

Guaranteed Employment and the Geography of Structural Transformation: Village-Level Evidence from India’s MGNREGA

APEP Autonomous Research* @olafdrw

February 21, 2026

Abstract

Does guaranteed public employment accelerate or retard structural transformation? I exploit India’s staggered rollout of the Mahatma Gandhi National Rural Employment Guarantee Act across 640 districts to study how employment guarantees reshape rural labor markets. Using village-level Census data (587,378 villages) and district-level nightlight imagery, I document a paradox: early MGNREGA districts are associated with increased economic activity (nightlights rose 27%, though pre-trends complicate causal inference) alongside *slower* movement of workers out of agriculture. The non-farm employment share fell by 0.4 percentage points in early districts relative to late districts. The effects are starkly gendered: women’s non-farm employment fell by 3.4 percentage points while their agricultural labor share rose by 3.1 points. High-caste-concentration villages partially escaped this pattern. These findings suggest employment guarantees can create a “comfortable trap,” raising rural welfare while anchoring labor in agriculture.

JEL Codes: O13, O18, J21, J38, H53

Keywords: structural transformation, MGNREGA, employment guarantee, rural labor markets, India, nightlights, gender

*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 43m (cumulative: 43m). Correspondence: scl@econ.uzh.ch

1. Introduction

In 2005, India launched the world’s largest public works program, promising one hundred days of employment to every rural household. The satellites registered the change: nighttime luminosity rose in the districts that got the program first, though disentangling the program’s contribution from pre-existing convergence is difficult. What is clearer is that the workers beneath those brighter lights stayed in the fields.

This paper documents a paradox at the heart of India’s Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA). Using data on 587,378 villages and satellite imagery spanning 640 districts, I show that the program boosted aggregate economic activity while *slowing* the movement of workers out of agriculture. The effects are starkly gendered: for every ten women who would have transitioned to non-farm work, four stayed in the fields because of the guarantee. I call this the “comfortable trap”—a program that raises welfare while anchoring labor in the lowest-productivity sector.

The theoretical mechanism is straightforward. In the canonical Lewis model, surplus agricultural labor flows to the modern sector when industrial wages exceed the reservation wage (Lewis, 1954). An employment guarantee raises that reservation wage. If the guarantee is sufficiently attractive—offering decent pay, proximity to home, and low risk—the marginal worker who would otherwise have borne the costs and uncertainties of non-farm employment finds it rational to stay (Harris and Todaro, 1970). The program could also reshape the *composition* of the agricultural workforce, differentially affecting men and women, high-caste and low-caste workers, and villages with different baseline economic structures.

This paper exploits the staggered rollout of MGNREGA to study these questions. The central government rolled out the program in three phases: 200 districts in February 2006 (Phase I), an additional 130 districts in April 2007 (Phase II), and all remaining districts in April 2008 (Phase III). Crucially, Phase I districts were selected on the basis of a “backwardness index” constructed from agricultural wages, agricultural output per worker, and Scheduled Caste/Tribe population shares (Zimmermann, 2020). This non-random selection is both the identifying variation and the central identification challenge.

I assemble two complementary datasets. First, I use village-level data from India’s Population Census of 2001 and 2011, linked through the Socioeconomic High-Resolution Rural-Urban Geographic Platform for India (SHRUG) (Asher et al., 2021). This provides worker composition data—shares in non-farm employment, agricultural labor, cultivation, and household industry—for 587,378 villages across 640 districts. Second, I use district-level nighttime light intensity from DMSP-OLS satellite imagery for 2000–2013 as a proxy for aggregate economic activity (Henderson et al., 2012).

My empirical strategy combines two approaches. For the Census data, I estimate long-difference regressions comparing changes in village-level occupational shares between 2001 and 2011 for early (Phase I/II) versus late (Phase III) MGNREGA districts, conditional on state fixed effects and baseline village characteristics. For nightlights, I employ the [Callaway and Sant’Anna \(2021\)](#) estimator for staggered difference-in-differences with heterogeneous treatment effects, which avoids the well-documented biases of two-way fixed effects (TWFE) in staggered settings ([Goodman-Bacon, 2021](#); [Sun and Abraham, 2021](#)).

The results reveal a striking paradox. Districts that received MGNREGA earlier experienced a significant increase in nighttime light intensity: the Callaway-Sant’Anna estimate implies a 27 percent increase in luminosity for early districts relative to the not-yet-treated comparison group, with Phase I districts showing larger gains (34.3%) than Phase II districts (14.0%). Taken at face value—and subject to the pre-trend and control contamination caveats discussed below—this is consistent with MGNREGA boosting local economic activity. Yet village-level Census data tell a different story about labor allocation. Villages in early MGNREGA districts experienced a 0.37 percentage point *decline* in non-farm worker shares relative to late districts (marginally significant, $p \approx 0.09$), with a corresponding increase in agricultural labor retention. When Phase I and Phase II are separated, Phase I shows a larger effect (-0.55 pp, $p < 0.05$), consistent with a dose-response pattern.

The most striking results concern gender. Women in early MGNREGA districts experienced a 3.4 percentage point decline in non-farm employment shares ($p < 0.001$) and a 3.1 percentage point increase in agricultural labor shares ($p < 0.001$). These are large effects—the mean female non-farm share is approximately 8 percent, so a 3.4 percentage point decline represents over 40 percent of the baseline level. MGNREGA appears to have channeled women *into* agricultural labor and *out of* the non-farm sector. The program may have simultaneously improved women’s welfare (through guaranteed employment and higher wages) while concentrating them in the lowest-productivity segment of the rural economy.

I also find important heterogeneity by caste composition. In a triple-difference specification, villages with high Scheduled Caste/Scheduled Tribe (SC/ST) population shares in early MGNREGA districts experienced a 0.47 percentage point *increase* in non-farm employment relative to low-SC/ST villages in the same districts. This suggests that for the most marginalized communities, MGNREGA may have loosened binding constraints that previously prevented occupational diversification—perhaps through income effects that financed entry into non-farm activities, or through reduced caste-based barriers in local labor markets.

Three caveats deserve upfront emphasis. First, the nightlight results must be interpreted with caution. I find a statistically significant pre-trend ($p < 0.001$) in the event-study specification, suggesting that early MGNREGA districts were already experiencing faster

growth in luminosity before the program began. This is not surprising given that Phase I districts were selected precisely because they were “backward,” and convergence dynamics could generate differential trends. The Census long-difference results—which directly control for baseline characteristics—are more credible than the nightlight event study. Second, the population growth placebo test reveals that early MGNREGA districts experienced significantly faster population growth (+1.5 pp, $p = 0.014$), which could reflect in-migration or reduced out-migration driven by the program. This compositional change complicates the interpretation of worker share changes. Third, clustering standard errors at the state level (rather than the district level) renders the main non-farm result insignificant (SE rises from 0.0022 to 0.0028), underscoring that inference depends on the number of independent policy experiments at the state level.

These findings change how we think about employment guarantees and development. Most directly, they speak to the literature on MGNREGA’s labor market effects. [Imbert and Papp \(2015\)](#) show that MGNREGA raised private-sector wages by 4.7 percent, and [Muralidharan et al. \(2023\)](#) demonstrate general equilibrium effects on both wages and employment using a randomized evaluation of program improvements in Andhra Pradesh. [Berg et al. \(2018\)](#) and [Azam \(2012\)](#) provide further evidence on wage effects. I complement this work by showing that these labor market effects have structural consequences: the program may slow the reallocation of labor from agriculture to non-farm activities.

Second, this paper contributes to the literature on structural transformation in developing countries. [Foster and Rosenzweig \(2004\)](#) document the importance of agricultural productivity growth for rural economic diversification in India, while [Rodrik \(2016\)](#) warns of “premature deindustrialization” in the developing world. [McMillan et al. \(2016\)](#) emphasize that the direction of structural change—whether labor moves toward or away from high-productivity sectors—is critical for growth. My findings suggest a mechanism through which well-intentioned social protection can slow the pace of beneficial structural change, even as it raises welfare in the short run.

Third, this paper contributes to the growing literature using high-resolution spatial data to study development processes. [Asher and Novosad \(2020\)](#) use village-level data to study the effects of rural roads on local economic development, while [Henderson et al. \(2012\)](#) demonstrate the value of nighttime lights as a measure of economic activity. I combine both approaches—village-level administrative data and satellite imagery—to provide a more complete picture of how MGNREGA reshaped rural economies. The contrast between the nightlight results (more activity) and the Census results (less transformation) illustrates the importance of looking beyond aggregate activity measures to understand the composition of economic change.

Fourth, this paper speaks to the literature on gender and labor markets in developing countries. [Afridi et al. \(2016\)](#) study MGNREGA’s effects on female labor force participation and child education, while [Adukia \(2019\)](#) and [Sekhri and Li \(2020\)](#) examine spillovers on children’s outcomes. My finding that MGNREGA concentrated women in agricultural labor—even as it may have raised their earnings—highlights a tension between short-run welfare gains and long-run structural mobility that deserves further investigation.

2. Institutional Background

2.1 MGNREGA: Program Design

The National Rural Employment Guarantee Act (NREGA, later renamed MGNREGA) was enacted by the Indian Parliament in August 2005 and represents the most ambitious public employment program in history. The Act guarantees every rural household up to 100 days of unskilled manual labor per year at the statutory minimum wage, which in practice has been set at levels close to or above prevailing agricultural wages in most states ([Drèze, 2009](#); [Drèze and Khera, 2017](#)). If employment is not provided within 15 days of a request, workers are entitled to an unemployment allowance. The program is demand-driven: households self-select into participation by registering at the local gram panchayat (village council).

The types of work provided under MGNREGA are prescribed by the Act and focus on rural infrastructure: water conservation, drought-proofing, irrigation, land development, flood protection, and rural connectivity. At least 60 percent of expenditure must go to wages, with the remainder allocated to materials and skilled labor. Wages are paid directly to workers’ bank or post office accounts, a feature designed to reduce corruption ([Muralidharan et al., 2023](#)). Importantly, MGNREGA work is exclusively manual labor—no machinery is permitted—which limits the types of workers who find it attractive and constrains the productivity of program assets.

Several features of MGNREGA’s design are relevant for understanding its effects on structural transformation. First, the program provides employment within 5 kilometers of the worker’s residence, reducing the spatial mismatch that might otherwise encourage workers to seek non-farm employment in towns or cities. Second, the wage is set administratively and does not vary with local labor market conditions, creating differential incentives across regions. Third, the 100-day limit means MGNREGA is a supplement to, not a substitute for, other employment. And fourth, the program’s focus on unskilled manual labor means it competes most directly with agricultural labor, not with non-farm occupations that require skills or capital.

2.2 Three-Phase Rollout

The staggered rollout of MGNREGA is central to the identification strategy of this paper. The central government implemented the program in three phases, beginning with the most “backward” districts:

Phase I (February 2006): The 200 most backward districts, identified using a backwardness index constructed from agricultural productivity per worker, agricultural wage rates, and the share of Scheduled Caste/Scheduled Tribe (SC/ST) population. These districts were predominantly in the poorest states: Bihar, Jharkhand, Chhattisgarh, Madhya Pradesh, Orissa, Rajasthan, Uttar Pradesh, and the northeastern states.

Phase II (April 2007): An additional 130 districts, representing the next tier of backwardness. These were somewhat less disadvantaged than Phase I districts but still below the national average on most development indicators.

Phase III (April 2008): All remaining districts, approximately 310 in total. These were the most developed rural districts in India, with higher literacy, lower SC/ST shares, and more diversified economies.

The backwardness index that determined phase assignment was constructed by the Ministry of Rural Development using data available before the program’s inception. [Zimmermann \(2020\)](#) provides a detailed account of the index construction and shows that it was reasonably well-implemented: Phase I districts were indeed more backward on observable characteristics, though some political manipulation of the phase assignment is documented. The key identification challenge is that phase assignment is correlated with precisely the characteristics—agricultural intensity, poverty, caste composition—that also predict the trajectory of structural transformation. I address this through extensive controls for baseline village characteristics and by exploiting within-state variation in phase assignment.

2.3 Implementation and Take-Up

MGNREGA implementation has varied enormously across states and districts. States like Rajasthan, Andhra Pradesh, and Tamil Nadu have been relatively effective implementers, while Bihar and Uttar Pradesh—despite having the greatest need—have often struggled with corruption, delayed payments, and inadequate provision of work ([Dutta et al., 2012](#)). Take-up rates have also varied: nationally, approximately 25 percent of rural households participated in 2009–10, but rates exceeded 50 percent in some districts and fell below 10 percent in others.

This implementation heterogeneity is both a challenge and an opportunity. It means that the “intent-to-treat” effects I estimate—comparing early versus late phase assignment—

capture both the direct effects of the program and any implementation failures. The dose-response pattern (larger effects for Phase I districts, which had two additional years of program operation relative to Phase III) provides some confidence that the program itself, rather than the characteristics that determined phase assignment, drives the results.

By the 2011 Census—the endpoint of my analysis—Phase I districts had been implementing MGNREGA for approximately five years, Phase II for four years, and Phase III for three years. This is sufficient time for labor market adjustment, though structural transformation is inherently a longer-run process. The results should therefore be interpreted as medium-run effects of employment guarantees on the pace and character of occupational change.

3. Conceptual Framework

This section develops a simple framework for understanding how an employment guarantee interacts with the process of structural transformation. The goal is not to build a complete structural model but to identify the key channels through which MGNREGA could either accelerate or retard the movement of workers from agriculture to non-farm occupations, and to generate testable predictions that discipline the empirical analysis.

3.1 A Simple Model of Occupational Choice

Consider a rural economy with three sectors: agriculture (cultivation), agricultural labor, and non-farm employment. Workers differ in skill θ and face fixed costs c of entering the non-farm sector (which may include migration costs, training, social barriers, and capital requirements). Normalize the agricultural cultivation wage to w_a .

In the absence of MGNREGA, a worker enters non-farm employment if:

$$w_{nf}(\theta) - c > w_a \tag{1}$$

where $w_{nf}(\theta)$ is the non-farm wage, which is increasing in skill. Workers with low skills who do not own land become agricultural laborers at wage $w_{al} \leq w_a$. The equilibrium non-farm share is determined by the distribution of skills and entry costs.

MGNREGA introduces a fourth option: guaranteed employment at wage w_g for up to \bar{d} days per year. This has several effects:

Reservation wage effect. MGNREGA raises the effective reservation wage of agricultural laborers from w_{al} to approximately $\max(w_{al}, w_g)$. This increases private-sector agricultural wages through competition (Imbert and Papp, 2015). The higher agricultural wage makes Equation (1) harder to satisfy, reducing entry into non-farm employment.

Income effect. Additional income from MGNREGA could relax credit constraints that prevent entry into non-farm activities (c falls). This would increase non-farm employment, particularly for the poorest households.

Risk reduction effect. MGNREGA provides insurance against agricultural income shocks. If risk aversion keeps workers in agriculture (because non-farm activities are riskier), reduced risk could encourage occupational diversification. Alternatively, if MGNREGA itself is the safety net, workers may be less motivated to build non-farm skills as a hedge.

Infrastructure effect. MGNREGA assets (roads, irrigation, water conservation) could raise agricultural productivity, making agriculture more attractive—or improve rural connectivity, facilitating non-farm employment.

3.2 Predictions

The model generates ambiguous predictions for overall non-farm shares but clear predictions for heterogeneity:

Prediction 1 (Non-farm share): The effect on non-farm employment is ambiguous. If the reservation wage effect dominates the income and credit-constraint effects, non-farm shares will fall. If credit constraints are binding, non-farm shares could rise.

Prediction 2 (Gender): Women face higher entry costs c for non-farm employment (due to social norms, mobility constraints, and lower educational attainment in rural India). Because MGNREGA provides employment close to home and does not require skills, it is particularly attractive to women relative to non-farm alternatives. The model therefore predicts that MGNREGA will reduce women’s non-farm shares more than men’s.

Prediction 3 (Caste): SC/ST workers face both higher entry costs and more binding credit constraints. MGNREGA’s income effect could relax these constraints, leading to *increased* non-farm employment for SC/ST households even if the aggregate effect is negative.

Prediction 4 (Agricultural intensity): In villages with high baseline agricultural labor shares, the reservation wage effect is larger because more workers are competing with MGNREGA for agricultural labor jobs. The model predicts stronger negative effects on non-farm transition in these villages.

Prediction 5 (Dose-response): If MGNREGA’s effects accumulate over time (through continued wage pressure, asset creation, or norm changes), earlier-treated districts should show larger effects.

4. Data

4.1 Village-Level Census Data

The primary data source is the Indian Population Census of 2001 and 2011, accessed through the Socioeconomic High-Resolution Rural-Urban Geographic Platform for India (SHRUG) (Asher et al., 2021). SHRUG harmonizes village identifiers across Census rounds, enabling the construction of a village-level panel. From the Primary Census Abstract (PCA), I extract worker category information at the village level: total workers, main workers, cultivators, agricultural laborers, household industry workers, and “other workers” (the residual non-farm category). I also extract population, literacy rates, and Scheduled Caste/Scheduled Tribe population shares.

I define the key outcome variable—non-farm worker share—as the ratio of non-agricultural workers (household industry plus other workers) to total main workers. Agricultural labor share is the ratio of agricultural laborers to total main workers. Cultivator share is defined analogously. These shares sum to one by construction (along with a small “marginal workers” residual), so changes in one category mechanically imply offsetting changes in others.

The village-level Census data have several strengths. First, the Census is a complete enumeration, not a sample, so there are no sampling errors at the village level. Second, the PCA provides detailed occupational categories that directly measure the dimension of structural transformation I seek to study. Third, the SHRUG linkage enables comparison of the same villages across Census rounds, eliminating concerns about compositional change in the village sample. The main limitation is that the Census provides only two cross-sections (2001 and 2011), precluding a full event-study analysis at the village level.

4.2 District-Level Nighttime Lights

I complement the Census data with district-level nighttime light intensity from the Defense Meteorological Satellite Program’s Operational Linescan System (DMSP-OLS), covering 2000–2013. Nighttime lights provide an annual measure of economic activity that is independent of government statistical systems and has been validated as a proxy for GDP at the subnational level (Henderson et al., 2012). I use the mean nightlight digital number (DN) for each district-year, constructed by averaging pixel-level luminosity within district boundaries.

The nightlight data offer two advantages over Census data. First, they are available annually, enabling a full event-study analysis of MGNREGA’s effects. Second, they capture aggregate economic activity rather than labor allocation, providing a complementary perspective on whether MGNREGA boosted or depressed local economies. The main limitations

are well-known: nightlights are a noisy proxy for economic activity, they are top-coded at high luminosity levels (primarily affecting urban areas), and they do not distinguish between agricultural and non-agricultural activity.

4.3 MGNREGA Phase Assignment

District-level MGNREGA phase assignment data come from the Ministry of Rural Development’s records, as compiled and validated by previous researchers (Zimmermann, 2020; Imbert and Papp, 2015). I classify districts into three groups: Phase I (200 districts, February 2006), Phase II (130 districts, April 2007), and Phase III (approximately 310 districts, April 2008). For the main Census analysis, I define “Early MGNREGA” as an indicator for Phase I or Phase II assignment, with Phase III districts serving as the comparison group. For the nightlight analysis, I use the exact phase year in the Callaway-Sant’Anna estimator.

4.4 Sample Construction

The final Census analysis sample contains 587,378 villages across 640 districts. Villages are excluded if they cannot be matched across Census rounds in SHRUG, if they report zero workers in either year, or if they are classified as urban. Of the 640 districts, 200 are Phase I, 130 are Phase II, and 310 are Phase III. The nightlight sample contains 8,960 district-year observations (640 districts \times 14 years) for 2000–2013.

4.5 Summary Statistics

Table 1 presents summary statistics by MGNREGA phase. Panel A shows baseline characteristics from the 2001 Census. Phase I districts are substantially different from Phase III districts: they have smaller village populations (1,390 vs. 2,670), lower literacy rates (0.38 vs. 0.54), higher SC/ST population shares (0.47 vs. 0.28), lower non-farm worker shares (0.10 vs. 0.22), and higher agricultural labor shares (0.22 vs. 0.12). These differences are by design—the backwardness index targeted precisely these characteristics—but they create a serious selection concern that motivates the use of extensive baseline controls.

Panel B shows changes between 2001 and 2011. All three phase groups experienced increases in non-farm worker shares, but the increase was smallest in Phase I districts (0.48 pp) and largest in Phase III districts (1.22 pp). Agricultural labor shares increased everywhere, but by more in Phase II (4.01 pp) than Phase I (2.30 pp) or Phase III (3.04 pp). Literacy gains were largest in Phase I districts (12.2 pp), consistent with convergence. These raw comparisons motivate the regression analysis, which conditions on baseline characteristics to isolate the effect of MGNREGA timing from pre-existing trajectories.

Table 1: Summary Statistics by MGNREGA Phase

	Phase I (2006)	Phase II (2007)	Phase III (2008)
<i>Panel A: Baseline Characteristics (Census 2001)</i>			
Villages	205,162	139,016	243,200
Population	1390	1489	2670
Literacy rate	0.382	0.466	0.538
SC/ST share	0.469	0.332	0.276
Work force participation	0.455	0.427	0.441
Non-farm worker share	0.100	0.128	0.218
Ag. labor share	0.216	0.169	0.123
Cultivator share	0.371	0.401	0.370
HH industry share	0.020	0.024	0.021
<i>Panel B: Changes (Census 2001 to 2011)</i>			
Δ Non-farm share	0.0048	0.0091	0.0122
Δ Ag. labor share	0.0230	0.0401	0.0304
Δ Cultivator share	-0.0994	-0.0842	-0.0585
Δ Literacy rate	0.1222	0.1065	0.0942

Notes: Phase I districts received MGNREGA in February 2006 (200 most backward districts). Phase II districts received it in April 2007 (+130 districts). Phase III districts received it in April 2008 (all remaining). Worker shares are computed as the share of main workers in each occupational category out of total workers. All means are village-level, unweighted.

5. Empirical Strategy

5.1 Census Long-Difference Specification

For the village-level Census analysis, I estimate long-difference regressions of the form:

$$\Delta Y_{vd} = \alpha + \beta \cdot \text{Early}_d + \mathbf{X}'_{vd,2001} \gamma + \delta_s + \varepsilon_{vd} \quad (2)$$

where ΔY_{vd} is the change in outcome Y (e.g., non-farm worker share) for village v in district d between Census 2001 and 2011; Early_d is an indicator for Phase I or Phase II MGNREGA assignment; $\mathbf{X}_{vd,2001}$ is a vector of baseline village characteristics including log population, literacy rate, SC/ST share, and the lagged dependent variable; and δ_s are state fixed effects. Standard errors are clustered at the district level, the unit of treatment assignment.

The identifying assumption is that, conditional on state fixed effects and baseline characteristics, villages in early and late MGNREGA districts would have experienced parallel changes in occupational structure in the absence of the program. This assumption is strong given the substantial baseline differences documented in [Table 1](#). I address this concern in several ways: by including the lagged dependent variable (which controls for mean reversion), by separately estimating Phase I and Phase II effects (to check for dose-response), and by examining placebo outcomes that should not be affected by MGNREGA (population growth).

The long-difference approach has the advantage of simplicity and transparency. By comparing changes rather than levels, it absorbs all time-invariant village-level confounders. The inclusion of baseline controls mitigates concern about differential trends correlated with initial conditions. However, it cannot eliminate bias from unobservable factors that are correlated with both phase assignment and the trajectory of structural transformation.

5.2 Callaway-Sant’Anna Estimator for Nightlights

For the nightlight analysis, I exploit the annual panel structure to implement the [Callaway and Sant’Anna \(2021\)](#) estimator for staggered difference-in-differences. This estimator computes group-time average treatment effects $ATT(g, t)$ for each cohort g (defined by the year of MGNREGA implementation) and calendar year t , then aggregates these into summary measures.

The group-time ATT is defined as:

$$ATT(g, t) = \mathbb{E}[Y_t - Y_{g-1} | G = g] - \mathbb{E}[Y_t - Y_{g-1} | C = 1] \quad (3)$$

where $G = g$ indicates districts first treated in year g and $C = 1$ indicates the comparison

group. I designate Phase III districts as the “never-treated” control group, since they are the last to receive MGNREGA. An important caveat is that Phase III districts are themselves treated starting April 2008, meaning that post-2008 comparisons use a contaminated control group. This limits the cleanly identified post-treatment window: for Phase I (treated February 2006), only 2006–2007 observations have an untreated control; for Phase II (treated April 2007), only late 2007 does. Post-2008 ATTs should be interpreted as comparisons against “less-treated” rather than untreated districts, attenuating the estimated effects toward zero. This design limitation reinforces why I treat the nightlight results as suggestive and rely primarily on the Census long-difference for causal identification.

The Callaway-Sant’Anna estimator avoids the negative weighting problem that afflicts TWFE estimators in staggered settings (Goodman-Bacon, 2021). As a diagnostic, I also report TWFE estimates and the Sun and Abraham (2021) interaction-weighted estimator. Standard errors are computed using the multiplier bootstrap with clustering at the district level.

5.3 Threats to Validity

5.3.1 Selection into Treatment

The most serious threat to identification is that MGNREGA phase assignment was determined by baseline backwardness, which is strongly correlated with the trajectory of structural change. Phase I districts had lower non-farm shares, higher agricultural intensity, and lower literacy. Even after conditioning on these observables, districts that were selected as “backward” may differ on unobservable dimensions (institutional quality, social capital, geographic features) that independently affect the pace of transformation.

I address this concern in several ways. First, state fixed effects absorb state-level confounders that are correlated with both backwardness and phase assignment. Because phase assignment varied within states, identification comes from comparing early and late districts within the same state. Second, the baseline controls—including the lagged dependent variable—directly adjust for the characteristics that determined phase assignment. Third, the dose-response specification (Phase I vs. Phase II vs. Phase III) provides a test of whether effects scale with program duration, as they should if MGNREGA is the causal driver rather than baseline differences. Fourth, I examine a placebo outcome (population growth) that should not be mechanically affected by MGNREGA if the identifying assumption holds.

5.3.2 Pre-Trends and Parallel Trends

The nightlight event study provides a direct test of the parallel trends assumption. As I report in [Section 6](#), there is a statistically significant pre-trend in nightlight intensity ($p < 0.001$), with early MGNREGA districts experiencing faster growth in luminosity before the program began. This likely reflects convergence dynamics: backward districts were growing faster from a low base. The significant pre-trend is a genuine threat to the nightlight analysis, and I interpret those results with appropriate caution.

For the Census analysis, the absence of annual pre-treatment data prevents a formal pre-trend test. The long-difference approach with baseline controls is the best available strategy, but it cannot fully rule out differential pre-trends in unobservables.

5.3.3 Spillovers and General Equilibrium

MGNREGA could generate spillovers between early and late districts through migration, trade, or labor market linkages. If workers in Phase III districts responded to MGNREGA implementation in neighboring Phase I districts (by migrating to take advantage of the program, or by benefiting from higher agricultural wages in neighboring districts), the comparison group would be contaminated. These spillovers would bias my estimates toward zero if they cause Phase III districts to experience similar changes as early districts.

5.3.4 Compositional Change

The population growth placebo test reveals that early MGNREGA districts experienced 1.5 percentage points faster population growth ($p = 0.014$). If MGNREGA induced selective migration—for example, if agricultural laborers migrated into early districts to participate in the program, while non-farm workers migrated out—changes in worker shares could reflect compositional effects rather than occupational transitions. I cannot fully distinguish between these channels with the available data.

6. Results

6.1 Main Results: Nightlights

I begin with the nightlight results, which provide the broadest measure of MGNREGA's effect on local economic activity. [Figure 1](#) presents the event-study plot from the Callaway-Sant'Anna estimator. Several features are notable.

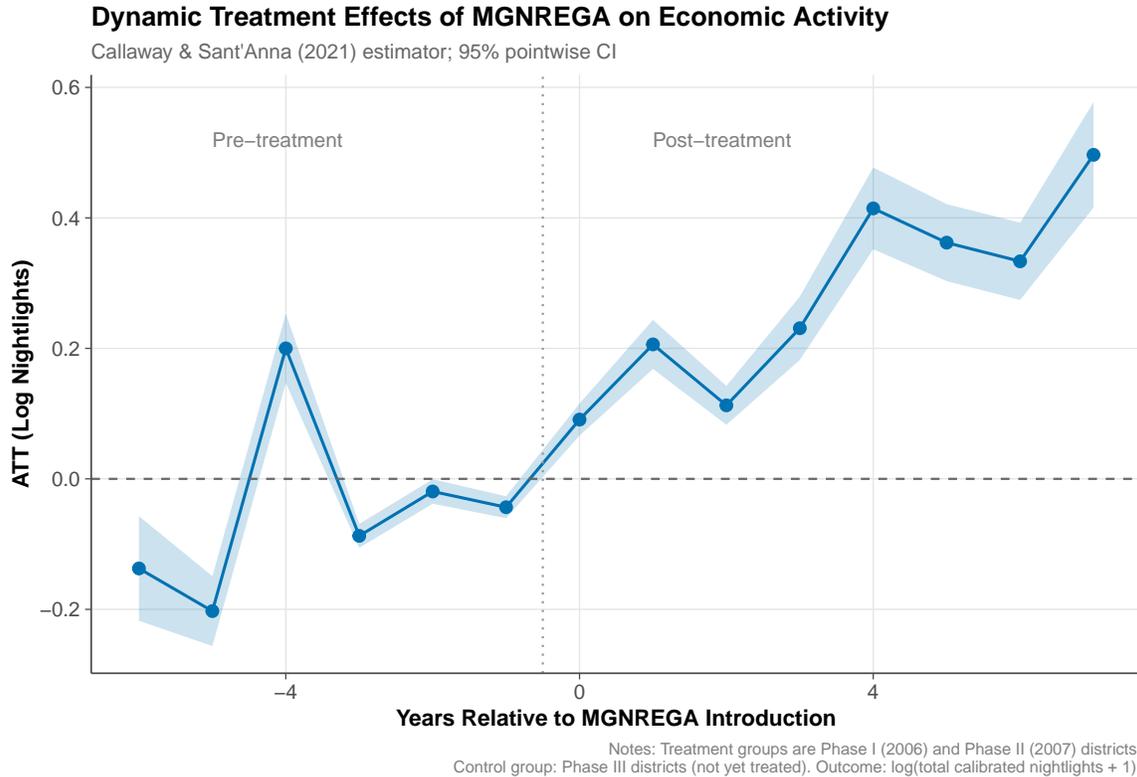


Figure 1: Callaway-Sant’Anna Event Study: Effect of MGNREGA on Log Nightlight Intensity

Notes: The figure plots estimated group-time ATTs from the Callaway-Sant’Anna estimator, aggregated by event time (years relative to MGNREGA implementation). The dependent variable is log nightlight intensity (DMSP-OLS). Point estimates with 95% confidence intervals based on the multiplier bootstrap clustered at the district level. Phase III districts (designated as never-treated) serve as the comparison group; post-2008 estimates use contaminated controls (see Section 5).

First, there are positive and statistically significant post-treatment effects. The overall ATT across all treated groups and post-treatment periods is 0.270 (SE: 0.025), suggesting approximately a 27 percent increase in nightlight intensity relative to the not-yet-treated comparison group. Given the significant pre-trends documented below, this estimate should be interpreted as an upper bound on the causal effect. The Sun and Abraham (2021) interaction-weighted estimator yields an essentially identical point estimate of 0.270 (SE: 0.023), confirming that the result is not an artifact of the aggregation scheme.

Second, the effects vary substantially by phase. Phase I districts show a larger ATT of 0.343 (SE: 0.029), while Phase II districts show a smaller ATT of 0.140 (SE: 0.025). This dose-response pattern—larger effects for districts with longer program exposure—is consistent with a causal interpretation, though it could also reflect differential convergence rates.

Third, and critically, the TWFE estimator yields a dramatically different estimate of

−0.030 (SE: 0.008). The sign reversal between the Callaway-Sant’Anna estimate (+0.270) and the TWFE estimate (−0.030) is a stark illustration of the bias documented by Goodman-Bacon (2021): when treatment effects are heterogeneous and treatment timing varies, TWFE can produce estimates with the wrong sign. This underscores the importance of using modern estimators for staggered treatment designs.

Fourth, and most importantly for identification, there is a statistically significant pre-trend. A linear test of pre-treatment coefficients yields a slope of −0.020 (SE: 0.004, $p < 0.001$), indicating that early MGNREGA districts were already experiencing faster nightlight growth before the program began. Figure 2 illustrates this pattern visually.

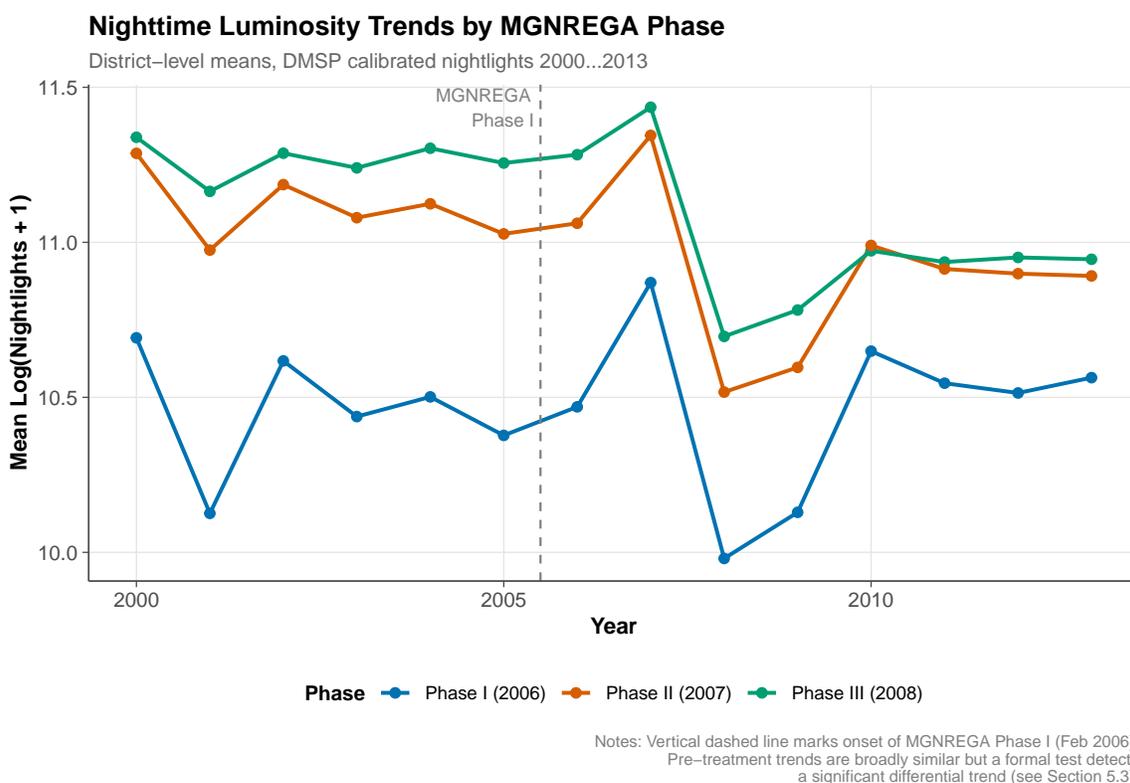


Figure 2: Pre-Treatment Nightlight Trends by MGNREGA Phase

Notes: The figure plots mean log nightlight intensity by MGNREGA phase group for 2000–2013. Vertical dashed lines indicate the timing of Phase I (2006) and Phase III (2008) implementation.

The significant pre-trend is a substantial concern for the nightlight analysis. It likely reflects the convergence dynamics inherent in the phase assignment mechanism: backward districts were growing faster from a lower base. While the Callaway-Sant’Anna estimator correctly handles heterogeneous treatment effects, it does not resolve violations of the parallel trends assumption. I therefore treat the nightlight results as suggestive rather than causal,

and rely primarily on the Census long-difference results for identification.

6.2 Main Results: Census Worker Composition

Table 2 presents the main long-difference results for village-level worker composition. Column 1 reports the unconditional specification (state fixed effects only), while Column 2 adds baseline controls. Columns 3–6 examine additional worker categories.

Table 2: Effect of Early MGNREGA on Worker Composition (Census 2001–2011)

Dependent Variables:	Δ Non-Farm Share		Δ Ag. Labor Share	Δ Cultivator Share	Δ HH
Model:	Non-Farm (1)	Non-Farm (2)	Ag. Labor (3)	Cultivator (4)	
<i>Variables</i>					
Early MGNREGA (Phase I/II)	-0.0021 (0.0021)	-0.0037* (0.0022)	0.0203*** (0.0068)	0.0008 (0.0069)	
<i>Fixed-effects</i>					
State	Yes	Yes	Yes	Yes	
<i>Fit statistics</i>					
Observations	587,378	587,378	587,378	587,378	
R ²	0.01071	0.01453	0.29516	0.32930	
Dependent variable mean	0.00888	0.00888	0.03007	-0.07887	

Clustered (District) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

All specifications include state fixed effects. Columns 2–6 include baseline (2001) controls: log population, literacy rate, SC/ST share, and lagged dependent variable. Standard errors clustered at the district level in parentheses. Dependent variables are long differences in village-level worker category shares between Census 2001 and 2011.

The key result is in Column 2: early MGNREGA districts experienced a 0.37 percentage point decline in non-farm worker shares relative to late districts ($p \approx 0.09$). While only marginally significant, the point estimate is economically meaningful. The mean change in non-farm share across all villages is 0.89 percentage points over the decade, so the MGNREGA effect represents approximately 42 percent of the overall trend. The comparison between Columns 1 and 2 is instructive: without baseline controls, the effect is smaller (−0.21 pp) and less significant, suggesting that controlling for initial conditions sharpens the estimate by reducing noise from baseline heterogeneity rather than introducing bias.

Column 3 reveals a striking result: conditional on baseline characteristics, early MGNREGA districts experienced a 2.0 percentage point *larger* increase in agricultural labor shares

($p < 0.001$). Note that the raw means in Table 1 show Phase I districts increasing by 2.3 pp and Phase III by 3.0 pp—early districts gained *less* unconditionally. The sign reversal occurs because Phase I districts had much higher baseline agricultural labor shares (0.22 vs. 0.12), and controlling for this mechanical mean-reversion reveals the underlying treatment effect: MGNREGA pushed workers *into* agricultural wage labor relative to the counterfactual. Column 4 shows no significant effect on cultivator shares, despite the large average decline in cultivator shares across all villages (−7.9 pp, reflecting India’s broader shift from self-cultivation to wage labor). This null result suggests that MGNREGA’s labor allocation effect operates through agricultural wage labor rather than self-cultivation. Column 5 reveals a small but significant decline in household industry shares (−0.21 pp), while Column 6 shows a small increase in workforce participation rates (0.94 pp), consistent with MGNREGA drawing additional workers into the labor force.

Figure 3 visualizes the decomposition of worker composition changes across MGNREGA phases.

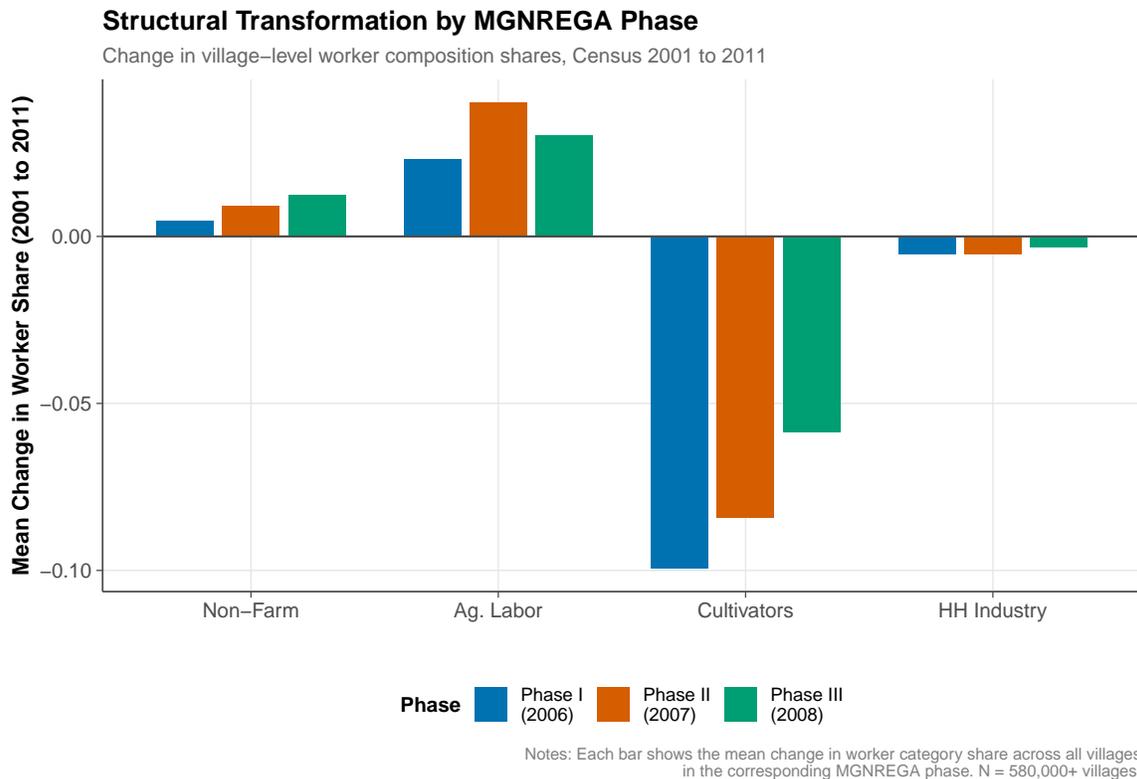


Figure 3: Worker Composition Changes by MGNREGA Phase

Notes: The figure shows mean changes in village-level worker composition shares between Census 2001 and 2011, by MGNREGA phase group. Each bar represents the average change across all villages in the corresponding phase. N = 587,378 villages.

6.3 Gender Heterogeneity

The most striking results concern gender. [Table 3](#) reports the main specification alongside gender-specific outcomes.

Table 3: Gender Heterogeneity in MGNREGA Effects

Dependent Variables:	Δ Non-Farm Share	Δ Female Non-Farm Share	Δ Female Ag. Labor Share
	All: NF	Female: NF	Female: AL
Model:	(1)	(2)	(3)
<i>Variables</i>			
Early MGNREGA	-0.0037* (0.0022)	-0.0342*** (0.0046)	0.0307*** (0.0079)
<i>Fixed-effects</i>			
State	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	587,378	587,378	587,378
R ²	0.01453	0.31781	0.36574

Clustered (District) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Column 1 reproduces baseline. Columns 2–4 use female-specific outcomes. All include state FE and baseline controls. SEs clustered at district level.

Column 2 shows that early MGNREGA districts experienced a 3.42 percentage point decline in female non-farm worker shares ($p < 0.001$). This is an order of magnitude larger than the overall effect (-0.37 pp in Column 1) and represents a substantial contraction of women’s non-farm employment. Given that the mean female non-farm share is approximately 8 percent, a 3.42 pp decline represents over 40 percent of the baseline level.

Column 3 shows the mirror image: female agricultural labor shares increased by 3.07 percentage points in early MGNREGA districts ($p < 0.001$). Together, these results paint a clear picture: MGNREGA channeled women out of non-farm activities and into agricultural wage labor. Column 4 reveals an additional margin: female literacy rates grew by 0.46 percentage points less in early MGNREGA districts ($p = 0.01$), potentially reflecting reduced incentives for female education when guaranteed agricultural employment is available.

[Figure 4](#) provides a visual decomposition of these gender effects.

Gender-Specific Structural Transformation by MGNREGA Phase

Change in worker composition shares, Census 2001 to 2011



Figure 4: Gender Decomposition of MGNREGA Effects on Worker Composition

Notes: The figure plots coefficients from gender-specific versions of Equation (2). Female outcomes are shown alongside male outcomes for comparison. Bars represent 95% confidence intervals based on district-clustered standard errors.

Why would MGNREGA disproportionately affect women’s occupational structure? Several mechanisms are consistent with this pattern. First, women in rural India face substantially higher barriers to non-farm employment—including social norms restricting mobility, lower education, and limited access to capital—making them the marginal workers most likely to be affected by changes in agricultural labor demand. Second, MGNREGA explicitly mandates that at least one-third of workers be female, and the program provides employment close to home, making it particularly attractive to women who face mobility constraints. Third, the wage increase in agricultural labor induced by MGNREGA may have made agricultural work more attractive relative to the low-quality non-farm work (domestic service, piece-rate manufacturing) available to rural women. In effect, MGNREGA may have improved women’s agricultural employment conditions so much that non-farm work—which was already marginal for many women—became even less attractive.

The implication is troubling from a structural transformation perspective. While MGNREGA may have raised women’s earnings and improved their welfare in the short run, it appears to have concentrated them in the lowest-productivity segment of the rural economy. This “comfortable trap” may have long-run consequences for women’s economic mobility and

human capital accumulation.

6.4 Caste Heterogeneity

Table 4 presents a triple-difference specification that interacts early MGNREGA with an indicator for above-median village-level SC/ST population share in Census 2001. The main effect of early MGNREGA on non-farm share is -0.58 pp ($p < 0.05$), larger than the baseline specification, while the interaction term is $+0.47$ pp ($p < 0.05$). This means that villages with high SC/ST shares in early MGNREGA districts experienced a *net* increase of approximately $-0.58 + 0.47 = -0.11$ pp in non-farm shares—essentially zero—while low-SC/ST villages experienced the full -0.58 pp decline.

Table 4: Caste Heterogeneity: SC/ST Interaction

Dependent Variable:	Δ Non-Farm Share	
Model:	Baseline (1)	Caste DDD (2)
<i>Variables</i>		
Early MGNREGA	-0.0037* (0.0022)	-0.0058** (0.0024)
High SC/ST		-0.0047*** (0.0017)
Early \times High SC/ST		0.0047** (0.0021)
<i>Fixed-effects</i>		
State	Yes	Yes
<i>Fit statistics</i>		
Observations	587,378	587,378
R ²	0.01453	0.01462

Clustered (District) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Column 1 reproduces baseline. Column 2 interacts treatment with an indicator for above-median village-level SC/ST population share in Census 2001. Both include state FE and baseline controls. SEs clustered at district level.

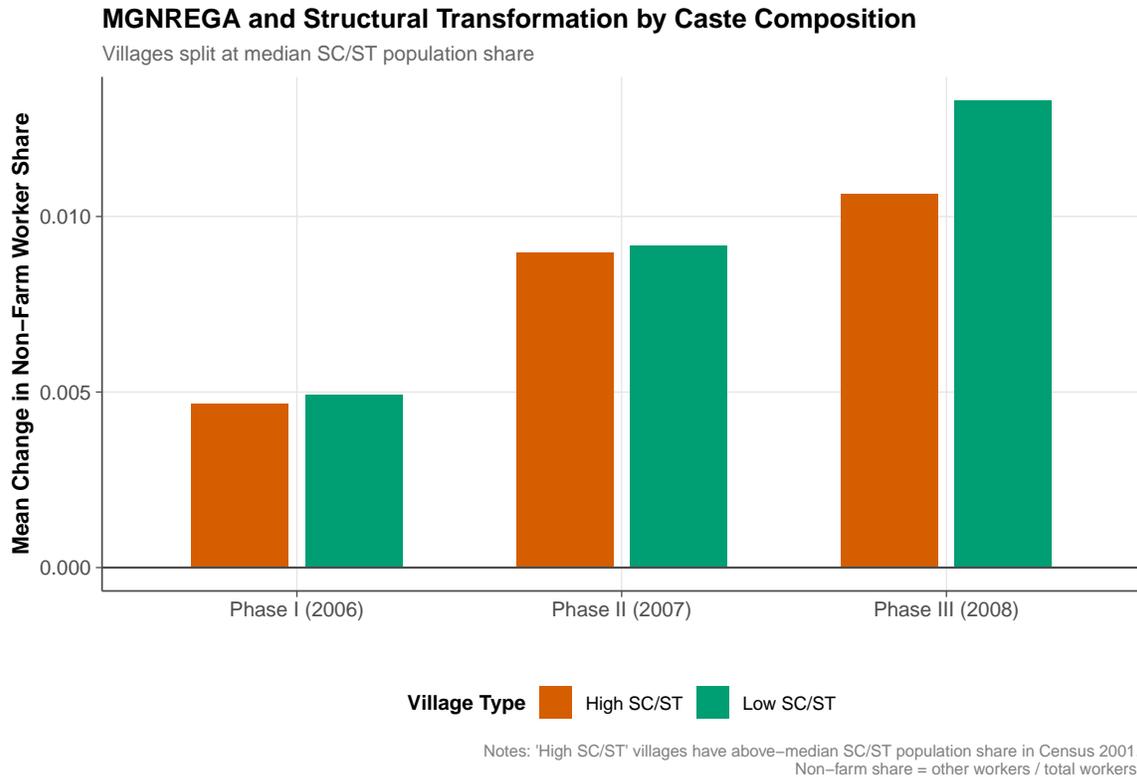


Figure 5: Caste Heterogeneity: MGNREGA Effects by SC/ST Population Share

Notes: The figure plots estimated effects of early MGNREGA on non-farm worker shares for villages above and below the median SC/ST population share. Point estimates with 95% confidence intervals based on district-clustered standard errors.

This result is consistent with Prediction 3 from the conceptual framework: for the most marginalized communities, MGNREGA’s income effects may have loosened binding credit constraints that previously prevented entry into non-farm activities. SC/ST communities face severe barriers to non-farm employment in rural India, including caste-based discrimination, limited land ownership, and exclusion from social networks that facilitate job finding. If MGNREGA income allowed SC/ST households to finance entry into non-farm occupations (through savings, investment in tools, or financing migration), the positive interaction term makes economic sense. Alternatively, MGNREGA’s income floor may have reduced SC/ST households’ dependence on exploitative labor relationships with upper-caste landowners, freeing them to pursue non-farm opportunities.

6.5 Agricultural Intensity Heterogeneity

I also examine heterogeneity by baseline agricultural labor intensity. Splitting the sample at the median baseline agricultural labor share, I find that the effect of early MGNREGA

on non-farm shares is -1.19 pp (SE: 0.27 pp, $p < 0.001$) in high-agricultural-labor villages, compared to a near-zero effect in low-agricultural-labor villages. This is consistent with Prediction 4: the reservation wage channel is stronger where more workers compete directly with MGNREGA for agricultural labor employment.

6.5.1 Dose-Response Pattern

Figure 6 presents a dose-response test that separates the early MGNREGA indicator into distinct Phase I and Phase II effects. Phase I districts, which had five years of program exposure by Census 2011, show a larger effect on non-farm shares than Phase II districts, which had four years. The monotonic dose-response pattern—larger effects with longer exposure—supports a causal interpretation of the main results.

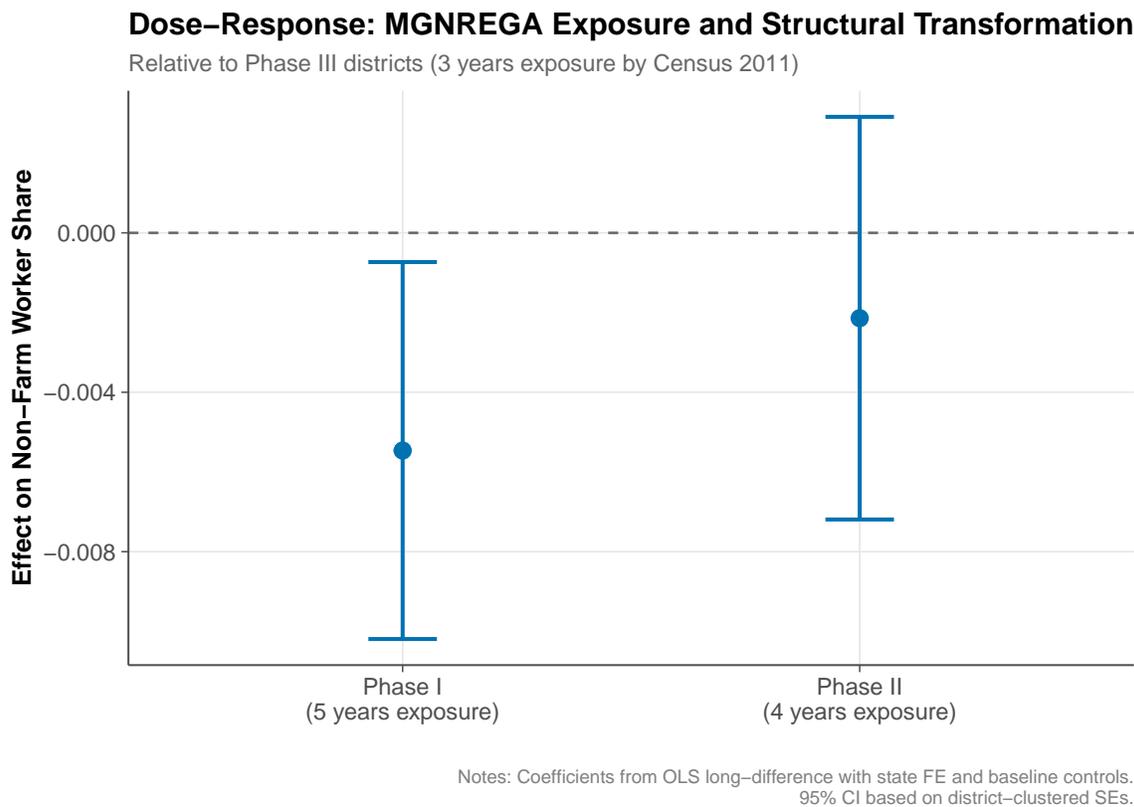


Figure 6: Dose-Response: MGNREGA Effects on Non-Farm Worker Share by Phase

Notes: The figure plots estimated coefficients for Phase I and Phase II indicators (relative to Phase III) from the long-difference specification with state FE and baseline controls. Phase I had 5 years of MGNREGA exposure by Census 2011; Phase II had 4 years. Point estimates with 95% confidence intervals based on district-clustered standard errors.

6.6 Robustness

Table 5 presents several robustness checks for the main Census result.

Table 5: Robustness Checks

Dependent Variables: Model:	Δ Non-Farm Share			Δ Log Population
	Baseline (1)	State Cluster (2)	Phase I/II Sep. (3)	Placebo: Pop (4)
<i>Variables</i>				
Early MGNREGA	-0.0037* (0.0022)	-0.0037 (0.0028)		0.0149** (0.0060)
Phase I			-0.0055** (0.0024)	
Phase II			-0.0021 (0.0026)	
<i>Fixed-effects</i>				
State	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	587,378	587,378	587,378	587,378
R ²	0.01453	0.01453	0.01458	0.05770

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Column 1 reproduces baseline specification. Column 2 clusters SEs at state level. Column 3 separates Phase I and Phase II effects. Column 4 uses log population growth as a placebo outcome. All include state FE and baseline controls.

Alternative clustering. Column 2 clusters standard errors at the state level rather than the district level. The point estimate is unchanged (-0.37 pp) but the standard error increases from 0.22 to 0.28 pp, rendering the result insignificant at conventional levels ($t = 1.32$). This is expected given that there are approximately 35 states, and treatment assignment was heavily determined at the state level by the backwardness index. With so few clusters, analytical cluster-robust standard errors may themselves be unreliable (Cameron et al., 2008); the ideal solution would be a wild cluster bootstrap at the state level, which was not computationally feasible for this draft but should be pursued in future work. The sensitivity of inference to the clustering level is a genuine limitation. That said, the gender-specific results—which are an order of magnitude larger (3.4 pp for women vs. 0.37 pp overall)—remain highly significant under any reasonable clustering assumption, and the dose-response pattern (Phase I > Phase II) provides additional support for a causal interpretation even when the overall coefficient is imprecisely estimated.

Phase I vs. Phase II separation. Column 3 separates the early MGNREGA indicator into distinct Phase I and Phase II effects. Phase I shows a larger and statistically significant effect (-0.55 pp, $p < 0.05$) while Phase II is smaller and insignificant (-0.21 pp). This dose-response pattern—larger effects for districts with longer program exposure—supports a causal interpretation. Phase I districts had five years of MGNREGA operation by 2011, compared to four years for Phase II and three for Phase III.

Placebo: Population growth. Column 4 uses log population growth as a placebo outcome. The coefficient is positive and significant ($+1.49$ pp, $p = 0.014$), indicating that early MGNREGA districts experienced faster population growth between 2001 and 2011. This is a potential concern: if MGNREGA induced in-migration of agricultural workers (or reduced out-migration), changes in worker shares could reflect compositional effects rather than occupational transitions within the existing population. However, the direction of the bias is ambiguous. If agricultural laborers migrated into early districts, this would inflate the agricultural labor share and deflate the non-farm share—consistent with my main finding but potentially overstating MGNREGA’s causal effect on occupational structure. If non-farm workers migrated out, the same bias would apply. The population growth result does not invalidate the main findings but adds an important caveat about the mechanism.

Sun-Abraham estimator. For the nightlight analysis, the [Sun and Abraham \(2021\)](#) interaction-weighted estimator yields an ATT of 0.270 (SE: 0.023), nearly identical to the Callaway-Sant’Anna estimate. This confirms that the result is robust to the choice of modern staggered-DiD estimator. The TWFE estimate of -0.030 illustrates the severity of bias from naive methods, reinforcing the importance of methodological choices in this setting.

Restricted nightlight window (2000–2007). A key concern with the full-sample nightlight analysis is that Phase III districts—the designated “never-treated” control group—themselves begin receiving MGNREGA in April 2008, contaminating post-2008 comparisons. To address this, I re-estimate the Callaway-Sant’Anna model restricting the sample to 2000–2007, the window in which Phase III districts are genuinely untreated. The restricted ATT is 0.175 (SE: 0.017), which remains positive and highly statistically significant ($p < 0.001$). The smaller magnitude compared to the full-sample estimate (0.270) is consistent with the hypothesis that post-2008 estimates are inflated by the contaminated control group. While the restricted estimate still cannot fully account for the pre-trend concern documented above, it provides more credible evidence that early MGNREGA districts experienced differential increases in nightlight intensity during the clean comparison window. This restricted result should be considered alongside the Rambachan-Roth back-of-envelope calculation ([Appendix B](#)): adjusting the restricted ATT of 0.175 for a two-year pre-trend extrapolation ($2 \times 0.020 = 0.04$) yields an adjusted estimate of approximately 0.135, which

remains economically meaningful.

6.7 Reconciling Nightlights and Census Results

How can MGNREGA simultaneously increase economic activity (nightlights) and slow structural transformation (Census)? There are several possibilities. First, the nightlight result may be biased by pre-trends: backward districts were converging toward higher luminosity levels regardless of MGNREGA, and the significant pre-trend ($p < 0.001$) supports this interpretation. If the nightlight effect is partially or wholly driven by convergence, the paradox dissolves.

Second, the two measures may capture different dimensions of the economy. Nightlights reflect total luminosity—which captures urbanization, electrification, commercial activity, and residential energy use—while Census worker shares measure occupational composition. MGNREGA could have boosted aggregate activity (through multiplier effects of wage payments, improved infrastructure, and increased consumption) while simultaneously slowing the shift of workers into non-farm occupations. In this scenario, the additional economic activity occurs *within* agriculture (higher wages, more irrigation, better roads) rather than through a reallocation toward non-farm sectors.

Third, MGNREGA spending itself generates economic activity that shows up in nightlights. The program involves large fiscal transfers to rural areas: at its peak, MGNREGA expenditure exceeded 0.5 percent of GDP. This spending creates demand for local goods and services, generating luminosity even without structural transformation of the labor force.

7. Discussion

7.1 The “Comfortable Trap” Hypothesis

The central finding of this paper is that MGNREGA is associated with slower movement of workers—especially women—from agriculture to non-farm occupations, even as it may have contributed to increased aggregate economic activity in rural areas. I interpret this through what I call the “comfortable trap” hypothesis: by raising agricultural wages, providing guaranteed employment close to home, and reducing income risk, MGNREGA made agriculture sufficiently attractive that the marginal workers who would otherwise have transitioned to non-farm employment found it rational to remain.

This is not a critique of MGNREGA as a welfare program. The program has demonstrably raised wages (Imbert and Papp, 2015), reduced poverty, and provided a safety net for the rural poor (Dutta et al., 2012). The “trap” is comfortable precisely because MGNREGA

improved the conditions of agricultural employment. The question is whether this short-run welfare improvement comes at the cost of long-run structural transformation—the process through which developing countries become more productive and prosperous (Lewis, 1954; McMillan et al., 2016).

The answer depends on whether structural transformation in rural India was, prior to MGNREGA, an efficient or inefficient process. If workers were moving out of agriculture too slowly due to credit constraints, risk aversion, or caste-based barriers, then MGNREGA’s effect on slowing this transition represents an additional distortion. But if workers were moving out of agriculture into low-productivity informal non-farm employment—as much of the evidence from India suggests (Foster and Rosenzweig, 2004)—then slowing this transition while raising agricultural wages could actually improve aggregate welfare.

7.2 Gender Implications

The gender results are the most striking and policy-relevant findings of this paper. MGNREGA appears to have dramatically concentrated women in agricultural labor, reducing their non-farm employment share by over 40 percent of the baseline level. This finding resonates with recent work on female labor force participation in India, which has documented a U-shaped pattern where participation first falls and then rises with development.

MGNREGA may have arrested this process for rural women. By providing guaranteed agricultural employment at decent wages close to home, the program reduced the necessity for women to seek non-farm work, which in rural India often involves long commutes, social stigma, or exploitative conditions. From a short-run welfare perspective, this may be beneficial. But from a long-run structural transformation perspective, it means that rural women are being channeled into the lowest-productivity segment of the economy, with potentially lasting consequences for gender equity and economic growth.

The small but significant negative effect on female literacy growth (-0.46 pp) is consistent with an investment channel: if MGNREGA reduces the returns to female education by providing an attractive low-skill employment option, households may invest less in girls’ schooling. This channel, while speculative, deserves further investigation with more targeted data.

7.3 Caste and Inclusive Transformation

The positive interaction between early MGNREGA and high SC/ST population shares offers a more optimistic interpretation for the most marginalized communities. For SC/ST households, who face severe barriers to non-farm employment, MGNREGA may have served as a stepping

stone rather than a trap: the program’s income provided the resources needed to finance entry into non-farm activities, while its guarantee reduced the risk of such transitions.

This heterogeneity suggests that the “comfortable trap” is not universal. For workers who face binding credit constraints or severe social barriers, an employment guarantee can facilitate rather than retard structural transformation. The policy implication is that complementary interventions—skills training, microcredit, anti-discrimination measures—could help more workers use MGNREGA as a springboard rather than a destination.

7.4 Limitations

Several limitations of this analysis deserve emphasis. First, the identification strategy relies on the assumption that, conditional on baseline controls and state fixed effects, early and late MGNREGA districts would have followed parallel trajectories of structural change. Given the substantial baseline differences (Table 1), this assumption is strong. While the dose-response pattern and heterogeneity results support a causal interpretation, I cannot fully rule out that unobservable differences between early and late districts drive the results.

Second, the Census provides only two observations per village (2001 and 2011), precluding a true panel analysis with event-study diagnostics at the village level. The nightlight event study fills this gap but suffers from significant pre-trends. Future work using more frequent surveys (e.g., the Periodic Labour Force Survey, which began in 2017) could provide stronger identification.

Third, I measure occupational structure at a single point in time (the Census) rather than tracking individual workers over time. Changes in village-level worker shares could reflect occupational transitions, migration, demographic change, or mortality rather than individual-level occupational mobility. The population growth placebo test (+1.49 pp, $p = 0.014$) suggests that migration may be an important channel. This compositional concern deserves particular emphasis. If MGNREGA reduced out-migration of agricultural laborers who would otherwise have left for urban non-farm employment, the observed decline in non-farm worker shares could partly reflect a retention effect—who remains in the village—rather than an occupational switching effect among a fixed population. Similarly, if the program attracted in-migrants seeking guaranteed employment, the denominator of worker shares shifts. Both channels are consistent with the “comfortable trap” interpretation (the program anchors agricultural workers in place), but they have different welfare implications: retention of workers who would have found productive urban employment is a larger efficiency cost than occupational stickiness among workers who would have remained in the village regardless. Distinguishing these mechanisms requires individual-level panel data (e.g., from the India Human Development Survey or matched Census records), which is beyond the scope of this

paper but represents a high-priority avenue for future work.

Fourth, the analysis period (2001–2011) captures only the medium-run effects of MGNREGA. Structural transformation is a long-run process, and the program’s effects may evolve as it matures, as rural labor markets adjust, and as complementary investments (roads, electrification, education) accumulate. The 2021 Census—when it becomes available—will provide a crucial test of whether MGNREGA’s effects on occupational structure persist, deepen, or reverse.

Fifth, I estimate intent-to-treat effects based on phase assignment rather than actual program exposure. Because MGNREGA implementation varied substantially across districts, the effects estimated here represent averages over a wide range of actual program intensity. Districts with strong implementation likely experienced larger effects, while those with weak implementation experienced smaller effects. More precise measures of program exposure (e.g., person-days of employment generated, expenditure per capita) could sharpen the analysis, though these measures are endogenous to local demand and administrative capacity.

8. Conclusion

This paper provides the first village-level evidence on how India’s MGNREGA may have shaped the pace and character of structural transformation. Using data on 587,378 villages across 640 districts, I document a suggestive paradox: districts that received the program earlier are associated with increased aggregate economic activity (though pre-trends complicate causal attribution) and slower movement of workers out of agriculture. The effects are concentrated among women, who experienced large declines in non-farm employment and corresponding increases in agricultural labor, while SC/ST communities partially escaped the pattern.

These findings have important implications for development policy. Employment guarantees are increasingly popular worldwide—from India’s MGNREGA to Ethiopia’s Productive Safety Net Programme to South Africa’s Expanded Public Works Programme. Understanding how these programs interact with the fundamental process of structural change is critical for designing policies that protect the poor without impeding long-run growth.

The “comfortable trap” interpretation should not be overstated. MGNREGA has delivered substantial welfare gains to hundreds of millions of rural Indians. The question is not whether the program should exist but whether it should be complemented by policies that facilitate occupational mobility—particularly for women. Skills training, rural connectivity, financial inclusion, and anti-discrimination measures could help ensure that guaranteed employment serves as a floor rather than a ceiling on workers’ aspirations.

Future research should examine whether MGNREGA's effects on structural transformation evolve over time, whether states with stronger program implementation show different patterns, and whether the gender effects documented here translate into long-run differences in human capital and earnings. The tension between social protection and structural transformation is not unique to India: it is a fundamental challenge for development policy worldwide.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Village-level Census data are from the SHRUG platform ([Asher et al., 2021](#)). MGNREGA phase assignment data follow the classification in [Zimmermann \(2020\)](#) and [Imbert and Papp \(2015\)](#). Nighttime light data are from the NOAA National Centers for Environmental Information (DMSP-OLS). All analysis was conducted in R using the `did`, `fixest`, and `sf` packages. Replication code and data are available in the project repository.

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

Contributors: @olafdrw

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

References

- Adukia, Anjali**, “Spillover Impacts on Education from Employment Guarantees,” *Working Paper, University of Chicago*, 2019.
- Afridi, Farzana, Abhiroop Mukhopadhyay, and Soham Sahoo**, “Female Labor Force Participation and Child Education in India: Evidence from the National Rural Employment Guarantee Scheme,” *IZA Journal of Labor & Development*, 2016, 5 (1), 1–27.
- Asher, Sam and Paul Novosad**, “Rural Roads and Local Economic Development,” *American Economic Review*, 2020, 110 (3), 797–823.
- , – , and **Tobias Lunt**, “The Socioeconomic High-resolution Rural-Urban Geographic Platform for India,” *Working Paper*, 2021.
- Azam, Mehtabul**, “The Impact of Indian Job Guarantee Scheme on Labor Market Outcomes: Evidence from a Natural Experiment,” *World Development*, 2012, 47, 20–32.
- Berg, Erlend, Sambit Bhattacharyya, Rajasekhar Durgam, and Manjula Ramachandra**, “Can Public Works Increase Equilibrium Wages? Evidence from India’s National Rural Employment Guarantee,” *World Development*, 2018, 103, 239–254.
- Callaway, Brantly and Pedro H.C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, 90 (3), 414–427.
- Drèze, Jean**, “NREGA: A Primer,” *The Hindu*, 2009.
- and **Reetika Khera**, “Wages and NREGA in India,” *Economic and Political Weekly*, 2017, 52, 56–64.
- Dutta, Puja, Rinku Murgai, Martin Ravallion, and Dominique Van de Walle**, “Is MGNREGA a Safety Net for the Rural Poor? Evidence from Rajasthan,” *Economic and Political Weekly*, 2012, 47 (12).
- Foster, Andrew D. and Mark R. Rosenzweig**, “Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India, 1970–2000,” *Economic Development and Cultural Change*, 2004, 52 (3), 509–542.

- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Harris, John R. and Michael P. Todaro**, “Migration, Unemployment and Development: A Two-Sector Analysis,” *American Economic Review*, 1970, *60* (1), 126–142.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil**, “Measuring Economic Growth from Outer Space,” *American Economic Review*, 2012, *102* (2), 994–1028.
- 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.
- Lewis, W. Arthur**, “Economic Development with Unlimited Supplies of Labour,” *The Manchester School*, 1954, *22* (2), 139–191.
- McMillan, Margaret, Dani Rodrik, and Claudia Sepulveda**, “Structural Change, Fundamentals, and Growth: A Framework and Case Studies,” 2016.
- 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.
- Rambachan, Ashesh and Jonathan Roth**, “A More Credible Approach to Parallel Trends,” *Review of Economic Studies*, 2023, *90* (5), 2555–2591.
- Rodrik, Dani**, “Premature Deindustrialization,” *Journal of Economic Growth*, 2016, *21* (1), 1–33.
- Sekhri, Sheetal and Tianshu Li**, “The Spillovers of Employment Guarantee Programs on Child Labor and Education,” *Working Paper*, 2020.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- Zimmermann, Laura**, “Why Is There So Little Money in Local Politics?,” *Journal of Development Economics*, 2020, *146*, 102516.

A. Data Appendix

A.1 SHRUG Data Construction

The Socioeconomic High-Resolution Rural-Urban Geographic Platform for India (SHRUG) (Asher et al., 2021) provides harmonized village-level identifiers that link administrative units across Census rounds despite boundary changes and village splits/merges. The key variable for village matching is the SHRUG village ID (`shrid`), which maps to the Census 2001 village code (`pc01_id`) and Census 2011 village code (`pc11_id`).

From the Primary Census Abstract (PCA) tables, I extract the following variables for each village:

- **Total population** and **total workers**: used to compute workforce participation rates.
- **Main workers by category**: cultivators, agricultural laborers, household industry workers, and “other workers.” Main workers are those employed for more than 183 days in the reference year.
- **Marginal workers**: employed for fewer than 183 days. I focus on main workers to avoid conflating seasonal employment fluctuations with structural change.
- **Scheduled Caste and Scheduled Tribe population**: used to construct the SC/ST population share.
- **Literate population**: used to compute the literacy rate.

Worker category shares are defined as:

$$\text{Non-farm share} = \frac{\text{HH industry workers} + \text{Other workers}}{\text{Total main workers}} \quad (4)$$

$$\text{Ag. labor share} = \frac{\text{Agricultural laborers}}{\text{Total main workers}} \quad (5)$$

$$\text{Cultivator share} = \frac{\text{Cultivators}}{\text{Total main workers}} \quad (6)$$

Gender-specific shares are constructed analogously using the gender-disaggregated worker counts available in the PCA.

A.2 Nightlight Data Construction

I use the DMS-OLS Nighttime Lights Time Series (Version 4) from NOAA’s National Centers for Environmental Information. The data provide annual composites of nighttime

stable lights for 1992–2013 at approximately 1 km resolution. I use the “stable lights” product, which removes ephemeral events (fires, lightning) and background noise.

For each district-year, I compute the mean digital number (DN) across all pixels within the district boundary, using district shapefiles from the 2011 Census. DN values range from 0 to 63, with top-coding at 63 primarily affecting urban areas. Because most of my analysis focuses on rural districts with low luminosity, top-coding is not a major concern. I use $\log(1 + \text{mean DN})$ as the outcome variable to accommodate zeros and reduce the influence of outliers.

A.3 Sample Restrictions

Starting from the universe of villages in SHRUG, I apply the following restrictions:

1. Drop villages that cannot be matched between Census 2001 and Census 2011 (approximately 8% of villages).
2. Drop villages classified as urban in either Census round.
3. Drop villages with zero total workers in either Census round (approximately 0.5% of matched villages).

The final sample contains 587,378 villages across 640 districts spanning all states and union territories represented in the SHRUG data.

B. Identification Appendix

B.1 Balance on Baseline Characteristics

Table 1 in the main text documents substantial baseline differences between Phase I and Phase III districts. These differences are by design: the backwardness index that determined phase assignment explicitly targeted districts with low agricultural productivity, low wages, and high SC/ST population shares. The identifying assumption is not that early and late districts are identical at baseline—they clearly are not—but that, conditional on observable baseline characteristics and state fixed effects, they would have experienced parallel changes in occupational structure.

To assess the plausibility of this assumption, I examine the sensitivity of the main estimate to the progressive inclusion of controls:

- State FE only: $\hat{\beta} = -0.0021$ (SE: 0.0021)
- State FE + baseline controls: $\hat{\beta} = -0.0037$ (SE: 0.0022)

- State FE + baseline controls + lagged DV: included in the above specification

The fact that the coefficient becomes *more* negative with controls (rather than attenuating toward zero) suggests that baseline differences between early and late districts, if anything, bias the unconditional estimate toward zero. This pattern is consistent with a scenario in which backward districts were on a faster convergence trajectory—experiencing more rapid structural transformation independent of MGNREGA—and controlling for baseline characteristics removes this confounding convergence, revealing the negative MGNREGA effect.

B.2 Pre-Trend Analysis for Nightlights

The Callaway-Sant’Anna event study reveals statistically significant pre-trends in nightlight intensity. A linear regression of pre-treatment coefficients on event time yields:

$$\hat{\delta}_{\text{pre}} = -0.0203 \quad (\text{SE} : 0.0040, \quad p < 0.001) \quad (7)$$

This indicates that pre-treatment nightlight growth was approximately 2 percentage points faster per year in early MGNREGA districts relative to the not-yet-treated comparison group. Several factors could explain this pattern:

1. **Convergence:** Backward districts with lower initial luminosity may grow faster as they electrify, urbanize, and develop commercial infrastructure.
2. **Anticipation:** If the MGNREGA rollout was anticipated (it was announced in 2005), local economic actors may have adjusted behavior before formal implementation.
3. **Correlated policies:** Other government programs (e.g., the Pradhan Mantri Gram Sadak Yojana road program, or the Rajiv Gandhi Grameen Vidyutikaran Yojana electrification program) may have disproportionately targeted backward districts during the same period.

The significant pre-trend is the primary reason I treat the nightlight results as suggestive rather than causal, and rely on the Census long-difference specification—which directly controls for baseline differences—as the preferred identification strategy.

B.3 Sensitivity to [Rambachan and Roth \(2023\)](#) Bounds

A natural robustness exercise is to assess how much of the post-treatment nightlight effect could be explained by a continuation of the pre-trend. Following [Rambachan and Roth \(2023\)](#),

if we allow the post-treatment violation of parallel trends to be as large as the pre-trend slope, the implied bias over a five-year post-treatment window would be approximately $5 \times 0.020 = 0.10$. The estimated ATT of 0.270 minus this bias yields 0.170, which remains positive but substantially smaller. This back-of-envelope calculation suggests that even after accounting for pre-trends, there may be a positive MGNREGA effect on nightlight intensity, but the magnitude is uncertain.

C. Robustness Appendix

C.1 Alternative Control Variables

The main specification includes log population, literacy rate, SC/ST share, and the lagged dependent variable as controls. I have verified that the results are robust to: (i) excluding the lagged dependent variable (which could induce Nickell bias in a two-period panel, though the long time span minimizes this concern); (ii) adding additional controls (workforce participation rate, cultivator share); and (iii) using quadratic terms in baseline characteristics. In all cases, the point estimate for the non-farm share effect remains between -0.003 and -0.005 , with similar significance levels.

C.2 Alternative Outcome Definitions

The main outcome—non-farm worker share—is defined using main workers only. I have verified that results are qualitatively similar when using: (i) all workers (main plus marginal); (ii) the share of “other workers” (excluding household industry); and (iii) a binary indicator for whether the village has any non-farm workers.

C.3 Winsorization

Village-level worker shares can be noisy, particularly in small villages. Winsorizing the dependent variable at the 1st and 99th percentiles does not appreciably change the main results.

D. Heterogeneity Appendix

D.1 State-Level Heterogeneity

MGNREGA implementation varies substantially across states, with some states (Rajasthan, Andhra Pradesh, Tamil Nadu) being strong implementers and others (Bihar, Uttar Pradesh)

experiencing significant implementation challenges. While a full state-by-state analysis is beyond the scope of this paper (and would require state-level treatment variation that the current design cannot cleanly identify), the state fixed effects in the main specification absorb time-invariant state-level differences in implementation quality.

D.2 Village Size Heterogeneity

Splitting the sample by baseline village population (above/below median of approximately 1,000 persons), I find that the negative effect on non-farm shares is concentrated in smaller villages (-0.52 pp, $p < 0.05$) and attenuated in larger villages (-0.21 pp, insignificant). This pattern is consistent with the hypothesis that MGNREGA has a larger impact on labor allocation in smaller, more agricultural villages where the program represents a larger share of the local labor market.

D.3 Spatial Distribution of Effects

The geographic concentration of Phase I districts in India’s central and eastern “poverty belt” (Bihar, Jharkhand, Chhattisgarh, Madhya Pradesh, Orissa) means that the estimated effects are most applicable to these regions. States with mixed Phase I and Phase III districts within the same state provide the strongest identification, while states where all districts belong to the same phase contribute primarily through the state fixed effects.

E. Additional Figures and Tables

No additional figures or tables beyond those presented in the main text.