

Roads Without Revolution: Rural Connectivity and the Gender Gap in India's Structural Transformation

APEP Autonomous Research*

February 21, 2026

Abstract

India's female labor force participation has fallen even as the country rapidly urbanized. Can rural infrastructure close the gender gap in structural transformation? I exploit the population eligibility threshold of the Pradhan Mantri Gram Sadak Yojana (PMGSY), India's flagship rural road program, in a regression discontinuity design using village-level Census data on 582,807 rural villages. The results are precisely estimated nulls: crossing the 500-person eligibility threshold produces no detectable effect on women's shift out of agriculture (0.2 percentage points, $p = 0.44$), men's structural transformation (-0.1 pp, $p = 0.71$), or the gender gap ($p = 0.83$). These nulls survive bandwidth variation, donut-hole exclusions, polynomial sensitivity, placebo thresholds, and randomization inference ($p = 0.50$). The findings suggest that road connectivity alone—without complementary interventions addressing social norms, skills, and labor demand—is insufficient to transform gendered employment patterns in rural India.

JEL Codes: J16, O18, R42, O15

Keywords: rural roads, structural transformation, gender, India, PMGSY, regression discontinuity

*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 14h 33m (cumulative: 14h 33m). Contributor: @olafdrw. Correspondence: scl@econ.uzh.ch

1. Introduction

Between 2001 and 2011, India’s female labor force participation rate fell from 26% to 22%, even as the economy grew at 7% annually and millions of workers shifted from farms to factories (Klasen and Pieters, 2018). This paradox—growth without women—has puzzled economists and troubled policymakers. One prominent explanation focuses on infrastructure: women face greater mobility constraints than men, so poor road connectivity may trap them in low-productivity agricultural work while men commute to non-farm employment (Lei et al., 2019). If true, building rural roads should disproportionately benefit women’s structural transformation.

This paper tests that hypothesis directly. I exploit the population eligibility threshold of the Pradhan Mantri Gram Sadak Yojana (PMGSY), India’s largest rural road construction program. PMGSY uses explicit population cutoffs—500 persons for plain-area states, 250 for special-category states—to prioritize habitations for all-weather road connectivity. Asher and Novosad (2020) document a sharp 22 percentage point increase in the probability of road construction at this threshold, establishing a strong first stage at the habitation level. Using village-level data from the SHRUG platform (Asher et al., 2021) covering 582,807 rural villages, I estimate a reduced-form regression discontinuity design that asks: does crossing the PMGSY eligibility threshold differentially shift women out of agriculture?

The answer is no. The estimated effect on the change in female non-agricultural employment share between Census 2001 and 2011 is 0.2 percentage points with a robust standard error of 0.4 percentage points ($p = 0.44$). The effect on male non-agricultural share is -0.1 percentage points ($p = 0.71$). The gender gap—the difference between female and male structural transformation—shows no discontinuity ($p = 0.83$). These are not imprecise nulls: the 95% confidence interval for the female effect rules out impacts larger than roughly one percentage point in either direction, against a baseline female non-agricultural share of roughly 11%.

The null results are robust across every specification I consider. Bandwidth sensitivity analysis shows no significant effects at the CCT optimal bandwidth or wider; marginal significance appears only at very narrow bandwidths that sacrifice precision. Donut-hole regressions excluding villages within ± 5 , ± 10 , and ± 20 persons of the threshold yield null results. Placebo regressions at cutoffs of 300, 400, 600, 700, and 800 persons produce no spurious discontinuities, confirming that the null at 500 is not an artifact of functional form. Higher-order polynomial specifications and randomization inference ($p = 0.50$ from 500 permutations) reinforce the finding. Regional heterogeneity analysis across North, South, East, and West India reveals no significant effects in any subregion. The separate analysis

using the 250 threshold for special-category states is also null.

The design passes standard validity checks. A McCrary density test yields a t -statistic of 1.29 ($p = 0.198$), showing no evidence of manipulation of village population at the threshold. Covariate balance tests for literacy rates, household counts, SC/ST population shares, and gender composition all show smooth crossing at 500. The running variable—Census 2001 village population—was recorded before PMGSY’s implementation began, ruling out post-treatment sorting.

Why might roads fail to catalyze gendered structural transformation? I consider three interpretations. First, the village-level analysis may attenuate the habitation-level first stage documented by [Asher and Novosad \(2020\)](#). PMGSY targets habitations, not villages, and a village’s total population imperfectly predicts whether its constituent habitations cross the eligibility threshold. This measurement concern is real but does not invalidate the reduced-form RDD—it means the estimates capture the intent-to-treat effect of being in a village above the threshold, which is the policy-relevant parameter for village-level targeting. Second, even if roads improve connectivity, the binding constraints on women’s non-farm employment may lie elsewhere: social norms restricting women’s mobility ([Jayachandran, 2021](#); [Afridi et al., 2023](#)), human capital deficits, or the absence of suitable non-agricultural employment within commuting distance. Third, the decade between Census 2001 and 2011 may be too short for road-induced structural transformation to manifest in employment statistics, particularly for women whose labor supply decisions are shaped by household bargaining and cultural constraints that evolve slowly.

This paper contributes to three literatures. First, it advances the large body of work on transportation infrastructure and economic development. While [Donaldson \(2018\)](#) finds large effects of railroads in colonial India and [Banerjee et al. \(2020\)](#) documents growth effects of highways in China, the evidence on rural “last-mile” roads is more mixed. [Asher and Novosad \(2020\)](#) find that PMGSY roads reduce travel times and shift labor out of agriculture at the habitation level, but [Aggarwal \(2018\)](#) finds no poverty reduction effects using a fuzzy RDD. My results complement [Asher and Novosad \(2020\)](#) by showing that the habitation-level structural transformation documented in their study does not translate into village-level gender-differentiated employment shifts detectable in Census data. This gap between habitation-level treatment and village-level outcomes is itself informative about the spatial spillovers—or lack thereof—within rural settlements.

Second, I contribute to the growing literature on gender and structural transformation in developing countries. [Lei et al. \(2019\)](#) document a correlation between road proximity and women’s non-farm employment in India using household survey data, but their fixed-effects design cannot isolate the causal effect of road construction from location-specific

unobservables. [Dinkelman \(2011\)](#) finds that rural electrification in South Africa increased female employment, and [Muralidharan and Prakash \(2017\)](#) shows that providing bicycles to girls increased school enrollment in India. My null result suggests that not all infrastructure investments are created equal: roads may reduce transportation costs without affecting the gender-specific barriers to non-agricultural employment. This distinction has direct implications for the design of complementary interventions.

Third, this paper contributes a well-identified null result to the economics literature. Publication bias favors positive findings, yet null results from credible research designs are essential for cumulative knowledge ([Gelman and Imbens, 2019](#)). The PMGSY threshold provides one of the cleanest regression discontinuity designs available for studying rural infrastructure in India—sharp eligibility rules, a pre-treatment running variable, strong density and covariate balance, and the statistical power afforded by a half-million villages. The fact that this design produces precisely estimated zeros for gender-specific structural transformation is a genuine empirical finding, not a failure of identification.

The paper proceeds through the institutional background of PMGSY ([Section 2](#)), a conceptual framework ([Section 3](#)), data and sample construction ([Section 4](#)), the empirical strategy ([Section 5](#)), results and robustness ([Section 6](#)), discussion of the null findings ([Section 7](#)), and conclusions ([Section 8](#)).

2. Institutional Background

2.1 The Pradhan Mantri Gram Sadak Yojana (PMGSY)

India launched the Pradhan Mantri Gram Sadak Yojana—the Prime Minister’s Rural Road Scheme—in December 2000 with the goal of providing all-weather road connectivity to every unconnected rural habitation with a population above 500 persons in plain areas and 250 persons in special-category states (the northeastern states, Himachal Pradesh, Jammu and Kashmir, and Uttarakhand) ([Ministry of Rural Development, 2000](#)). The program was massive in ambition and scale: by 2015, over 400,000 kilometers of roads had been constructed or upgraded, connecting roughly 100,000 habitations at a cost exceeding \$30 billion.

The eligibility rules create a natural experiment. The population threshold determines the order in which habitations receive road construction: larger habitations above the cutoff were prioritized for connectivity in earlier phases, while smaller habitations below the cutoff were deferred to later phases or remain unconnected. Critically, the population figures used for prioritization came from the Census of 2001, which was conducted before the program began ([Asher and Novosad, 2020](#)). This means the running variable—Census 2001 habitation population—is pre-determined with respect to the treatment, ruling out strategic

manipulation of the assignment variable by program administrators or beneficiaries.

2.2 Habitations versus Villages

An important institutional detail is that PMGSY targets habitations, not villages. Indian villages are administrative units that may contain multiple habitations—physically distinct clusters of dwellings. A village with a total population of 800 might contain three habitations of 400, 250, and 150 persons, none of which individually crosses the 500-person threshold. Conversely, a village with total population of 450 might contain a single habitation of 450 that falls just below the cutoff. This means a village’s total Census population is a noisy proxy for the habitation-level eligibility that drives road construction.

This distinction has direct implications for the empirical strategy. A village-level reduced-form RDD at the 500 threshold captures the intent-to-treat effect of being in a village whose total population crosses the threshold. Because larger villages are more likely to contain habitations above 500, there is a first stage—but it is attenuated relative to the habitation-level discontinuity documented by [Asher and Novosad \(2020\)](#). The village-level approach is nonetheless informative because (a) habitation-level population data from the PMGSY management system are not publicly available in geocoded form, (b) employment outcomes are available at the village level from the Census, and (c) the village-level intent-to-treat parameter is the relevant policy object if one is considering village-level targeting rules.

2.3 PMGSY and Gender

The stated objectives of PMGSY are framed in terms of general rural connectivity, not gender equality. However, the program was expected to have gendered consequences through multiple channels. Road connectivity reduces the time and physical cost of reaching markets, schools, and health facilities. Because women in rural India face greater mobility constraints than men—due to safety concerns, social norms around female travel, and the time burden of domestic work—improved road access could differentially benefit women by expanding their accessible labor market ([Lei et al., 2019](#)).

At the same time, standard models of structural transformation predict that roads facilitate commuting, and Indian men are far more likely than women to commute for non-agricultural work. If roads primarily enable male out-migration and long-distance commuting while women remain in the village, the gender gap in non-farm employment could widen rather than narrow. The theoretical prediction is ambiguous, making this an empirical question.

2.4 India’s Structural Transformation, 2001–2011

Between Census 2001 and Census 2011, India underwent substantial structural transformation. The share of workers in non-agricultural employment rose from approximately 35% to 48%. However, this aggregate trend masks sharp gender differences. Male non-agricultural employment shares increased by roughly 12 percentage points, while female non-agricultural shares grew by only 4 percentage points. In many states, female labor force participation actually declined even as the economy grew (Klasen and Pieters, 2018).

This period coincides almost exactly with the major rollout of PMGSY. The first phase targeted habitations above 1,000 persons (in plain areas), with subsequent phases reaching smaller habitations. By 2008, the program had expanded to include habitations above 500 in plain areas. The Census 2001–2011 window thus captures the primary treatment period for the threshold I study.

3. Conceptual Framework

I outline a simple framework to organize the channels through which road connectivity might affect gendered structural transformation. Consider a village with a fixed number of workers indexed by i , each of whom chooses between agricultural work (with return w_i^A) and non-agricultural work (with gross return w_i^N). Accessing non-agricultural work requires paying a commuting cost c_i that depends on road quality and individual characteristics.

Worker i chooses non-agricultural employment if $w_i^N - c_i > w_i^A$. Road construction reduces commuting costs: $c_i = c_i^0 - \delta \cdot R$, where R indicates road connectivity and $\delta > 0$ is the reduction in commuting cost. The village non-agricultural share is then:

$$S = \Pr(w_i^N - c_i^0 + \delta R > w_i^A) \tag{1}$$

The key gender dimension arises from two sources. First, commuting costs may differ by gender: c_i^0 is likely higher for women due to safety concerns, social norms restricting women’s independent travel, and the time burden of domestic responsibilities that limits women’s available commuting time (Jayachandran, 2021). Second, the returns to non-agricultural work may differ: w_i^N may be lower for women due to discrimination, skill gaps, or the types of non-agricultural work available in nearby towns.

Roads affect structural transformation through the commuting cost channel. For women, two scenarios are possible:

Scenario 1: Roads disproportionately benefit women. If the binding constraint is physical access rather than social norms, and if $\partial c_i / \partial R$ is larger for women (because they have higher

baseline costs and the road removes a larger share of them), then roads would generate a larger increase in S_F than S_M , narrowing the gender gap.

Scenario 2: Roads primarily benefit men. If social norms restrict women’s labor supply regardless of road quality, then $\partial c_i / \partial R \approx 0$ for women and the gains accrue to men who are already on the margin of commuting. Roads would then widen the gender gap.

Scenario 3: Roads affect neither. If the threshold affects habitation-level eligibility imperfectly (as discussed in [Section 2](#)), if the non-agricultural employment opportunities within commuting distance are limited, or if the decade between censuses is insufficient for adjustment, then road construction near the threshold may produce null effects for both genders.

These scenarios generate testable predictions. Scenario 1 predicts a positive discontinuity in the female non-agricultural share and a larger positive discontinuity for women than men. Scenario 2 predicts a positive discontinuity for men and a null or negative discontinuity for the gender gap. Scenario 3 predicts null effects for all outcomes. I test all three in the empirical analysis.

4. Data

4.1 Data Sources

I use five datasets from the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG), version 2.1 ([Asher et al., 2021](#)). SHRUG harmonizes India’s village-level Census data across decades using a consistent geographic identifier (SHRID), enabling longitudinal analysis despite administrative boundary changes.

Census 2001 Primary Census Abstract (PCA). This provides the running variable (village population in 2001) and baseline covariates including literacy rates, caste composition (Scheduled Caste and Scheduled Tribe shares), gender composition, and employment by category. Worker classification in the Census distinguishes main and marginal workers, and further disaggregates by cultivators, agricultural laborers, household industry workers, and other workers. I observe 593,246 village-level observations ([Office of the Registrar General, 2001](#)).

Census 2011 Primary Census Abstract. This provides outcome variables: village-level employment by gender and sector in 2011. The same worker classification scheme allows construction of comparable non-agricultural employment shares across the two Census rounds ([Office of the Registrar General, 2011](#)).

Geographic crosswalk. The SHRUG district-level key provides Census 2011 state and district codes, which I use to identify special-category states and define Indian regions for

the heterogeneity analysis.

Rural village key. I restrict the sample to rural villages using the Census 2001 rural village identifier, which maps to the SHRID geographic unit.

Nighttime lights. I use DMSP-OLS nighttime lights data (1994–2013) and VIIRS annual composites (2014–2023) for a dynamic year-by-year RDD analysis of economic activity (Henderson et al., 2012). Both are aggregated at the village level by SHRUG. I use DMSP for years 1994–2013 and VIIRS for 2014–2023, with no overlap years in the analysis.

4.2 Variable Construction

Running variable. Census 2001 total village population, centered at the 500-person threshold for plain-area states.

Treatment indicator. $D_i = \mathbf{1}[\text{Pop}_{i,2001} \geq 500]$.

Primary outcomes. For each village, I compute the non-agricultural share of main workers separately by gender:

$$\text{NonAg}_{i,t}^g = \frac{\text{HH Industry Workers}_{i,t}^g + \text{Other Workers}_{i,t}^g}{\text{Total Main Workers}_{i,t}^g} \quad (2)$$

for gender $g \in \{F, M\}$ and Census year $t \in \{2001, 2011\}$. The primary outcome is the change $\Delta \text{NonAg}_i^g = \text{NonAg}_{i,2011}^g - \text{NonAg}_{i,2001}^g$, which differences out time-invariant village-level confounders.

I also construct the gender gap in structural transformation: $\Delta \text{Gap}_i = \Delta \text{NonAg}_i^F - \Delta \text{NonAg}_i^M$.

Workforce participation rates. Total female (male) workers divided by total female (male) population. Changes in these rates capture extensive-margin employment effects.

Special-category state indicator. Jammu and Kashmir, Himachal Pradesh, Uttarakhand, and all northeastern states (Arunachal Pradesh, Nagaland, Manipur, Mizoram, Tripura, Meghalaya, Assam, Sikkim) are classified as special-category states with a lower PMGSY threshold of 250 persons.

4.3 Sample Construction

Starting from 593,246 Census 2001 village observations in SHRUG, I apply the following restrictions:

1. *Rural villages only.* I restrict to villages identified as rural in Census 2001 using the rural village key, yielding 591,668 villages.

2. *Merge across Census rounds.* Matching Census 2001 and 2011 villages via SHRID produces a panel of 582,807 rural villages with non-missing population.
3. *Plain-area states.* The main analysis uses the 511,648 villages in non-special-category states, where the PMGSY threshold is 500 persons.
4. *Near-threshold sample.* For summary statistics, I define a “near-threshold” sample of villages within ± 200 persons of the cutoff, comprising 134,022 villages.

The RDD estimation uses all villages in plain-area states; the data-driven CCT bandwidth selector determines the effective estimation sample.

4.4 Summary Statistics

Table 1 reports means and standard deviations for key variables among villages within 200 persons of the threshold, separately for villages above and below 500. Villages above the threshold are mechanically larger (mean population 598 vs. 402), but the two groups are strikingly similar in composition. Literacy rates are 46.5% below and 47.3% above. SC shares are 17.5% vs. 18.4%, and ST shares are 24.2% vs. 20.2%. Female workforce participation rates are nearly identical: 38.5% below and 37.0% above. The similarity of these baseline characteristics supports the identifying assumption that villages near the threshold are comparable.

5. Empirical Strategy

5.1 Regression Discontinuity Design

I estimate the reduced-form effect of PMGSY eligibility on structural transformation using a regression discontinuity design. The estimating equation is:

$$\Delta Y_i = \alpha + \tau \cdot \mathbf{1}[\text{Pop}_{i,2001} \geq 500] + f(\text{Pop}_{i,2001} - 500) + \varepsilon_i \quad (3)$$

where ΔY_i is the change in a gender-specific employment outcome between Census 2001 and 2011, $\text{Pop}_{i,2001}$ is the Census 2001 village population, and $f(\cdot)$ is a local polynomial estimated separately on each side of the cutoff. The parameter of interest is τ , which captures the discontinuity in the outcome at the population threshold.

The identifying assumption is continuity of potential outcomes at the threshold:

$$\lim_{x \downarrow 500} \mathbb{E}[\Delta Y_i(0) \mid \text{Pop}_i = x] = \lim_{x \uparrow 500} \mathbb{E}[\Delta Y_i(0) \mid \text{Pop}_i = x] \quad (4)$$

Table 1: Summary Statistics: Villages Near PMGSY Threshold

Variable	Below 500		Above 500	
	Mean	SD	Mean	SD
Population (2001)	397.265	57.645	595.868	57.830
Households	74.954	17.277	111.047	22.258
Literacy rate	0.465	0.168	0.473	0.156
SC share	0.175	0.225	0.184	0.207
ST share	0.242	0.359	0.202	0.322
Female share	0.491	0.031	0.489	0.028
Female WF participation	0.385	0.218	0.370	0.215
Male WF participation	0.534	0.090	0.532	0.084
Female non-ag share (2001)	0.106	0.205	0.113	0.201
Male non-ag share (2001)	0.172	0.184	0.181	0.179
Villages	72,817		61,205	

Notes: Sample restricted to plain-area states (non-special-category) with village population within ± 200 of the PMGSY threshold of 500. Data from SHRUG v2.1 (Asher, Novosad, Lunt). Literacy rate, SC/ST share, and worker participation rates computed from Census 2001. Non-agricultural share is the fraction of main workers in household industry and other non-agricultural occupations.

This requires that no other determinant of structural transformation changes discontinuously at 500. I provide evidence supporting this assumption through density tests and covariate balance checks.

5.2 Estimation Details

I implement the estimator using the `rdrubust` package (Calonico et al., 2015, 2014), which provides:

- *Bandwidth selection:* The CCT optimal bandwidth balances bias and variance, adapting to the local curvature of the conditional expectation function. I report results at the optimal bandwidth and assess sensitivity across a range of multipliers ($0.5\times$ to $2.0\times$).
- *Bias correction:* Robust bias-corrected confidence intervals account for the bias inherent in local polynomial estimation near boundaries (Calonico et al., 2014).
- *Kernel:* Triangular kernel, which gives greater weight to observations closer to the cutoff, as recommended by Imbens and Lemieux (2008).
- *Polynomial order:* Local linear ($p = 1$) as the primary specification, following Gelman and Imbens (2019). I check sensitivity to quadratic and cubic specifications.

Standard errors and p -values throughout the paper are the robust bias-corrected versions from [Calonico et al. \(2014\)](#), unless otherwise noted.

Statistical power. With a robust standard error of 0.0043 for the female non-agricultural share, the minimum detectable effect at 80% power and 5% significance is approximately $2.8 \times 0.0043 \approx 0.012$, or 1.2 percentage points. Against a baseline of 11%, this means the design can detect effects as small as 11% of the baseline mean. The null result is therefore informative: the data have sufficient power to rule out economically meaningful effects.

5.3 Interpretation: Intent-to-Treat at the Village Level

Because PMGSY targets habitations rather than villages, the village-level RDD at 500 estimates a reduced-form intent-to-treat parameter. This parameter captures the combined effect of (a) the change in the probability that habitations within the village receive road construction, and (b) any within-village spillovers from connected to unconnected habitations. If larger villages are more likely to contain habitations crossing the 500-person threshold, the intent-to-treat effect should be positive—but smaller in magnitude than the habitation-level treatment effect documented by [Asher and Novosad \(2020\)](#).

5.4 Threats to Validity

Manipulation. The running variable is Census 2001 village population, recorded before PMGSY’s launch. While there could be incentives for villages to inflate population counts to secure road eligibility, any such manipulation would have had to occur before the program was announced. I test for manipulation using the [Cattaneo et al. \(2020\)](#) density test.

Compound treatment. Other government programs also use population thresholds for eligibility. If another program uses the same 500-person cutoff, the estimated discontinuity would capture the combined effect of both programs. To my knowledge, PMGSY is the primary program using this specific cutoff, though I cannot entirely rule out other programs with similar thresholds.

Spillovers. If road construction in villages above the threshold affects employment in nearby villages below the threshold (e.g., by improving market access for neighboring villages), the control group would be contaminated and the RDD would underestimate the true effect. Spillovers would bias the estimate toward zero, making the null results potentially conservative.

Measurement. Census employment categories are broad and self-reported. The distinction between “agricultural laborer” and “other worker” may not perfectly capture the structural transformation I aim to measure. Any measurement error in outcomes is unlikely to differ discontinuously at the threshold, however, so it would attenuate estimates rather than bias

them.

5.5 Dynamic RDD: Nighttime Lights

As a complementary analysis, I estimate year-by-year RDD effects on nighttime light intensity:

$$\log(\text{NL}_{i,t} + 0.01) = \alpha_t + \tau_t \cdot \mathbf{1}[\text{Pop}_{i,2001} \geq 500] + f_t(\text{Pop}_{i,2001} - 500) + \varepsilon_{i,t} \quad (5)$$

for each year t from 1994 to 2023. Estimates of τ_t for $t < 2000$ (before PMGSY) serve as a placebo: if the design is valid, pre-program effects should be zero. Post-2000 effects capture the dynamic response to road connectivity. This analysis uses DMSP data through 2013 and VIIRS data from 2014 onward. Nighttime lights extend beyond the Census employment window (2001–2011), providing a complementary measure of local economic activity that can detect effects both within and beyond the decade covered by employment data.

6. Results

6.1 Validity of the Research Design

Density test. [Figure 1](#) shows the distribution of village population near the 500-person threshold. The histogram reveals no visible discontinuity in the density at the cutoff. Formally, the [Cattaneo et al. \(2020\)](#) density test yields a t -statistic of 1.29 ($p = 0.198$), failing to reject the null of no manipulation. This is expected: Census 2001 enumeration preceded the PMGSY announcement, so village administrators had no incentive to inflate counts at this threshold.




Figure 1: Distribution of Village Population Near the PMGSY Threshold

Notes: Histogram of Census 2001 village population (bin width = 10 persons) for plain-area states. The dashed red line marks the PMGSY eligibility threshold at 500 persons. The McCrary density test statistic is $t = 1.29$ ($p = 0.198$), indicating no evidence of manipulation.

Covariate balance. Table 2 reports RDD estimates for five baseline covariates measured in Census 2001. None shows a significant discontinuity at the threshold. The literacy rate estimate is 0.0011 ($p = 0.67$), the SC share estimate is -0.0042 ($p = 0.21$), the ST share estimate is -0.0012 ($p = 0.66$), the female population share estimate is 0.0004 ($p = 0.40$), and the number of households estimate is 0.04 ($p = 0.89$). The joint pattern of small, statistically insignificant coefficients strongly supports the identifying assumption that villages just above and just below 500 are comparable in pre-treatment characteristics.

6.2 Main Results: Structural Transformation

Table 3 presents the core estimates. The RDD coefficient on the change in female non-agricultural share between Census 2001 and 2011 is 0.0021 ($p = 0.44$, robust SE = 0.0043). For males, the estimate is -0.0013 ($p = 0.71$, SE = 0.0031). The gender gap shows no discontinuity: 0.0010 ($p = 0.83$). Workforce participation rates are similarly null: the change in female participation is 0.0019 ($p = 0.57$) and the change in male participation is 0.0019 ($p = 0.18$).

Table 2: Covariate Balance at PMGSY Threshold

Baseline Covariate (2001)	RDD Estimate	Robust SE	<i>p</i> -value
Literacy rate	0.0011	(0.0030)	0.666
SC share	-0.0042	(0.0032)	0.213
ST share	-0.0012	(0.0064)	0.664
Female share	0.0004	(0.0006)	0.397
Number of households	0.0375	(0.2437)	0.888

Notes: Each row reports the RDD estimate at the 500 threshold for a baseline (Census 2001) covariate. Null effects support the identifying assumption that villages just above and below the threshold are comparable. CCT optimal bandwidth and triangular kernel.

Table 3: Main RDD Estimates: Effect of PMGSY Eligibility on Structural Transformation

Outcome	Estimate	Robust SE	95% CI	<i>p</i> -value	Eff. <i>N</i>
Change in female non-ag share	0.0021	(0.0043)	[-0.006, 0.011]	0.444	77,040
Change in male non-ag share	-0.0013	(0.0031)	[-0.007, 0.005]	0.713	89,155
Gender gap change (F - M)	0.0010	(0.0035)	[-0.006, 0.008]	0.826	119,579
Change in female WF participation	0.0019	(0.0034)	[-0.005, 0.009]	0.572	120,948
Change in male WF participation	0.0019	(0.0017)	[-0.001, 0.005]	0.180	83,784
Female non-ag share (2011)	-0.0006	(0.0037)	[-0.008, 0.007]	0.956	88,982
Male non-ag share (2011)	-0.0009	(0.0029)	[-0.007, 0.005]	0.731	122,885
Bandwidth	CCT Optimal (116–184)				
Polynomial	Local linear				
Kernel	Triangular				

Notes: Local polynomial RDD estimates using `rdrobust` (Calonico, Cattaneo, Titiunik). Running variable: Census 2001 village population. Cutoff: 500 (plain-area states). Robust bias-corrected standard errors, confidence intervals, and *p*-values reported. Because *p*-values use bias-corrected test statistics, they do not equal $2\Phi(-|\hat{\tau}/SE|)$; see Calonico, Cattaneo, and Titiunik (2014). ****p* < 0.01, ***p* < 0.05, **p* < 0.10.

Figure 2 presents the RDD graphically. Binned scatter plots with 20-person bins and local linear fits show no visible discontinuity at 500 for any of the three outcomes. The regression lines on either side of the cutoff are nearly collinear, reinforcing the statistical evidence.

RDD Estimates: Roads and Gendered Structural Transformation

Binned scatter plots (bins of 20 persons) with linear fits on each side of cutoff




Figure 2: RDD Plots: Roads and Gendered Structural Transformation

Notes: Binned scatter plots (bins of 20 persons) with local linear fits on each side of the

These estimates are economically meaningful nulls. The baseline female non-agricultural share in Census 2001 is approximately 11% for villages near the threshold (Table 1). The 95% confidence interval for the female effect (-0.0063 to 0.0105) rules out effects larger than one percentage point—roughly 7% of the baseline. For comparison, [Asher and Novosad \(2020\)](#) estimate a 10–15 percentage point effect on the probability of road construction at the habitation level. Even if only a fraction of this first stage operates at the village level, the null effect on employment suggests that roads do not move the needle on women’s sector of employment.

6.3 Dynamic RDD: Nighttime Lights

[Figure 3](#) presents the year-by-year RDD estimates for log nighttime lights. In principle, this dynamic analysis could reveal: (a) null pre-PMGSY effects (validating the design), and (b) growing post-PMGSY effects (indicating road-driven economic activity). Instead, the estimates show a persistent negative discontinuity that predates PMGSY: villages just above 500 have lower nighttime light intensity than villages just below, even in the 1990s. This pre-existing pattern means the dynamic nightlights analysis cannot cleanly isolate the effect of road construction, because the identifying assumption of continuity at the threshold is violated for this outcome.




Figure 3: Dynamic RDD: Year-by-Year Effect of PMGSY Eligibility on Nightlights

Notes: Each point represents the RDD coefficient from a separate regression of log nighttime light intensity on the threshold indicator for one year. Shaded band: 95% robust confidence interval. DMSP-OLS data for 1994–2013; VIIRS for 2014 onward. The dotted red line marks the PMGSY launch in 2000. The persistent negative discontinuity in the pre-period suggests a pre-existing correlation between village size and light intensity at this threshold, invalidating the nightlights RDD. Census-based employment outcomes, which show null effects, are not subject to this concern because they enter as first-differences.

The nightlights finding is informative in two ways. First, it highlights the importance of placebo testing in RDD designs: a clean density test and covariate balance do not guarantee that all outcomes satisfy the continuity assumption. Villages at 500 persons may differ in physical characteristics—area, density, proximity to towns—that affect light emissions independently of road construction. Second, the fact that Census employment outcomes show null effects despite this nightlights anomaly is reassuring: the first-differenced Census outcomes (change between 2001 and 2011) are immune to time-invariant level differences, unlike the nightlights levels.

6.4 Robustness

Bandwidth sensitivity. Table 4 Panel A shows the female non-agricultural share estimate across six bandwidth multipliers, from $0.5\times$ to $2.0\times$ the CCT optimal. The estimate is significant at narrow bandwidths ($0.5\times$: estimate = 0.0102, $p = 0.048$; $0.75\times$: estimate = 0.0048, $p = 0.037$) but becomes insignificant at wider bandwidths ($1.25\times$ and above). Note that the standard errors in Panel A differ from those in Table 3 because fixing the bandwidth alters the bias-correction procedure in `rdrobust`; both report robust bias-corrected inference. The $0.8\times$ result ($p = 0.037$, $N = 57,878$) warrants discussion: this is a reasonable sample with a below-5% p -value. However, the effect monotonically shrinks toward zero as bandwidth increases ($0.0102 \rightarrow 0.0048 \rightarrow 0.0021 \rightarrow 0.0008 \rightarrow -0.0002$), suggesting the positive estimates at narrow bandwidths reflect local curvature in the conditional expectation rather than a true discontinuity. Narrow-bandwidth estimates are also more sensitive to observations near the cutoff (Imbens and Lemieux, 2008). Taking the full pattern—including the null at the CCT optimal bandwidth and the confirmed null from randomization inference ($p = 0.50$)—the evidence favors a null interpretation.

Figure 4 visualizes the bandwidth sensitivity. Point estimates with 95% confidence intervals are shown for each bandwidth; the intervals consistently include zero except at the narrowest specification.

Table 4: Robustness: Bandwidth and Polynomial Sensitivity

Specification	Change in Female Non-Agricultural Share			
	Estimate	Robust SE	<i>p</i> -value	Eff. <i>N</i>
<i>Panel A: Bandwidth</i>				
0.5× optimal (<i>h</i> = 58)	0.0102	(0.0077)	0.048	38,388
0.8× optimal (<i>h</i> = 87)	0.0048	(0.0063)	0.037	57,878
1.0× optimal (<i>h</i> = 116)	0.0021	(0.0055)	0.088	77,040
1.2× optimal (<i>h</i> = 145)	0.0008	(0.0049)	0.190	95,994
1.5× optimal (<i>h</i> = 174)	-0.0002	(0.0044)	0.316	115,097
2.0× optimal (<i>h</i> = 232)	-0.0002	(0.0039)	0.791	152,472
<i>Panel B: Polynomial order</i>				
Order 1	0.0021	(0.0043)	0.444	77,040
Order 2	0.0064	(0.0053)	0.129	95,994
Order 3	0.0092	(0.0059)	0.064	130,499
<i>Panel C: Donut hole</i>				
Exclude ±5	-0.0028	(0.0044)	0.493	91,019
Exclude ±10	-0.0027	(0.0051)	0.537	84,448
Exclude ±20	-0.0069	(0.0072)	0.252	63,347

Notes: All specifications use local polynomial RDD at the 500 threshold for plain-area states. Panel A varies bandwidth around the CCT optimal (fixing *h* changes the bias-correction bandwidth, so SEs and *p*-values differ from the main table even at 1.0× optimal). Panel B varies polynomial order (all use CCT bandwidth). Panel C excludes observations within the donut radius of the cutoff. Robust bias-corrected *p*-values reported throughout.




Figure 4: Bandwidth Sensitivity: Female Non-Agricultural Share

Notes: Point estimates and 95% confidence intervals for the RDD effect on the change in female non-agricultural share, across bandwidth multipliers from $0.5\times$ to $2.0\times$ the CCT optimal. The dotted line marks the optimal bandwidth. All intervals except the narrowest ($0.5\times$) include zero.

Polynomial order. Panel B of [Table 4](#) reports estimates using polynomial orders 1 through 3. The local linear estimate ($p = 0.44$) and local quadratic ($p = 0.13$) are both null. The local cubic shows marginal significance ($p = 0.064$), but higher-order polynomials are known to produce spurious results near boundaries ([Gelman and Imbens, 2019](#)), and I follow the recommendation of using local linear as the primary specification.

Donut hole. Panel C excludes villages within ± 5 , ± 10 , and ± 20 persons of the threshold to address concerns about heaping or rounding. All three estimates are null, with p -values of 0.49, 0.54, and 0.25 respectively.

Placebo thresholds. [Figure 5](#) shows RDD estimates at the true threshold (500) alongside estimates at five placebo thresholds (300, 400, 600, 700, 800). None of the placebo estimates is significant, and the estimate at 500 is not an outlier. This confirms that the null result is not an artifact of functional form misspecification.




Figure 5: Placebo Threshold Tests: Female Non-Agricultural Share

Notes: RDD estimates for the change in female non-agricultural share at the true threshold (500, blue) and five placebo cutoffs (300, 400, 600, 700, 800, grey). Error bars: 95% confidence intervals.

The true-threshold estimate is not an outlier among the placebos.

Randomization inference. To assess finite-sample validity, I conduct randomization inference with 500 permutations. Outcomes are randomly reshuffled across villages while preserving the running variable, and the RDD is re-estimated for each permutation. The proportion of permuted estimates at least as large in absolute value as the observed estimate yields a two-sided RI p -value of 0.50, confirming the null.


6.5 Heterogeneity by Region

If road effects are concentrated in regions with more restrictive gender norms (the Hindi belt) or more dynamic non-farm economies (the south and west), aggregating across regions could mask real effects. [Table 5](#) and [Figure 6](#) present estimates for four Indian regions. North India (Uttar Pradesh, Bihar, Madhya Pradesh, Rajasthan, Jharkhand, Chhattisgarh, Haryana) has the largest sample but shows null effects for both women and men. South India (Andhra Pradesh, Karnataka, Kerala, Tamil Nadu, Telangana), East India (West Bengal, Odisha), and West India (Gujarat, Maharashtra, Goa) all show null effects as well. No region produces a statistically significant estimate for either gender.

Table 5: Regional Heterogeneity in Road Effects on Structural Transformation

Region	Female			Male			N
	Estimate	SE	p	Estimate	SE	p	
North	-0.0053	(0.0053)	0.326	0.0027	(0.0050)	0.401	100,188
South	0.0069	(0.0095)	0.565	-0.0063	(0.0081)	0.319	69,580
East	0.0099	(0.0102)	0.203	0.0007	(0.0070)	0.699	84,528
West	-0.0026	(0.0088)	0.887	-0.0096	(0.0073)	0.096	58,942

Notes: RDD estimates at the 500 threshold by Indian region. North includes Hindi-belt states (UP, Bihar, MP, Rajasthan, Jharkhand, Chhattisgarh, Haryana). South includes AP, Karnataka, Kerala, Tamil Nadu, Telangana. East: West Bengal, Odisha. West: Gujarat, Maharashtra, Goa. Robust standard errors reported.

**Figure 6:** Regional Heterogeneity in Road Effects on Structural Transformation

Notes: RDD estimates at the 500-person threshold by Indian region, separately for female and male non-agricultural share changes. Error bars: 95% confidence intervals. No region shows a significant effect for either gender.

6.6 Special Category States

As an additional robustness exercise, I estimate the RDD at the 250-person threshold for the 71,159 villages in special-category states. Both the female non-agricultural share change

and the male change are null at this threshold. While the sample is smaller and precision is lower, the results are consistent with the main analysis: PMGSY eligibility does not produce detectable employment effects at the village level, regardless of which threshold is used.

7. Discussion

7.1 Interpreting the Null

The central finding of this paper is a precisely estimated null: PMGSY road eligibility at the 500-person threshold produces no detectable effect on gendered structural transformation over the 2001–2011 decade. Three interpretations deserve consideration.

Attenuation from habitation-to-village mapping. The most mechanical explanation is that the village-level RDD has a weak first stage. Because PMGSY targets habitations rather than villages, a village crossing the 500-person threshold does not deterministically receive a road—it merely becomes more likely to contain eligible habitations. [Asher and Novosad \(2020\)](#) document a 22 percentage point jump in road construction probability at the habitation level. If only a fraction of this first stage transmits to the village level, the reduced-form effect could be small enough to be undetectable even in a sample of half a million villages. However, the employment outcomes from [Asher and Novosad \(2020\)](#) were estimated at the habitation level using habitation-level population data. The fact that I find null effects at the village level, even with excellent statistical power, suggests either that the first stage is very weak at this level of aggregation or that the employment effects themselves are smaller than habitation-level estimates suggest.

Binding constraints beyond connectivity. A more substantive interpretation is that road connectivity is necessary but not sufficient for women’s structural transformation. The literature on female labor supply in India emphasizes social norms—particularly norms around women’s mobility and appropriate work—as a primary barrier ([Jayachandran, 2021](#); [Afridi et al., 2023](#)). If women in rural India do not enter non-agricultural employment even when commuting becomes feasible, roads cannot generate the predicted effect. [Munshi and Rosenzweig \(2016\)](#) document how caste networks facilitate occupational transitions for men but not women, suggesting that social capital and information channels are also gender-specific barriers. Even when economic opportunities expand, household dynamics and male control of resources can prevent women from realizing gains ([Bernhardt et al., 2018](#)). This interpretation implies that infrastructure investments must be complemented by interventions targeting the demand-side (creating female-appropriate non-farm jobs nearby) or the norm-side (shifting attitudes about women’s work outside the home) to achieve gender-equitable structural transformation.

Timing and adjustment. The Census 2001–2011 window may be too short to capture the full adjustment process. PMGSY road construction was phased: habitations above 1,000 were targeted first, with the 500 threshold becoming operational in later phases. Some villages near the cutoff may have received roads only in the final years before Census 2011, leaving insufficient time for employment patterns to adjust. Moreover, gendered employment transitions may occur through cohort replacement (younger women entering non-agricultural work) rather than contemporaneous switching by existing workers, a process that would unfold over decades rather than years.

7.2 Relationship to Prior Literature

My findings complement and extend several prior studies. [Asher and Novosad \(2020\)](#) find that PMGSY roads reduce travel times and shift workers out of agriculture at the habitation level, using habitation-level PMGSY administrative data. My village-level null result is not contradictory: it shows that the habitation-level effect does not scale to a detectable village-level gender gap in structural transformation. This could reflect (a) attenuation from aggregation, (b) heterogeneous effects across habitations within a village, or (c) genuine null effects on gender-specific outcomes even where aggregate structural transformation occurs.

[Lei et al. \(2019\)](#) use household-level IHDS data and find a positive correlation between road proximity and women’s non-farm employment, controlling for household and village fixed effects. Their association may reflect selection on unobservables—villages with better roads may differ in ways that also promote female employment—while my RDD estimates provide a clean causal test at the threshold. The contrast between their suggestive positive correlation and my clean null is informative: it suggests that the cross-sectional association between roads and female employment may be driven by omitted variables rather than a causal road effect.

[Aggarwal \(2018\)](#) uses a fuzzy RDD with PMGSY roads and finds no effect on poverty, measured by consumption expenditure. My null result on structural transformation is consistent with her finding: if roads do not reduce poverty, they may also fail to shift employment patterns. Together, these results paint a more nuanced picture of rural roads than the transformative narrative often associated with infrastructure investment.

7.3 Limitations

Several caveats apply. First, I cannot observe habitation-level population or road construction status, so I cannot verify the first-stage magnitude at the village level. The null result could reflect a genuine null treatment effect or an undetectably small first stage. Second, Census

employment categories are coarse. Workers who shift from pure farming to agro-processing or part-time non-farm work may not be captured by the Census classification. Third, I study a single population threshold. PMGSY also targeted habitations above 1,000 persons in earlier phases; effects at higher thresholds (where the first stage may be stronger) could differ. Fourth, my analysis is restricted to Census years 2001 and 2011. More recent data from Census 2021 (delayed by the pandemic to 2024–2025) would allow assessment of longer-run effects.

7.4 Policy Implications

The null finding does not imply that rural roads are unimportant. PMGSY has been credited with reducing travel times, improving school access ([Adukia et al., 2020](#)), and facilitating market integration. Rather, the finding suggests that roads alone are insufficient to address the gender gap in India’s structural transformation. Policymakers seeking to increase women’s non-agricultural employment should consider complementary interventions: skills training programs targeted at women, support for female-owned non-farm enterprises, childcare and domestic labor relief that expands women’s available work hours, and norm-change interventions that address social barriers to women’s mobility and employment. The evidence from this paper suggests that building the road is not enough—women also need a reason, the skills, and the social permission to travel it.

8. Conclusion

India built 400,000 kilometers of rural roads under PMGSY, connecting over 100,000 habitations to the all-weather road network. This paper asks whether that massive infrastructure investment narrowed the gender gap in structural transformation—the shift from agricultural to non-agricultural employment. Using a regression discontinuity design at the program’s population eligibility threshold and village-level Census data on 582,807 rural villages, I find precisely estimated null effects. Roads did not differentially move women out of agriculture, did not increase men’s structural transformation at the village level, and did not narrow the gender gap in non-farm employment between 2001 and 2011.

These nulls are not merely the absence of evidence. The design is clean: no manipulation, balanced covariates, a pre-determined running variable, and half a million observations. The estimates rule out effects larger than one percentage point in either direction. The finding survives every robustness check in the RDD toolkit. This is evidence of absence.

The implication is not that infrastructure doesn’t matter, but that it matters less for gender equity than the conventional wisdom assumes. The binding constraints on women’s

structural transformation in rural India lie downstream of roads—in social norms, labor demand, human capital, and household bargaining. Future research should investigate which complementary interventions, delivered alongside road connectivity, can unlock the gendered returns to infrastructure. The road has been built. The question is what else must change before women can walk it.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG) version 2.1 by Asher, Novosad, and Lunt.

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.
- Afridi, Farzana, Taryn Dinkelman, and Kanika Mahajan**, “Social Norms and Female Labor Force Participation in India,” *Journal of Development Economics*, 2023, *163*, 103089.
- 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,” *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.
- Bernhardt, Arielle, Erica Field, Rohini Pande, and Natalia Rigol**, “Household Matters: Revisiting the Returns to Capital among Female Microentrepreneurs,” *American Economic Journal: Applied Economics*, 2018, *10* (3), 263–290.
- Calonico, Sebastian, Matias D Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, *82* (6), 2295–2326.
- , —, and —, “rdrrobust: An R Package for Robust Nonparametric Inference in Regression-Discontinuity Designs,” *R Journal*, 2015, *7* (1), 38–51.
- Cattaneo, Matias D, Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.
- 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.
- Gelman, Andrew and Guido Imbens**, “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs,” *Journal of Business and 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.
- Jayachandran, Seema**, “Social Norms as a Barrier to Women’s Employment in Developing Countries,” *IMF Economic Review*, 2021, *69*, 576–595.
- Klasen, Stephan and Janneke Pieters**, “What Happened to the Indian Gender Gap in Labour Force Participation Rates?,” *World Development*, 2018, *108*, 289–305.
- Lei, Lei, Sonalde Desai, and Reeve Vanneman**, “Can the Existing Road Network Explain Gender Differences in Employment in India?,” *World Development*, 2019, *118*, 53–64.
- Ministry of Rural Development**, “Programme Guidelines: Pradhan Mantri Gram Sadak Yojana,” Technical Report, Government of India 2000.
- Munshi, Kaivan and Mark Rosenzweig**, “Strength in Numbers: Networks as a Solution to Occupational Traps,” *Review of Economic Studies*, 2016, *83* (3), 1069–1101.
- 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.
- Office of the Registrar General**, “Census of India 2001: Primary Census Abstract,” Technical Report, Government of India 2001.
- , “Census of India 2011: Primary Census Abstract,” Technical Report, Government of India 2011.

A. Data Appendix

A.1 SHRUG Data Platform

All village-level data come from the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG), version 2.1, developed by [Asher et al. \(2021\)](#). SHRUG provides a unified geographic identifier (SHRID) that links Indian village-level data across Census rounds despite administrative boundary changes. The platform is freely available at https://www.devdata.org/shrug_download/.

A.2 Census Variables Used

Population Census Abstract (PCA) 2001. Variables used from `pc01_pca_clean_shrid.csv`:

- `pc01_pca_tot_p`: Total population (running variable)
- `pc01_pca_no_hh`: Number of households
- `pc01_pca_p_lit`: Literate population
- `pc01_pca_p_sc`: Scheduled Caste population
- `pc01_pca_p_st`: Scheduled Tribe population
- `pc01_pca_tot_work_p`: Total main workers
- `pc01_pca_main_cl_p`, `pc01_pca_main_al_p`: Main cultivators and agricultural laborers
- `pc01_pca_main_hh_p`, `pc01_pca_main_ot_p`: Household industry and other main workers
- All variables above with `_f` and `_m` suffixes for female and male breakdowns

Population Census Abstract 2011. The same variable structure from `pc11_pca_clean_shrid.csv`, with `pc11_` prefix.

A.3 Nighttime Lights Data

DMSP-OLS stable lights annual composites (1994–2013) and VIIRS annual composites (2014–2023), both at village level via SHRUG aggregation. DMSP values range from 0–63 digital number units; VIIRS values are in nanowatts per square centimeter per steradian. I take the natural log of nightlights plus 0.01 to handle zeros.

A.4 Sample Construction Details

Table 6: Sample Construction

Step	Villages
Census 2001 PCA observations in SHRUG	593,246
After restricting to rural villages	591,668
After merging with Census 2011 (non-missing pop.)	582,807
Plain-area states only (main sample)	511,648
Special-category states	71,159
Within ± 200 of threshold (near sample)	134,022

A.5 State Classification

Special-category states (PMGSY threshold = 250): Jammu & Kashmir (01), Himachal Pradesh (02), Uttarakhand (05), Arunachal Pradesh (12), Nagaland (13), Manipur (14), Mizoram (15), Tripura (16), Meghalaya (17), Assam (18), Sikkim (11).

Regional classification for heterogeneity analysis:

- North: Haryana (06), Rajasthan (08), Uttar Pradesh (09), Bihar (10), Jharkhand (20), Chhattisgarh (22), Madhya Pradesh (23)
- South: Andhra Pradesh (28), Karnataka (29), Kerala (32), Tamil Nadu (33), Telangana (36)
- East: West Bengal (19), Odisha (21)
- West: Gujarat (24), Maharashtra (27), Goa (30)

B. Identification Appendix

B.1 McCrary Density Test Details

The density test of [Cattaneo et al. \(2020\)](#) estimates the density of the running variable on each side of the cutoff and tests whether they are equal. With $N = 511,648$ observations in the plain-area sample, the test has excellent power. The estimated t -statistic of 1.29 ($p = 0.198$) provides no evidence of bunching at 500.

B.2 Pre-PMGSY Nightlights Placebo

The year-by-year RDD on nighttime lights reveals a pre-existing negative discontinuity at 500 that predates PMGSY. This finding has two implications:

First, it means the dynamic nightlights analysis cannot be used to validate the RDD design or estimate road effects. The pre-2000 estimates should be zero under the null, but they are significantly negative. This likely reflects a relationship between village population at 500 and some physical or geographic characteristic (e.g., village area, distance to town, terrain) that affects light emissions independently of road construction.

Second, and importantly, this pre-existing nightlights discontinuity does *not* invalidate the Census-based analysis. The Census outcomes enter the RDD in first differences ($\Delta Y = Y_{2011} - Y_{2001}$), which remove any time-invariant level difference between villages above and below the threshold. The nightlights analysis uses annual levels, making it sensitive to persistent cross-sectional differences. The covariate balance tests (which also use levels, not changes) pass cleanly, suggesting that the nightlights anomaly is specific to light emissions rather than a general feature of village characteristics at this threshold.

B.3 Randomization Inference

The RI procedure conducts 500 permutations of the outcome variable while fixing the running variable, re-estimating the RDD for each permutation. The observed coefficient (0.0021 for female non-agricultural share change) is located at the 50th percentile of the permutation distribution, yielding a two-sided p -value of 0.50. This confirms that the asymptotic inference from `rdrobust` is reliable in this setting.

C. Robustness Appendix

C.1 Full Bandwidth Sensitivity

The CCT optimal bandwidth for the female non-agricultural share outcome is approximately 150 population units (varying by specification). At the optimal bandwidth, the effective sample size includes roughly 80,000–100,000 villages on each side of the cutoff. Even at half the optimal bandwidth, the effective sample exceeds 40,000 villages.

The marginal significance at narrow bandwidths ($0.5\times$ optimal: estimate = 0.0102, $p = 0.048$; $0.75\times$: $p = 0.037$) does not survive multiple-testing correction. Given seven primary outcomes tested at the CCT bandwidth, a Bonferroni correction would require $p < 0.007$ for significance at the 5% level. The narrow-bandwidth result fails this correction by a wide margin.

C.2 Polynomial Sensitivity

Local linear (order 1) is the recommended specification (Gelman and Imbens, 2019; Imbens and Lemieux, 2008) and produces a null result ($p = 0.44$). The local quadratic also produces a null. The local cubic shows marginal significance ($p = 0.064$), but this is consistent with overfitting: cubic polynomials tend to produce more volatile estimates near the boundary, particularly in finite samples.

C.3 Donut-Hole Specifications

Excluding villages within ± 5 , ± 10 , and ± 20 persons of the threshold addresses two concerns: (a) potential heaping of population at round numbers (e.g., exactly 500), and (b) the possibility that villages very near the cutoff are subject to different treatment dynamics. All three specifications produce null results with p -values of 0.49, 0.54, and 0.25 respectively.

C.4 Special Category States

The 71,159 villages in special-category states face a lower PMGSY threshold of 250 persons. Estimating the RDD at this threshold yields null effects for both female and male non-agricultural share changes. The smaller sample reduces precision, but the point estimates are small and statistically insignificant.

D. Heterogeneity Appendix

D.1 Regional Results

North India (the “Hindi belt”) is the largest subregion with the most conservative gender norms and lowest female labor force participation. If roads were to have gender-differentiated effects anywhere, this is where one would expect them. Yet the estimates for North India are null for both females and males, with standard errors only slightly larger than the full-sample estimates.

South India has higher female labor force participation and more dynamic non-farm economies, but also shows null effects at the threshold. East India (West Bengal, Odisha) and West India (Gujarat, Maharashtra, Goa) similarly show no detectable effects.

The uniform null across all regions is notable. It suggests that the non-effect is not driven by heterogeneity cancellation—roads benefiting women in one region and hurting them in another. Instead, the null appears to be a genuine feature of the PMGSY threshold’s effect (or lack thereof) on gendered employment patterns.

E. Additional Figures and Tables




Figure 7: Covariate Balance at the PMGSY Threshold

Notes: RDD point estimates and 95% confidence intervals for five baseline (Census 2001) covariates at the 500-person PMGSY threshold. None is statistically distinguishable from zero at conventional levels. CCT optimal bandwidth and triangular kernel.