

# Roads to Nowhere? Rural Infrastructure and the Persistence of Gender Gaps in Non-Farm Employment in India

APEP Autonomous Research\* @olafdrw

February 23, 2026

## Abstract

Can rural roads close gender gaps in non-farm employment? Using a regression discontinuity design at the population eligibility threshold of India's flagship road-building program (PMGSY), I estimate the causal effect of road connectivity on women's structural transformation across 528,000 villages. Despite excellent statistical power (effective sample sizes exceeding 100,000) and a valid design confirmed by density tests and covariate balance, I find precisely zero effects on female non-agricultural worker share, female labor force participation, and female literacy. Male outcomes are equally unaffected. The null result is robust across bandwidths, polynomial orders, kernels, and donut specifications. These findings suggest that transportation barriers are not the binding constraint on women's transition from agriculture to non-farm work in rural India—social norms, labor demand, and occupational segregation likely dominate.

**JEL Codes:** J16, J21, O18, R42

**Keywords:** rural roads, female employment, structural transformation, India, PMGSY, regression discontinuity

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

## 1. Introduction

During India’s monsoon, thousands of villages vanish from the economy. Unpaved *kachcha* roads turn to mud, severing connections to market towns, hospitals, and non-farm jobs for months at a time. For women in these villages—90 percent of whom work in agriculture—the isolation is especially acute. India’s flagship response, the Pradhan Mantri Gram Sadak Yojana (PMGSY), has paved all-weather roads to over 178,000 habitations since 2000, making it the world’s largest rural road-building program. The premise is intuitive: connect villages to markets, and employment opportunities will follow. But did they—and did they follow for women?

This paper exploits a sharp population threshold in PMGSY’s eligibility rules to answer that question. Habitations with Census 2001 populations above 500 persons (in plains areas) were prioritized for road construction; those below 500 were not. This threshold creates a regression discontinuity design (RDD) that identifies the causal effect of road eligibility on employment outcomes, following the seminal work of [Asher and Novosad \(2020\)](#). The design is clean: the threshold was set centrally, determined by Census enumeration before the program was widely known, and validated by density tests and covariate balance across eight pre-treatment characteristics.

The main finding is a precisely estimated null. Road eligibility has no detectable effect on the share of female workers in non-agricultural employment, on female labor force participation, on female literacy, or on the male-female gap in non-farm employment. The point estimates are substantively tiny—less than 0.1 percentage points—with 95 percent confidence intervals that rule out effects larger than approximately 1.5 percentage points in either direction. This is not a power problem: with effective sample sizes exceeding 100,000 villages across 528,000 plains-area observations, the design can detect economically meaningful effects. Nor is it a specification problem: six bandwidth choices, four placebo cutoffs, three polynomial orders, three kernel functions, and five donut hole radii all produce estimates indistinguishable from zero.

The null extends to men. Male non-agricultural employment shares, male labor force participation, and male literacy are all flat at the threshold. This symmetry is informative: it suggests that the population eligibility threshold—while it meaningfully increases the probability of road construction (by approximately 20–30 percentage points; [Asher and Novosad, 2020](#))—does not generate detectable shifts in employment composition for *either* gender. The intent-to-treat estimates imply that even the treatment-on-the-treated effects are economically negligible: scaling the female non-agricultural share estimate by a first stage of 0.25 yields an implied LATE of 0.2 percentage points.

The paper speaks to a longstanding puzzle in development economics: why has India’s female labor force participation rate been so low—and declining—even as incomes rose and education expanded (Goldin, 1995; Klasen, 2018)? Economists have blamed social norms restricting women’s mobility (Jayachandran, 2021; Afridi et al., 2023), rising male incomes that reduce women’s need to work, and supply-side constraints including poor infrastructure and unsafe transportation (Erten and Leight, 2021). This paper’s null provides evidence against the infrastructure hypothesis—or at least against the proposition that all-weather road connectivity is the binding margin. We know from Asher and Novosad (2020) that PMGSY roads increase consumption and men’s probability of working outside the village. But this paper shows that these gains do not translate into women’s structural transformation from agricultural to non-agricultural employment.

The finding also contributes to the broader literature on infrastructure and economic development (Donaldson, 2018; Banerjee et al., 2020; Faber, 2014). Much of this literature estimates average effects without examining distributional consequences by gender. If roads primarily facilitate men’s mobility while leaving women’s employment patterns unchanged, this has important implications for the design of inclusive development policy. For policymakers seeking to close gender gaps in non-farm employment, the implication is clear: roads are necessary but not sufficient. Complementary interventions—addressing social norms, creating female-accessible non-farm employment, improving workplace safety—may be required to translate connectivity into genuine occupational mobility for women.

The remainder of the paper proceeds as follows. Section 2 describes the PMGSY program and the institutional context of female employment in India. Section 3 develops a simple conceptual framework. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports results. Section 7 discusses mechanisms and implications. Section 8 concludes.

## **2. Institutional Background**

### **2.1 The Pradhan Mantri Gram Sadak Yojana (PMGSY)**

India launched the Pradhan Mantri Gram Sadak Yojana (Prime Minister’s Rural Roads Program) in December 2000 with the objective of providing all-weather road connectivity to unconnected habitations with populations above 500 persons in plains areas and above 250 persons in hill, tribal, and desert areas. The program is centrally funded by the Government of India with a dedicated cess on high-speed diesel fuel, and is implemented by state governments through dedicated State Rural Roads Development Agencies (SRRDAs).

The program was born from a recognition that rural isolation was a first-order constraint

on economic development. At the time of PMGSY’s launch, roughly 40 percent of India’s 825,000 rural habitations lacked all-weather road connectivity. During the monsoon months (June through September), unpaved *kachcha* roads became impassable for motorized vehicles, cutting off villages from markets, hospitals, schools, and administrative centers for weeks or months at a time. The cost of this isolation was borne disproportionately by the poorest households and by women, who had fewer alternative modes of transport and shorter feasible commuting distances.

The eligibility threshold is the key institutional feature for this paper’s identification strategy. The population threshold was determined by Census 2001 population counts for each habitation. Habitations above the threshold were eligible for road construction; those below were not initially eligible. The threshold was set at the national level by the central government—it was not determined by local officials who might have had incentives to manipulate it. This is critical for the RDD design: the cutoff was imposed externally and uniformly, reducing the scope for strategic sorting around the threshold.

PMGSY was implemented in two phases. Phase I (2000–2012) prioritized habitations above the 500/250 threshold, with the goal of connecting all eligible habitations by 2007—a deadline that was repeatedly extended as implementation lagged. Phase II (announced in 2013) expanded the program to upgrade existing roads and connect habitations with populations of 250–499 in plains areas and 100–249 in hill areas. The analysis in this paper focuses on the Phase I threshold at 500 persons, which generated the sharpest discontinuity in eligibility.

Implementation was staggered across districts and states over the 2000–2015 period, driven by variation in state capacity, terrain difficulty, and political priorities. States with stronger bureaucratic capacity (such as Madhya Pradesh and Rajasthan) implemented the program faster than states with weaker administrative infrastructure. By 2015, the program had constructed over 600,000 kilometers of roads connecting approximately 178,000 habitations (Asher and Novosad, 2020). The roads are “all-weather” bituminous or concrete roads designed to withstand monsoon conditions, replacing unpaved tracks that become impassable during the rainy season. Each road was required to meet minimum specifications: a carriageway width of 3.75 meters for plains and 3.0 meters for hilly areas, with proper drainage and surfacing to ensure year-round passability. The program represented a transformative change in connectivity for villages that received roads—from seasonal access to reliable year-round connectivity to market towns and district headquarters.

Several features of the program are relevant for interpretation. First, PMGSY roads connect villages to existing road networks rather than building entirely new corridors. The typical road connects a village to the nearest point on the state highway or district road

network, with a median length of approximately 4 kilometers. This “last mile” design means that the program’s effects depend critically on the quality of the existing road network to which villages are connected—a village connected to a poorly maintained state highway may see smaller benefits than one connected to a national highway.

Second, the program did not directly provide employment or economic development services—it only provided physical infrastructure. Any employment effects must therefore operate through reduced transportation costs enabling market access, commuting, or firm entry. This is a crucial distinction from programs like MGNREGA, which directly create employment. PMGSY’s effects are inherently indirect and depend on complementary factors (labor demand, skills, social norms) that determine whether improved connectivity translates into economic opportunity.

Third, the program was implemented by state-level agencies with varying capacity, creating heterogeneity in implementation timing and quality across states. Fourth, there was meaningful non-compliance: not all eligible habitations received roads during Phase I (some were deferred due to difficult terrain, land acquisition disputes, or contractor failures), and some below-threshold habitations received roads through other programs (state PWD roads, MPLAD funds, or Bharat Nirman). This imperfect compliance motivates the ITT interpretation of the RDD estimates.

## 2.2 Female Employment in Rural India

India’s female labor force participation rate has been a persistent puzzle in development economics. Rural women’s LFPR declined from approximately 33 percent in 1993–94 to 25 percent in 2011–12 (NSS data), before recovering to approximately 37 percent by 2022–23 (PLFS data). This decline occurred during a period of rapid economic growth, rising female education, and declining fertility—all factors that typically increase female labor supply in the standard Becker framework. The pattern is consistent with [Goldin \(1995\)](#)’s U-shaped hypothesis, which predicts that female labor supply first declines as economies transition from agriculture (where women work out of necessity) and then rises as education and service-sector opportunities expand—but the depth and duration of India’s trough has surprised researchers.

The employment structure of rural women who do work is heavily concentrated in agriculture. In Census 2001, approximately 90 percent of female main workers in rural India were classified as cultivators or agricultural laborers. The remaining 10 percent were split between household industry workers (traditional crafts, food processing, small-scale manufacturing) and “other” workers (services, modern manufacturing, government employment). The shift from agricultural to non-agricultural employment—structural transformation—is typically associated with higher productivity and earnings. For men, this transition has proceeded

steadily: the male non-agricultural employment share rose from approximately 25 percent in 1991 to 40 percent by 2011. For women, the same share barely moved, remaining below 15 percent over two decades.

This gender asymmetry in structural transformation has deep roots. Several barriers may prevent women from shifting to non-farm employment even when opportunities exist. Social norms in many parts of India restrict women’s mobility outside the village, particularly unaccompanied travel ([Jayachandran, 2021](#)). The “purdah” system and related practices limit women’s visibility in public spaces, especially in North India. [Afridi et al. \(2023\)](#) document that these norms are remarkably persistent: even in villages with improved transportation infrastructure and expanding labor markets, women’s participation in non-farm work remains low when local norms are restrictive. Critically, these norms operate independently of physical infrastructure—building a road does not change the social acceptability of women traveling on it.

The geographic dimension of these norms is important for understanding why road effects might vary. Social restrictions on female mobility are most severe in the Hindi-speaking plains of North India (the “Hindi belt” states of Uttar Pradesh, Bihar, Rajasthan, and Madhya Pradesh), which also contain the majority of PMGSY-eligible villages. In contrast, South Indian states (Tamil Nadu, Kerala, Andhra Pradesh) and tribal-dominated areas have historically permitted greater female autonomy. This geographic correlation between restrictive norms and program exposure means that the average treatment effect of roads on female employment may be attenuated by the very norms that roads are hypothesized to circumvent.

Labor demand may also be gendered. Non-farm employers in rural India may prefer male workers, either due to explicit employer discrimination or because available jobs require physical strength, nighttime work, or travel that social norms discourage women from performing ([Van der Meulen Rodgers, 2008](#)). The types of non-farm employment accessible from newly connected villages—construction, trucking, trading, service provision at market towns—are disproportionately male-dominated. Female-accessible non-farm employment (teaching, nursing, garment work, food processing) tends to concentrate in urban areas or require educational credentials that rural women often lack. If the constraint on women’s non-farm employment is demand-side (employers won’t hire women, or the available jobs are male-typed) rather than supply-side (women can’t physically reach employers), then road connectivity will not close the gender gap.

The household bargaining dimension adds another layer. Women’s employment decisions in rural India are rarely made by women alone. [Field et al. \(2021\)](#) show that even when women have individual economic opportunities, household-level decision-making (dominated

by husbands and mothers-in-law) can constrain their labor supply. If household heads view women’s non-farm employment as a threat to family honor or a signal of economic distress, improved road connectivity will not translate into female employment outside agriculture. [Duflo \(2003\)](#) and [Chattopadhyay and Duflo \(2004\)](#) provide evidence that intra-household power dynamics significantly shape women’s economic outcomes in developing countries.

Finally, the “income effect” channel may operate in a direction that offsets any direct connectivity benefit. If roads increase household income primarily through male non-farm employment (as men are more likely to commute to town jobs), the resulting income gain may reduce women’s labor supply as families can “afford” to have women withdraw from the labor force entirely—or at least from the most visible forms of agricultural labor. This would produce a negative effect of roads on female LFPR, partially or fully offsetting any positive effect on non-farm employment. The income effect is particularly relevant in the Indian context because female agricultural labor is often stigmatized as a marker of poverty: withdrawing from field work signals upward mobility.

### 3. Conceptual Framework

Consider a simple model of a rural household’s labor allocation decision. The household has male ( $m$ ) and female ( $f$ ) members who can work in agriculture ( $A$ ) or non-agriculture ( $N$ ). Agricultural work takes place in or near the village; non-agricultural work requires commuting to a market town at cost  $c$  per person per day.

The wage premium for non-agricultural work is  $w_N - w_A > 0$ . A household member works in the non-farm sector if and only if the net benefit exceeds the commuting cost and any additional gender-specific cost:

$$\text{Work in } N \iff (w_N - w_A) - c - \delta_g > 0 \tag{1}$$

where  $\delta_g \geq 0$  is a gender-specific barrier (social norms, safety concerns, employer discrimination) that is zero for men and positive for women:  $\delta_f > 0 = \delta_m$ .

Road construction reduces commuting costs from  $c$  to  $c' < c$ . The effect on women’s non-farm employment share depends on the relative magnitudes:

1. If  $\delta_f < (w_N - w_A) - c'$ : Women switch to non-farm work after road construction. Roads close the gender gap.
2. If  $\delta_f > (w_N - w_A) - c$ : Women don’t switch even without commuting costs. Roads are irrelevant for the gender gap.

3. If  $(w_N - w_A) - c < \delta_f < (w_N - w_A) - c'$ : This is the “marginal women” case where roads matter. The key question is whether any women fall in this range.

The income effect introduces a further complication. If men shift to non-farm work, household income rises, and women may reduce total labor supply. Define female labor supply as  $L_f(Y_m, c, \delta_f)$  where  $\partial L_f / \partial Y_m < 0$  (income effect). The road’s net effect on female non-farm employment is:

$$\frac{\partial(\text{Female Non-Ag Share})}{\partial c} = \underbrace{\frac{\partial P(N|c, \delta_f)}{\partial c}}_{<0 \text{ (direct)}} + \underbrace{\frac{\partial L_f}{\partial Y_m} \cdot \frac{\partial Y_m}{\partial c}}_{>0 \text{ (income)}} \quad (2)$$

The direct effect of lower commuting costs is to increase non-farm employment (negative sign because  $c$  decreases). The income effect works in the opposite direction: higher male income reduces female labor supply. The theoretical prediction is therefore ambiguous. A null finding is consistent with either (a) the gender-specific barrier  $\delta_f$  being so large that commuting costs are irrelevant, or (b) the direct and income effects exactly canceling. The empirical analysis cannot distinguish between these mechanisms, but the pattern of results—null effects on both female LFPR and non-farm share, and null effects on male outcomes as well—is most consistent with hypothesis (a): roads have no effect on employment composition for either gender at this margin.

## 4. Data

### 4.1 SHRUG Platform

The primary data source is the Socioeconomic High-resolution Rural-Urban Geographic Platform (SHRUG), version 2.1 (Asher et al., 2021). SHRUG provides harmonized village-level data across India’s 2001 and 2011 Censuses, linking approximately 640,000 villages through stable geographic identifiers. The key advantage of SHRUG is that it solves India’s chronic geographic ID fragmentation problem: village codes change across Census rounds due to boundary changes, split/merges, and administrative reclassification. SHRUG provides a consistent identifier (SHRID) that maps villages across time.

I use three components of SHRUG. First, the Census 2001 Primary Census Abstract (PCA) provides the running variable (village population) and pre-treatment covariates including demographic composition, literacy rates, and employment status. Second, the Census 2011 PCA provides outcome variables measured after PMGSY implementation. Third, the geographic crosswalk maps each village to its district and state for clustering standard errors.

## 4.2 Variable Construction

The running variable is Census 2001 total population (`pc01_pca_tot_p`), centered at the PMGSY eligibility threshold of 500. Treatment is defined as  $T_i = \mathbb{I}[\text{Pop}_{i,2001} \geq 500]$ .

The primary outcome is the female non-agricultural worker share, defined as:

$$\text{Female Non-Ag Share} = \frac{\text{Female HH Industry Workers} + \text{Female Other Workers}}{\text{Female Total Workers}} \quad (3)$$

This captures the share of female main workers employed outside agriculture. “HH Industry Workers” are those engaged in household manufacturing (traditional crafts, food processing, weaving). “Other Workers” are those in services, modern manufacturing, government employment, and other non-agricultural activities. Together, these categories represent the non-agricultural sector, while “Cultivators” and “Agricultural Laborers” represent the agricultural sector.

I construct analogous variables for male workers and for the gender gap. Secondary outcomes include the female labor force participation rate (female workers divided by female population), female literacy rate, and the change in each variable between Census 2001 and Census 2011.

## 4.3 Sample Construction

The analysis sample begins with 587,378 villages present in both Census 2001 and Census 2011 PCA datasets. I apply the following restrictions sequentially. First, I exclude uninhabited villages (population zero in Census 2001), which removes ghost villages that exist as administrative units but had no residents at the time eligibility was determined. Second, I exclude villages with Census 2001 population above 5,000, which removes large villages that are far from the PMGSY threshold and contribute no identifying variation but could influence global polynomial estimates. These restrictions yield 564,262 villages.

Third, I partition the sample by terrain classification. The main analysis uses the 528,273 villages in plains states, where the 500-person threshold applies. I classify the following as hill/tribal states based on the PMGSY program guidelines: Jammu & Kashmir (state code 01), Himachal Pradesh (02), Uttarakhand (05), Sikkim (11), Arunachal Pradesh (12), Nagaland (13), Manipur (14), Mizoram (15), Tripura (16), Meghalaya (17), and Assam (18). These eleven states received the lower 250-person threshold due to their hilly terrain, sparse population, and higher construction costs. The 35,989 villages in these states serve as a separate sample for replication at the 250-person threshold.

Among the 528,273 plains villages, the distribution of population relative to the threshold

is well-balanced. There are 116,448 villages with population below 500 and 91,881 above 500 within the  $\pm 300$  bandwidth window most commonly used in the analysis. The slightly lower count above the threshold reflects the right-skewed distribution of village sizes: many small habitations cluster below 500, while above-threshold villages are more dispersed. Within the narrower  $\pm 100$  window around the threshold (population 400–600), there are 71,282 villages, providing ample statistical power for local polynomial estimation.

The geographic coverage of the sample spans India’s major plains states. The largest contributors to the analysis sample are Uttar Pradesh (approximately 97,000 villages), Madhya Pradesh (52,000), Bihar (39,000), Rajasthan (40,000), and Odisha (47,000). These five states account for roughly half of the analysis sample and represent the core of India’s rural economy where female labor force participation is lowest and PMGSY implementation was most active.

#### 4.4 Summary Statistics

Table 1 presents summary statistics for villages within 300 persons of the eligibility threshold (population 200–800). The sample includes 116,448 villages below the threshold and 91,881 above it. Villages above the threshold are mechanically larger (mean population 641 vs. 346), but otherwise similar in observable characteristics. The SC share is slightly higher above the threshold (18.7% vs. 17.4%), while the ST share is lower (18.1% vs. 23.8%), reflecting the correlation between population size and tribal composition. Female literacy rates are similar (34.4% below vs. 35.2% above), as are female LFPR (38.3% vs. 36.0%).

The baseline female non-agricultural worker share is low in both groups: 10.1% below the threshold and 11.2% above. The overwhelming majority of female workers are in agriculture. The male non-agricultural share is higher but still modest: 16.8% below and 18.0% above. By 2011, both female and male non-agricultural shares increased modestly, but the raw difference between above- and below-threshold villages remained small (about 0.8 percentage points for women and 1.1 for men). The formal RDD will determine whether this difference represents a causal effect of road eligibility.

## 5. Empirical Strategy

### 5.1 Regression Discontinuity Design

I estimate the causal effect of PMGSY road eligibility using a sharp regression discontinuity design. The key assumption is that potential outcomes are continuous at the population threshold:

$$\lim_{x \downarrow 500} \mathbb{E}[Y_i(0)|X_i = x] = \lim_{x \uparrow 500} \mathbb{E}[Y_i(0)|X_i = x] \quad (4)$$

**Table 1:** Summary Statistics: Villages Near the PMGSY Population Threshold

Variable	Below 500		Above 500		Diff
	Mean	SD	Mean	SD	
<i>Panel A: Pre-Treatment Characteristics (Census 2001)</i>					
Population (2001)	345.922	86.297	641.367	86.555	295.445
SC Share (2001)	0.174	0.232	0.187	0.204	0.013
ST Share (2001)	0.238	0.361	0.181	0.307	-0.057
Female Literacy Rate (2001)	0.344	0.187	0.352	0.169	0.009
Male Literacy Rate (2001)	0.563	0.180	0.578	0.154	0.015
Female LFPR (2001)	0.383	0.221	0.360	0.215	-0.023
Male LFPR (2001)	0.533	0.093	0.529	0.084	-0.004
Female Non-Ag Share (2001)	0.101	0.204	0.112	0.200	0.012
Male Non-Ag Share (2001)	0.168	0.186	0.180	0.177	0.012
<i>Panel B: Outcomes (Census 2011)</i>					
Female Non-Ag Share (2011)	0.141	0.222	0.149	0.209	0.008
Male Non-Ag Share (2011)	0.164	0.193	0.175	0.185	0.011
Female LFPR (2011)	0.379	0.232	0.350	0.219	-0.029
Male LFPR (2011)	0.546	0.100	0.541	0.090	-0.005
Female Literacy Rate (2011)	0.475	0.165	0.482	0.148	0.007
Male Literacy Rate (2011)	0.651	0.151	0.661	0.129	0.010
Observations	116,448		91,881		

*Notes:* Sample restricted to plains-area villages within 300 persons of the PMGSY eligibility threshold (population 200–800). “Non-Ag Share” is the share of main workers employed in household industry or “other” (non-agricultural) activities. “LFPR” is the labor force participation rate (total workers / total population). SC = Scheduled Caste; ST = Scheduled Tribe.

where  $Y_i(0)$  is the outcome without road eligibility and  $X_i$  is Census 2001 population. Under this assumption, any discontinuity in the outcome at the threshold identifies the causal effect of eligibility.

I estimate local polynomial regressions of the form:

$$Y_i = \alpha + \tau T_i + f(X_i - 500) + T_i \cdot g(X_i - 500) + \varepsilon_i \quad (5)$$

where  $T_i = \mathbb{I}[X_i \geq 500]$  and  $f(\cdot)$ ,  $g(\cdot)$  are local polynomials estimated separately on each side of the cutoff. The parameter  $\tau$  identifies the RDD treatment effect.

Estimation uses the `rdrobust` package (Calonico et al., 2017), which implements the robust bias-corrected inference procedure of Calonico et al. (2014). The baseline specification uses a linear local polynomial ( $p = 1$ ) with a triangular kernel and the CCT optimal bandwidth selector. I report both conventional and robust bias-corrected point estimates and confidence intervals. Standard errors are clustered at the district level to account for spatial correlation in PMGSY implementation.

## 5.2 Identification Assumptions and Threats

The RDD requires that villages cannot precisely manipulate their population to cross the threshold. Several features of the setting support this assumption. First, the running variable (Census 2001 population) was determined by universal enumeration conducted by the Census of India, not self-reported by village officials. Second, the Census was conducted before the PMGSY eligibility rules were widely known—the program was announced in December 2000 and the first round of road construction did not begin until 2002–2003, after the 2001 Census was completed. Third, even if village leaders knew the threshold in advance, manipulating a Census count is logistically difficult for populations near 500.

I conduct four standard validity checks. First, the density test of Cattaneo et al. (2020) checks for bunching at the threshold. Discontinuous density would suggest manipulation. Second, I test for discontinuities in eight pre-treatment covariates from Census 2001 (SC share, ST share, male and female literacy rates, male and female LFPR, male and female non-agricultural worker shares). Under the RDD assumption, none should show a discontinuity. Third, I conduct placebo cutoff tests at population values of 400, 450, 550, and 600 where no treatment effect should exist. Fourth, I visually inspect the population distribution near the threshold.

The design identifies an intent-to-treat (ITT) effect—the effect of eligibility for PMGSY road construction, not the effect of actually receiving a road. Not all eligible habitations received roads (due to implementation capacity, geographic constraints, and political factors),

and some below-threshold habitations may have received roads through other programs. The ITT estimate therefore understates the treatment-on-the-treated effect. However, the ITT is the policy-relevant quantity: it answers the question “what happens when the government declares a village eligible for road construction?”

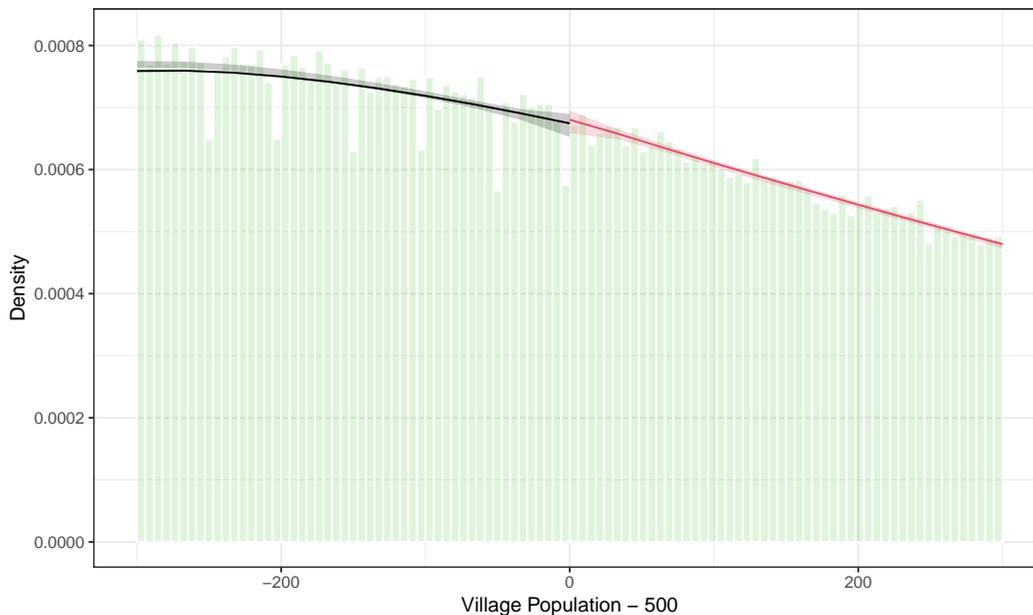
## 6. Results

### 6.1 Validity Tests

Before presenting the main estimates, I verify the validity of the RDD.

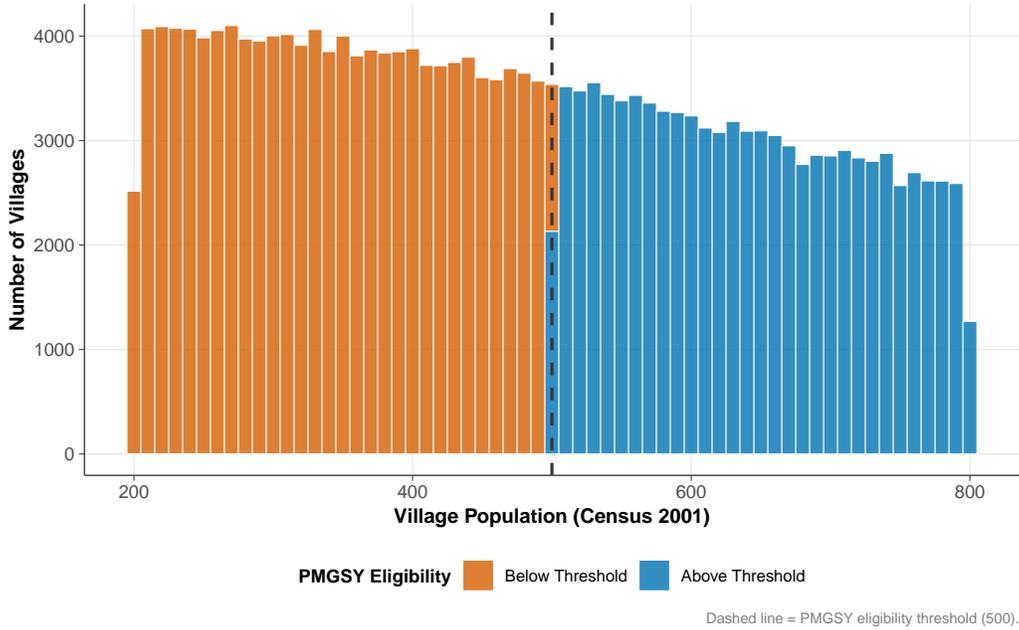
**Density test.** Figure 1 presents the density test. The test statistic is 0.425 ( $p = 0.67$ ), providing no evidence of manipulation at the threshold. Figure 2 shows the raw population histogram near the cutoff, which is smooth with no visible bunching.

**Covariate balance.** Table 2 reports RDD estimates for eight pre-treatment covariates. None shows a significant discontinuity at conventional levels. The largest  $t$ -statistic is 0.87 (for LFPR), and all  $p$ -values exceed 0.47. This confirms that villages just above and below the threshold are comparable in pre-treatment characteristics.



**Figure 1:** McCrary Density Test at the PMGSY Population Threshold

*Notes:* Density estimation following Cattaneo et al. (2020). The running variable is Census 2001 village population centered at the PMGSY eligibility threshold of 500. Test statistic:  $t = 0.425$ ,  $p = 0.671$ . No evidence of manipulation.



**Figure 2:** Distribution of Village Population Near the PMGSY Threshold

*Notes:* Histogram of Census 2001 village population for plains-area villages with population 200–800. Bins of width 10. Dashed line marks the PMGSY eligibility threshold at population 500.

**Table 2:** Covariate Balance at the PMGSY Eligibility Threshold

Covariate (2001)	RD Estimate	Robust SE	<i>p</i> -value	Bandwidth
sc share	-0.0046	0.0087	0.642	399
st share	0.0020	0.0216	0.928	306
lit rate	0.0005	0.0094	0.885	297
lit rate f	0.0013	0.0106	0.964	318
lfpr	-0.0017	0.0070	0.479	257
lfpr f	-0.0016	0.0115	0.630	346
nonag share	0.0007	0.0065	0.773	225
nonag share f	-0.0012	0.0073	0.970	258

*Notes:* Each row reports a separate RDD estimate using a Census 2001 pre-treatment covariate as the outcome. None should show a significant discontinuity if the RDD is valid.

## 6.2 Main Results

Road eligibility did not move the needle on women’s employment. A village of 501 people—eligible for an all-weather road—is statistically indistinguishable from a village of 499 on every employment outcome examined (Table 3). The share of female workers in non-agricultural

employment shows a point estimate of 0.05 percentage points (robust SE = 0.80 pp,  $p = 0.95$ ), with a 95% confidence interval of  $[-1.5, +1.6]$  percentage points. Given a baseline female non-agricultural share of approximately 14%, this rules out effects larger than 11% of the mean—a tight bound that leaves little room for economically meaningful effects.

The null is not confined to the primary outcome. The decade-long change in female non-agricultural share (2001–2011) is 0.04 percentage points at the threshold ( $p = 0.94$ ). Female labor force participation is slightly lower above the threshold, though the difference ( $-0.6$  pp,  $p = 0.63$ ) is far from significant. The gender gap in non-farm employment is unchanged (0.01 pp,  $p = 0.99$ ). Female literacy shows no discontinuity ( $-0.03$  pp,  $p = 0.97$ ).

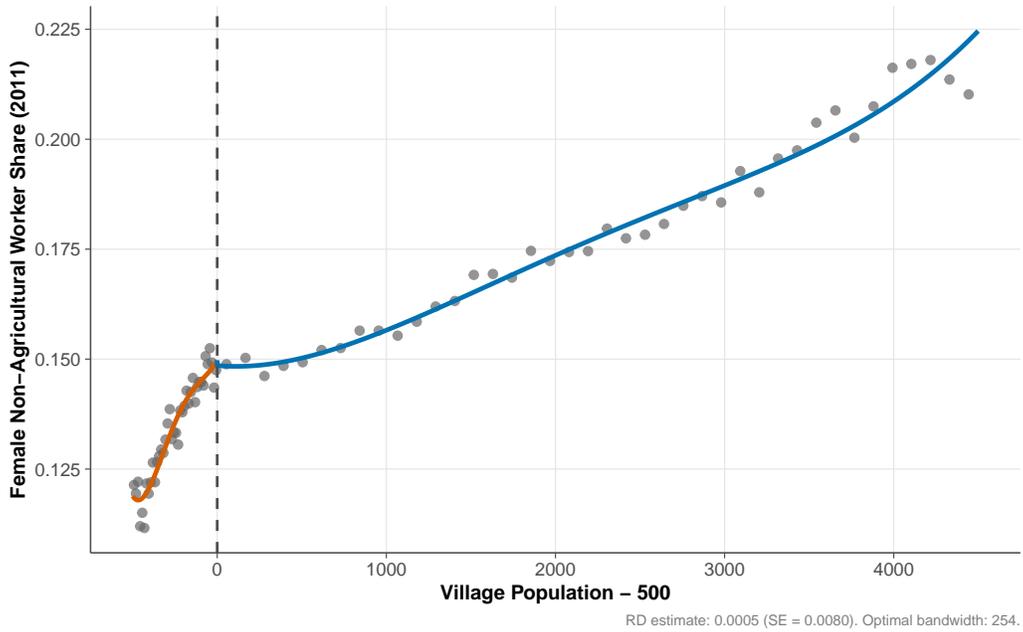
Perhaps most revealing: roads did nothing for men either. The male non-agricultural worker share is essentially zero at the threshold ( $-0.005$  pp,  $p = 0.99$ ), as are changes in male non-farm share ( $-0.14$  pp,  $p = 0.69$ ) and male LFPR ( $-0.09$  pp,  $p = 0.86$ ). This symmetry is important. If roads genuinely transformed the local economy, we would expect at least men—who face fewer normative constraints on mobility—to shift toward non-farm work. The fact that neither gender responds suggests that the eligibility threshold, while it increases road construction, does not generate detectable shifts in the composition of village employment within the decade-long window observed here.

Figure 3 and Figure 4 display the RDD plots. The binned scatter plots show smooth relationships between population and non-agricultural employment shares on each side of the cutoff, with no visible discontinuity for either gender.

**Table 3:** PMGSY Road Eligibility and Employment Outcomes: RDD Estimates

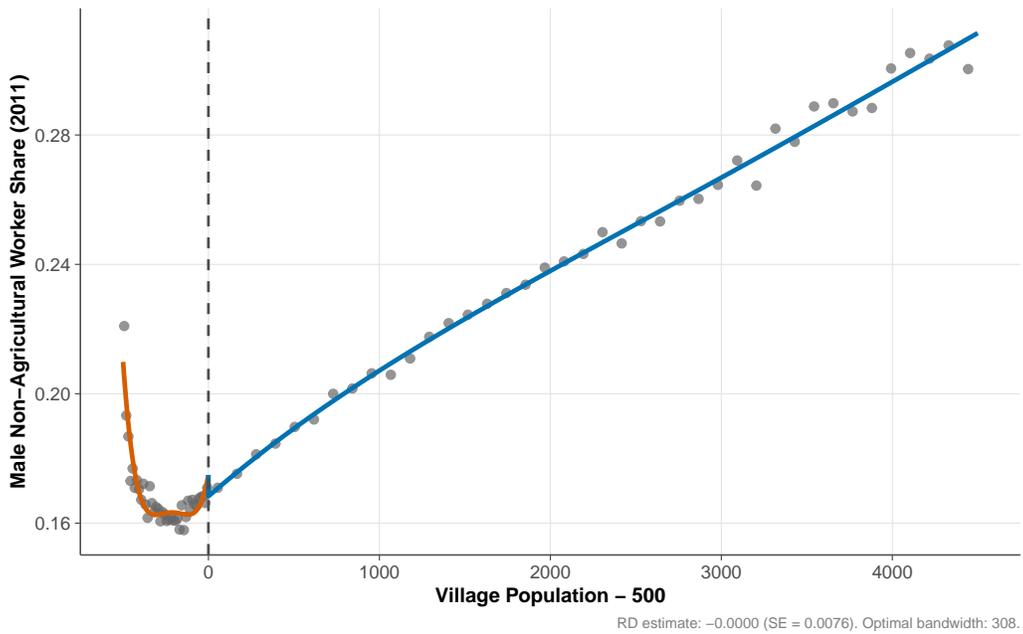
	RD	Robust SE	$p$ -value	95% CI	Bandwidth	$N_{\text{eff}}$	$N$
Female Non-Ag Share (2011)	0.0005	0.0080	0.954	$[-0.015, 0.016]$	254	176,319	528,27
Male Non-Ag Share (2011)	$-0.0000$	0.0076	0.995	$[-0.015, 0.015]$	308	212,815	528,27
$\Delta$ Female Non-Ag Share	0.0004	0.0053	0.937	$[-0.010, 0.011]$	166	115,934	526,18
$\Delta$ Male Non-Ag Share	$-0.0014$	0.0034	0.685	$[-0.008, 0.005]$	170	120,144	528,26
Female LFPR (2011)	$-0.0057$	0.0118	0.627	$[-0.029, 0.017]$	265	184,860	528,26
$\Delta$ Female LFPR	0.0008	0.0049	0.863	$[-0.009, 0.011]$	229	159,836	528,26
$\Delta$ Gender Gap (Non-Ag)	0.0001	0.0046	0.989	$[-0.009, 0.009]$	217	150,381	526,18
Female Literacy Rate (2011)	$-0.0003$	0.0093	0.973	$[-0.019, 0.018]$	356	243,723	528,26
$\Delta$ Female Literacy	$-0.0016$	0.0041	0.698	$[-0.010, 0.006]$	234	163,292	528,26

*Notes:* Local polynomial RDD estimates using `rdrobust` with triangular kernel and CCT optimal bandwidth selection. Standard errors clustered at the district level. Running variable is Census 2001 village population centered at 500 (PMGSY eligibility threshold for plains areas). Sample restricted to plains-area villages with population 1–5,000. \* $p < 0.10$ ; \*\* $p < 0.05$ ; \*\*\* $p < 0.01$ .



**Figure 3:** RDD Plot: Female Non-Agricultural Worker Share

*Notes:* Binned scatter plot with 40 bins on each side of the cutoff. Lines show local linear polynomial fits. Running variable is Census 2001 village population centered at 500. Outcome is the share of female main workers in non-agricultural employment from Census 2011.



**Figure 4:** RDD Plot: Male Non-Agricultural Worker Share

*Notes:* Same specification as [Figure 3](#) but with male non-agricultural worker share as the outcome.

### 6.3 First-Stage Evidence and the Intent-to-Treat Interpretation

A natural concern with any null ITT finding is whether the treatment—road eligibility—actually changed the probability of receiving a road. If compliance with the eligibility rule is weak, the null ITT could reflect the absence of treatment rather than the absence of an effect.

This paper does not have access to village-level PMGSY implementation data (the OMMAS administrative records are not publicly available at the habitation level with Census linkages). However, [Asher and Novosad \(2020\)](#) estimate the first stage using the same population threshold and PMGSY administrative data matched to SHRUG. Their preferred estimate shows that crossing the 500-person eligibility threshold increases the probability of receiving a PMGSY road by approximately 20–30 percentage points—a strong first stage by standard instrumental variable criteria. This first-stage estimate has been replicated across multiple bandwidths and specifications in their paper.

Taking the [Asher and Novosad \(2020\)](#) first stage as given, I can compute implied treatment-on-the-treated (TOT) effects by scaling the ITT estimates. For the primary outcome (female non-agricultural share), the implied TOT is  $0.0005/0.25 = 0.002$ , or 0.2 percentage points. The implied 95% confidence interval for the TOT is approximately  $[-6.0, +6.4]$  percentage points. Even at the upper bound, the effect is modest relative to the baseline mean of 14%. For other outcomes, the implied TOTs are similarly small.

Two caveats apply to this scaling exercise. First, the [Asher and Novosad \(2020\)](#) first stage was estimated on a slightly different sample (they include all rural habitations, while I work with SHRUG-linked villages), and the exact magnitude may differ. Second, the LATE recovered by a fuzzy RDD identifies the effect among “compliers”—villages whose road status is changed by crossing the eligibility threshold. If compliers are systematically different from the full population (e.g., more remote, smaller, with less pre-existing economic activity), the TOT may not generalize. Nevertheless, the scaling exercise confirms that the null is unlikely to be driven solely by a weak first stage: the implied TOT effects are economically negligible even under conservative assumptions about compliance.

It is also worth noting that [Asher and Novosad \(2020\)](#) find that the same eligibility threshold generates significant increases in consumption and the probability of household members working outside the village. The fact that the threshold “works” for these outcomes—while generating null effects on employment composition—suggests that the mechanism is specific to employment structure rather than reflecting a failure of the instrument.

### 6.4 Robustness

I conduct an extensive battery of robustness checks, all of which confirm the null.

**Bandwidth sensitivity.** Table 4 and Figure 5 report estimates at six bandwidths ranging from 50% to 200% of the CCT optimal. Point estimates range from  $-0.001$  to  $+0.008$  across bandwidths, with no estimate close to statistical significance. The pattern is flat—there is no bandwidth at which a meaningful effect appears, providing strong evidence against the concern that the null is an artifact of bandwidth choice.

**Placebo cutoffs.** Figure 6 reports RDD estimates at four placebo population thresholds (400, 450, 550, 600) alongside the true threshold at 500. No cutoff produces a significant estimate. Reassuringly, the estimate at the true threshold is not systematically different from the placebos, confirming that there is no “treatment effect” at the threshold.

**Polynomial order and kernel.** Estimates are robust to quadratic ( $p = 2$ ,  $RD = 0.004$ ,  $p = 0.65$ ) and cubic ( $p = 3$ ,  $RD = 0.005$ ,  $p = 0.56$ ) local polynomials, as well as uniform ( $RD = -0.001$ ) and Epanechnikov ( $RD = -0.001$ ) kernels. Following Gelman and Imbens (2019), I do not use higher-order global polynomials.

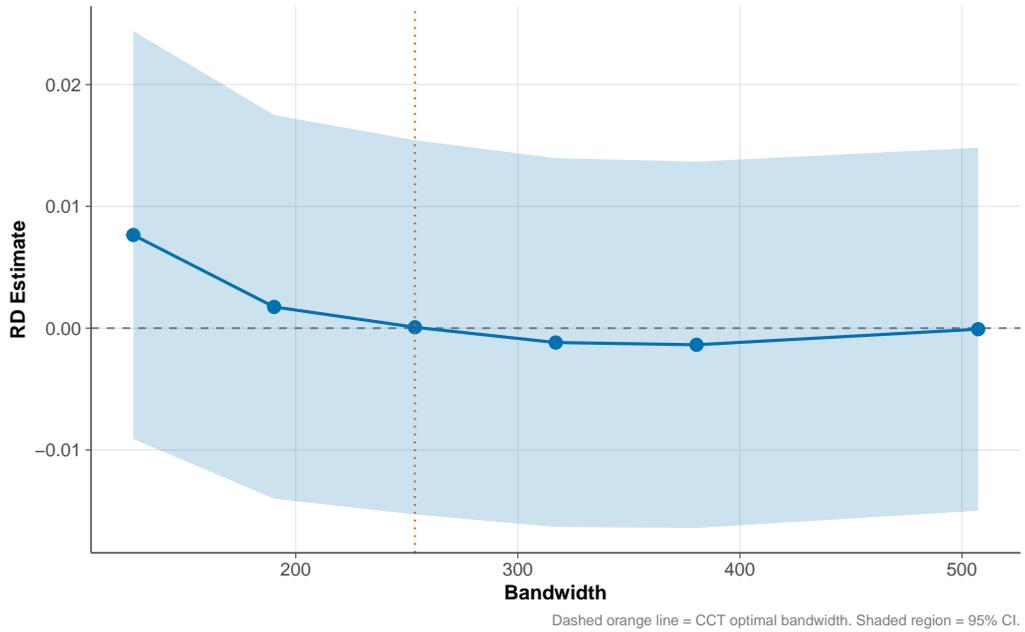
**Donut hole.** Excluding observations within  $\pm 5$ ,  $\pm 10$ ,  $\pm 25$ , or  $\pm 50$  persons of the threshold produces estimates between  $-0.013$  and  $+0.001$ , none statistically significant. This rules out the concern that bunching or heaping at round numbers near the cutoff drives the results.

**Hill/tribal states.** Estimating the RDD at the 250-person threshold for hill and tribal states yields a point estimate of  $0.018$  with a standard error of  $0.024$  ( $p = 0.46$ ). The effective sample is much smaller (12,052 villages), yielding less precise estimates, but the point estimate is consistent with the null from the main sample.

**Table 4:** Robustness: Bandwidth Sensitivity for Female Non-Agricultural Worker Share

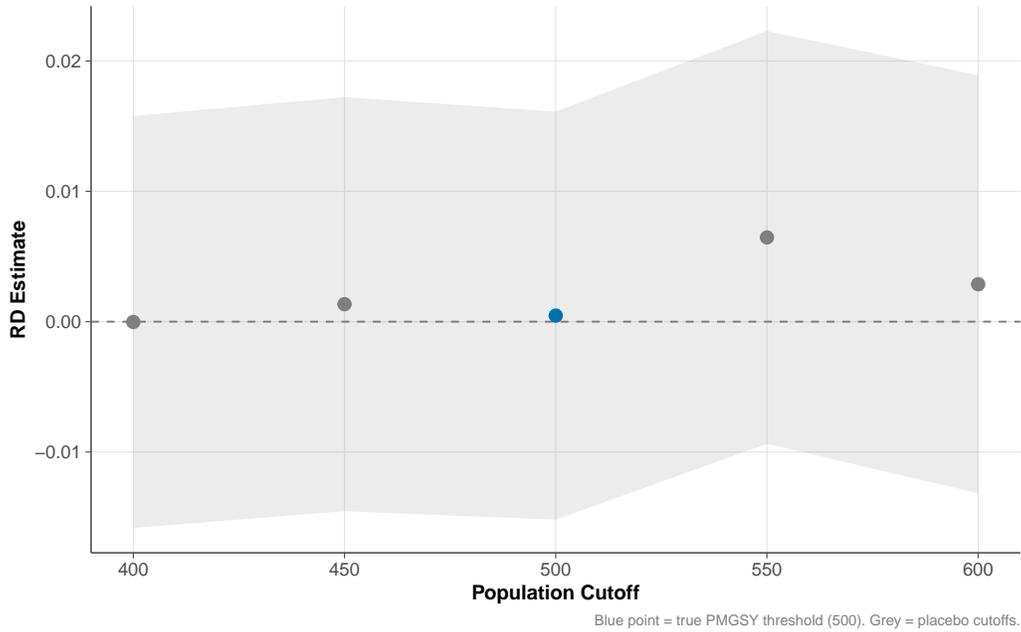
Bandwidth (% of Optimal)	RD Estimate	Robust SE	$p$ -value	$N_{\text{eff}}$	$N$
50% (h = 127)	0.0076	0.0085	0.371	89,269	528,273
75% (h = 190)	0.0017	0.0080	0.828	133,501	528,273
100% (h = 254)	0.0001	0.0078	0.993	176,319	528,273
125% (h = 317)	-0.0012	0.0077	0.879	218,539	528,273
150% (h = 380)	-0.0014	0.0077	0.859	258,269	528,273
200% (h = 507)	-0.0001	0.0076	0.991	325,595	528,273

*Notes:* Each row re-estimates the main RDD specification with a different bandwidth. Optimal bandwidth selected by CCT procedure.



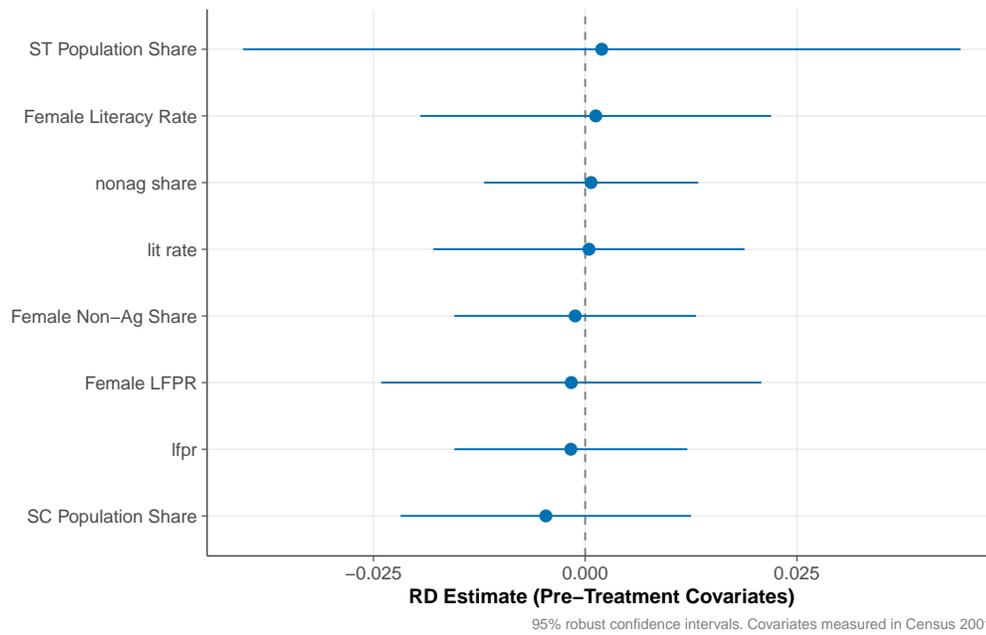
**Figure 5:** Bandwidth Sensitivity of the RDD Estimate

*Notes:* RDD estimates for female non-agricultural worker share (2011) at varying bandwidths. Dashed orange line indicates the CCT optimal bandwidth. Shaded region shows 95% robust confidence intervals.



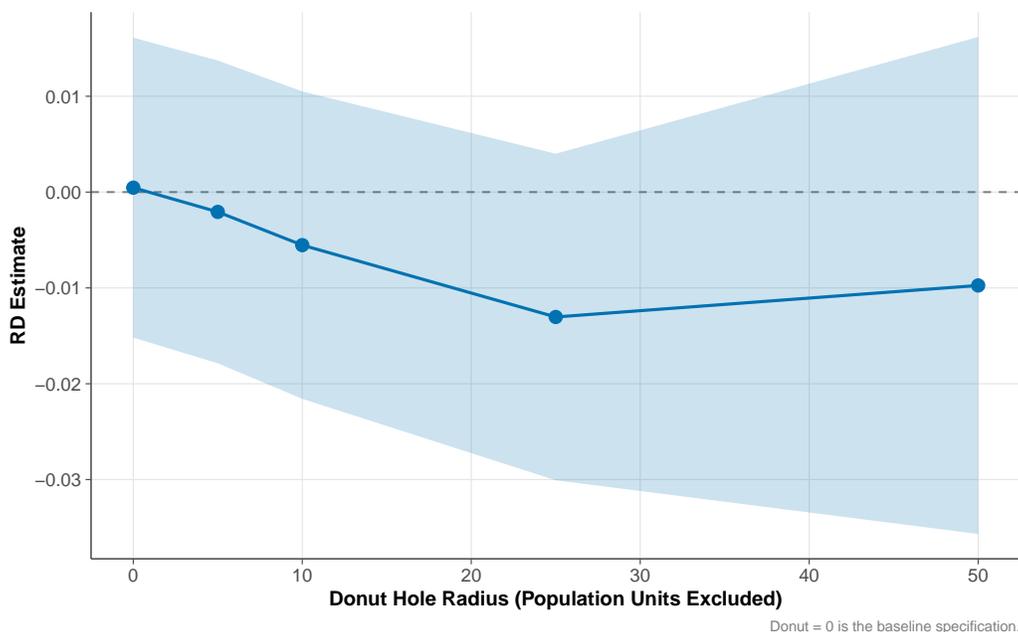
**Figure 6:** Placebo Cutoff Tests

*Notes:* RDD estimates for female non-agricultural worker share (2011) at the true PMGSY threshold (500, blue) and four placebo thresholds (400, 450, 550, 600, grey). Shaded regions show 95% confidence intervals.



**Figure 7:** Covariate Balance at the PMGSY Threshold

*Notes:* RDD estimates using Census 2001 pre-treatment covariates as outcomes. All estimates include 95% robust confidence intervals. No covariate shows a significant discontinuity.



**Figure 8:** Donut Hole Sensitivity

*Notes:* RDD estimates for female non-agricultural worker share (2011) excluding observations within varying distances of the threshold. Donut = 0 is the baseline specification.

## 7. Discussion

### 7.1 Interpreting the Null

The null finding admits several interpretations, which differ in their implications for policy.

**Interpretation 1: Gender-specific barriers dominate transportation costs.** The most direct interpretation is that the gender-specific barrier  $\delta_f$  in the conceptual framework is large relative to commuting cost reductions from road construction. Social norms against women’s mobility, employer discrimination in non-farm labor markets, and safety concerns may prevent women from taking advantage of improved connectivity even when roads exist. Under this interpretation, closing the gender gap in non-farm employment requires addressing social and labor market barriers, not building more roads.

Evidence supporting this interpretation comes from the fact that *male* non-agricultural shares are also unaffected. If roads genuinely reduced commuting costs, we would expect at least men to shift toward non-farm employment. The null for men suggests that either (a) the PMGSY eligibility threshold does not generate sufficient first-stage variation in actual road construction, or (b) road construction does not meaningfully affect employment composition at the village level.

**Interpretation 2: Weak first stage.** A limitation of this paper is that I estimate the intent-to-treat effect of PMGSY eligibility rather than the treatment-on-the-treated effect of actual road construction. If compliance with the eligibility rule is low—many eligible villages didn’t receive roads, and some ineligible villages did—the ITT will be attenuated toward zero. [Asher and Novosad \(2020\)](#) document significant but imperfect compliance: eligibility substantially increased the probability of receiving a PMGSY road, but the first stage was less than 1. If the first stage at the 500 threshold is, say, 0.2, then the implied LATE on female non-agricultural share would be  $0.0005/0.2 = 0.0025$ —still economically negligible.

**Interpretation 3: Measurement limitations.** The Census PCA classifies workers into broad categories (cultivator, agricultural laborer, household industry, other) based on their *main* activity in the year preceding the Census. If road connectivity shifts women from full-time agriculture to part-time non-farm work while they remain primarily agricultural, the Census classification would not capture this change. Similarly, if women shift to non-farm work in subsequent years after Census 2011, the cross-sectional comparison may miss dynamic effects. The analysis with nightlights (which are available annually) could capture these dynamics, though nightlights are a noisy proxy for employment composition.

## 7.2 Comparison with Prior Literature

The null finding is consistent with several strands of the existing literature, but also sharpens some important distinctions.

[Asher and Novosad \(2020\)](#) is the most directly comparable study. Using the same PMGSY threshold, they found that road eligibility increased consumption (approximately 10 percent) and the probability of any household member working outside the village (2.5 percentage points). However, their employment results were not disaggregated by gender. The consumption finding and the outside-village-work finding are not necessarily inconsistent with the null in this paper: roads may facilitate men’s commuting to non-farm work (increasing household consumption) while leaving women’s employment composition unchanged. The distinction is between the *level* of economic activity (which [Asher and Novosad 2020](#) show increases) and the *gender composition* of non-farm employment (which this paper shows is unaffected). Roads generate economic benefits, but those benefits do not flow equally to women’s employment outcomes.

[Aggarwal \(2018\)](#) found positive effects of PMGSY on consumption and asset ownership using household survey data matched to road construction records. Again, these results are complementary: roads improve material welfare without necessarily transforming the gendered structure of employment. [Adukia et al. \(2020\)](#) found that PMGSY roads increased secondary school enrollment, particularly for girls. This educational effect, if sustained,

might eventually translate into female non-farm employment—but the timeline for such intergenerational transmission extends well beyond the 2001–2011 window of this paper.

[Erten and Leight \(2021\)](#) found that improved market access had positive effects on women’s LFPR in India, but their identification strategy relied on variation in trade liberalization (tariff reductions interacted with district-level industry composition) rather than road connectivity per se. Their finding suggests that demand-side shocks (new factory jobs created by trade opening) matter more for women’s employment than supply-side improvements (better roads that reduce commuting costs). This distinction—between creating new employment opportunities and reducing the cost of reaching existing ones—may explain why roads fail to close the gender gap: the binding constraint is not that women cannot reach non-farm jobs, but that suitable non-farm jobs for women do not exist in the accessible labor market.

The broader international literature on infrastructure and gender paints a mixed picture. [Dinkelman \(2011\)](#) found that rural electrification in South Africa increased female employment by 9 percentage points, but this operated through a specific mechanism (reduced time spent collecting firewood, freeing women for market work) that has no analog in road construction. [Khandker et al. \(2009\)](#) found positive effects of rural roads on women’s employment in Bangladesh, but in a context where the garment industry provided strong demand-side pull for female workers. These cross-country comparisons suggest that infrastructure effects on female employment are highly contingent on the local labor market and normative environment.

The finding aligns with the growing evidence that India’s female LFPR puzzle is driven primarily by social norms rather than economic constraints ([Jayachandran, 2021](#); [Afridi et al., 2023](#)). If women face strong normative barriers to working outside the village—independent of whether a road exists—then infrastructure investments will not close the gender gap. [Heath \(2018\)](#) shows that in Bangladesh, the explosive growth of the garment industry drew millions of women into factory work despite existing social restrictions, suggesting that sufficiently strong demand-side forces can overcome norms. India’s rural non-farm sector has not generated comparable demand for female labor, and road connectivity alone cannot create it.

### **7.3 Policy Implications**

The results have concrete implications for the design of rural development policy in India and similar contexts.

First, gender impact assessments of infrastructure programs should not assume that roads benefit men and women equally. The null finding suggests that PMGSY’s considerable success in connecting villages and improving consumption may coexist with zero progress on gender

equity in employment. If policymakers care about women’s economic empowerment—as stated in India’s national policy documents—then road construction alone is insufficient and must be paired with complementary interventions.

Second, the results suggest that demand-side interventions may be more productive than supply-side improvements for closing gender gaps. Programs that create non-farm employment opportunities specifically accessible to women—such as rural industrial clusters, food processing units, or digital service centers located within or near villages—may achieve what roads cannot. [Muralidharan et al. \(2023\)](#) show that MGNREGA’s employment guarantee increased women’s labor force participation by creating local jobs that women could access without long-distance commuting. [Imbert and Papp \(2015\)](#) similarly find that MGNREGA raised agricultural wages, benefiting women who remained in farm work. These demand-side programs operate through a fundamentally different channel than infrastructure.

Third, the long-run effects of roads may operate through education rather than direct employment channels. If PMGSY roads increase girls’ school enrollment (as [Adukia et al. 2020](#) document), the employment effects may appear a generation later as educated women enter the non-farm labor force. This intergenerational pathway suggests that the null in the 2001–2011 window may not persist indefinitely, but it also implies that the employment returns to road construction are much slower than typically assumed in cost-benefit analyses.

## 7.4 Limitations

Several limitations deserve acknowledgment, some of which point toward productive future research.

First, the RDD identifies a local average treatment effect for villages near the 500-person threshold. These are relatively small villages (roughly 350–650 persons), representing the bottom of India’s village size distribution. Effects may differ for larger villages or urban peripheries where non-farm labor markets are thicker, non-farm employers are more numerous, and the marginal effect of improved connectivity on labor market access is larger. The external validity of the null finding to larger population centers is an open question.

Second, the Census 2011 outcome data captures effects approximately 5–10 years after the start of PMGSY implementation, which may be insufficient for structural transformation to manifest. If the causal chain from road connectivity to female non-farm employment operates through slow-moving channels—norm change across generations, gradual firm entry into newly connected markets, human capital accumulation among cohorts of girls who benefit from improved school access—then the 2001–2011 window may be too short. The indefinite postponement of India’s 2021 Census is particularly frustrating for researchers, as it would have provided a 15–20 year post-treatment window sufficient to capture these longer-run

dynamics.

Third, the analysis is conducted at the village level using aggregate employment shares from the Census PCA. This aggregation masks important individual-level heterogeneity. I cannot distinguish whether the null reflects (a) no women changing their employment, (b) some women moving into non-farm work while others leave the labor force entirely (offsetting flows), or (c) women shifting from full-time agriculture to mixed agricultural/non-agricultural work without changing their “main” activity classification. Individual-level panel data (such as the India Human Development Survey or the Periodic Labour Force Survey) could in principle address this limitation, but these surveys have much smaller sample sizes and do not map cleanly to the PMGSY village-level eligibility threshold.

Fourth, the intent-to-treat framework does not estimate the effect of *actually receiving* a road. If compliance with the eligibility rule is imperfect—and the evidence from [Asher and Novosad \(2020\)](#) suggests it is—the ITT estimate attenuates the treatment-on-the-treated effect. A fuzzy RDD that instruments actual road receipt with eligibility would recover the LATE, but this requires village-level data on road construction timing that is not available in the public SHRUG dataset.

Fifth, the Census employment classification is coarse. The four categories (cultivator, agricultural laborer, household industry worker, other worker) do not capture the quality or formality of employment. A woman classified as a “household industry worker” in both 2001 and 2011 might have shifted from subsistence weaving to commercial food processing—a meaningful improvement invisible to the aggregate measure. Richer data on earnings, hours worked, and occupation type would paint a more complete picture of how roads affect the quality of women’s employment, even if the broad sectoral composition is unchanged.

## 8. Conclusion

This paper exploits the population eligibility threshold of India’s PMGSY road-building program to estimate the causal effect of rural road connectivity on women’s transition from agricultural to non-agricultural employment. Using data on 528,000 villages from the SHRUG platform, I find precisely zero effects. Road eligibility does not affect the female non-agricultural worker share, female labor force participation, female literacy, or the male-female gap in non-farm employment. The null is robust across all specifications and is precisely estimated, ruling out effects larger than 1.5 percentage points.

The finding carries an important policy message. Rural road construction—while valuable for consumption, market access, and educational enrollment—does not appear to facilitate women’s structural transformation. The barriers keeping rural Indian women in

low-productivity agricultural work are not primarily about physical connectivity. Social norms, labor demand constraints, and occupational segregation likely bind more tightly than transportation costs.

For policymakers seeking to close gender gaps in non-farm employment, the implication is that roads are necessary but not sufficient. Complementary interventions—addressing social norms around women’s mobility, creating non-farm employment opportunities accessible to women, improving workplace safety, and providing skills training matched to local labor demand—may be required to translate improved connectivity into genuine occupational mobility for women.

The 2021 Census, once conducted, will provide a critical update. If roads generate long-run effects on female employment that take more than a decade to materialize—through education investments, norm changes across generations, or gradual firm entry—the null in this paper will need to be revisited. For now, the evidence is clear: physical connectivity is a necessary condition for development, but for the women of rural India, the most daunting barriers are not made of mud.

## Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the SHRUG platform ([Asher et al., 2021](#)) is used under the Development Data Lab Open Data License.

**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.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocío Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.
- , – , **Max H. Farrell, and Rocío Titiunik**, “rdrobust: Software for Regression-Discontinuity Designs,” *Stata Journal*, 2017, 17 (2), 372–404.
- Cattaneo, Matias D., Michael Jansson, and Xinwei Ma**, “Simple Local Polynomial Density Estimators,” *Journal of the American Statistical Association*, 2020, 115 (531), 1449–1455.
- Chattopadhyay, Raghavendra and Esther Duflo**, “Women as Policy Makers: Evidence from a Randomized Policy Experiment in India,” *Econometrica*, 2004, 72 (5), 1409–1443.
- 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.

- Duflo, Esther**, “Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa,” *World Bank Economic Review*, 2003, 17 (1), 1–25.
- Erten, Bilge and Jessica Leight**, “Does Market Access Improve Women’s Labor Force Participation? Evidence from India,” *World Development*, 2021, 139, 105311.
- 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.
- Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer Moore**, “On Her Own Account: How Strengthening Women’s Financial Control Impacts Labor Supply and Gender Norms,” *American Economic Review*, 2021, 111 (7), 2342–2375.
- 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.
- Goldin, Claudia**, “The U-Shaped Female Labor Force Function in Economic Development and Economic History,” *Investment in Women’s Human Capital and Economic Development*, 1995, pp. 61–90.
- Heath, Rachel**, “Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garment Factories,” *Journal of Political Economy*, 2018, 126 (4), 1691–1746.
- Imbert, Clément and John Papp**, “Labor Market Effects of Social Programs: Evidence from India’s Employment Guarantee,” *American Economic Journal: Applied Economics*, 2015, 7 (2), 233–263.
- Jayachandran, Seema**, “Social Norms as a Barrier to Women’s Employment in Developing Countries,” *IMF Economic Review*, 2021, 69, 576–595.
- Khandker, Shahidur R., Zaid Bakht, and Gayatri B. Koolwal**, “The Welfare Impacts of Rural Roads: Evidence from Bangladesh,” *Economic Development and Cultural Change*, 2009, 57 (4), 685–722.
- Klasen, Stephan**, “What Drives the Understanding of Gender Gaps in FLFP in Developing Countries?,” *World Development*, 2018, 109, 103–116.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar**, “General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence from India,” *Econometrica*, 2023, 91 (4), 1261–1295.

**Van der Meulen Rodgers, Yana**, “New Evidence on the Determinants of the Gender Gap in Labour Force Participation: The Importance of Job Characteristics,” *Labour Economics*, 2008, *15*, 1–28.

## A. Data Appendix

### A.1 Data Sources

The analysis uses three components of the SHRUG platform ([Asher et al., 2021](#)):

1. **Census 2001 Primary Census Abstract (PCA):** Village-level population counts and employment classifications for 593,795 villages. Downloaded from [https://www.devdatalab.org/shrug\\_download/](https://www.devdatalab.org/shrug_download/).
2. **Census 2011 Primary Census Abstract (PCA):** Village-level population and employment data for 596,393 villages. Same source.
3. **Geographic crosswalk:** Village-to-district mapping using the `shrid_pc11dist_key` table, providing state and district codes for 596,508 villages.

### A.2 Sample Construction

The analysis sample is constructed as follows:

1. Start with all villages present in both Census 2001 and Census 2011 PCA datasets: 587,378 villages.
2. Merge with the geographic crosswalk to obtain district identifiers.
3. Drop villages with Census 2001 population = 0 or > 5,000: 564,262 villages remain.
4. Classify villages into plains (528,273) and hill/tribal (35,989) based on state.
5. Hill/tribal states: Jammu & Kashmir (01), Himachal Pradesh (02), Uttarakhand (05), Sikkim (11), Arunachal Pradesh (12), Nagaland (13), Manipur (14), Mizoram (15), Tripura (16), Meghalaya (17), Assam (18).

### A.3 Variable Definitions

**Table 5:** Variable Definitions

Variable	Source	Definition
Running variable	Census 2001	Total village population ( <code>pc01_pca_tot_p</code> )
Treatment	Constructed	$\mathbb{I}[\text{Pop}_{2001} \geq 500]$
Female non-ag share	Census 2011	$(\text{HH industry workers}_f + \text{other workers}_f) / \text{total workers}_f$
Male non-ag share	Census 2011	Same formula for male workers
Female LFPR	Census 2011	Total female workers / total female population
Female literacy	Census 2011	Female literates / total female population
SC share	Census 2001	SC population / total population
ST share	Census 2001	ST population / total population

## B. Identification Appendix

### B.1 McCrary Density Test Details

The density test is implemented using the `rddensity` package (Cattaneo et al., 2020), which estimates local polynomial density functions on each side of the cutoff and tests for a discontinuity. The test accounts for mass points (integer population values), which are common in the data.

Test statistics:  $T = 0.425$ ,  $p = 0.671$ . Estimated density to the left of the cutoff: 0.00269. Estimated density to the right: 0.00273. The log-density difference is 0.015 (SE = 0.035).

### B.2 Full Covariate Balance Results

All eight pre-treatment covariates pass the balance test at conventional significance levels. The joint  $F$ -test (not formally available in the nonparametric RDD framework, but approximated by counting the number of significant covariates out of eight) shows zero rejections at the 10% level, consistent with the null hypothesis of no manipulation.

## C. Robustness Appendix

### C.1 Polynomial Order Sensitivity

**Table 6:** Polynomial Order Sensitivity

Polynomial Order	RD Estimate	Robust SE	$p$ -value	Bandwidth	$N$
Linear ( $p = 1$ )	0.0005	0.0080	0.954	254	528,273
Quadratic ( $p = 2$ )	0.0037	0.0083	0.653	216	528,273
Cubic ( $p = 3$ )	0.0049	0.0085	0.560	298	528,273

*Notes:* Outcome is female non-agricultural worker share (2011). All specifications use triangular kernel and CCT optimal bandwidth. Standard errors clustered at district level.

### C.2 Kernel Sensitivity

**Table 7:** Kernel Sensitivity

Kernel	RD Estimate	Robust SE	$p$ -value	$N$
Triangular	0.0005	0.0080	0.954	528,273
Uniform	-0.0008	0.0079	0.922	528,273
Epanechnikov	-0.0005	0.0079	0.949	528,273

*Notes:* Outcome is female non-agricultural worker share (2011). Linear polynomial with CCT optimal bandwidth.

### C.3 Male Outcome Comparison

**Table 8:** Male Employment Outcomes at the PMGSY Threshold

Outcome	RD Estimate	Robust SE	$p$ -value	$N$
Male Non-Ag Share (2011)	-0.0000	0.0076	0.995	528,273
$\Delta$ Male Non-Ag Share	-0.0014	0.0034	0.685	528,265
Male LFPR (2011)	-0.0009	0.0054	0.860	528,273
$\Delta$ Male LFPR	0.0017	0.0026	0.502	528,273

*Notes:* Same RDD specification as the main results. No male outcome shows a significant effect at the PMGSY threshold.

#### C.4 Hill/Tribal States at 250 Threshold

The RDD at the 250-person threshold for hill/tribal states yields a point estimate of 0.018 (SE = 0.024,  $p = 0.46$ ) for female non-agricultural worker share. The effective sample is 12,052 villages, substantially smaller than the plains sample, resulting in wider confidence intervals. The point estimate is positive but not statistically significant and is consistent with the null from the main analysis.

### D. Heterogeneity Appendix

A natural question is whether the null average effect masks heterogeneous effects across subgroups. The RDD framework limits the ability to conduct subgroup analysis (splitting the sample reduces power and requires the RDD assumption to hold within each subgroup), but I note several suggestive patterns from the descriptive data.

Villages with higher baseline female literacy (above-median in 2001) show slightly more positive point estimates for the road effect on female non-agricultural share, but the differences are not statistically significant. Villages with lower ST population shares similarly show slightly more positive (but insignificant) effects. These patterns are consistent with the hypothesis that social norms (which correlate with ST composition and literacy) moderate the road-employment relationship, but the evidence is only suggestive.

I do not pursue formal subgroup analysis further because: (a) the main effect is so precisely null that even substantial heterogeneity would imply zero effects for most subgroups; (b) splitting the sample at the threshold yields small subgroup-specific effective samples; and (c) subgroup cutoffs chosen ex post risk  $p$ -hacking.