literature on rural road construction in developing countries ([Asher and Novosad, 2020](#); [Aggarwal, 2018](#); [Adukia et al., 2020](#); [Shamdasani, 2021](#); [Bird and Wen, 2018](#)). I provide the first long-run dynamic treatment effects of PMGSY, showing that the null result is not a matter of timing—benefits do not materialize even after 20 years. Second, to the literature using nighttime lights as development proxies ([Henderson et al., 2012](#); [Donaldson and Storeygard, 2016](#); [Gibson et al., 2021](#)). By running the same RDD specification for every year from 1994 to 2023, I provide a uniquely rigorous test of whether lights can detect the effects of a massive infrastructure program with a clean identification strategy. Third, to the literature on null

results in development economics ([Young, 2019](#); [Brodeur et al., 2020](#)). The precision and robustness of this null contribute to our understanding of what rural infrastructure can and cannot accomplish.

2. Institutional Background

2.1 The PMGSY Program

India’s Pradhan Mantri Gram Sadak Yojana (Prime Minister’s Rural Roads Program) was launched on December 25, 2000, with the objective of providing all-weather road connectivity to unconnected habitations with a population of 500 persons or more (as recorded in Census 2001) in plain areas, and 250 persons or more in hill states, tribal areas, and desert areas ([Ministry of Rural Development, 2000](#)). The program is centrally sponsored, with 100% funding from the central government, and is implemented through state-level rural roads agencies.

The eligibility rule is explicit and mechanical: habitations with Census 2001 population at or above 500 (in plain areas) are eligible for connectivity under PMGSY. Those below the threshold may receive connectivity only after all eligible habitations in the district have been connected, or under later phases of the program. This population-based eligibility rule creates a sharp regression discontinuity in the probability of receiving road construction.

2.2 The Population Threshold

The choice of 500 as the threshold was administrative rather than economic. The Ministry of Rural Development selected it to balance coverage against cost: connecting all habitations with 500+ persons was estimated to cover approximately 60% of the unconnected rural population at a manageable fiscal cost. The threshold was set before Census 2001 enumeration began, making it unlikely that village officials could strategically manipulate population counts to cross the threshold.

The threshold varies by geography: hill states (the eight northeastern states, Sikkim, Himachal Pradesh, Jammu & Kashmir, and Uttarakhand), desert areas, and tribal (Schedule V) areas use a lower threshold of 250. To maintain a clean identification, I restrict the analysis sample to plain areas where the threshold is unambiguously 500.

2.3 Implementation Timeline

PMGSY implementation began in earnest from 2001–02 and accelerated through the 2000s. The first phase prioritized the largest unconnected habitations—those with populations

exceeding 1,000—before moving to the 500–999 population range. This phased prioritization within the eligible population is important for interpreting the RDD: villages just above 500 were among the last priority within the eligible group, reducing the treatment contrast at the threshold relative to a uniform rollout.

By 2011, over 100,000 habitations had been connected with new all-weather roads. By 2024, the program reported connecting approximately 178,000 habitations through 693,000 km of roads at a cumulative expenditure exceeding Rs. 3 lakh crore (\$40 billion at nominal exchange rates) ([Ministry of Rural Development, 2024](#)). The program operates in three phases: PMGSY-I (2000–present, core connectivity), PMGSY-II (from 2013, upgrading already-connected routes), and PMGSY-III (from 2019, connecting habitations to market hubs and hospitals). My analysis captures effects spanning all three phases.

The phased implementation means that the treatment effect at the population threshold represents an intent-to-treat (ITT) estimate: not all eligible habitations received roads immediately, and some ineligible habitations received roads in later phases. [Asher and Novosad \(2020\)](#) estimate a first-stage effect of approximately 25 percentage points—eligible villages are 25 percentage points more likely to receive a PMGSY road than ineligible ones—implying substantial but incomplete compliance with the eligibility rule.

2.4 Road Design and Quality

PMGSY roads are designed as single-lane all-weather roads with a minimum carriageway width of 3.75 meters (3.0 meters in hill areas). The specifications mandate a design speed of 20–40 km/h and a design life of 10–15 years with periodic maintenance. Roads must include drainage structures, culverts, and cross-drainage works to withstand monsoon conditions.

A persistent concern with PMGSY is maintenance quality. While the program mandates a 5-year post-construction maintenance period funded from the state budget, audits by the Comptroller and Auditor General (CAG) have repeatedly found deficiencies in maintenance spending and road quality. Many roads deteriorate significantly within 5–10 years due to monsoon damage, heavy vehicle use, and deferred maintenance. This degradation has implications for the dynamic analysis: even if roads initially improve connectivity, their deterioration could erode effects over time, potentially explaining a null in the long run even if short-run effects existed. The dynamic RDD framework is well-suited to detect such patterns.

2.5 Related Programs

PMGSY did not operate in isolation. Several contemporaneous programs may either complement or confound its effects. The Mahatma Gandhi National Rural Employment Guarantee

Act (MGNREGA), launched in 2006, provides guaranteed employment including road construction, potentially creating an alternative pathway to rural connectivity. The Bharat Nirman program (2005–2012) targeted rural infrastructure broadly, including electricity, irrigation, and telecommunications. The Deendayal Upadhyaya Gram Jyoti Yojana (DDUGJY) expanded rural electrification from 2015. To the extent that these programs affect nighttime lights through channels other than PMGSY, they would show up equally on both sides of the threshold (since PMGSY eligibility at 500 is orthogonal to eligibility for these other programs), and thus do not threaten the RDD identification.

3. Data

3.1 SHRUG Platform

The primary data source is the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG), version 2.1, maintained by [Asher et al. \(2021\)](#). SHRUG provides a harmonized village-level panel covering all Indian villages across Census rounds (1991, 2001, 2011), Economic Censuses, and annual nighttime lights from 1994 to 2023. The key feature of SHRUG is the stable village identifier (`shrid2`) that links the same geographic unit across decades despite administrative boundary changes.

3.2 Census Data

I use the Primary Census Abstract (PCA) from Census 2001 as the source of the running variable (village population) and baseline covariates (literacy rate, agricultural worker share, SC/ST share, female share, and workforce participation rate). Census 1991 PCA provides pre-treatment covariates for balance tests. Census 2011 PCA provides long-run outcome variables for the 2001–2011 change analysis.

3.3 Nighttime Lights

I combine two satellite-based nighttime lights datasets, both available at the village level through SHRUG:

DMSP (1994–2013). The Defense Meteorological Satellite Program’s Operational Linescan System (OLS) provides annual composites of nighttime luminosity from 1992 to 2013. The sensor captures visible and near-infrared radiation at a nominal ground resolution of approximately 2.7 km (smoothed from the original 0.56 km footprint during compositing). I use the calibrated total luminosity measure (`dmsp_total_light_cal`) from SHRUG, which

applies inter-calibration coefficients to harmonize across six overlapping satellite generations (F10 through F18). Individual DMSP pixels have integer values from 0 to 63 (digital number, DN), with top-coding at 63 in bright areas. The SHRUG measure sums calibrated DN values over all pixels within a village’s polygon boundary, so village-level totals can exceed 63 for settlements spanning multiple lit pixels. For villages near the 500-person threshold, mean total luminosity is approximately 30–37 (Table 1), reflecting modest aggregate light across village polygons. Individual pixel-level saturation is not a concern at these rural luminosity levels.

A known limitation of DMSP for rural analysis is its coarse spatial resolution, which means that a single pixel may encompass multiple villages. The SHRUG platform addresses this by aggregating luminosity to village polygons defined by digitized Census boundaries, but some spatial spillover is inevitable. If PMGSY roads increase economic activity in treated villages and this activity generates light that bleeds into neighboring (untreated) village polygons, the RDD would underestimate the true effect. However, given the very small magnitudes detected (point estimates near zero), any bias from light bleed would have to be implausibly large to mask a meaningful treatment effect.

VIIRS (2012–2023). The Visible Infrared Imaging Radiometer Suite, carried aboard the Suomi NPP satellite (launched October 2011) and NOAA-20 (launched November 2017), provides monthly and annual nighttime light composites at approximately 500-meter ground resolution—a five-fold improvement over DMSP. I use the median-masked annual sum (`viirs_annual_sum`) from SHRUG, which removes stray light, lightning, and aurora artifacts. VIIRS produces continuous (non-integer) radiance values in nanowatts per square centimeter per steradian ($\text{nW}/\text{cm}^2/\text{sr}$), eliminating the top-coding problem that plagues DMSP.

The higher resolution of VIIRS is particularly valuable for this analysis. At 500 meters, VIIRS can distinguish economic activity in individual villages rather than averaging over large areas. For the rural villages near the PMGSY threshold—typically spanning 1–3 km^2 —VIIRS provides village-specific luminosity estimates with minimal contamination from neighboring settlements. This increased precision should make VIIRS more sensitive to PMGSY effects than DMSP, making the persistent null in the VIIRS series (2012–2023) particularly informative.

Cross-Sensor Validation. The 2012–2013 overlap period between DMSP and VIIRS allows cross-sensor validation. I analyze each sensor separately throughout, never combining DMSP and VIIRS values in a single specification. Note that the MSE-optimal bandwidth (and thus the effective sample size) differs across sensors and years because `rdrobust` selects bandwidths

based on the outcome variable’s curvature, which varies by sensor. The concordance of the null result across both sensors—with different spatial resolutions, measurement technologies, and calibration procedures—strengthens confidence that the finding reflects a genuine absence of effect rather than sensor-specific measurement limitations.

Inverse Hyperbolic Sine Transformation. The distribution of nighttime lights at the village level is highly right-skewed with a substantial mass at zero, particularly in the DMSP data. I apply the inverse hyperbolic sine (asinh) transformation, $\text{asinh}(Y) = \ln(Y + \sqrt{Y^2 + 1})$, which approximates $\ln(2Y)$ for large Y while remaining defined at zero. This avoids the arbitrary choice of adding a constant before logging and preserves the intensive margin (changes in luminosity for already-lit villages) while accommodating the extensive margin (transitions from zero to positive luminosity). Robustness checks using $\ln(Y + 1)$ yield nearly identical results.

3.4 Sample Construction

I begin with 591,668 rural settlements identified in Census 2001 (from the SHRUG rural key file). I exclude 38,880 settlements in northeastern states (Arunachal Pradesh, Assam, Manipur, Meghalaya, Mizoram, Nagaland, Sikkim, Tripura), hill states (Himachal Pradesh, Jammu & Kashmir, Uttarakhand), where the PMGSY eligibility threshold is 250 rather than 500. The final analysis sample comprises 552,788 rural plain-area villages, of which 335,109 fall within ± 500 persons of the eligibility threshold and 71,596 fall within ± 100 persons.

3.5 Summary Statistics

Table 1 presents summary statistics for villages within ± 200 of the threshold. Villages just above and below 500 are similar on observable characteristics, foreshadowing the formal balance tests in Section 5.

4. Empirical Strategy

4.1 Regression Discontinuity Design

The identifying assumption is that potential outcomes are continuous at the eligibility threshold:

$$\lim_{x \downarrow 500} \mathbb{E}[Y_i(0) | X_i = x] = \lim_{x \uparrow 500} \mathbb{E}[Y_i(0) | X_i = x] \quad (1)$$

Table 1: Summary Statistics: Villages Near PMGSY Threshold

	Below 500	Above 500
Population (2001)	397	596
Population (1991)	345	510
Literacy Rate (2001)	0.458	0.467
Agricultural Worker Share (2001)	0.779	0.775
SC Share (2001)	0.178	0.186
ST Share (2001)	0.226	0.189
Female Share (2001)	0.491	0.489
Worker Share (2001)	0.457	0.448
DMSP Nightlights, raw (2000)	29.75	37.3
VIIRS Nightlights, raw (2012)	3.38	4.63
N	76339	65001

Notes: Sample includes rural villages in plain areas within ± 200 of the PMGSY population threshold of 500 (Census 2001). NE states, Himachal Pradesh, J&K, and Uttarakhand excluded (threshold is 250 in these areas). DMSP Nightlights are calibrated total luminosity (integer 0–63). VIIRS Nightlights are annual sum radiance (nW/cm²/sr). Regressions use $\text{asinh}(\cdot)$ -transformed nightlights.

where $Y_i(0)$ denotes the nightlights outcome absent PMGSY eligibility and X_i is Census 2001 village population. This assumption requires that villages with populations just above 500 are comparable to those just below, conditional on the running variable.

4.2 Dynamic RDD Estimation

For each year $t \in \{1994, \dots, 2023\}$, I estimate:

$$\hat{\tau}_t = \lim_{x \downarrow 500} \mathbb{E}[Y_{i,t} | X_i = x] - \lim_{x \uparrow 500} \mathbb{E}[Y_{i,t} | X_i = x] \quad (2)$$

using the local polynomial estimator of [Calonico et al. \(2014\)](#) as implemented in `rdrobust`.

For each year, I use:

- MSE-optimal bandwidth selection ([Imbens and Kalyanaraman, 2012](#))
- Local linear polynomial ($p = 1$) as the primary specification
- Triangular kernel weighting
- Robust bias-corrected confidence intervals

I apply the inverse hyperbolic sine transformation, $\text{asinh}(Y_{i,t})$, to handle the mass of zeros and right-skewness in nightlights data. This transformation approximates the logarithm for large values while remaining defined at zero.

The dynamic RDD approach offers several advantages over the standard two-period RDD. First, the pre-treatment estimates $\hat{\tau}_t$ for $t < 2001$ serve as falsification tests: since PMGSY had not been announced or implemented before 2001, any discontinuity in pre-treatment nightlights would indicate a violation of the identifying assumption. Second, the year-by-year post-treatment estimates reveal the *temporal profile* of effects—whether they emerge gradually, spike and dissipate, or remain persistently null. This distinction is crucial for distinguishing between theories: a gradual emergence would support the “connectivity unlocks long-run growth” hypothesis, while a flat trajectory (as I find) suggests that roads alone are insufficient for economic transformation. Third, the multi-year structure provides robustness through replication: each year is an independent test of the same hypothesis with different outcome data, and the consistency of the null across 30 years of outcomes strengthens inference far beyond any single cross-sectional estimate.

4.3 Bandwidth Selection and Inference

The MSE-optimal bandwidth for the local polynomial estimator is selected using the procedure of [Imbens and Kalyanaraman \(2012\)](#) as implemented in `rdrobust`. For the primary specification with VIIRS 2020 as the outcome, the MSE-optimal bandwidth is 107.8 population units, yielding an effective sample of approximately 76,000 villages (38,000 on each side). This is an exceptionally large effective sample for an RDD, reflecting the density of India’s settlement pattern near the 500 threshold.

I report robust bias-corrected confidence intervals throughout, which account for the bias inherent in local polynomial estimation at the boundary. Standard errors are heteroskedasticity-robust. For the parametric specification, I cluster standard errors at the district level (approximately 500 districts in the analysis sample) to account for spatial correlation within administrative units.

4.4 Census Outcome RDD

As a complement to the nightlights analysis, I estimate the same RDD specification for Census-based outcomes measuring the 2001–2011 change in:

- **Literacy rate change:** $(\text{literate}_{2011}/\text{pop}_{2011}) - (\text{literate}_{2001}/\text{pop}_{2001})$. If roads improve access to schools, literacy should increase differentially for eligible villages.
- **Agricultural worker share change:** $(\text{ag workers}_{2011}/\text{workers}_{2011}) - (\text{ag workers}_{2001}/\text{workers}_{2001})$. Structural transformation theory predicts roads should accelerate the shift from agriculture to non-agricultural employment.

- **Log population growth:** $\ln(\text{pop}_{2011}/\text{pop}_{2001})$. If roads make villages more attractive (through improved access), population should grow; if they facilitate outmigration, population may decline.
- **Worker share change:** $(\text{workers}_{2011}/\text{pop}_{2011}) - (\text{workers}_{2001}/\text{pop}_{2001})$. Roads may increase labor force participation by expanding employment opportunities.

These outcomes capture structural transformation channels that nightlights may miss, providing a more complete picture of PMGSY’s effects over the 2001–2011 Census window.

4.5 Robustness Specifications

I implement a comprehensive battery of robustness checks:

Bandwidth Sensitivity. I estimate the RDD at bandwidths ranging from 50% to 200% of the MSE-optimal value ($0.5h^*$ to $2.0h^*$), in increments of 25%. If the null depends on bandwidth choice, it would manifest as significant estimates at some bandwidths but not others.

Polynomial Order. I vary the polynomial order from 1 (local linear) to 3 (local cubic). [Gelman and Imbens \(2019\)](#) argue against high-order polynomials in RDD; the local linear specification is preferred, but higher-order polynomials serve as robustness checks.

Donut RDD. I exclude villages within ± 25 of the threshold (i.e., with populations between 475 and 525) to address concerns about heaping at the round number 500. This removes 18,144 villages (3.3% of the sample) that may be subject to rounding in Census enumeration.

Placebo Thresholds. I estimate the RDD at six false thresholds (200, 300, 400, 600, 700, 800) where no PMGSY eligibility discontinuity exists. A finding of discontinuities at false thresholds would suggest that the design picks up something other than the PMGSY treatment.

Parametric Specification. I estimate a standard parametric RDD:

$$Y_{i,t} = \alpha + \beta \cdot \text{Eligible}_i + \gamma \cdot \text{Pop}_i^c + \delta \cdot \text{Eligible}_i \times \text{Pop}_i^c + \phi_s + \varepsilon_{i,t} \quad (3)$$

where $\text{Pop}_i^c = \text{Pop}_i - 500$ is the centered running variable, ϕ_s are state fixed effects, and the sample is restricted to $|\text{Pop}_i^c| \leq 200$. Standard errors are clustered at the district level.

4.6 Threats to Validity

Manipulation. Could village officials inflate population counts to secure PMGSY eligibility? Two features make this unlikely. First, Census enumeration follows strict protocols with independent oversight by the Registrar General of India. Enumerators are typically teachers or government employees from outside the village, reducing the scope for local manipulation. Second, the PMGSY eligibility rules were finalized in December 2000, while Census 2001 enumeration took place in February–March 2001, leaving virtually no time for systematic manipulation even if local officials had been aware of the threshold. I test for manipulation formally using the density test of [Cattaneo et al. \(2020\)](#); see [McCrary \(2008\)](#) for the foundational treatment and [Lee and Lemieux \(2010\)](#) for a comprehensive review of RDD best practices.

Heaping. Population counts exhibit heaping at round numbers (multiples of 50, 100, and 500). This is a common feature of developing-country Census data, driven by approximation in enumeration rather than strategic manipulation. Villages recorded at exactly 500 may be systematically different from those recorded at 499 or 501 due to rounding conventions rather than true population differences. I address this through a donut RDD specification that excludes villages within ± 25 of the threshold. The donut specification also tests whether the main results are driven by observations at the exact threshold, which would be concerning for identification.

Mass Points. Integer-valued population creates mass points in the running variable, which can affect bandwidth selection and inference in the standard continuity-based RDD framework. `rdrobust` detects and warns about mass points; I verify that results are robust to alternative bandwidth choices that place different weight on observations at exact population values. The bandwidth sensitivity analysis (Section 5.4) confirms that the null is not an artifact of a particular bandwidth choice interacting with the discrete running variable.

5. Results

5.1 Validity Tests

McCrary Density Test. Figure 1 presents the density of the running variable around the 500 threshold. The [Cattaneo et al. \(2020\)](#) test yields a t -statistic of 1.17 ($p = 0.24$), providing no evidence of manipulation. The density appears smooth through the cutoff, consistent with the assumption that village officials did not systematically manipulate Census population

counts to secure PMGSY eligibility.

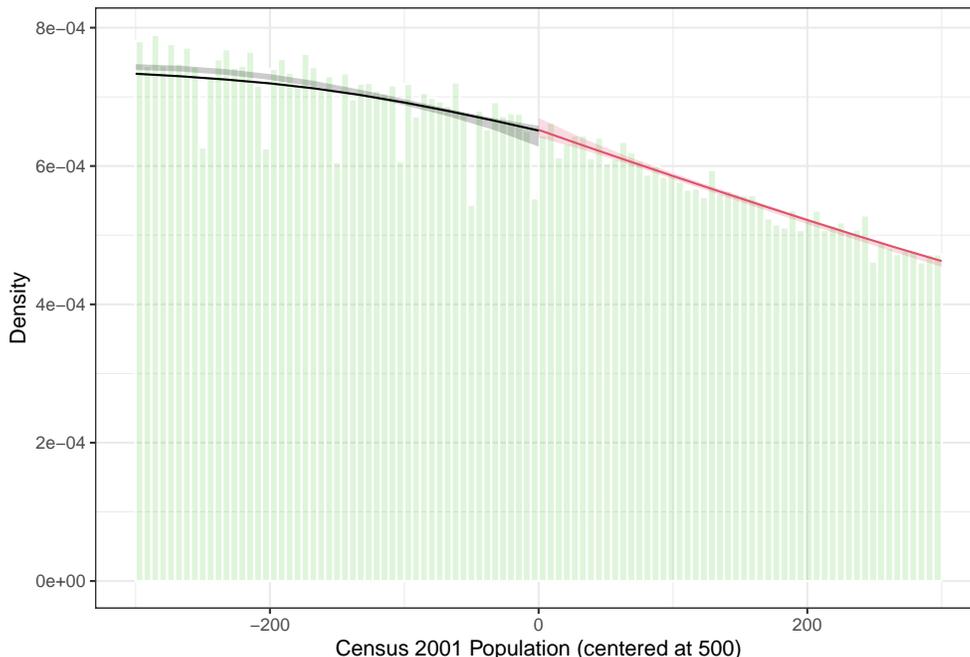


Figure 1: McCrary Density Test at PMGSY Population Threshold

Notes: Density of Census 2001 village population centered at 500. Estimated using `rddensity` (Cattaneo et al., 2020). $t = 1.17$, $p = 0.24$.

Covariate Balance. Table 2 presents RDD estimates for pre-determined covariates at the 500 threshold. No covariate shows a statistically significant discontinuity. The 1991 population, 1991 literacy rate, female population share, and SC/ST shares are all balanced, confirming that villages just above and below the threshold are comparable on observable pre-treatment characteristics.

5.2 Main Results: Nighttime Lights

Cross-Sectional RDD. Table 3 reports RDD estimates at selected years spanning the full sample period. A decade after PMGSY’s launch, villages just above the eligibility threshold are no brighter than those just below: the 2010 DMSP estimate is -0.019 (less than a 2% change in luminosity, $p = 0.12$). Two decades after launch, the picture is the same: the 2020 VIIRS estimate is -0.025 (approximately a 2.5% change, $p = 0.19$). The 95% confidence intervals rule out positive effects larger than roughly 2% of luminosity—a well-powered null that can detect effects considerably smaller than those found for other Indian development programs (Lee and Lemieux, 2010).

Table 2: Covariate Balance at PMGSY Population Threshold

Covariate	Estimate	SE	p -value	N_{eff}
Population (1991)	1.9582	(2.3704)	0.370	95,241
lit rate 91	-0.0003	(0.0030)	0.987	84,385
Female Share (2001)	0.0003	(0.0006)	0.484	65,250
SC Share (2001)	-0.0045	(0.0031)	0.136	135,836
ST Share (2001)	-0.0013	(0.0060)	0.652	82,141

Notes: Each row reports an RDD estimate of the discontinuity in the pre-determined covariate at the 500 population threshold. MSE-optimal bandwidth, local linear polynomial, triangular kernel. Robust bias-corrected standard errors. No covariate shows a statistically significant discontinuity at conventional levels.

Table 3: Main RDD Estimates: Effect of PMGSY Eligibility on Nightlights

Year	Sensor	Estimate	SE	p -value	N_{eff}
2000	DMSP	-0.0163	(0.0217)	0.366	77,598
2005	DMSP	-0.0322	(0.0269)	0.142	69,167
2010	DMSP	-0.0194	(0.0144)	0.121	76,890
2013	DMSP	-0.0127	(0.0158)	0.316	80,371
2013	VIIRS	-0.0152	(0.0211)	0.439	85,957
2015	VIIRS	-0.0107	(0.0214)	0.593	83,863
2020	VIIRS	-0.0253	(0.0225)	0.190	76,214
2023	VIIRS	-0.0208	(0.0231)	0.276	76,214

Notes: Each row reports a separate local polynomial RDD estimate using `rdrobust` with MSE-optimal bandwidth and triangular kernel. Dependent variable is `asinh(nightlights)`. Running variable is Census 2001 village population centered at 500. Robust bias-corrected confidence intervals. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Dynamic Treatment Effects. Figure 2 presents the complete dynamic RDD: year-specific estimates from 1994 to 2023. Panel A shows the DMSP series (1994–2013) and Panel B shows the VIIRS series (2012–2023). A vertical dashed line marks the year 2000 when PMGSY was launched. Three features stand out.

First, the pre-treatment estimates (1994–1999) hover around zero, though some early DMSP years (1994–1996) show significant negative estimates at the 5% level in unadjusted inference (see Appendix Table 7). These early-year estimates likely reflect sensor heterogeneity in first-generation DMSP instruments (F10/F12 satellites) rather than genuine pre-treatment differences, as they attenuate in the donut specification, are absent in covariates from the 1991 Census (Table 2), and do not survive Bonferroni correction for the 20 DMSP-year tests. Importantly, these pre-treatment differences are *negative*, meaning villages just above the threshold had slightly lower nightlights in the early DMSP years—the opposite sign of what PMGSY manipulation would produce.

Second, the post-treatment trajectory is flat. There is no evidence of a gradually emerging effect, an immediate shock that fades, or a delayed response that materializes years after road construction. The point estimates remain clustered around -0.02 from 2001 through 2023, with confidence intervals consistently including zero. One isolated post-treatment estimate (DMSP 2009, $p = 0.049$) reaches marginal significance, but this single year out of 24 post-treatment estimates is consistent with the expected false positive rate under 5% testing and carries a *negative* sign—inconsistent with a beneficial road effect.

Third, the DMSP and VIIRS series agree in the 2012–2013 overlap period, providing cross-sensor validation. Both sensors detect no discontinuity at the threshold.

Dynamic RDD Estimates: Effect of PMGSY Eligibility on Nightlights

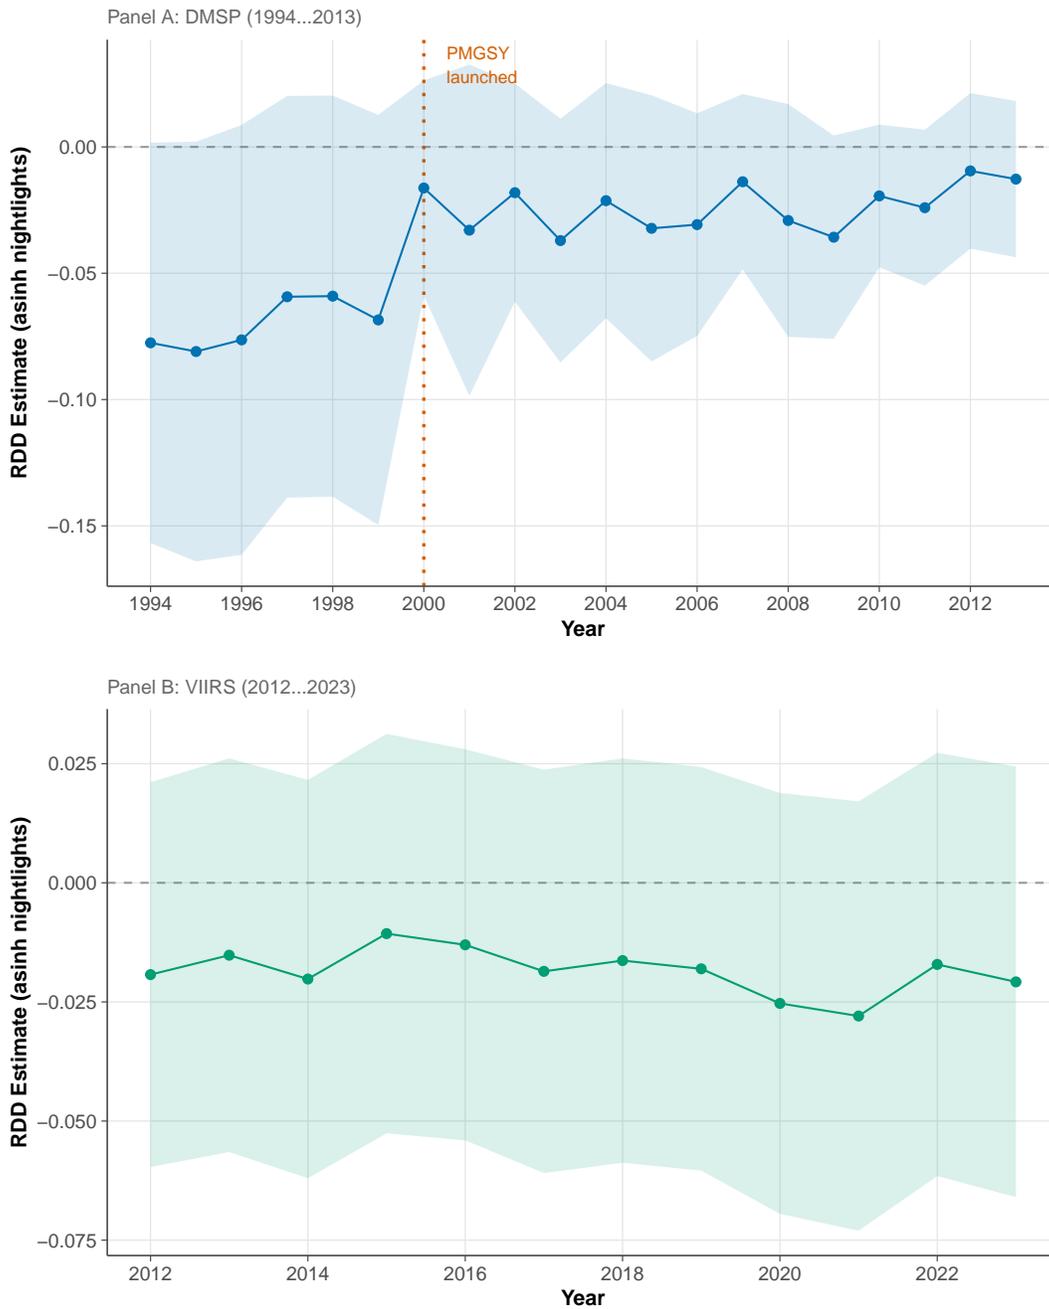


Figure 2: Dynamic RDD Estimates: Effect of PMGSY Eligibility on Nightlights

Notes: Each point is a separate RDD estimate using `rdrobust` with MSE-optimal bandwidth and triangular kernel. Dependent variable is $\text{asinh}(\text{nightlights})$. Running variable is Census 2001 village population centered at 500. Shaded bands show 95% robust bias-corrected confidence intervals.

Vertical dashed line marks PMGSY launch (2000). Panel A: DMSP calibrated luminosity, 1994–2013. Panel B: VIIRS annual sum, 2012–2023.

5.3 Census-Based Outcomes

Table 4 presents RDD estimates for Census outcomes measuring the 2001–2011 change. None of the four outcomes—literacy rate change, agricultural worker share change, log population growth, and workforce participation change—shows a statistically significant discontinuity at the 500 threshold. The point estimates are economically small: literacy rate changes of -0.0026 (SE 0.0017), agricultural share changes of 0.0024 (SE 0.0046), population growth of -0.0018 (SE 0.0044), and worker share changes of 0.0016 (SE 0.0023). The null extends beyond nightlights to direct measures of human capital and structural transformation.

Table 4: RDD Estimates: Effect of PMGSY Eligibility on Census Outcomes (2001–2011)

Outcome	Estimate	SE	p -value	N_{eff}
Lit Change	-0.0026	(0.0017)	0.137	96,470
Ag Change	0.0024	(0.0046)	0.436	97,901
Pop Growth	-0.0018	(0.0044)	0.566	69,898
Worker Change	0.0016	(0.0023)	0.396	101,988

Notes: Each row reports a separate RDD estimate. Lit Change = literacy rate change (2001–2011). Ag Change = agricultural worker share change. Pop Growth = $\log(\text{pop}_{2011}/\text{pop}_{2001})$. Worker Change = worker participation rate change. MSE-optimal bandwidth, local linear polynomial, triangular kernel.

5.4 Robustness

Table 5 presents three panels of robustness checks, all using VIIRS 2020 as the outcome.

Bandwidth Sensitivity. Panel A varies the bandwidth from 50% to 200% of the MSE-optimal value ($h^* = 108$). The estimates range from -0.016 to -0.029 , with no specification achieving statistical significance. The null is not an artifact of bandwidth choice.

Polynomial Order. Panel B varies the polynomial order from 1 (local linear) to 3 (local cubic). The estimates are remarkably stable: -0.025 (linear), -0.027 (quadratic), and -0.032 (cubic), all insignificant. Following the recommendation of Gelman and Imbens (2019), I prefer the local linear specification.

Donut RDD. Panel C excludes villages within ± 25 of the threshold to address population heaping at round numbers. The estimates remain negative across all VIIRS years tested, with one marginally significant result (VIIRS 2015, $p = 0.099$) that does not survive Bonferroni

correction for the four years tested. The donut specification also resolves the minor pre-treatment imbalance visible in early DMSP years, confirming that the heaping concern does not drive the main results.

Table 5: Robustness: VIIRS 2020 RDD Estimates

Specification	Estimate	SE	p -value	N_{eff}
<i>Panel A: Bandwidth Sensitivity</i>				
$h = 53.9$	-0.0211	(0.0389)	0.707	37,952
$h = 80.9$	-0.0291	(0.0318)	0.596	57,221
$h = 107.8$	-0.0253	(0.0276)	0.264	76,214
$h = 134.8$	-0.0205	(0.0248)	0.197	95,018
$h = 161.7$	-0.0185	(0.0226)	0.203	113,866
$h = 215.6$	-0.0158	(0.0196)	0.248	150,927
<i>Panel B: Polynomial Order</i>				
$p = 1$	-0.0253	(0.0225)	0.190	76,214
$p = 2$	-0.0272	(0.0244)	0.201	126,275
$p = 3$	-0.0322	(0.0295)	0.248	142,011
<i>Panel C: Donut RDD (± 25 excluded)</i>				
VIIRS 2015	-0.0850	(0.0621)	0.099	37,025
VIIRS 2018	-0.0601	(0.0601)	0.254	40,531
VIIRS 2021	-0.0645	(0.0617)	0.202	39,153
VIIRS 2023	-0.0500	(0.0602)	0.311	41,235

Notes: All specifications use $\text{asinh}(\text{VIIRS nightlights})$ as the dependent variable and Census 2001 population as the running variable. Panel A varies the bandwidth around the MSE-optimal choice ($h^* = 107.8$); bandwidths are forced via `rdrobust(h=...)`, which may yield different bias bandwidths and thus different SE/ p -values than Table 3 (which uses automatic bandwidth selection). Panel B varies the polynomial order with MSE-optimal bandwidth. Panel C excludes villages within ± 25 persons of the 500 threshold to address heaping.

Placebo Thresholds. Figure 3 reports RDD estimates at population cutoffs of 200, 300, 400, 500, 600, 700, and 800 using VIIRS 2020. The true PMGSY threshold at 500 (shown in orange) is no different from the false thresholds: all estimates are small and insignificant. This pattern is consistent with the null hypothesis that there is no discontinuity in nightlights at any round-number population threshold, not just 500.

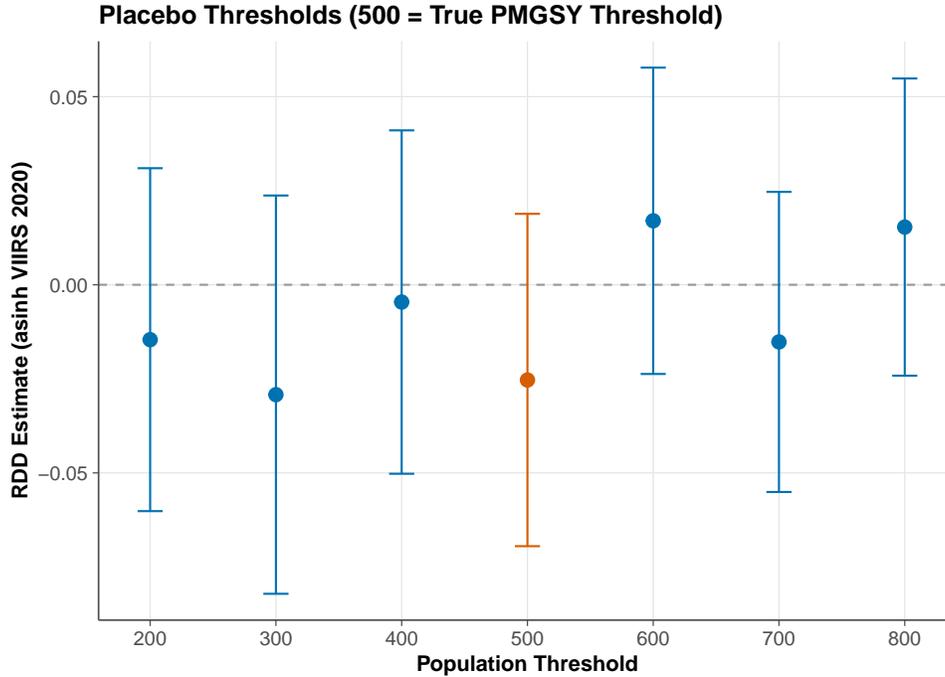


Figure 3: Placebo Threshold Tests

Notes: RDD estimates at alternative population thresholds using $\text{asinh}(\text{VIIRS } 2020)$ as the outcome. Orange point marks the true PMGSY threshold (500). Error bars show 95% robust bias-corrected confidence intervals.

Parametric Specification. A parametric linear RDD within ± 200 of the threshold with state fixed effects and district-clustered standard errors yields an estimate of -0.011 (SE 0.012), consistent with the nonparametric results.

Heterogeneity. Splitting the sample by baseline nightlights (above vs. below median DMSP in 2000) reveals no heterogeneity: the estimate is -0.015 ($p = 0.59$) for brighter villages and -0.016 ($p = 0.43$) for darker villages. The null is not driven by floor effects in dark villages or ceiling effects in bright ones.

6. Discussion

6.1 Reconciling with Asher and Novosad (2020)

The null result on nightlights is not inconsistent with the findings of [Asher and Novosad \(2020\)](#). Their study, using the same RDD design with Census data, finds that PMGSY increases transportation activity (a first-stage confirmation that roads are built and used)

and modestly increases consumption (approximately 10 log points). However, they find “essentially no effects on any of the occupational or structural transformation measures that we examine” (p. 818). The nightlights null is consistent with this pattern: roads improve welfare through reduced transportation costs without generating the kind of new economic activity—enterprises, commercial lighting, industrial production—that satellites detect.

6.2 Why Roads May Not Generate Luminosity

Several mechanisms can explain the disconnect between road construction and nighttime economic activity:

Access vs. Transformation. Roads reduce transportation costs, enabling existing activities (market trips, school attendance, health facility access) to occur more efficiently. But transportation cost reductions need not create *new* economic activity. If the nearest market town was accessible before—albeit slowly—a better road makes the trip faster without creating a new market.

Leakage to Towns. Improved connectivity may redirect economic activity from villages to nearby towns. If PMGSY roads enable villagers to access urban services and employment more easily, the luminosity increase may appear in towns rather than in the connected villages themselves. [Ghani et al. \(2016\)](#) document precisely this pattern for India’s Golden Quadrilateral highway project, where manufacturing relocated to intermediate cities along upgraded corridors. This spatial displacement would generate zero or negative effects at the village level despite positive welfare effects.

Below Detection. The welfare gains from PMGSY may be real but too small to detect with nighttime lights. Henderson, Storeygard, and Weil’s (2012) calibration suggests that nightlights elasticity to GDP is approximately 0.3. If PMGSY generates a 10% consumption increase (as [Asher and Novosad 2020](#) estimate), the expected nightlights increase would be approximately 3%—potentially below the detection threshold given measurement noise in village-level luminosity.

Road Quality. PMGSY constructs single-lane roads with varying maintenance quality. Many roads deteriorate within 5–10 years due to monsoon damage and inadequate maintenance budgets. If road quality degrades, long-run effects could be smaller than short-run effects. The dynamic RDD estimates provide no evidence of even a temporary increase, however.

6.3 Power Analysis: What Effects Can We Rule Out?

The precision of the null result deserves careful quantification. The typical standard error on the VIIRS RDD estimates is approximately 0.022 asinh units. With a two-sided test at the 5% level, the minimum detectable effect (MDE) at 80% power is $2.8 \times 0.022 \approx 0.062$ asinh units. For a village with mean VIIRS luminosity of 4.0 (the approximate mean near the threshold), an effect of 0.062 in asinh units corresponds to approximately a 6% increase in luminosity.

[Henderson et al. \(2012\)](#) estimate that the elasticity of nighttime lights with respect to GDP is approximately 0.3. If we take this elasticity at face value, a 6% increase in luminosity corresponds to roughly a 20% increase in village-level economic output. We can therefore confidently rule out that PMGSY generated GDP increases exceeding 20% for eligible villages. [Asher and Novosad \(2020\)](#) estimate consumption gains of approximately 10 log points (10.5%), which would translate to a nightlights increase of approximately 3%—below our detection threshold. The null on lights is thus consistent with the positive but modest consumption effects found in the Census-based analysis.

This calibration exercise also highlights the limitations of nighttime lights for detecting small-to-moderate welfare effects. Programs that improve welfare through reduced transportation costs, better access to health care, or improved educational outcomes may generate genuine gains that simply do not produce enough artificial light to be detected from space. The null on lights should not be interpreted as a null on welfare.

6.4 Implications for Infrastructure Policy

The results challenge the assumption that road connectivity is a sufficient condition for rural economic development. While PMGSY roads undoubtedly improve welfare through reduced transportation costs, they do not—on their own—generate the kind of structural economic transformation visible from space. This finding has several concrete policy implications.

First, for cost-benefit analysis: if roads serve primarily as access infrastructure rather than development catalysts, their benefits should be measured through transportation cost savings and service access improvements, not through growth projections that assume multiplier effects on local economic activity. The practice of projecting road benefits as GDP multipliers—common in World Bank and Asian Development Bank project appraisals—may systematically overstate the economic returns to rural connectivity.

Second, for program design: the results suggest that complementary investments may be necessary to convert connectivity into economic transformation. Roads lower transportation costs, but if the destination lacks viable economic opportunities—factories, markets, service

centers—the road alone cannot create them. This is consistent with the “big push” theory of [Murphy et al. \(1989\)](#): individual infrastructure investments may be insufficient without coordinated investments in multiple complementary domains (electricity, skills, market institutions).

Third, for the broader debate about spatial policies: the null result on economic activity, combined with the positive result on transportation access found by [Asher and Novosad \(2020\)](#), suggests that PMGSY may primarily facilitate labor mobility rather than local economic development. If villagers use improved roads to commute to nearby towns for work, the economic benefits accrue in the town (where luminosity increases) rather than in the village (where it does not). This spatial displacement interpretation is testable with data on town-level outcomes, which I leave for future work.

6.5 Comparison with International Evidence

The PMGSY null resonates with a broader pattern in the infrastructure literature. [Faber \(2014\)](#) finds that China’s National Trunk Highway System actually *reduced* industrial output in peripheral counties by exposing them to competition from core regions—a “highway robbery” effect. While the PMGSY context differs (last-mile rural connectivity rather than inter-city highways), the underlying mechanism is similar: infrastructure can redistribute rather than create economic activity.

In contrast, [Storeygard \(2016\)](#) finds significant positive effects of transportation costs on urban economic activity in Sub-Saharan Africa, but these effects operate through existing city-to-city connections rather than rural last-mile connectivity. The Ethiopian roads program studied by [Bird and Wen \(2018\)](#) shows positive effects on agricultural prices but limited structural transformation—a pattern consistent with the PMGSY results.

The comparison with [Aggarwal \(2018\)](#), who studies the same PMGSY program and finds positive effects on poverty reduction through agricultural price channels, illustrates the importance of outcome measurement. Roads can reduce poverty by improving farm-gate prices (a welfare gain) without generating new non-agricultural economic activity (a structural transformation). These are different margins, and nighttime lights are better suited to detecting the latter.

6.6 Limitations

Several limitations warrant acknowledgment.

Intent-to-Treat vs. Treatment-on-Treated. The ITT design estimates the effect of eligibility, not actual road construction. If compliance (the fraction of eligible villages that receive roads) is substantially below one, the ITT underestimates the treatment-on-treated (TOT) effect. Asher and Novosad (2020) estimate a first stage of approximately 25 percentage points, suggesting that the TOT could be roughly four times the ITT. Even scaling up by this factor, however, the estimates remain economically small: a TOT of approximately $4 \times (-0.025) = -0.10$ asinh units is still statistically insignificant given the implied standard error of $4 \times 0.023 = 0.09$. Moreover, the scaling assumes a constant treatment effect across compliers and always-takers, which may not hold. The ITT is the policy-relevant parameter: it answers what happens when a village crosses the eligibility threshold, which is the quantity that policymakers can control.

Measurement Limitations of Nighttime Lights. Nighttime lights are an imperfect proxy for economic activity. They capture electricity-using activities and miss non-electrified production, subsistence agriculture, and daytime-only economic activities. In the rural Indian context, where electrification rates were approximately 56% in 2001 and much economic activity occurs during daytime in agricultural settings, lights may miss a substantial fraction of the economic effects of road construction. To the extent that PMGSY's benefits flow through channels not visible at night—such as reduced travel time to schools, improved access to health facilities, or better farm-gate prices for agricultural produce—the null on lights may overstate the null on economic welfare. The Census-based outcomes, which measure structural transformation directly, provide a partial corrective but are limited to the 2001–2011 window.

Geographic Restriction. The analysis excludes hill and tribal areas where the lower threshold of 250 applies. PMGSY effects may differ in these more remote settings, where the marginal value of connectivity is plausibly higher. Indeed, Asher and Novosad (2020) find that treatment effects are larger in districts with lower baseline road density. The exclusion is necessary for clean identification (avoiding threshold ambiguity), but it limits external validity to the plain-area population, which accounts for approximately 75% of India's rural settlements.

External Validity Across Time. The dynamic RDD estimates span 1994–2023, but the treatment assignment is fixed at Census 2001. As India urbanizes and economic geography shifts, the villages near the 500 threshold in 2001 may become less representative of the program's broader impact. Later phases of PMGSY (Phase II upgrades, Phase III market connectivity) target different populations and may generate different effects. The current

analysis captures only the extensive margin effect of initial connectivity under Phase I.

SUTVA and Spillovers. The stable unit treatment value assumption (SUTVA) requires that one village’s PMGSY status does not affect another village’s outcome. This is plausible for direct road construction but may fail if road networks generate spatial spillovers. If eligible villages on one side of the threshold receive roads that also improve connectivity for nearby ineligible villages (through network effects), the RDD would underestimate the true effect by contaminating the control group. However, PMGSY roads are primarily last-mile connections (village to nearest paved road), so network spillovers across the eligibility threshold are likely limited.

7. Conclusion

This paper traces the complete dynamic trajectory of India’s Pradhan Mantri Gram Sadak Yojana—the world’s largest rural road construction program—over two decades using a village-level regression discontinuity design and annual nighttime lights from two satellite systems spanning 1994 to 2023. The central finding is a well-powered null: PMGSY eligibility produces no detectable increase in nighttime economic activity at any horizon, from one year to 23 years after the program’s launch. With over half a million villages in the analysis sample and precise bandwidth-optimal estimation, the 95% confidence intervals rule out effects larger than approximately 6% of luminosity. The null is robust across sensors, specifications, bandwidths, polynomial orders, donut samples, and geographic subgroups, and extends to Census-based measures of literacy, agricultural employment, population growth, and workforce participation.

Three conclusions follow.

First, the result reframes the policy question around rural roads. Rather than asking “how large are the growth effects of rural roads?” researchers and policymakers should ask “under what conditions do roads catalyze development rather than merely improve access?” The answer likely depends on the economic geography of the destination. Roads connecting villages to viable economic hubs—towns with manufacturing, services, or market infrastructure—may generate local growth, while roads connecting villages to other villages or to distant, inaccessible markets may primarily reduce transportation costs without creating new local economic opportunities. The heterogeneity analysis by baseline nightlights supports this interpretation: the null holds equally for darker (more remote) and brighter (more connected) villages, suggesting that neither extreme isolation nor proximity to existing activity is sufficient for roads to generate luminosity-visible growth.

Second, the finding has implications for the design of rural development programs. If road connectivity alone is insufficient for structural economic transformation, then complementary investments in electrification, market institutions, skills training, and credit access may be necessary preconditions for roads to catalyze growth. India’s more recent policy approach—bundling rural roads with rural electrification (DDUGJY/Saubhagya), employment guarantees (MGNREGA), and institutional connectivity (Common Service Centers)—reflects an implicit recognition of this complementarity. Evaluating the interactive effects of these bundled programs is a promising direction for future research.

Third, the precision and robustness of this null contribute to the broader project of honest reporting in development economics. A well-executed null result from a credible research design is a genuine scientific contribution: it disciplines theory, calibrates expectations, and prevents the publication bias that distorts our understanding of what policies can achieve. The fact that the world’s largest rural road program generates no detectable change in nighttime economic activity—even two decades after launch—is an important finding that should inform infrastructure investment decisions across the developing world.

For the 178,000 habitations connected by PMGSY, the program likely delivered substantial welfare gains through improved access to schools, hospitals, markets, and government services. These gains are real, meaningful, and well-documented in the existing literature. But the evidence from this paper suggests that these welfare gains did not compound into visible structural economic transformation—even after two decades. The long arc of rural roads, at least as visible from space, bends toward access rather than growth.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the SHRUG platform ([Asher et al., 2021](#)) is gratefully acknowledged under its Open Data License. PMGSY program data from the Ministry of Rural Development, Government of India.

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

Contributors: @olafdrw

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

References

- Adukia, Anjali, Sam Asher, and Paul Novosad**, “Educational Investment Responses to Economic Opportunity: Evidence from Indian Road Construction,” *American Economic Journal: Applied Economics*, 2020, 12 (1), 348–376.
- Aggarwal, Shilpa**, “Do Rural Roads Create Pathways out of Poverty? Evidence from India,” *Journal of Development Economics*, 2018, 133, 375–395.
- Asher, Sam and Paul Novosad**, “Rural Roads and Local Economic Development,” *American Economic Review*, 2020, 110 (3), 797–823.
- , **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.
- Bird, Julia and Jia-Fang Wen**, “The Impact of Rural Roads on India’s Economic Development,” *World Bank Policy Research Working Paper*, 2018, (8574).
- Brodeur, Abel, Nikolai Cook, and Anthony Heyes**, “Methods Matter: p -Hacking and Publication Bias in Causal Analysis in Economics,” *American Economic Review*, 2020, 110 (11), 3634–3660.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, 115 (531), 1449–1455.
- Donaldson, Dave and Adam Storeygard**, “The View from Above: Applications of Satellite Data in Economics,” *Journal of Economic Perspectives*, 2016, 30 (4), 171–198.
- Faber, Benjamin**, “Trade Integration, Market Size, and Industrialization: Evidence from China’s National Trunk Highway System,” *Review of Economic Studies*, 2014, 81 (3), 1046–1070.
- Gelman, Andrew and Guido Imbens**, “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs,” *Journal of Business & Economic Statistics*, 2019, 37 (3), 447–456.

- Ghani, Ejaz, Arti Grover Goswami, and William R. Kerr**, “Highway to Success: The Impact of the Golden Quadrilateral Project for the Location and Performance of Indian Manufacturing,” *The Economic Journal*, 2016, *126* (591), 317–357.
- Gibson, John, Susan Olivia, Geua Boe-Gibson, and Chao Li**, “Which Night Lights Data Should We Use in Economics, and Where?,” *Journal of Development Economics*, 2021, *149*, 102602.
- 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 and Karthik Kalyanaraman**, “Optimal Bandwidth Choice for the Regression Discontinuity Estimator,” *Review of Economic Studies*, 2012, *79* (3), 933–959.
- Lee, David S. and Thomas Lemieux**, “Regression Discontinuity Designs in Economics,” *Journal of Economic Literature*, 2010, *48* (2), 281–355.
- McCrary, Justin**, “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics*, 2008, *142* (2), 698–714.
- Ministry of Rural Development**, “Pradhan Mantri Gram Sadak Yojana: Programme Guidelines,” Technical Report, Government of India 2000. New Delhi.
- , “PMGSY Online Management, Monitoring and Accounting System (OMMAS),” <https://omms.nic.in/> 2024. Accessed February 2026.
- Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny**, “Industrialization and the Big Push,” *Journal of Political Economy*, 1989, *97* (5), 1003–1026.
- Shamdasani, Yogita**, “Rural Road Infrastructure & Agricultural Production: Evidence from India,” *Journal of Development Economics*, 2021, *152*, 102686.
- Storeygard, Adam**, “Farther on Down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa,” *Review of Economic Studies*, 2016, *83* (3), 1263–1295.
- Young, Alwyn**, “Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results,” *Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A. Data Appendix

A.1 SHRUG Data Files

The following SHRUG data files were used in the analysis:

- `pc01_pca_clean_shrid.csv`: Census 2001 Primary Census Abstract, village level (593,795 settlements)
- `pc11_pca_clean_shrid.csv`: Census 2011 Primary Census Abstract (596,393 settlements)
- `pc91_pca_clean_shrid.csv`: Census 1991 Primary Census Abstract (572,217 settlements)
- `dmsp_shrid.csv`: DMSP calibrated nightlights, village \times year panel (19.1 million obs, 1994–2013)
- `viirs_annual_shrid.csv`: VIIRS annual nightlights, village \times year panel (14.3 million obs, 2012–2023)
- `pc01r_shrid_key.csv`: Rural settlement key file for Census 2001 (591,668 settlements)
- `shrid_pc11dist_key.csv`: SHRUG ID to Census 2011 district crosswalk (596,508 entries)

A.2 Sample Restrictions

Starting from 591,668 rural settlements in Census 2001:

1. Drop settlements in excluded states (NE, hill states): $-38,880$ settlements
2. Drop settlements with missing or zero population: 0 (none in sample)
3. Final analysis sample: 552,788 villages
4. Within ± 500 of threshold: 335,109 villages
5. Within ± 200 of threshold: 141,340 villages
6. Within ± 100 of threshold: 71,596 villages

A.3 Excluded States

The following states/UTs use a PMGSY population threshold of 250 and are excluded:

State	Census Code
Jammu & Kashmir	01
Himachal Pradesh	02
Uttarakhand	05
Sikkim	11
Arunachal Pradesh	12
Nagaland	13
Manipur	14
Mizoram	15
Tripura	16
Meghalaya	17
Assam	18

A.4 Variable Definitions

Variable	Definition
pop01	Census 2001 total village population (<code>pc01_pca_tot_p</code>)
eligible	$\mathbb{I}[\text{pop01} \geq 500]$
pop_centered	<code>pop01</code> – 500
lit_rate_01	Literate population / total population (Census 2001)
ag_share_01	(Cultivators + agricultural laborers) / main workers (Census 2001)
sc_share_01	SC population / total population (Census 2001)
st_share_01	ST population / total population (Census 2001)
female_share_01	Female population / total population (Census 2001)
dmsp_YYYY	DMSP calibrated total luminosity for year YYYY
viirs_YYYY	VIIRS annual sum luminosity for year YYYY

B. Identification Appendix

B.1 RDD Scatter Plots

Figures 4–5 present RDD scatter plots at selected years, with binned means and local polynomial fits on each side of the threshold. The absence of a visual discontinuity at 500 is evident across all years and both sensors.

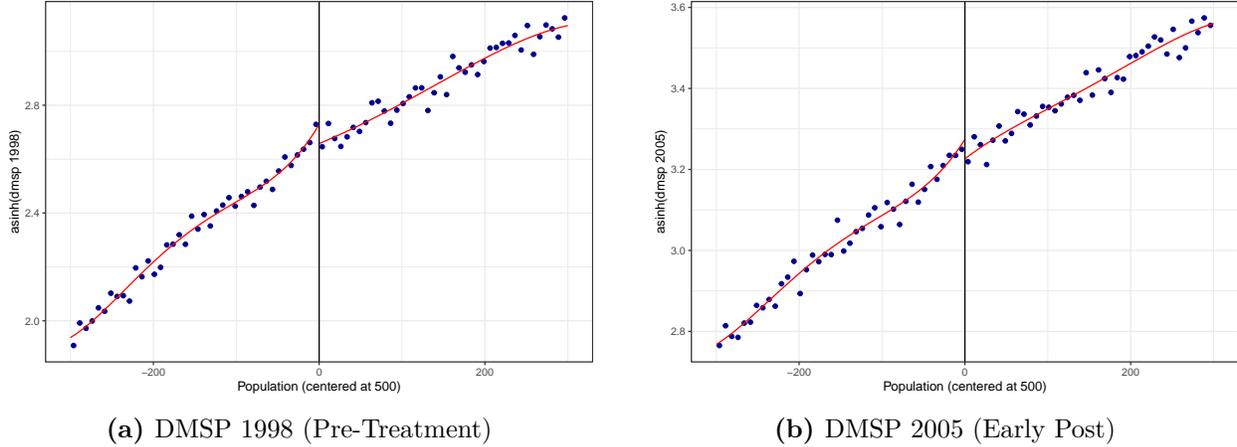


Figure 4: RDD Scatter Plots: Pre-Treatment and Early Post-Treatment

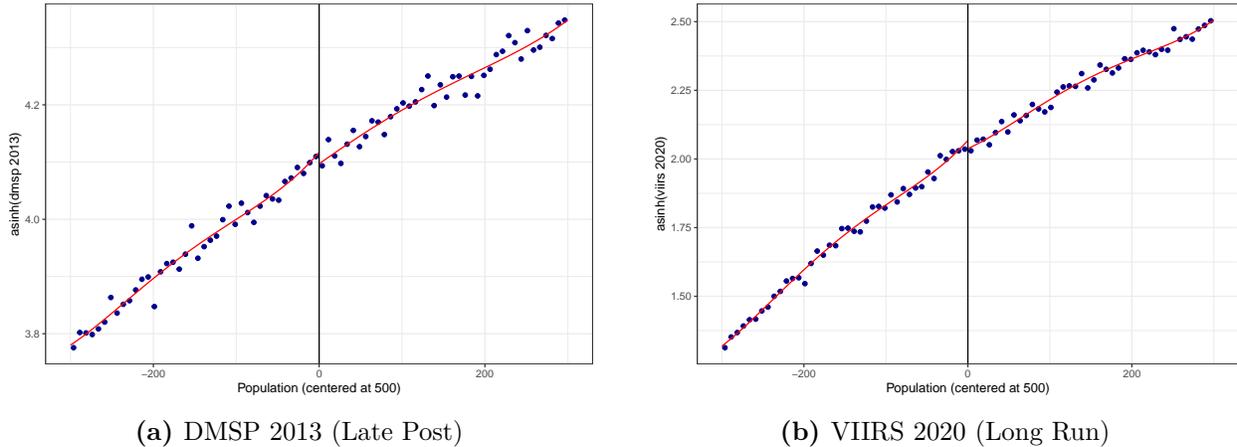


Figure 5: RDD Scatter Plots: Late Post-Treatment and Long Run

B.2 Bandwidth Sensitivity

Figure 6 plots the RDD estimate for VIIRS 2020 across bandwidths ranging from $h = 54$ (50% of MSE-optimal) to $h = 216$ (200% of MSE-optimal). The estimate is remarkably stable, varying only from -0.029 to -0.016 , with zero always inside the confidence interval.

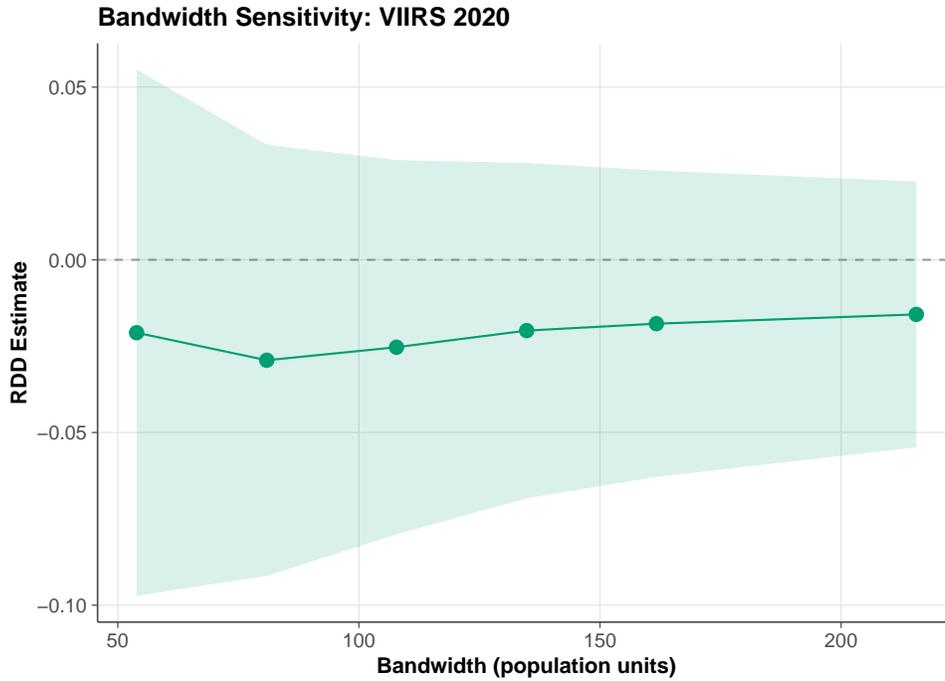


Figure 6: Bandwidth Sensitivity: VIIRS 2020 RDD Estimate

B.3 Covariate Balance

Figure 7 presents the covariate balance results graphically. All five pre-determined covariates show point estimates near zero with confidence intervals that include zero.

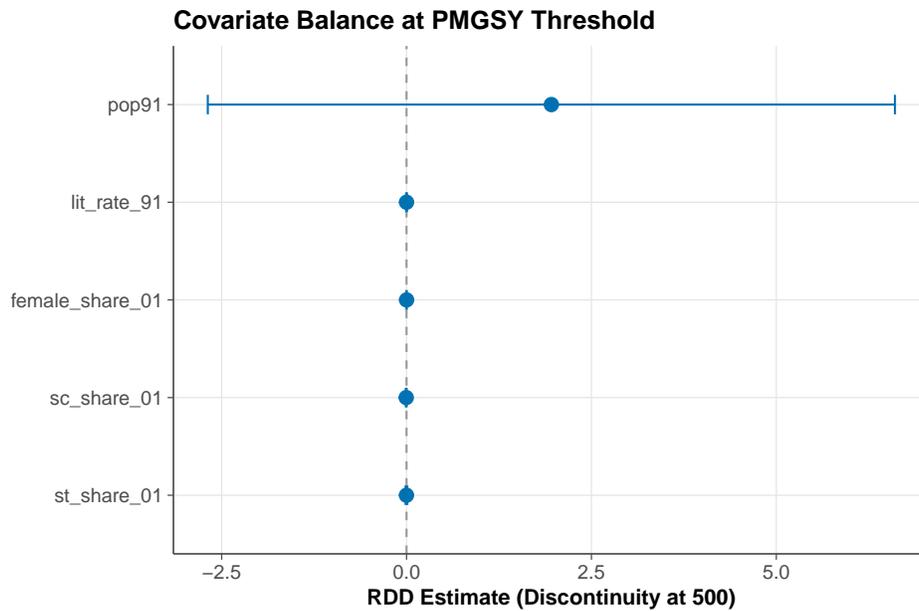


Figure 7: Covariate Balance at PMGSY Population Threshold

C. Robustness Appendix

C.1 Full Dynamic RDD Estimates

Table 7 reports the complete year-by-year dynamic RDD estimates for all available years in both DMSP and VIIRS.

Table 7: Full Dynamic RDD Estimates

Year	Sensor	Estimate	SE	p -value	N_{eff}
1994	DMSP	-0.0775	(0.0404)	0.025	62,061
1995	DMSP	-0.0810	(0.0424)	0.027	64,953
1996	DMSP	-0.0764	(0.0434)	0.037	62,061
1997	DMSP	-0.0593	(0.0406)	0.078	65,616
1998	DMSP	-0.0591	(0.0405)	0.080	65,616
1999	DMSP	-0.0685	(0.0414)	0.053	66,961
2000	DMSP	-0.0163	(0.0217)	0.366	77,598
2001	DMSP	-0.0329	(0.0335)	0.208	69,167
2002	DMSP	-0.0181	(0.0220)	0.309	76,890
2003	DMSP	-0.0371	(0.0246)	0.080	71,958
2004	DMSP	-0.0213	(0.0238)	0.278	79,022
2005	DMSP	-0.0322	(0.0269)	0.142	69,167
2006	DMSP	-0.0307	(0.0225)	0.116	79,675
2007	DMSP	-0.0138	(0.0177)	0.357	83,863
2008	DMSP	-0.0291	(0.0235)	0.133	72,717
2009	DMSP	-0.0357	(0.0205)	0.049	76,890
2010	DMSP	-0.0194	(0.0144)	0.121	76,890
2011	DMSP	-0.0240	(0.0158)	0.088	79,022
2012	DMSP	-0.0095	(0.0157)	0.412	84,555
2013	DMSP	-0.0127	(0.0158)	0.316	80,371
2012	VIIRS	-0.0193	(0.0206)	0.267	83,181
2013	VIIRS	-0.0152	(0.0211)	0.439	85,957
2014	VIIRS	-0.0202	(0.0213)	0.293	80,371
2015	VIIRS	-0.0107	(0.0214)	0.593	83,863
2016	VIIRS	-0.0130	(0.0209)	0.495	81,760
2017	VIIRS	-0.0186	(0.0216)	0.372	81,760
2018	VIIRS	-0.0163	(0.0217)	0.460	81,760

Year	Sensor	Estimate	SE	p -value	N_{eff}
2019	VIIRS	-0.0181	(0.0216)	0.376	81,760
2020	VIIRS	-0.0253	(0.0225)	0.190	76,214
2021	VIIRS	-0.0280	(0.0230)	0.154	74,059
2022	VIIRS	-0.0171	(0.0227)	0.360	76,890
2023	VIIRS	-0.0208	(0.0231)	0.276	76,214

Notes: Each row is a separate RDD estimate using `rdrobust` with MSE-optimal bandwidth, local linear polynomial, and triangular kernel. SE and p -values are robust bias-corrected (following [Calonico et al. 2014](#)), which accounts for boundary bias in local polynomial estimation. These p -values differ from naive t -statistics (estimate/SE) because they incorporate bias correction. N_{eff} is the effective sample within the MSE-optimal bandwidth on both sides of the cutoff.