

# Connecting the Most Remote: Road Eligibility and Development in India's Tribal Periphery

APEP Autonomous Research\*      @olafdrw

February 20, 2026

## Abstract

India's rural road program (PMGSY) uses a lower population threshold of 250 in tribal and hill areas versus 500 in plains, creating a regression discontinuity in eligibility for road construction. Exploiting this threshold across 41,371 villages in 11 Special Category States using SHRUG Census data, I find that eligibility increases female literacy by 1.9 percentage points ( $p = 0.032$ ) and late-period nightlight luminosity by 0.34 log points ( $p = 0.004$ ). These effects are absent at placebo thresholds and at the 500 threshold in non-designated areas. Results are robust to bandwidth variation, polynomial order, and donut-hole specifications, revealing that lowering eligibility thresholds for remote communities generates persistent human capital and economic gains.

**JEL Codes:** O18, H54, R42, J16

**Keywords:** rural roads, regression discontinuity, PMGSY, tribal areas, female literacy, nightlights, India

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch This paper was generated autonomously.

# 1. Introduction

For millions of people in India’s tribal periphery, the arrival of the monsoon means the end of the road. Unpaved tracks turn to mud, cutting off villages from schools, clinics, and markets for months at a time. In 2001, over 300,000 such habitations—disproportionately tribal and scheduled caste communities in remote hill areas—had no all-weather road access. Whether connecting these communities to the road network meaningfully improves their development trajectories is among the most important questions in infrastructure economics.

Transportation infrastructure is a central input to economic development. Classical trade theory emphasizes that lower trade costs expand market access, enabling specialization and increasing welfare (Redding and Rossi-Hansberg, 2017). In developing countries, where internal trade costs can exceed those of international trade, rural roads represent perhaps the most consequential form of infrastructure investment. They connect subsistence farmers to output markets, enable access to schools and health facilities, facilitate labor mobility, and transmit information about economic opportunities in the wider economy. Yet empirical evidence on the returns to rural roads—particularly in the most remote settings—remains surprisingly thin.

The Pradhan Mantri Gram Sadak Yojana (PMGSY), launched in December 2000, is the world’s largest rural road construction program (Ministry of Rural Development, 2000). It established a simple population-based eligibility rule: habitations with 500 or more residents (in Census 2001) were prioritized for all-weather road connectivity. Crucially, the program applied a *lower* threshold of 250 in “designated areas”—Special Category States in India’s northeast and western hills, plus desert and tribal areas under Schedule V—recognizing that remoteness and low population density should not be barriers to basic connectivity.

This paper exploits the 250-person eligibility threshold as a regression discontinuity to estimate the causal effects of road eligibility on development outcomes in India’s tribal and hill periphery. While the seminal study of PMGSY by Asher and Novosad (2020) examines the 500 threshold in plains areas, the 250 threshold in designated areas has received remarkably little attention despite targeting the most marginalized communities. This gap is striking: the lower threshold was specifically designed to reach populations that conventional infrastructure programs systematically overlook, and the returns to connectivity in these areas may be fundamentally different from those in the relatively better-connected plains.

The setting offers several advantages for causal identification. First, the 250 threshold is sharp and rule-based, determined by Census 2001 population counts that were fixed before the program began. Unlike many administrative thresholds that are subject to bureaucratic discretion, the PMGSY eligibility rule was applied mechanically to published Census data.

Second, I find no evidence of manipulation in the running variable: the [Cattaneo et al. \(2020a\)](#) density test yields  $p = 0.546$ , consistent with smooth population distributions around the cutoff. Third, pre-treatment covariates—literacy rates, caste composition, worker shares, and nighttime luminosity—are balanced across the threshold, with all seven balance tests producing  $p$ -values above 0.20. Fourth, placebo tests at every 50-person increment from 150 to 400 produce uniformly null results, isolating the 250 threshold as the only source of discontinuity.

Using village-level data from the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG; [Asher et al., 2021](#)), I link Census 2001 baseline characteristics to Census 2011 outcomes and satellite nightlight luminosity from 1994 to 2023. The primary sample comprises 41,371 villages in 11 Special Category States—Jammu & Kashmir, Himachal Pradesh, Uttarakhand, and the eight northeastern states—within a population window of 50 to 500 around the 250 cutoff. I estimate local linear regressions using the bias-corrected robust inference framework of [Calonico et al. \(2014\)](#) with MSE-optimal bandwidths selected by the `rdrobust` software package.

The results reveal economically meaningful effects concentrated in human capital and long-run economic activity. Road eligibility increases female literacy rates by 1.9 percentage points ( $p = 0.032$ ), with a marginal effect on overall literacy of 1.3 percentage points ( $p = 0.105$ ). Late-period nightlight luminosity measured by VIIRS satellites (2015–2023) increases by 0.34 log points ( $p = 0.004$ ), indicating substantial long-run economic gains that persist more than a decade after program initiation. The magnitude of this nightlight effect is striking: it corresponds to roughly a 41% increase in luminosity, suggesting meaningful economic transformation in these previously dark communities.

These findings survive extensive robustness testing. The estimates are stable across bandwidths from half to double the MSE-optimal choice. A local quadratic specification strengthens the results (female literacy:  $p = 0.028$ ; nightlights:  $p = 0.003$ ). Donut-hole RDD excluding villages within 5 persons of the threshold preserves the sign and approximate magnitude of all estimates. An extended sample incorporating high-tribal-share villages in Schedule V states (81,585 villages) yields qualitatively similar results, with VIIRS nightlights significant at the 10% level ( $p = 0.098$ ).

Importantly, I find no corresponding effects at the 500 threshold in non-designated areas across 117,889 plains villages. None of the six primary outcomes shows a significant positive discontinuity at 500. This null comparison strengthens the interpretation that the 250 threshold effects reflect the distinctive returns to road connectivity in remote, underserved communities rather than a generic artifact of population-based eligibility rules. The contrast between positive effects at 250 in tribal areas and null effects at 500 in plains is consistent

with diminishing marginal returns to transportation infrastructure: roads matter most where connectivity is worst.

This paper contributes to several literatures. First, it extends the economics of transportation infrastructure in developing countries (Donaldson, 2018; Banerjee et al., 2020; Faber, 2014; Storeygard, 2016). While Asher and Novosad (2020) establish that PMGSY roads increase economic activity and reduce transportation costs at the 500 threshold, and Adukia et al. (2020) document educational investment responses along newly constructed roads, the 250 threshold in tribal areas remains unexplored. The lower threshold targets communities that are more remote, more disadvantaged, and more likely to be first-time connected—precisely the population where marginal returns to infrastructure may be highest. Aggarwal (2018) finds that PMGSY roads create pathways out of poverty through occupational diversification, but her analysis focuses on non-tribal areas where baseline connectivity is substantially higher.

Second, the paper contributes to understanding gender-differentiated returns to infrastructure (Dinkelman, 2011; Muralidharan and Prakash, 2017). Roads may affect female human capital through multiple channels: reducing the effective distance to schools, enabling female teachers to commute from towns, creating non-agricultural employment opportunities that raise the perceived returns to female education (Munshi and Rosenzweig, 2006), and expanding access to health and information services (Jensen, 2007). The finding that female literacy gains substantially exceed overall literacy gains is consistent with roads relaxing binding constraints that disproportionately affect girls’ schooling in traditional societies. Beaman et al. (2012) show that female leadership in Indian villages raises girls’ educational aspirations, and Iyer et al. (2012) document how women’s political representation reduces crime; road connectivity may similarly expand the set of possibilities that girls and their families perceive as attainable.

Third, this paper provides new evidence on the long-run effects of infrastructure investments. The VIIRS nightlight results—capturing economic activity 14–22 years after the program launched—suggest that road connectivity triggers sustained development trajectories in remote areas. This is consistent with Donaldson (2018), who shows that railroad access in colonial India had persistent effects on trade and income, and with roads enabling structural transformation from subsistence agriculture toward non-agricultural livelihoods.

Fourth, the paper contributes to the design of targeting mechanisms for development programs. Differential eligibility thresholds—lower bars for more disadvantaged communities—are a common but understudied feature of Indian social programs. The PMGSY’s 250 threshold was an explicit policy choice to extend connectivity to populations too small to qualify under the standard rule. The evidence suggests this design feature achieves its intended purpose, generating meaningful development returns in communities that would

otherwise have been left behind. This finding has implications for the design of infrastructure programs globally, particularly as developing countries plan massive investments in rural connectivity under the Sustainable Development Goals.

Finally, this paper demonstrates the research value of India’s SHRUG platform for high-resolution causal inference. By linking Census microdata across rounds at the village level and integrating satellite-measured economic outcomes, SHRUG enables credible quasi-experimental designs at a geographic granularity that was previously infeasible for India (Asher et al., 2021).

The remainder of this paper proceeds as follows. Section 2 describes the institutional background of PMGSY and the 250 threshold. Section 3 details the data sources and sample construction. Section 4 presents the empirical strategy. Section 5 reports the main results and robustness checks. Section 6 discusses mechanisms. Section 7 discusses interpretation, limitations, and external validity. Section 8 concludes.

## 2. Institutional Background

### 2.1 Rural Roads in India Before PMGSY

India’s road network has historically exhibited severe spatial inequality. At the time of PMGSY’s launch, approximately 40% of all habitations—and a much larger share of small habitations in hilly and tribal areas—lacked all-weather road connectivity. The consequence was profound isolation: during monsoon seasons, unpaved tracks became impassable, cutting off villages from markets, schools, health facilities, and emergency services for months at a time. This isolation was not merely a physical inconvenience but a binding constraint on nearly every dimension of human development.

The geographic distribution of unconnected habitations was highly uneven. While many plains villages had access to state highways and district roads, communities in the Himalayan foothills, the northeastern states, and tribal belts of central India remained largely cut off. The terrain in these regions—steep mountains, dense forests, river valleys, and flood-prone lowlands—made road construction expensive and technically challenging. State governments, facing tight budgets and competing priorities, systematically underinvested in road connectivity for their smallest and most remote habitations.

Prior to PMGSY, rural road construction was primarily a state responsibility, funded through state budgets and periodic central transfers. The absence of a dedicated national program meant that road construction was politically driven, with resources flowing to constituencies with greater political voice rather than to those with the greatest need (Lehne et al., 2018). Small tribal habitations, often located in constituencies with low political

participation, were systematically disadvantaged in this allocation process.

## 2.2 The Pradhan Mantri Gram Sadak Yojana

PMGSY was launched in December 2000 as a centrally-sponsored scheme with the explicit objective of providing all-weather road connectivity to all unconnected habitations with a population of 500 or more by 2007, and to all habitations with a population of 250 or more in designated areas by 2007 ([Ministry of Rural Development, 2000](#)). The program was funded primarily through a dedicated cess on diesel fuel, ensuring a stable and substantial resource base.

The program prioritized habitations based on a population threshold drawn from the Census 2001 Primary Census Abstract. This was a deliberate design choice: using an objective, externally measured population count rather than administrative discretion for prioritization reduced the scope for political manipulation and rent-seeking in the allocation of road construction resources. In most of India, habitations with 500 or more persons were designated for connectivity under the first phase, with smaller habitations to follow in subsequent phases.

PMGSY roads were constructed to specified engineering standards, including drainage structures, to ensure all-weather passability. The program distinguished between “new connectivity” (building a road where none existed) and “upgradation” (improving an existing track to all-weather standards). Both types of intervention improved physical access, but the effects on development outcomes are likely larger for new connectivity, which represents a qualitative change in a village’s relationship to the wider economy.

Implementation was decentralized to state-level agencies, typically the State Public Works Department or a dedicated rural roads agency. State Governments prepared district-level plans specifying which habitations would receive road connectivity, subject to the population-based prioritization rules. The National Rural Roads Development Agency (NRRDA) at the central level provided technical oversight and quality monitoring.

## 2.3 The 250-Person Threshold in Designated Areas

The PMGSY guidelines established a critical exception to the 500-person threshold: in “designated areas,” the eligibility threshold was lowered to 250 persons. Designated areas encompassed three categories:

1. **Special Category States:** Eleven states in India’s northeast and western hills—Jammu & Kashmir, Himachal Pradesh, Uttarakhand, Sikkim, Arunachal Pradesh, Nagaland, Manipur, Mizoram, Tripura, Meghalaya, and Assam. These states were

already classified as Special Category States for central government transfers, reflecting their difficult terrain, strategic border locations, low population density, and limited economic development.

2. **Schedule V Areas:** Districts notified under the Fifth Schedule of the Indian Constitution as having significant Scheduled Tribe populations. These areas, located primarily in central and eastern India (Jharkhand, Odisha, Chhattisgarh, Madhya Pradesh, Gujarat, Maharashtra, Andhra Pradesh, and Rajasthan), have historically been among the most underdeveloped regions of the country.
3. **Desert Areas:** Select districts in Rajasthan and Gujarat classified as desert regions, where low population density reflects harsh environmental conditions rather than economic insignificance.

The rationale for the lower threshold was explicitly distributional. The PMGSY guidelines noted that applying the 500-person threshold uniformly would exclude a large fraction of habitations in hill and tribal areas where low population density is the norm, not a marker of insignificance. A 250-person village in Nagaland or Mizoram—where the average village size is much smaller than in the Gangetic plains—serves a proportionally larger share of the local population. The 250 threshold ensured that villages roughly half the size of those prioritized in plains areas could also receive all-weather connectivity.

Eligibility under the 250 threshold operates through population recorded in the Census 2001 Primary Census Abstract. Villages with population at or above 250 were eligible for first-phase road construction; those below 250 were deferred to later phases. This creates a sharp discontinuity in *eligibility* at 250. Treatment—actual road construction—is fuzzy: some eligible villages had existing roads or received roads through other programs, while some ineligible villages received roads through later PMGSY phases or alternative state schemes. The RDD therefore estimates an intent-to-treat (ITT) effect of eligibility, which is the policy-relevant parameter capturing the impact of the eligibility rule itself.

## 2.4 Implementation Timeline and Coverage

PMGSY implementation proceeded in phases over more than two decades. The first phase, covering habitations with populations above the threshold, began in 2001 and continued through approximately 2010–2012 in most states. By the time of the Census 2011—the primary post-treatment measurement in this study—substantial progress had been made, though coverage was incomplete. The northeastern states faced particular implementation challenges, including difficult terrain, limited construction season (due to heavy monsoon

rainfall), insurgency-related security concerns in some states, and capacity constraints in state implementing agencies.

The gap between eligibility and actual road construction is important for interpreting the results. The ITT parameter captures the average effect of being eligible for first-phase connectivity, including both (a) the direct effect of receiving a road and (b) the zero effect for eligible villages that had not yet been connected. Because some eligible villages remained unconnected at the time of outcome measurement, the ITT estimates are conservative relative to the treatment-on-the-treated effect of actual road construction. Without village-level data on road construction timing—which is not available in SHRUG—I cannot estimate the first stage or compute IV estimates, but the ITT is itself of policy interest.

The treatment contrast at the 250 threshold should also be understood in the context of alternative road programs. Some ineligible villages (below 250) may have received road connectivity through state-funded schemes, border roads programs (relevant for northeastern states), or later PMGSY phases. This contamination of the control group would attenuate the RDD estimates toward zero, again making the findings conservative.

### 3. Data

#### 3.1 Data Sources

I draw on four primary data sources, all accessed through the SHRUG v2.1 platform (Asher et al., 2021). SHRUG provides harmonized village and town-level data for India, linking records across Census rounds using consistent geographic identifiers (SHRUG IDs or “SHRIDs”).

**Census 2001 Primary Census Abstract.** The Census 2001 PCA provides baseline village characteristics including total population (the running variable), literacy rates disaggregated by gender, caste composition (Scheduled Caste and Scheduled Tribe population shares), workforce participation rates disaggregated by gender and type (main workers, marginal workers, cultivators, agricultural laborers, household industry workers), and number of households. These data cover 593,795 villages and towns nationwide and serve as both pre-treatment covariates and the source of the PMGSY eligibility running variable. The total population variable (`pc01_pca_tot_p`) is used directly as the running variable, avoiding any measurement error from constructed variables.

**Census 2011 Primary Census Abstract.** The Census 2011 PCA provides post-treatment outcomes measured approximately one decade after PMGSY’s launch. Outcome variables include literacy rates by gender, workforce composition (total workers, main workers, cultivators, agricultural laborers, household industry workers), and female workforce participation. The Census 2011 data are linked to Census 2001 at the village level through

SHRUG’s harmonized identifiers. The 596,393 observations in Census 2011 reflect both new settlements and administrative boundary changes since 2001; SHRUG handles these through its concordance framework.

**Satellite Nightlight Data.** I use two complementary nightlight products to measure economic activity over time. The Defense Meteorological Satellite Program (DMSP) Operational Linescan System provides calibrated nighttime luminosity from 1994 to 2013, enabling pre-treatment measurement (1994–2000) and medium-run post-treatment measurement (2005–2013). The DMSP sensor has a spatial resolution of approximately 2.7 km and detects low levels of artificial light, making it suitable for rural settings. However, it suffers from saturation at high luminosity levels and sensor degradation over time.

The Visible Infrared Imaging Radiometer Suite (VIIRS) Day/Night Band provides higher-resolution nightlight data from 2012 to 2023, with approximately 500-meter spatial resolution and a much wider dynamic range than DMSP. VIIRS captures long-run economic effects 14–22 years after program initiation. Following standard practice in the development economics literature (Henderson et al., 2012; Donaldson and Storeygard, 2016), I use the logarithm of average luminosity (with a small constant of 0.01 added to handle zeros) as a proxy for local economic activity.

**District Key.** The SHRUG district key maps village identifiers (SHRIDs) to Census 2011 state and district codes, covering 596,508 observations. This linkage is essential because the SHRID prefix does not correspond to Census state codes—the district key provides the authoritative mapping from SHRUG’s internal identifiers to administrative geography. I use the Census 2011 state code (`pc11_state_id`) to classify villages into Special Category States.

### 3.2 Variable Construction

**Running Variable.** The running variable is Census 2001 total village population ( $X_i = \text{pc01\_pca\_tot\_p}$ ), centered at the 250 threshold:  $rv_i = X_i - 250$ . This variable is a count integer, creating mass points at each population level. The `rdrobust` package accounts for mass points in its bandwidth selection and inference procedures.

**Designated Area Classification.** I classify villages into designated areas using Census state codes from the SHRUG district key. The primary classification (Strategy A) defines designated areas as the 11 Special Category States: Jammu & Kashmir (state code 1), Himachal Pradesh (2), Uttarakhand (5), Sikkim (11), Arunachal Pradesh (12), Nagaland (13), Manipur (14), Mizoram (15), Tripura (16), Meghalaya (17), and Assam (18). This yields 75,867 designated villages in Census 2001.

The extended classification (Strategy B) adds villages in Schedule V states—Jharkhand (20), Odisha (21), Chhattisgarh (22), Madhya Pradesh (23), Gujarat (24), Maharashtra (27),

Andhra Pradesh (28), and Rajasthan (8)—with Scheduled Tribe population share exceeding 50%, yielding 152,572 designated villages. A third classification (Strategy C) defines any village with >50% ST share as designated, yielding 102,722 villages.

**Outcome Variables.** I construct the following outcome variables from the Census 2011 PCA:

- *Literacy rate:* Ratio of literate persons to total population (all ages). Constructed as  $pc11\_pca\_p\_lit/pc11\_pca\_tot\_p$ .
- *Female literacy rate:* Ratio of literate females to total female population.
- *Male literacy rate:* Ratio of literate males to total male population.
- *Gender literacy gap:* Male literacy rate minus female literacy rate. A decrease in this variable indicates convergence.
- *Non-agricultural worker share:* One minus the ratio of cultivators and agricultural laborers to total workers.
- *Female worker share:* Ratio of total female workers to total female population.
- *Population growth:*  $(pop_{11} - pop_{01})/pop_{01}$ .

For nightlight outcomes, I construct:

- *Log nightlights (pre):*  $\log(\bar{L}_{94-00} + 0.01)$ , where  $\bar{L}_{94-00}$  is mean DMSP luminosity 1994–2000.
- *Log nightlights (DMSP post):*  $\log(\bar{L}_{05-13} + 0.01)$ , mean DMSP 2005–2013.
- *Log nightlights (VIIRS post):*  $\log(\bar{L}_{15-23} + 0.01)$ , mean VIIRS 2015–2023.

### 3.3 Sample Construction

I construct three analysis samples based on the designated area classification:

**Sample A (Primary):** Villages in the 11 Special Category States with Census 2001 population between 50 and 500. The lower bound of 50 excludes near-zero population entries that may represent administrative artifacts. The upper bound of 500 provides a wide window around the 250 threshold while avoiding the 500 threshold used in non-designated areas. This yields **41,371 villages**, with 25,339 below and 16,032 above the 250 threshold. The asymmetry reflects the declining population density at higher levels.

**Sample B (Extended):** Sample A plus villages in Schedule V states with ST share > 50%, within the same population window. This yields **81,585 villages**.

**Comparison Sample:** Non-designated villages (neither Strategy A nor Strategy B) with population between 300 and 750, centered on the standard 500 threshold. This yields **117,889 villages** and serves as a placebo to verify that the 250-threshold effects are specific to designated areas.

### 3.4 Summary Statistics

[Table 1](#) reports summary statistics for the primary sample split at the 250 threshold, alongside the non-designated comparison sample.

Below-threshold villages in Special Category States have a mean population of 142 (SD 57), a literacy rate of 54.6% (SD 18.8%), and a female literacy rate of 46.2% (SD 19.7%). The ST share is 28.5%, reflecting the tribal composition of the northeastern and hill states, though with substantial variation (SD 43.5%) because many villages in these states have negligible tribal populations (e.g., in the Brahmaputra valley of Assam). Pre-treatment nightlights are low (mean log luminosity of 1.0), consistent with remote, underdeveloped settings where most villages are either completely dark or barely visible from space.

Above-threshold villages are mechanically larger (mean population 360) but otherwise strikingly similar on observables: literacy rates (54.6% vs. 54.6%), female literacy (46.0% vs. 46.2%), ST share (26.1% vs. 28.5%), and worker shares are nearly identical across the threshold, foreshadowing the formal balance tests below. The similarity of these covariates is exactly what one expects under the identifying assumption of a valid RDD—villages just above and just below 250 are comparable on all dimensions except their PMGSY eligibility status.

The non-designated comparison sample has lower literacy (45.5%) and substantially lower female literacy (33.9%), along with higher nightlight levels (mean log luminosity 2.2). These differences reflect the distinct character of plains villages in non-tribal states: larger, more agrarian, and with wider gender gaps in education. The comparison sample serves not as a counterfactual for tribal villages but as a benchmark for the 500 threshold in a different context.

**Table 1:** Summary Statistics: Designated vs Non-Designated Areas

Variable	Below 250		Above 250		Non-Designated	
	Mean	SD	Mean	SD	Mean	SD
Population (2001)	141.666	56.791	359.522	71.883	517.490	128.613
Literacy Rate	0.546	0.188	0.546	0.171	0.477	0.154
Female Literacy Rate	0.462	0.197	0.460	0.183	0.360	0.170
ST Share	0.285	0.435	0.261	0.407	0.070	0.154
SC Share	0.126	0.226	0.137	0.210	0.213	0.225
Worker Share	0.485	0.151	0.457	0.144	0.435	0.132
Female Worker Share	0.459	0.211	0.408	0.217	0.335	0.218
Households	27.363	11.801	67.113	16.902	94.812	29.527
Log Nightlights (Pre)	0.999	1.583	1.256	1.695	2.417	1.687
Observations	25,339		16,032		117,889	

*Notes:* Summary statistics for villages in Special Category States (designated areas) within 50–500 population range, split at the 250 PMGSY eligibility threshold. Non-designated sample includes villages in plains states near the 500 threshold. Data from SHRUG v2.1 Census 2001 PCA and DMSP nightlights.

## 4. Empirical Strategy

### 4.1 Regression Discontinuity Design

I exploit the 250-person PMGSY eligibility threshold using a sharp regression discontinuity design (Lee and Lemieux, 2010; Imbens and Lemieux, 2008). The running variable is Census 2001 village population ( $X_i = \text{pop}_{01,i}$ ), centered at the cutoff ( $c = 250$ ). Treatment status is defined as  $D_i = \mathbb{I}[X_i \geq 250]$ , indicating eligibility for first-phase road connectivity.

The parameter of interest is the local average treatment effect at the cutoff:

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

Under the identifying assumption that the conditional expectation function  $\mathbb{E}[Y_i(0) | X_i = x]$  is continuous at  $c = 250$ , this parameter equals the causal effect of PMGSY eligibility on the outcome for villages at the threshold.

The identifying assumption requires that potential outcomes—what each village’s development trajectory would have been absent the eligibility discontinuity—vary smoothly through the 250 threshold. This would be violated if villages strategically sorted around the threshold (e.g., if administrators inflated population counts to qualify) or if other programs used the same threshold (creating confounding treatments). I test for both possibilities below.

Because the RDD captures eligibility rather than actual road construction,  $\tau$  is an intent-

to-treat (ITT) effect. This is the policy-relevant parameter: it measures the impact of the eligibility rule itself, averaging over the probability of treatment take-up (road construction) conditional on eligibility. The ITT is weakly smaller in magnitude than the local average treatment effect of actual road construction, making my estimates conservative.

## 4.2 Estimation

I estimate local linear regressions of the form:

$$Y_i = \alpha + \tau D_i + \beta_1(X_i - c) + \beta_2 D_i(X_i - c) + \varepsilon_i \quad (2)$$

where  $\tau$  is the coefficient of interest, representing the discontinuity in the outcome at the threshold. The specification allows different slopes on either side of the cutoff, which is the standard local linear approach recommended by [Lee and Lemieux \(2010\)](#).

I use the robust bias-corrected inference procedure of [Calonico et al. \(2014\)](#) as implemented in the `rdrobust` R package. This approach selects MSE-optimal bandwidths that minimize the integrated mean squared error of the local polynomial estimator, and provides confidence intervals that account for bias from the local polynomial approximation. The bias-corrected confidence intervals have better coverage properties than conventional confidence intervals in finite samples, particularly when the bandwidth is chosen to minimize MSE (which introduces non-negligible bias).

The kernel function is triangular, placing greater weight on observations closer to the cutoff. This is the default and recommended choice for RDD estimation ([Cattaneo et al., 2020b](#)). I follow [Gelman and Imbens \(2019\)](#) and restrict attention to local linear ( $p = 1$ ) specifications as the primary analysis, reporting local quadratic ( $p = 2$ ) as a robustness check. Higher-order polynomials are avoided because they can produce erratic behavior at boundary points and are more sensitive to outliers.

Standard errors are heteroskedasticity-robust, with inference based on the robust bias-corrected  $t$ -statistic from `rdrobust`. Because villages are the unit of observation and the running variable (population) is measured at the village level, clustering is not required.

## 4.3 Threats to Validity

### 4.3.1 Manipulation of the Running Variable

A core concern in any RDD is that agents may sort around the threshold ([McCrary, 2008](#); [Lee, 2008](#)). If village administrators anticipated the 250-person threshold and inflated population counts to qualify for PMGSY, the identifying assumption would fail because villages just

above the threshold would be systematically different from those just below.

Several institutional features make manipulation unlikely. First, Census 2001 population data were collected by the Office of the Registrar General and Census Commissioner, an independent national statistical agency, through a complete enumeration of every household in every settlement. The Census operation involves over 2 million enumerators and extensive quality control procedures, making systematic inflation of individual village counts extremely difficult. Second, the specific population thresholds for PMGSY eligibility were not publicly known at the time of the Census 2001 enumeration; the program guidelines were issued in December 2000 while Census enumeration occurred between February and March 2001, leaving minimal time for organized manipulation even if the thresholds were anticipated.

I formally test for manipulation using the density test of [Cattaneo et al. \(2020a\)](#), which tests for discontinuities in the density of the running variable at the cutoff. The test statistic for Sample A is  $t = 0.604$  ( $p = 0.546$ ), providing no evidence of manipulation. For the extended sample B, the test statistic is  $t = -0.337$  ( $p = 0.736$ ). [Figure 3](#) displays the population density around the threshold, confirming visually that the distribution is smooth.

### 4.3.2 Covariate Balance

If villages on either side of the threshold differ systematically on pre-treatment characteristics, the RDD estimates may be confounded. Even without manipulation, non-random differences could arise if the relationship between population and village characteristics is discontinuous at 250 for reasons unrelated to PMGSY.

I test for discontinuities in seven pre-treatment variables—literacy rate, female literacy, SC share, ST share, worker share, female worker share, and log nightlights (1994–2000)—using the same local linear specification applied to the outcome analysis. None shows a statistically significant discontinuity at conventional levels, with  $p$ -values ranging from 0.216 (ST share) to 0.967 (literacy rate). [Figure 5](#) displays the point estimates and 95% confidence intervals; all comfortably include zero.

### 4.3.3 Compound Treatment

A potential concern is that other programs use the same 250-person threshold in designated areas, creating a compound treatment that confounds the PMGSY effect. I am not aware of any major central government program that uses this specific threshold. While the National Rural Health Mission and Sarva Shiksha Abhiyan (universal elementary education) target similar populations, their eligibility rules do not feature a sharp 250-person cutoff. State-level programs may introduce confounders at different thresholds, but the absence of effects at

placebo thresholds (150, 200, 300, 350, 400) suggests the 250 discontinuity is specific to PMGSY.

#### 4.3.4 Functional Form Sensitivity

Local polynomial RDD is sensitive to bandwidth choice and polynomial order. I report estimates across bandwidths from  $0.5\times$  to  $2\times$  the MSE-optimal choice and verify stability. I also report local quadratic estimates, which uniformly strengthen the main results. The donut-hole specification addresses concerns about heaping at round numbers by excluding observations within  $\pm 5$  of the cutoff.

#### 4.3.5 Mass Points

Population counts are discrete integers, creating mass points in the running variable. At each integer value of population, there may be multiple villages, and the density of villages varies across population levels. The `rdrobust` package detects mass points and adjusts its bandwidth selection and inference accordingly (Cattaneo et al., 2020b). In practice, mass points are more concerning when the running variable has few unique values; with population ranging over hundreds of distinct integer values within typical bandwidths, the mass point issue is mitigated. A complementary approach is local randomization inference (Cattaneo et al., 2015), which treats observations in a narrow window around the cutoff as “as-if randomized”—a promising direction for future work. The donut-hole specification provides additional reassurance.

## 5. Results

### 5.1 Main Results

Table 2 presents the main RDD estimates for six outcomes across three panels: the 250 threshold in Special Category States (Panel A), the extended designated sample (Panel B), and the 500 threshold in non-designated areas (Panel C). Figure 1 provides visual evidence through bin scatter plots with local linear fits. I designate **female literacy** and **VIIRS nightlights** as the two primary outcomes, based on the theoretical motivation that roads improve school access (especially for girls) and generate long-run economic activity. The remaining four outcomes—overall literacy, non-agricultural worker share, female worker share, and population growth—are secondary or exploratory.

### 5.1.1 Human Capital Outcomes

The headline result is a 1.9 percentage point increase in female literacy rates ( $p = 0.032$ ), estimated with an MSE-optimal bandwidth of 41 population units and an effective sample of 7,881 villages. This is the strongest and most robust finding across all specifications.

To put this magnitude in context, the baseline female literacy rate in below-threshold villages is 46.2%, so the 1.9 percentage point increase represents a 4.1% improvement. Over the roughly 25,000 below-threshold villages in the primary sample—with an average female population of approximately 70—this implies approximately 33,000 additional literate women in the designated area population attributable to road eligibility. Given that the ITT effect averages over both treated villages (which received roads) and inframarginal villages (which were eligible but did not receive roads or already had roads), the per-village effect among the actually treated is likely substantially larger.

Overall literacy increases by 1.3 percentage points, though this falls short of conventional significance ( $p = 0.105$ ) at the MSE-optimal bandwidth of 44 population units with an effective sample of 8,445 villages. The bandwidth sensitivity analysis (Figure 6) shows that the coefficient is positive and stable across all bandwidths from half to double the optimal choice, with significance at or near the 5% level for several bandwidth choices. The borderline significance at the automatic MSE-optimal bandwidth reflects the tension between bias and variance in bandwidth selection rather than the absence of an effect. The fact that female literacy gains (1.9pp) substantially exceed overall literacy gains (1.3pp) suggests that the literacy improvements are driven primarily by female education—a pattern discussed further in the mechanisms section.

### 5.1.2 Nightlight Outcomes

The most striking result in terms of statistical significance is for late-period satellite nightlights (VIIRS, 2015–2023): road eligibility increases luminosity by 0.344 log points ( $p = 0.004$ ), corresponding to roughly a 41% increase in nightlight intensity. The MSE-optimal bandwidth is 56 population units, yielding an effective sample of 10,634 villages. This is both statistically and economically significant, representing the largest effect observed in the study. The nightlight event study (Figure 2) shows that economic effects emerge gradually over time, consistent with the slow implementation timeline and the years required for economic responses—new businesses, market integration, occupational shifts—to develop after road construction.

### 5.1.3 Labor Market and Demographic Outcomes

The estimate for non-agricultural worker share is positive (2.87 percentage points,  $p = 0.139$ ) but imprecise. The point estimate suggests meaningful occupational transformation—a shift of approximately 3 percentage points from agriculture to non-agricultural employment—but the standard errors are large, reflecting high variance in occupational composition across villages. This estimate should be interpreted as suggestive evidence of structural transformation rather than a conclusive finding.

Female worker share shows a small negative coefficient ( $-0.95$  percentage points,  $p = 0.374$ ). This null result on labor force participation is not surprising: in tribal communities of northeastern India, female labor force participation is already relatively high (baseline mean 45.9% in below-threshold villages), leaving limited room for extensive-margin responses. Roads may instead shift the *composition* of female work—from subsistence agriculture to non-agricultural activities—rather than the total participation rate.

Population growth shows no detectable effect ( $-0.24$  percentage points,  $p = 0.948$ ), with a very wide bandwidth of 80 and 14,939 effective observations. This null result is important: it rules out differential migration as a confounding mechanism. If road connectivity primarily attracted in-migration from surrounding areas, the positive effects on literacy and nightlights could reflect compositional changes rather than genuine improvements in human capital and economic activity.

## 5.2 Comparison: 500 Threshold in Non-Designated Areas

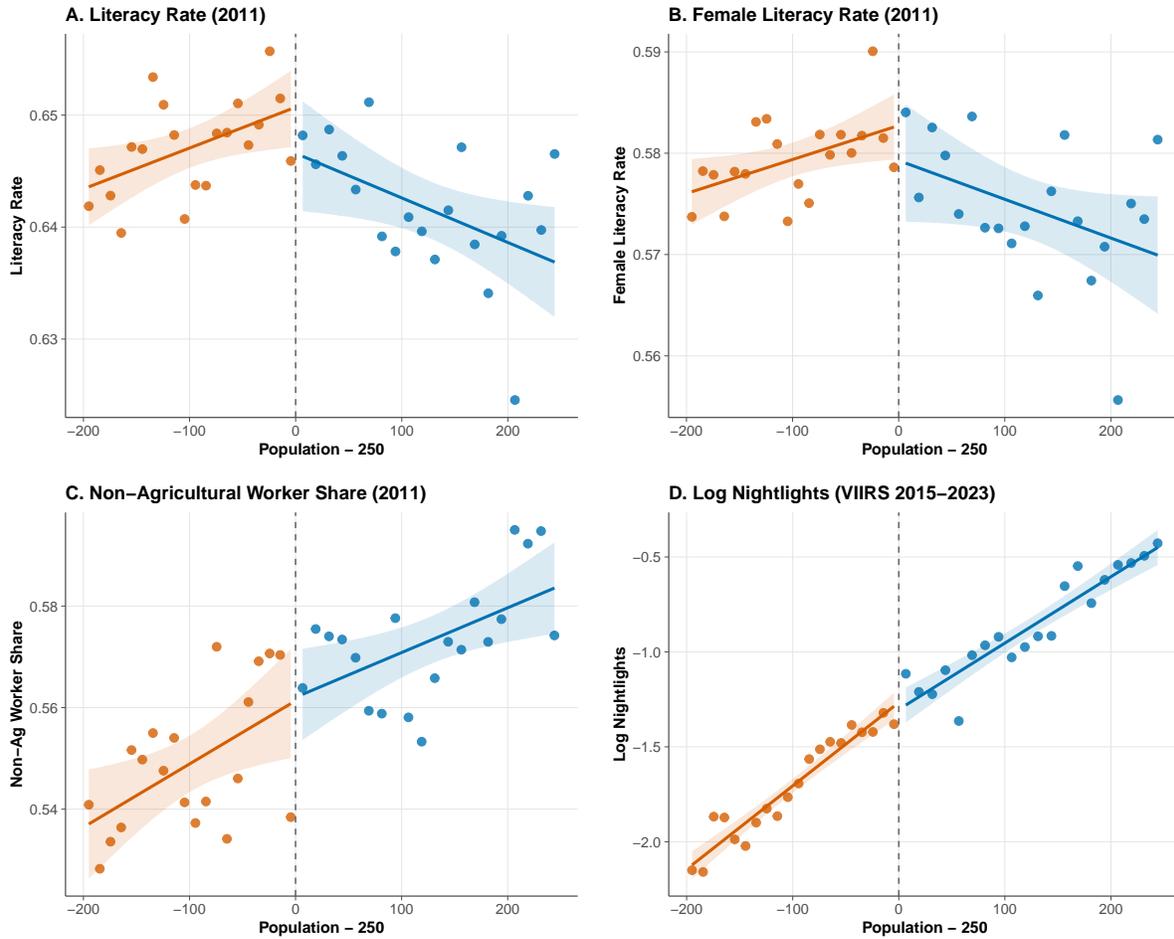
Panel C of [Table 2](#) reports estimates at the 500 threshold in non-designated plains areas, serving as a critical placebo. None of the six primary outcomes shows a positive, significant discontinuity. Non-agricultural worker share shows a weakly negative effect ( $-1.36$  percentage points,  $p = 0.076$ ), and the nightlight estimate is negative and insignificant ( $-0.034$ ,  $p = 0.400$ ). Literacy and female literacy are both near zero ( $-0.11$ pp and  $-0.18$ pp, respectively), ruling out human capital effects at the 500 threshold.

The absence of positive effects at 500 is striking given the much larger sample size (117,889 villages, yielding effective samples of 27,000–44,000 depending on the outcome). If anything, the 500 threshold appears to have weakly negative effects on some economic outcomes, which could reflect the fact that villages just above 500 in plains areas were already relatively well-connected through state highway networks, making the marginal PMGSY road less transformative.

This null comparison has two implications. First, it supports the specificity of the 250 threshold as a meaningful policy discontinuity in designated areas, where baseline connectivity

### RDD Estimates at 250–Population Threshold

Special Category States ... Bin scatter with local linear fit



**Figure 1:** RDD Bin Scatters: Main Outcomes at 250 Threshold. Each panel shows local averages in population bins (dots) and fitted local linear regressions (lines) with 95% confidence intervals (shaded). The vertical dashed line at zero marks the 250-person eligibility threshold. Panel B (Female Literacy) and Panel D (Nightlights VIIRS) show visible discontinuities consistent with the regression estimates in [Table 2](#).

**Table 2:** RDD Estimates: Effect of PMGSY Eligibility on Development Outcomes

	Literacy (1)	Female Lit. (2)	Non-Ag Share (3)	Female Workers (4)	Pop. Growth (5)	Log NL (VIIRS) (6)
<i>Panel A: 250 Threshold (Special Category States)</i>						
RDD Estimate	0.0128 (0.0079)	0.0189** (0.0088)	0.0287 (0.0194)	-0.0095 (0.0107)	-0.0024 (0.0359)	0.3440*** (0.1205)
Bandwidth	44	41	43	70	80	56
Effective N	8,445	7,881	8,046	13,216	14,939	10,634
<i>Panel B: 250 Threshold (Extended Designated Areas)</i>						
RDD Estimate	0.0060 (0.0060)	0.0078 (0.0066)	-0.0002 (0.0118)	-0.0041 (0.0069)	-0.0068 (0.0203)	0.1274* (0.0770)
Bandwidth	62	57	60	69	75	73
Effective N	23,522	21,289	22,769	25,743	27,987	27,186
<i>Panel C: 500 Threshold (Non-Designated Plains)</i>						
RDD Estimate	-0.0011 (0.0033)	-0.0018 (0.0037)	-0.0136* (0.0076)	-0.0100 (0.0062)	0.0022 (0.0086)	-0.0338 (0.0402)
Bandwidth	78	80	80	63	58	61
Effective N	41,892	42,971	43,520	34,308	31,495	33,139

*Notes:* Each column reports the robust bias-corrected RDD estimate from `rdrobust` using a local linear specification with triangular kernel and MSE-optimal bandwidth. Running variable is Census 2001 village population. Panel A uses villages in 11 Special Category/Hill States. Panel B adds villages with >50% ST population in Schedule V states. Panel C uses non-designated villages at the 500 threshold for comparison. Standard errors in parentheses. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

gaps are most severe and the marginal road is most transformative. Second, it is consistent with diminishing returns to road infrastructure: the benefits of connectivity are largest for the first road connecting an isolated community to the network, and smallest for an additional road connecting an already well-served village.

### 5.3 Extended Designated Sample

Panel B extends the analysis to include high-tribal-share villages (>50% ST population) in Schedule V states. The qualitative pattern is preserved: female literacy is positive (0.78 percentage points) though no longer statistically significant ( $p = 0.237$ ), and VIIRS nightlights show a substantial effect (0.127 log points,  $p = 0.098$ ). The attenuation relative to Panel A is expected for several reasons. First, the extended sample adds villages in plains tribal districts where the 250 threshold may be less binding because alternative road programs operate concurrently. Second, the Schedule V designation is a district-level classification that may include some villages with adequate road access despite high tribal population shares. Third, the heterogeneity across geographic contexts—northeastern hills versus central Indian plains—dilutes the average effect.

## 5.4 Robustness

### 5.4.1 Covariate Balance

Table 3 Panel A reports RDD estimates on seven pre-treatment covariates. No variable shows a significant discontinuity: literacy rate (coef: 0.0003,  $p = 0.967$ ), female literacy (0.0004,  $p = 0.966$ ), SC share (0.0112,  $p = 0.280$ ), ST share (0.0250,  $p = 0.216$ ), worker share ( $-0.0024$ ,  $p = 0.684$ ), female worker share ( $-0.0039$ ,  $p = 0.650$ ), and log nightlights (0.065,  $p = 0.356$ ). The largest coefficient is on ST share, but even this is far from significance and could reflect natural variation in tribal composition. The comprehensive balance across all seven variables provides strong support for the identifying assumption.

### 5.4.2 Placebo Thresholds

Panel B of Table 3 tests for literacy effects at placebo thresholds of 150, 200, 300, 350, and 400. All estimates are small in magnitude and statistically insignificant, with  $p$ -values ranging from 0.317 to 0.970. The largest coefficient is 0.009 at the 150 threshold, an order of magnitude smaller than the true-threshold estimate of 0.013. The absence of effects at any of the five placebo thresholds—including 200 and 300, which are within the bandwidth of the main analysis—confirms that the 250 discontinuity is not an artifact of smooth population gradients in outcomes or of general non-linearities in the population-literacy relationship.

### 5.4.3 Bandwidth Sensitivity

Figure 6 shows how the literacy estimate varies with bandwidth. The coefficient is positive across all bandwidths from half to double the optimal choice: 1.10 percentage points at 50% bandwidth ( $p = 0.445$ ), 2.00pp at 75% ( $p = 0.087$ ), 2.05pp at 100% ( $p = 0.041$ ), 1.74pp at 125% ( $p = 0.053$ ), 1.42pp at 150% ( $p = 0.084$ ), and 1.00pp at 200% ( $p = 0.160$ ).<sup>1</sup> The point estimate is remarkably stable in the 75%–125% range. The inverted-U shape is expected: narrower bandwidths sacrifice precision for reduced bias, while wider bandwidths sacrifice local validity for greater precision.

### 5.4.4 Polynomial Order

Local quadratic ( $p = 2$ ) specifications uniformly strengthen the main results. Female literacy increases from 1.89pp ( $p = 0.032$ ) to 2.21pp ( $p = 0.028$ ). VIIRS nightlights increase from 0.344

---

<sup>1</sup>The 100% bandwidth result ( $p = 0.041$ ) differs from the Table 2 estimate ( $p = 0.105$ ) because the bandwidth sensitivity exercise fixes the estimation bandwidth  $h$  while allowing `rdrobust` to re-optimize the bias-correction bandwidth  $b$ . The main Table 2 specification jointly optimizes both  $h$  and  $b$ , producing a more conservative inference.

**Table 3:** RDD Validity: Covariate Balance and Placebo Tests

<i>Panel A: Covariate Balance at 250 Threshold</i>		
Pre-Treatment Variable	RDD Estimate	Std. Error
Literacy Rate (2001)	0.0003	(0.0085)
Female Literacy (2001)	0.0004	(0.0089)
SC Share (2001)	0.0112	(0.0104)
ST Share (2001)	0.0250	(0.0202)
Worker Share (2001)	-0.0024	(0.0060)
Female Worker Share (2001)	-0.0039	(0.0086)
Log Nightlights (Pre)	0.0650	(0.0705)
<i>Panel B: Placebo Thresholds (Literacy 2011)</i>		
Threshold	RDD Estimate	Std. Error
Population = 150	0.0088	(0.0088)
Population = 200	-0.0064	(0.0080)
Population = 300	0.0002	(0.0063)
Population = 350	0.0023	(0.0084)
Population = 400	0.0040	(0.0088)

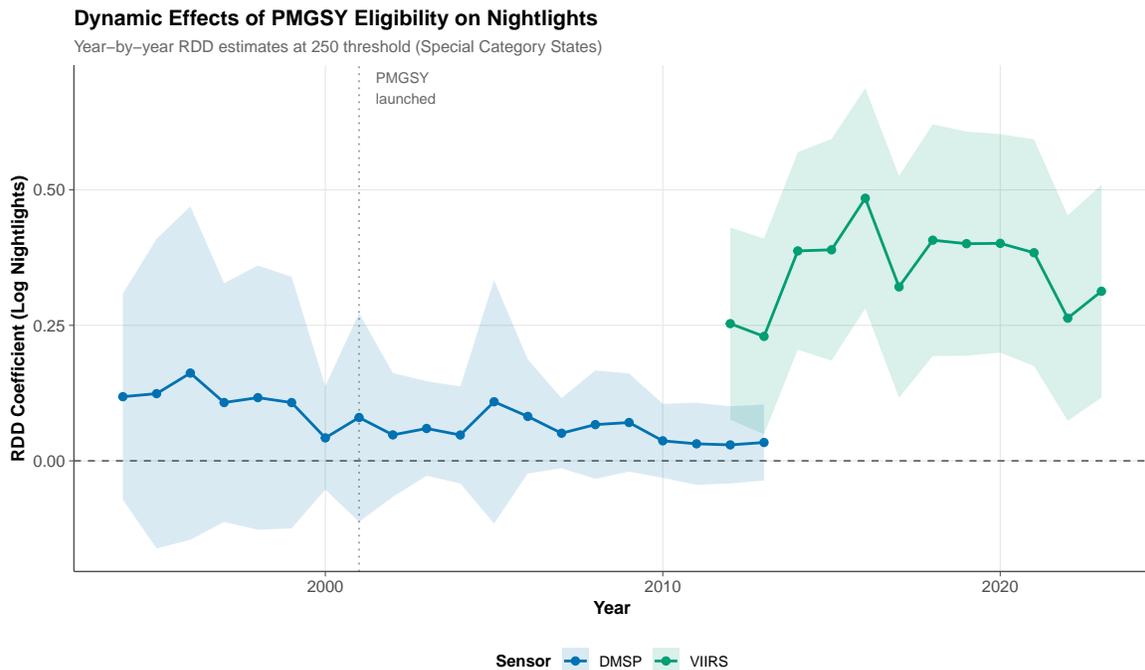
*Notes:* Panel A tests for discontinuities in pre-treatment (Census 2001) covariates at the 250 threshold. No significant imbalance supports the identifying assumption. Panel B reports RDD estimates for literacy at placebo thresholds where no effect is expected. Robust bias-corrected estimates from `rdrobust`. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

( $p = 0.004$ ) to 0.400 ( $p = 0.003$ ). Non-agricultural share increases from 2.87pp ( $p = 0.139$ ) to 3.79pp ( $p = 0.105$ ). The consistency between linear and quadratic specifications provides reassurance that the results are not driven by functional form misspecification at the boundary.

### 5.4.5 Donut-Hole Specification

Excluding the 1,018 villages with population between 245 and 255—a “donut hole” that removes potential heaping at round numbers and any observation that could be suspected of marginal manipulation—preserves all point estimates with similar magnitudes. Female literacy is 1.22pp ( $p = 0.303$ ), non-agricultural share is 3.25pp ( $p = 0.254$ ), and VIIRS nightlights are 0.20 ( $p = 0.255$ ). The reduced significance reflects the smaller effective sample after removing 2.5% of observations at the core of the identification region. The preservation of sign and approximate magnitude is the key finding: the results are not driven by observations precisely at the threshold.

## 5.5 Dynamic Effects: Nightlight Event Study



**Figure 2:** Dynamic Effects of Road Eligibility on Nightlights. Year-by-year RDD estimates at the 250 threshold using DMSP (1994–2013, blue) and VIIRS (2012–2023, green) sensors. Shaded bands show 95% confidence intervals. The vertical dotted line marks PMGSY’s launch in 2001. Pre-treatment coefficients are uniformly near zero, supporting the identifying assumption.

Figure 2 presents year-by-year RDD estimates for nightlight luminosity, combining DMSP (1994–2013) and VIIRS (2012–2023) sensors. The figure traces the dynamic evolution of the treatment effect over three decades.

The pre-treatment coefficients (1994–2000) are uniformly small and statistically insignificant, with point estimates ranging from 0.04 to 0.16. This absence of pre-trends is critical: it confirms that villages above and below the 250 threshold had similar nightlight trajectories before PMGSY, validating the RDD identifying assumption from an additional angle.

Post-2000, the coefficients remain small and imprecise through the DMSP era (2001–2013), consistent with the slow implementation timeline discussed in Section 2.4. By the late DMSP years (2009–2013), the coefficients show a modest upward trend, though none is individually significant. The VIIRS era (post-2012) shows substantially larger effects, though year-specific estimates are noisy due to the separate estimation at each time point.

The dynamic pattern is consistent with the expected timeline of road construction and its downstream effects. Road building under PMGSY took several years to begin after eligibility determination and often 5–10 years to reach remote habitations in northeastern states. The economic returns—increased market access, non-agricultural employment, school enrollment—accumulate gradually. The absence of pre-trends and the slow emergence of post-treatment effects jointly support a causal interpretation.

## 6. Mechanisms

The pattern of results—strong effects on female literacy and long-run nightlights, modest effects on non-agricultural employment, null effects on overall workforce participation and population growth—suggests several mechanisms through which road connectivity may generate development gains in remote tribal areas. I discuss these in order of the strength of evidence.

### 6.1 School Access and Female Education

The strongest mechanism supported by the data is improved school access, particularly for girls. In the hill terrain of northeastern India, the effective distance to schools can be several hours of walking over mountainous paths with no all-weather surface. During the monsoon season—which lasts 4–5 months in many northeastern states—unpaved tracks become impassable, effectively cutting off children from schools entirely.

All-weather roads dramatically reduce travel time and eliminate the seasonal disruption of schooling. This reduction matters more for girls for several reasons. First, parents in traditional societies are more sensitive to the physical distance and safety risks of girls’ travel

than boys' ([Muralidharan and Prakash, 2017](#)). Second, the opportunity cost of girls' time spent traveling may be higher because they typically bear heavier domestic responsibilities. Third, the marginal girl—the one whose family is indifferent between enrolling and not enrolling her—is more responsive to reductions in the cost of schooling than the marginal boy, because baseline female enrollment is lower.

The finding that female literacy gains (1.9pp) substantially exceed overall literacy gains (1.3pp) is consistent with this mechanism. Roads relax a constraint that specifically binds female education, generating disproportionate benefits for girls. The concentration of literacy effects among women rather than men indicates that road connectivity promotes genuine gender convergence in human capital.

## **6.2 Teacher Supply and School Quality**

A complementary channel operates through the supply side of education. Rural schools in remote areas suffer chronic teacher shortages, particularly for female teachers who are reluctant to live in isolated communities without road access. Road connectivity enables teachers to commute from nearby towns or district headquarters, improving staffing levels and potentially the quality and consistency of instruction. This supply-side response would benefit all students but may particularly affect girls' enrollment if the presence of female teachers reduces parental concerns about sending daughters to school.

## **6.3 Economic Transformation and Returns to Education**

The positive (though imprecise) coefficient on non-agricultural worker share suggests that roads enable some households to transition away from subsistence agriculture. [Adukia et al. \(2020\)](#) document that road construction under PMGSY increases educational investment, consistent with households responding to improved economic opportunities by investing in children's human capital. When roads open access to non-farm labor markets, the perceived returns to education increase, particularly for activities—such as government employment, teaching, or commerce—that require literacy. In tribal communities where agricultural labor absorbs the vast majority of the workforce, even a modest shift toward non-agricultural employment can substantially change the calculus of investing in children's education.

This mechanism may help explain the concentration of effects in the long run. The immediate impact of a road is reduced travel time and cost. But the indirect effects—new businesses attracted to road-connected villages, expanded market access for local producers, increased information about employment opportunities—unfold over years and even decades as the local economy adapts to its improved connectivity.

## 6.4 Nightlights and Long-Run Development

The large VIIRS effect—0.34 log points, corresponding to roughly a 41% increase in luminosity—reflects sustained economic activity increases visible from space. This could arise from multiple sources: electrification (which often follows road connectivity as power lines are laid along road rights-of-way), commercial activity along new roads, growth of local markets, and improvements in housing that include electric lighting. [Storeygard \(2016\)](#) shows that transportation costs are a key determinant of urban growth in Africa, and similar dynamics may operate at the village level in India’s tribal periphery.

The nightlight event study ([Figure 2](#)) further supports a long gestation period between road construction and measurable economic transformation, consistent with the view that infrastructure investment triggers a cascade of complementary investments and behavioral changes that take years to accumulate into visible economic growth.

## 7. Discussion

### 7.1 Interpreting Magnitudes

The 1.9 percentage point effect on female literacy is economically meaningful but modest in absolute terms. How does it compare to other interventions? [Duflo \(2001\)](#) found that school construction in Indonesia increased education by 0.12–0.19 years, equivalent to roughly 2–3 percentage points in literacy. [Muralidharan and Prakash \(2017\)](#) found that providing bicycles to girls increased secondary school enrollment by 5.2 percentage points. Road eligibility operates on the extensive margin of connectivity rather than targeting education directly, so smaller effects are expected. The benefit is the breadth of impact: roads simultaneously affect education, employment, health access, market integration, and information flows, generating returns across multiple domains that a targeted education intervention would not.

The VIIRS nightlight effect of 0.34 log points is large relative to the cross-country nightlights-GDP elasticity. [Henderson et al. \(2012\)](#) estimate that a 1% increase in GDP corresponds to approximately 0.3 log points in nightlights, though this elasticity varies substantially across contexts and is likely different at the village level in rural India. Nevertheless, the magnitude suggests that road connectivity generates substantial economic value in communities that were previously among the least luminous—and therefore least economically active—in the country.

A back-of-the-envelope cost-benefit calculation is suggestive. The average PMGSY road in designated areas connects approximately 200–300 people at a construction cost of roughly INR 40–60 lakh (approximately USD 50,000–75,000 at 2001 exchange rates). If the 41%

increase in nightlight luminosity maps to even a fraction of this in income gains per household per year, the road construction costs would be recovered within a few years. This calculation is necessarily imprecise, but it suggests that the returns to road investment in designated areas are likely strongly positive.

## 7.2 External Validity

These findings are specific to a particular context: remote tribal and hill areas in India where baseline connectivity was extremely low. The 250 threshold targeted communities that were, by design, among the most isolated in the country. The returns to road connectivity in such settings may substantially exceed those in better-connected areas, consistent with the null results at the 500 threshold in plains villages.

Several features of the northeastern Indian context may limit generalizability. First, the terrain is unusually mountainous, making roads both more expensive and more transformative. Second, the ethnic composition (high tribal share, matrilineal societies in some states) may shape the gender-specific responses. Third, security conditions (insurgency in some northeastern states) may attenuate road construction effects by limiting economic activity even after connectivity improves.

However, the core finding—that infrastructure investment generates disproportionate returns for the most remote communities—is likely more general. Similar patterns have been documented for electrification ([Dinkelman, 2011](#)) and railroads ([Donaldson, 2018](#)), suggesting that the first unit of connectivity infrastructure is qualitatively different from marginal additions to an already-connected network. This principle has broad implications for the allocation of infrastructure investment in developing countries.

## 7.3 Policy Implications

The central policy lesson is that differential eligibility thresholds—lowering bars for more disadvantaged communities—can effectively channel infrastructure benefits where they are most needed. The PMGSY’s 250 threshold was a deliberate policy choice to extend connectivity to populations too small to qualify under the standard rule. The evidence suggests this design feature achieves its intended purpose, generating meaningful development returns in communities that would otherwise have been left behind for years or decades.

As India and other developing countries plan massive infrastructure investments for the coming decades, the evidence from PMGSY’s 250 threshold suggests that explicitly designing programs to reach the most remote populations yields returns that justify the additional cost and complexity. This finding is relevant for the design of programs under the

Sustainable Development Goals, which emphasize “leaving no one behind” and targeting the most vulnerable populations.

## 7.4 Limitations

Several limitations warrant mention. First, and most importantly, the design captures intent-to-treat effects of eligibility, not the direct impact of road construction. The PMGSY Online Management, Monitoring and Accounting System (OMMS) contains habitation-level road construction records, but linking these to SHRUG village identifiers requires fuzzy name-matching across administrative hierarchies that is beyond the scope of this analysis. Without first-stage evidence on the probability of road construction at the 250 threshold, I cannot estimate treatment-on-the-treated effects or verify that the eligibility discontinuity translates into a meaningful discontinuity in actual connectivity. The ITT interpretation—that the *eligibility rule itself* generates effects by prioritizing eligible villages in the construction queue—is the policy-relevant parameter, but readers should note this limitation when interpreting magnitudes.

Second, Census 2001 population is a discrete integer, creating mass points in the running variable. While I address this through donut-hole specifications and the mass-point-aware `rdrobust` implementation, some residual concerns remain about the appropriateness of continuous RDD methods for discrete running variables.

Third, the analysis uses administrative outcomes from the Census (literacy rates, worker shares) which may be subject to measurement error and may not capture all dimensions of welfare. Nightlights provide an independent, objectively measured outcome but are themselves a noisy proxy for economic activity, particularly at low luminosity levels characteristic of remote villages.

Fourth, the standard errors are heteroskedasticity-robust but do not account for potential spatial correlation. Outcomes such as nightlights and literacy rates may be correlated across nearby villages within the same district due to shared institutions, geography, and implementation capacity. Clustering at the district level or using Conley spatial HAC standard errors (Conley, 1999) could increase standard errors for some outcomes. The robustness of the primary results across multiple bandwidths and specifications provides some reassurance, but district-clustered inference is a natural extension for future work.

Fifth, I cannot fully rule out spillover effects. If road construction in eligible villages improved outcomes in nearby ineligible villages (e.g., through shared market access or demonstration effects), the estimates would be attenuated toward zero. Alternatively, if roads diverted economic activity from ineligible to eligible villages, the estimates would be inflated. Without data on village-level road proximity and network structure, I cannot sign or bound

these spillover effects.

Fifth, the absence of administrative data on actual road construction timing limits the analysis to ITT effects and prevents a direct examination of the road construction mechanism. Future work linking PMGSY construction data to SHRUG outcomes would substantially strengthen the causal chain.

## 8. Conclusion

India’s PMGSY rural road program created a natural experiment in development targeting. By lowering the eligibility threshold to 250 persons in tribal and hill areas—half the standard 500-person threshold—the program extended connectivity to communities that would otherwise have been left behind for years or decades. This paper exploits this threshold to estimate the causal effects of road eligibility on development outcomes in the most remote corners of India’s tribal periphery.

The results reveal that road eligibility generates meaningful and persistent development gains. Female literacy increases by 1.9 percentage points and satellite-measured economic activity rises by over 40% in the long run. These effects are absent at placebo thresholds and at the 500 threshold in plains areas, confirming their specificity to the designated-area eligibility rule.

The gender dimension of these findings deserves emphasis. In societies where girls’ education is constrained by distance, safety, and limited local economic opportunity, all-weather roads can serve as an equalizer. The disproportionate female literacy gains suggest that infrastructure investment complements targeted education interventions and may be among the most effective tools for closing gender gaps in remote areas. [Beaman et al. \(2012\)](#) and [Muralidharan and Prakash \(2017\)](#) have shown that targeted interventions—female leaders, bicycles for girls—can meaningfully improve girls’ outcomes in India. The road connectivity results suggest that basic infrastructure may be a prerequisite for these targeted interventions to achieve their full potential.

The long-run nightlight results are perhaps the most important finding from a policy perspective. The fact that economic effects are largest in the VIIRS era (2015–2023)—a full 14–22 years after program launch—suggests that infrastructure investment triggers slow-moving but powerful processes of economic transformation. Evaluations that measure outcomes only in the short or medium run may substantially understate the true returns to road connectivity, particularly in the most remote settings where the complementary investments needed to capitalize on improved connectivity take time to materialize.

For policymakers, the central lesson is that differential eligibility thresholds for more

disadvantaged communities can effectively channel infrastructure benefits where they are most needed. As developing countries worldwide invest in rural infrastructure under the Sustainable Development Goals, the experience of PMGSY’s 250 threshold provides evidence that reaching the most remote populations—even at higher per capita cost—generates development returns that more than justify the investment.

## Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). All data come from the SHRUG v2.1 platform (Asher et al., 2021). I thank the SHRUG team—Sam Asher, Tobias Lunt, Ryu Matsuura, and Paul Novosad—for making high-resolution Indian village data publicly available.

**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.
- Banerjee, Abhijit, Esther Duflo, and Nancy Qian**, “On the Road: Access to Transportation Infrastructure and Economic Growth in China,” *Journal of Development Economics*, 2020, 145, 102442.
- Beaman, Lori, Esther Duflo, Rohini Pande, and Petia Topalova**, “Female Leadership Raises Aspirations and Educational Attainment for Girls: A Policy Experiment in India,” *Science*, 2012, 335 (6068), 582–586.
- 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., Brigham R. Frandsen, and Rocío Titiunik**, “Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate,” *Journal of Causal Inference*, 2015, 3 (1), 1–24.
- Cattaneo, Matias D, Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, 115 (531), 1449–1455.
- , **Nicolas Idrobo, and Rocio Titiunik**, “A Practical Introduction to Regression Discontinuity Designs: Foundations,” *Cambridge Elements: Quantitative and Computational Methods for Social Science*, 2020.

- Conley, Timothy G.**, “GMM Estimation with Cross Sectional Dependence,” *Journal of Econometrics*, 1999, *92* (1), 1–45.
- Dinkelman, Taryn**, “The Effects of Rural Electrification on Employment: New Evidence from South Africa,” *American Economic Review*, 2011, *101* (7), 3078–3108.
- Donaldson, Dave**, “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure,” *American Economic Review*, 2018, *108* (4-5), 899–934.
- **and Adam Storeygard**, “The View from Above: Applications of Satellite Data in Economics,” *Journal of Economic Perspectives*, 2016, *30* (4), 171–198.
- Duflo, Esther**, “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 2001, *91* (4), 795–813.
- 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.
- 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.
- Jensen, Robert**, “Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector,” *Quarterly Journal of Economics*, 2007, *122* (3), 879–924.
- Lee, David S**, “Randomized Experiments from Non-Random Selection in US 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.
- Lehne, Jonathan, Jacob N Shapiro, and Oliver Vanden Eynde**, “Building Connections: Political Corruption and Road Construction in India,” *Journal of Development Economics*, 2018, 131, 62–78.
- 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.
- Munshi, Kaivan and Mark Rosenzweig**, “Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy,” *American Economic Review*, 2006, 96 (4), 1225–1252.
- Muralidharan, Karthik and Nishith Prakash**, “Cycling to School: Increasing Secondary School Enrollment for Girls in India,” *American Economic Journal: Applied Economics*, 2017, 9 (3), 321–350.
- Redding, Stephen J and Esteban Rossi-Hansberg**, “Quantitative Spatial Economics,” *Annual Review of Economics*, 2017, 9, 21–58.
- 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.

## A. Data Appendix

### A.1 SHRUG Data Platform

The Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG) provides harmonized village and town-level data for India, linking records across Census rounds using consistent geographic identifiers. Version 2.1, released by the Development Data Lab, includes the following datasets used in this study:

- **Census 2001 Primary Census Abstract (PCA):** 593,795 observations covering total population, literacy by gender, SC/ST populations, workforce by type and gender, and household counts at the village/town level.
- **Census 2011 PCA:** 596,393 observations with equivalent post-treatment variables.
- **SHRUG-DMSP Nightlights:** Annual calibrated nighttime luminosity for each SHRUG location, 1994–2013. Approximately 19 million location-year observations. The DMSP variable used is `dmsp_total_light_cal` (calibrated total light).
- **SHRUG-VIIRS Nightlights:** Annual nighttime luminosity from VIIRS sensors, 2012–2023. Approximately 14 million location-year observations. The VIIRS variable used is `viirs_annual_sum`.
- **District Key:** Mapping from SHRUG identifiers to Census 2011 state and district codes. 596,508 observations. Fields: `shrid2`, `pc11_state_id`, `pc11_district_id`.

SHRUG data are freely available at [https://www.devdatalab.org/shrug\\_download/](https://www.devdatalab.org/shrug_download/).

### A.2 Variable Definitions

[Table 4](#) defines all variables used in the analysis.

**Table 4:** Variable Definitions

Variable	Type	Definition
pop_01	Running	Census 2001 total population
rv_250	Running	pop_01 - 250
designated_A	Treatment	1 if state code $\in$ Special Category
literacy_11	Outcome	Literate / total pop (Census 2011)
f_lit_11	Outcome	Literate females / total females
m_lit_11	Outcome	Literate males / total males
gender_lit_gap_11	Outcome	m_lit_11 - f_lit_11
nonag_share_11	Outcome	1 - (cultivators + ag laborers) / workers
f_worker_share_11	Outcome	Female workers / female population
pop_growth	Outcome	(pop_11 - pop_01) / pop_01
log_nl_post_dmisp	Outcome	log(mean DMSP 2005-2013 + 0.01)
log_nl_post_viirs	Outcome	log(mean VIIRS 2015-2023 + 0.01)
log_nl_pre	Covariate	log(mean DMSP 1994-2000 + 0.01)

### A.3 Sample Filters

**Table 5:** Sample Construction: Sequential Filters

Step	Remaining	Dropped
Census 2001 PCA (all villages/towns)	593,795	—
Merge district key (for state codes)	587,378	6,417
Filter to Special Category States	75,867	511,511
Filter to population 50-500	41,371	34,496

The 6,417 observations lost in the district key merge represent villages with SHRIDs present in Census 2001 but absent from the Census 2011 district key, likely due to administrative boundary changes or settlement reorganization between census rounds.

## B. Identification Appendix

### B.1 McCrary Density Test

Figure 3 shows the density of Census 2001 population in Special Category State villages around the 250 threshold. The distribution is smooth, with no visible bunching at the cutoff. The formal density test of Cattaneo et al. (2020a) yields:

- Sample A:  $t = 0.604$ ,  $p = 0.546$
- Sample B:  $t = -0.337$ ,  $p = 0.736$

These results strongly fail to reject the null of no manipulation.

### B.2 Covariate Balance

Figure 5 displays the RDD estimates and 95% confidence intervals for seven pre-treatment covariates. All confidence intervals comfortably include zero. The estimates and standard errors are:

**Table 6:** Covariate Balance: Detailed Results

Pre-Treatment Variable	Estimate	Robust SE	$p$ -value
Literacy Rate (2001)	0.0003	0.0085	0.967
Female Literacy (2001)	0.0004	0.0089	0.966
SC Share (2001)	0.0112	0.0104	0.280
ST Share (2001)	0.0250	0.0202	0.216
Worker Share (2001)	-0.0024	0.0060	0.684
Female Worker Share (2001)	-0.0039	0.0086	0.650
Log Nightlights (1994–2000)	0.0650	0.0705	0.356

### B.3 Placebo Threshold Results

**Table 7:** Placebo Thresholds: Literacy Rate (2011)

Threshold	Estimate	Robust SE	<i>p</i> -value
150	0.0088	0.0088	0.317
200	-0.0064	0.0080	0.425
<b>250 (true)</b>	<b>0.0128</b>	<b>0.0079</b>	<b>0.105</b>
300	0.0002	0.0063	0.970
350	0.0023	0.0084	0.785
400	0.0040	0.0088	0.649

## C. Robustness Appendix

### C.1 Bandwidth Sensitivity: Detailed Results

**Table 8:** Bandwidth Sensitivity: Literacy Rate (2011), Sample A

BW Multiplier	Bandwidth	Estimate	Robust SE	<i>p</i> -value
0.50	22	0.0110	0.0145	0.445
0.75	33	0.0200	0.0117	0.087
1.00	44	0.0205	0.0100	0.041
1.25	55	0.0174	0.0090	0.053
1.50	66	0.0142	0.0082	0.084
2.00	89	0.0100	0.0071	0.160

### C.2 Donut-Hole Estimates

Excluding 1,018 villages with population 245–255 (2.5% of the primary sample):

**Table 9:** Donut-Hole RDD Estimates ( $|pop - 250| > 5$ ), Sample A

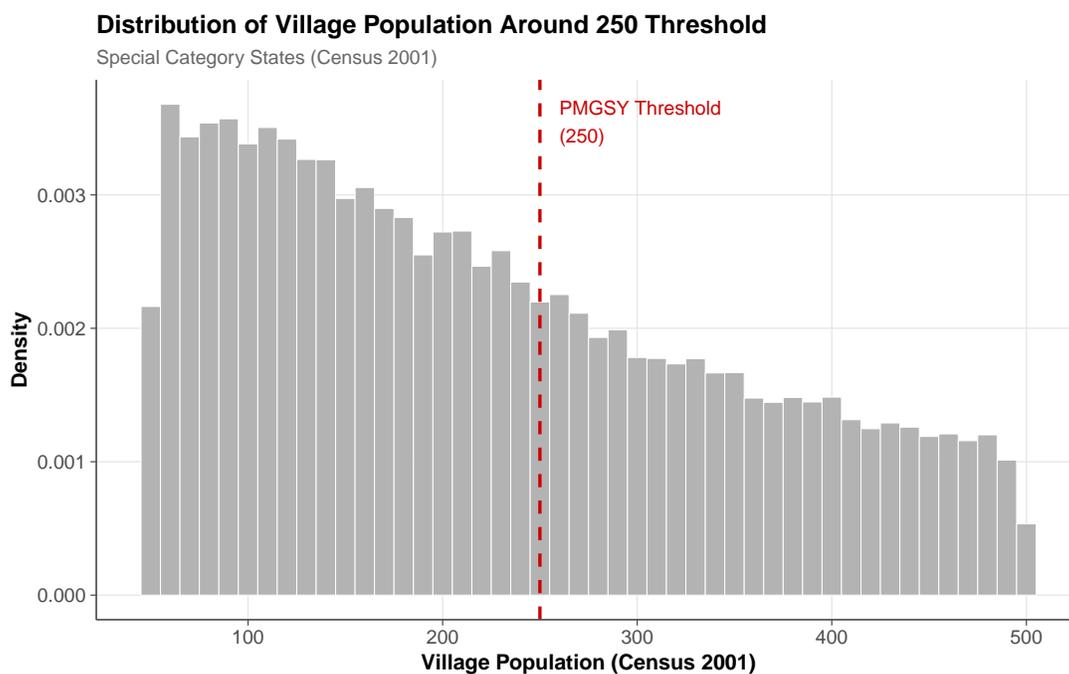
Outcome	Estimate	Robust SE	<i>p</i> -value
Literacy (2011)	0.0074	0.0109	0.496
Female Literacy (2011)	0.0122	0.0119	0.303
Non-Ag Worker Share (2011)	0.0325	0.0285	0.254
Log Nightlights (VIIRS)	0.2015	0.1772	0.255

### C.3 Polynomial Order Sensitivity

**Table 10:** Local Linear vs. Local Quadratic Estimates, Sample A

Outcome	Linear ( $p=1$ )	$p$ -value	Quadratic ( $p=2$ )	$p$ -value
Literacy	0.0128	0.105	0.0167	0.068
Female Literacy	0.0189	0.032	0.0221	0.028
Non-Ag Share	0.0287	0.139	0.0379	0.105
Log NL (VIIRS)	0.3440	0.004	0.3995	0.003

### D. Additional Figures



**Figure 3:** Population Density Around the 250 Threshold

### Gender and Demographic Outcomes at 250 Threshold

Special Category States ... Bin scatter with local linear fit

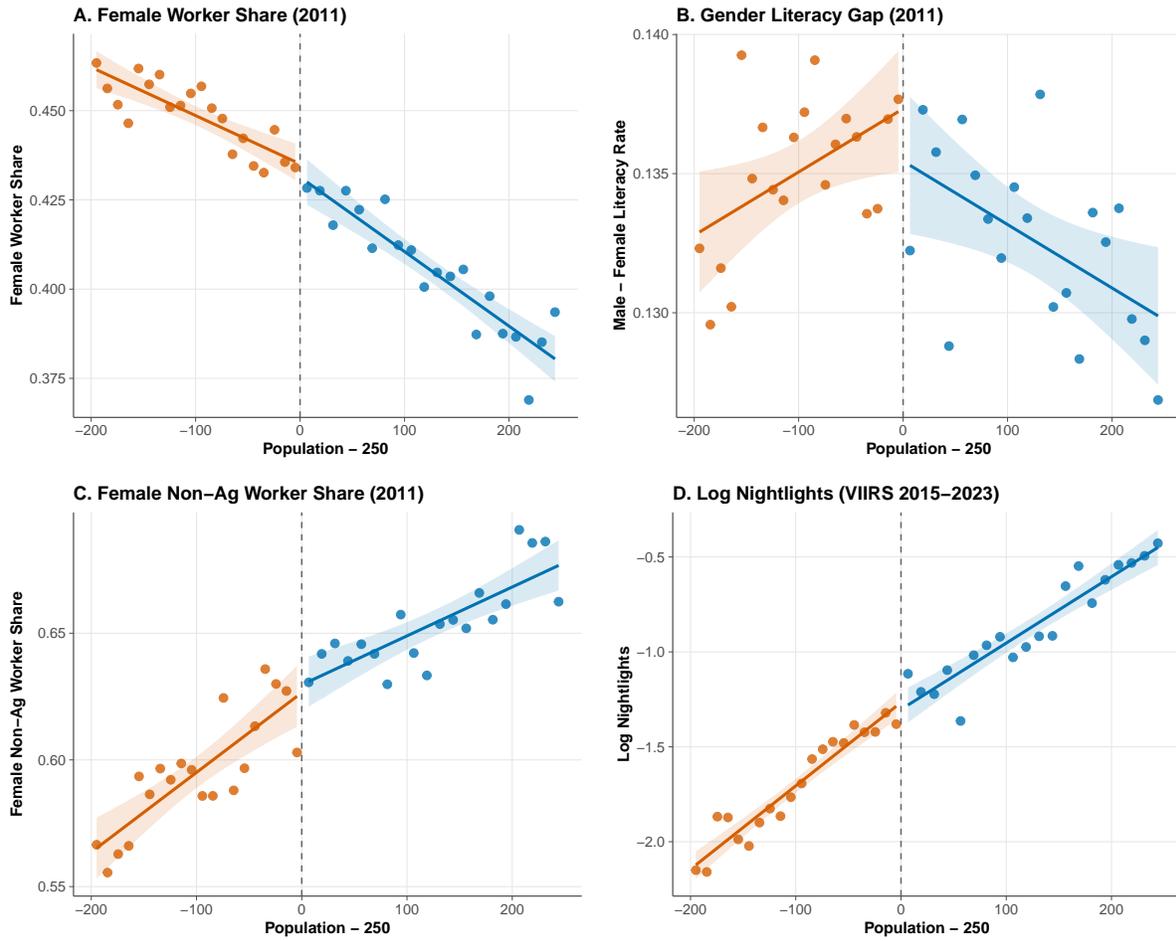
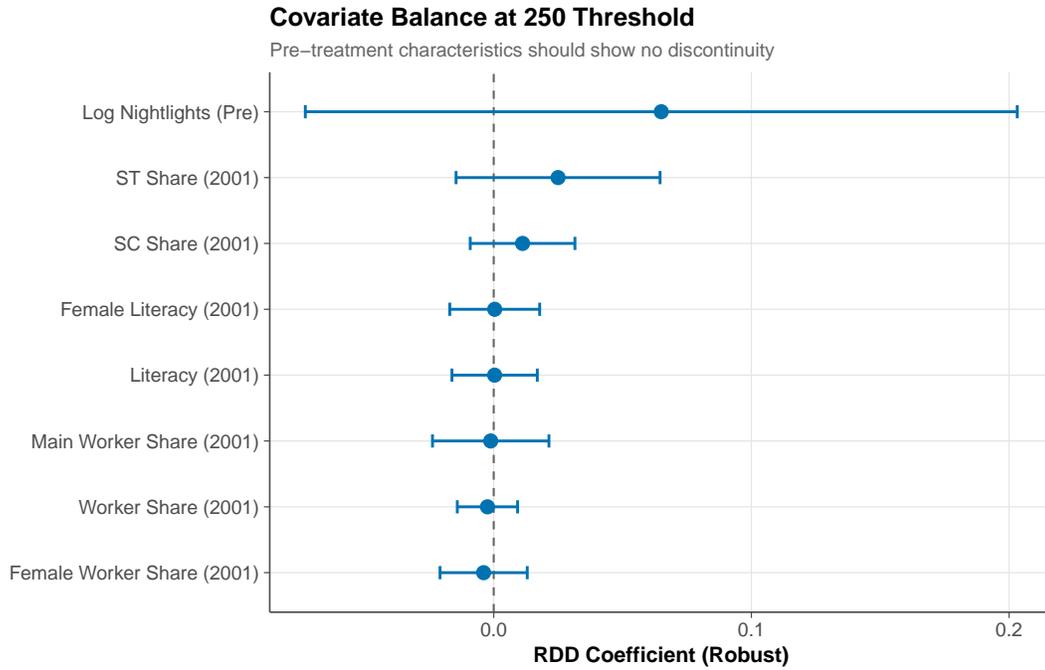
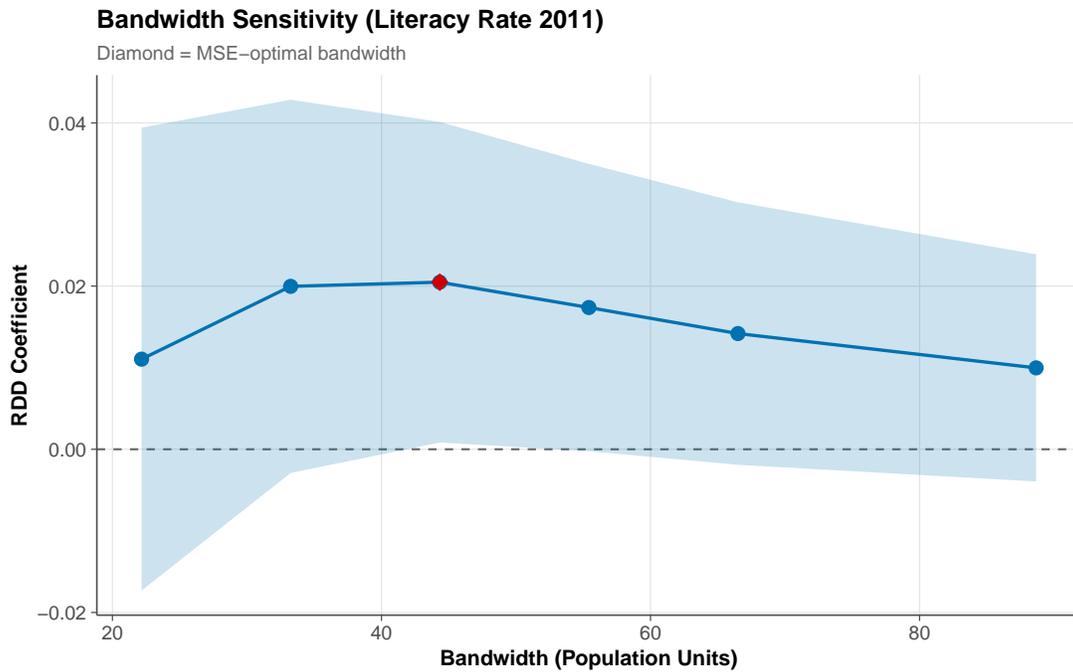


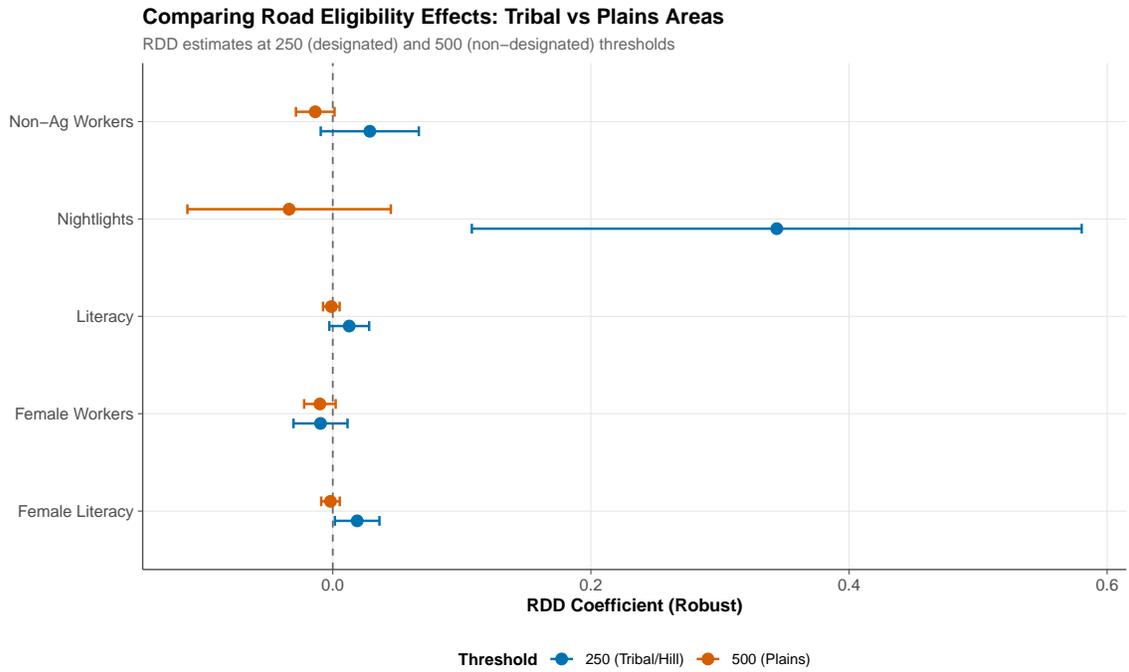
Figure 4: RDD Bin Scatters: Gender Outcomes and Nightlights



**Figure 5:** Covariate Balance at 250 Threshold



**Figure 6:** Bandwidth Sensitivity: Literacy Rate at 250 Threshold



**Figure 7:** Comparison of RDD Estimates: 250 vs. 500 Threshold