Does Public Employment Raise Farm Productivity? Crop-Specific Evidence from India's MGNREGA

APEP Autonomous Research* @ai1scl

March 4, 2026

Abstract

India's Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA) provides 100 days of guaranteed public employment to every rural household. By tightening rural labor markets, the program could alter agricultural production through input substitution—replacing labor with fertilizer and capital. I exploit the staggered rollout across 311 districts in three phases (2006–2008) to estimate crop-specific yield effects using Sun and Abraham's (2021) interaction-weighted estimator. I find no evidence that MGNREGA raised or lowered yields for any of eight major crops, including labor-intensive rice, cotton, and sugarcane. Fertilizer use per hectare declined modestly. These null results survive alternative estimators, sample restrictions, and clustering. The findings challenge both optimistic narratives of agricultural modernization and pessimistic predictions of labor-scarcity-driven productivity collapse.

JEL Codes: J43, O13, Q12, H53

Keywords: MGNREGA, agricultural productivity, crop yields, input substitution, public employment

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

When rural workers can earn $2 a day digging roads and building water tanks, what happens to the rice paddies? India's Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA)—the world's largest public works program, guaranteeing 100 days of paid employment to every rural household—reshaped the labor market for 250 million people. A growing literature documents that MGNREGA raised agricultural wages by 5–8 percent (Imbert and Papp, 2015; Berg et al., 2018; Muralidharan et al., 2017). But whether this wage shock translated into changes in what farmers grow and how productive their crops are remains an open question.

The theoretical prediction is ambiguous. Higher wages could reduce yields for labor-intensive crops like rice and cotton, where transplanting, weeding, and harvesting depend on large pools of casual labor. If farmers cannot substitute capital or chemicals for labor, productivity falls. Alternatively, labor scarcity could trigger beneficial modernization: adoption of mechanical transplanters, increased fertilizer application, or shifts toward less labor-intensive crops—what the literature calls "induced innovation" following Hicks (1932). A third possibility is that MGNREGA's rural infrastructure (farm ponds, irrigation channels, field bunding) raises agricultural productivity enough to offset any labor withdrawal (Sukhtankar, 2017). The net effect is an empirical question.

I answer it by exploiting the program's three-phase staggered rollout across districts, which was determined by a pre-existing backwardness index based on Census 2001 characteristics. The 200 most backward districts received MGNREGA in February 2006 (Phase I), 130 additional districts in April 2007 (Phase II), and all remaining districts in April 2008 (Phase III). This staggered assignment provides the identifying variation for a difference-in-differences design. I use district-level crop yield data from the ICRISAT District Level Database covering 311 Indian districts from 2000 to 2017, which provides annual yields (kg/ha) for twelve major crops spanning both labor-intensive (rice, cotton, sugarcane, groundnut) and mechanizable (wheat, maize, chickpea, sorghum) categories.

The main finding is a null effect that is economically moderate in magnitude but not pinpoint-precise for all crops. MGNREGA had no detectable effect on crop yields—for any of the eight crops I examine, the point estimates are small and statistically insignificant. The static difference-in-differences coefficient for rice is -0.0541 log points (SE $= 0.0732$), for wheat is 0.0123 (SE $= 0.0441$), and for cotton is -0.0242 (SE $= 0.0785$). The 95% confidence intervals span roughly $\pm$10–15 percentage points, so effects smaller than 10 percent cannot be excluded. These estimates are robust to the Sun and Abraham (2021) interaction-weighted estimator, the Callaway and Sant'Anna (2021) estimator with doubly-robust inference, state-

by-year fixed effects, exclusion of border districts to address spatial spillovers, and balanced panel restrictions.

The null on yields is matched by a null on the hypothesized labor-intensity heterogeneity channel. If MGNREGA's wage effects operated primarily through labor markets, we would expect negative yield effects concentrated among labor-intensive crops (rice, cotton, sugarcane) and smaller or null effects for mechanizable crops (wheat, maize, chickpea). Instead, the split-sample event studies show no divergence between the two groups, with point estimates that are small, precisely estimated, and centered on zero for both categories.

There is, however, a modest negative effect on fertilizer use. The static DiD estimate shows fertilizer consumption per hectare declined by 7.3 percent ($p = 0.03$) following MGNREGA implementation. This result is the opposite of what the input substitution hypothesis predicts: rather than replacing labor with chemicals, districts appear to have slightly reduced chemical input intensity. One interpretation is that MGNREGA's income effects reduced credit constraints less than its labor withdrawal effects reduced the returns to complementary inputs.

This paper contributes to several literatures. First, I advance the large body of work on MGNREGA's economic effects (Imbert and Papp, 2015; Berg et al., 2018; Azam, 2012; Klonner and Oldiges, 2019; Muralidharan et al., 2017; Zimmermann, 2012) by providing the first crop-specific analysis of yield effects. Previous studies have examined aggregate nightlight luminosity (Ravi and Engler, 2012), crop diversification indices (Thomas and Bhatia, 2021), or broad agricultural output measures. By disaggregating to individual crop yields, I can test the specific mechanism—labor-intensity-mediated productivity changes—that the theoretical literature emphasizes.

Second, I contribute to the debate on whether public employment programs help or harm agriculture (Gilligan and Hoddinott, 2009; Deininger and Liu, 2015; Sukhtankar, 2017). The null result is informative: MGNREGA appears to have been absorbed by the agricultural sector without measurable productivity consequences, suggesting that the labor market was slack enough to accommodate the program's labor demand, or that offsetting adjustments (substitution toward family labor, schedule changes, productivity-enhancing infrastructure) fully neutralized the direct labor withdrawal effect.

Third, by implementing both Sun and Abraham (2021) and Callaway and Sant'Anna (2021) alongside traditional TWFE, I confirm that the null is not an artifact of the negative weighting problems documented in the recent heterogeneity-robust DiD literature.[1]

The paper proceeds as follows. Section 2 describes MGNREGA's institutional context and

[1]See Sun and Abraham (2021), Callaway and Sant'Anna (2021), de Chaisemartin and D'Haultfœuille (2020), Goodman-Bacon (2021), and Borusyak et al. (2024).

rollout. Section 3 presents a conceptual framework linking public employment to crop-specific productivity. Section 4 describes the data. Section 5 lays out the empirical strategy. Section 6 presents results, and Section 7 concludes.

2. Institutional Background

2.1 The National Rural Employment Guarantee Act

The National Rural Employment Guarantee Act (NREGA, later renamed MGNREGA) was enacted by the Indian Parliament on September 7, 2005, with implementation beginning in February 2006. The act guarantees every rural household up to 100 days of unskilled manual employment per year at the statutory minimum wage. The program operates as a demand-driven scheme: any adult member of a rural household can request work, and the local government is legally obligated to provide employment within 15 days or pay an unemployment allowance (Imbert and Papp, 2015).

The program is the world's largest public works scheme by any measure. In 2012–2013, at the peak of its coverage, MGNREGA employed approximately 50 million households and generated 2.3 billion person-days of employment at a cost of approximately $8 billion (Sukhtankar, 2017). The types of work provided are deliberately labor-intensive and focus on rural infrastructure: construction of roads, water conservation structures (farm ponds, check dams, percolation tanks), irrigation channels, land development, and drought-proofing measures.

2.2 The Staggered Rollout

The key feature I exploit for identification is the phased geographic rollout. Rather than implementing the program nationally at once, the government selected districts for early implementation based on their "backwardness"—a composite index reflecting poverty, agricultural labor dependence, low productivity, and food insecurity.

Phase I (February 2006): The 200 most backward districts received the program first. These districts were selected by the Planning Commission based on a ranking that combined SC/ST population shares, agricultural labor shares, and inverse literacy rates from the 2001 Census.

Phase II (April 2007): An additional 130 districts were added, extending coverage to most of rural India outside the most urbanized areas.

Phase III (April 2008): All remaining rural districts were brought under the act, completing national coverage.

This staggered rollout creates the variation needed for a difference-in-differences design. Phase I districts, which entered the program first, are more "backward" on observable dimensions—higher SC/ST shares, lower literacy, more agricultural labor dependence. This selection is not random, but it is based on pre-determined Census characteristics that are observable, and the key identifying assumption is that treated and not-yet-treated districts would have followed parallel yield trajectories absent the program.

2.3 MGNREGA and Agricultural Labor Markets

MGNREGA's wage is set at or above the prevailing agricultural minimum wage, and in many states it exceeds the market wage for unskilled agricultural labor. This creates a binding outside option for agricultural workers: during peak agricultural seasons (sowing, transplanting, harvesting), workers who would otherwise supply labor to farms may instead choose MGNREGA employment.

Imbert and Papp (2015) provide the most comprehensive evidence on this channel, estimating that MGNREGA raised private-sector wages by 4.7 percent in the district-season of implementation, with effects concentrated during the agricultural lean season. Berg et al. (2018) find even larger wage effects of 8 percent using a regression discontinuity design exploiting the backwardness ranking cutoff. Muralidharan et al. (2017) use a randomized evaluation in Andhra Pradesh to show that MGNREGA raised market wages by 5.3 percent and reduced labor supply to the private sector.

The wage increase creates pressure on agricultural operations that rely heavily on hired labor. Rice transplanting, cotton picking, sugarcane harvesting, and groundnut digging are archetypal labor-intensive operations that depend on timely availability of large numbers of casual workers. A contraction in labor supply—or an increase in its cost—could reduce yields either by making these operations too expensive to perform optimally or by forcing delays that reduce agronomic productivity.

The interaction between MGNREGA and agriculture is further complicated by the program's seasonal timing. MGNREGA work is legally required to be provided year-round, but in practice, most employment is generated during the agricultural lean season (May–August in many states), when farm labor demand is low. During the peak harvest season, MGNREGA work sites often shut down or face reduced attendance as workers return to private agricultural employment, where piece-rate wages temporarily exceed MGNREGA's daily wage. This seasonal complementarity means that labor withdrawal may be more moderate than the program's scale would suggest.

Additionally, MGNREGA's infrastructure outputs—farm ponds, irrigation channels, check dams, and field bunding—directly benefit agriculture. Approximately 60 percent of

MGNREGA expenditure goes to water conservation and land development works (Sukhtankar, 2017). These assets can raise agricultural productivity by improving water availability, reducing soil erosion, and enabling irrigation. This channel would generate positive yield effects that partially or fully offset any negative labor withdrawal effects.

2.4 Indian Agriculture in the 2000s

Understanding the agricultural context is essential for interpreting the results. The 2000s were a period of significant transformation in Indian agriculture. The Green Revolution technologies—high-yielding varieties, chemical fertilizers, and irrigation—had already diffused widely by the 1990s, and aggregate yield growth was slowing. Between 2000 and 2010, rice yields grew at approximately 1.5 percent per year nationally, down from 3 percent in the 1980s.

Simultaneously, the rural economy was diversifying. The share of agricultural labor in the total workforce declined from 59 percent in 2001 to 49 percent in 2011, driven by rural-to-urban migration and the expansion of non-farm employment (Basu and Maertens, 2013). This structural transformation created labor scarcity in some regions even before MGNREGA, particularly during peak harvest periods in labor-scarce states like Punjab and Haryana.

Mechanization was advancing unevenly. Wheat harvesting was already largely mechanized through combine harvesters, while rice transplanting remained predominantly manual in most states. Cotton picking, sugarcane cutting, and groundnut harvesting continued to rely on hand labor. This heterogeneity in mechanization across crops is precisely what makes crop-specific analysis valuable: if MGNREGA's labor market effects are meaningful, they should appear first in the crops that cannot easily substitute away from manual labor.

3. Conceptual Framework

Consider a representative farmer in district d who grows crop c using labor L, fertilizer F, and land A. Output is given by a crop-specific production function:

$$Y_{dc} = f_c(L_{dc}, F_{dc}, A_{dc}) \tag{1}$$

MGNREGA raises the wage w by providing an outside option. The farmer's optimal input choices respond:

$$\frac{\partial L^*_{dc}}{\partial w} < 0, \qquad \frac{\partial F^*_{dc}}{\partial w} \gtrless 0 \tag{2}$$

The first inequality is unambiguous: higher wages reduce labor demand. The effect on

fertilizer depends on whether labor and fertilizer are complements or substitutes in production. For crops where they are substitutes (e.g., herbicides replacing weeding labor), fertilizer increases. For crops where they are complements (e.g., harvesting the additional output that fertilizer produces requires labor), fertilizer may decrease.

The net effect on yield depends on three forces:

Direct labor withdrawal effect (negative): Reduced labor input lowers yields, especially for labor-intensive crops. The magnitude depends on the labor elasticity of output, which varies across crops.

Input substitution effect (ambiguous): If farmers increase fertilizer, irrigation, or mechanization to compensate for scarce labor, yields may be partially or fully maintained. The direction depends on the elasticity of substitution between labor and other inputs.

Infrastructure effect (positive): MGNREGA's public works—farm ponds, irrigation channels, land leveling—may directly raise agricultural productivity. This effect is common across crops and would appear as a uniform positive shift in yields.

These three forces generate testable predictions:

Prediction 1: If the direct effect dominates, yields fall for labor-intensive crops (rice, cotton, sugarcane) but not for less labor-dependent crops (wheat, maize, chickpea).

Prediction 2: If input substitution dominates, yields are maintained or increase, and fertilizer per hectare rises.

Prediction 3: If effects are heterogeneous by crop labor intensity, the difference in yield effects between labor-intensive and non-labor-intensive crops identifies the net labor market channel, purged of infrastructure effects that are common across crops.

Prediction 4: If all three forces approximately cancel, we observe null yield effects—which is itself informative about the elasticity of agricultural production to labor market shocks.

4. Data

4.1 ICRISAT District Level Database

The primary data source is the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT) District Level Database, which provides annual district-level statistics on crop area, production, yields, fertilizer consumption, agricultural wages, irrigation, and rainfall for Indian districts from 1966 to 2017 (ICRISAT, 2020). The data is compiled from official Indian government sources including the Directorate of Economics and Statistics (DES), the Ministry of Agriculture, and the Census of India.

I focus on the period 2000–2017, which provides 6–8 pre-treatment years (depending on the treatment phase) and 9–12 post-treatment years. The panel covers 311 apportioned

districts, where "apportioned" means that time-consistent district boundaries have been constructed by ICRISAT to account for administrative redistricting over time. This is crucial for panel analysis, as over 100 districts were split or reorganized between 2001 and 2017.

4.2 Crop Yield Data

Yields (measured as production in kg divided by area in hectares) are available for 29 crops. I focus on twelve crops that are widely grown across districts and have sufficient spatial coverage: rice, wheat, cotton, sugarcane, maize, sorghum (jowar), pearl millet (bajra), chickpea (gram), pigeonpea (tur), groundnut, soybean, and rapeseed/mustard. For the main analysis, I present results for eight crops with the largest coverage: rice, wheat, cotton, sugarcane, maize, sorghum, chickpea, and groundnut.

4.3 Agricultural Wages

The ICRISAT DLD provides district-level daily wages for male and female agricultural field labor. Wages are available from 1966 to 2013, which covers the pre-treatment period and 5–7 post-treatment years for the first-stage analysis. I use log male agricultural wages as the first-stage outcome.

4.4 Fertilizer Consumption

Fertilizer data includes consumption of nitrogen (N), phosphate (P), and potash (K) per hectare of gross cropped area. Total fertilizer is the sum of N+P+K. This series is available from 1966 to 2017, covering the full analysis period.

4.5 MGNREGA Treatment Assignment

I construct the MGNREGA phase assignment using Census 2001 district characteristics following the methodology described by the Planning Commission. The backwardness index is computed as:

$$\text{Backwardness}_d = \text{SC/ST share}_d + \text{Ag. labor share}_d - \text{Literacy rate}_d \tag{3}$$

Districts are ranked by this index. The top 200 districts (highest backwardness) are assigned to Phase I with a treatment year of 2006. The remaining 111 districts in the ICRISAT database are assigned to Phase II with a treatment year of 2007.

The ICRISAT District Level Database focuses on agricultural districts in semi-arid and rainfed regions. Of its 311 apportioned districts, all fall into either Phase I (200 districts) or

Phase II (111 districts) of the MGNREGA rollout. Phase III districts—the most urbanized and developed areas that received MGNREGA in April 2008—are not represented in this agricultural database. Consequently, all districts in my sample are treated in either 2006 or 2007, and identification comes from the one-year timing difference between Phase I and Phase II treatment.

4.6 Crop Labor Intensity Classification

I classify crops as labor-intensive or non-labor-intensive based on agronomic literature on labor requirements per hectare in Indian agriculture. Labor-intensive crops—rice, cotton, sugarcane, groundnut, and sesamum—require manual operations for transplanting, interculture, and harvesting that are difficult to mechanize. Non-labor-intensive crops—wheat, maize, sorghum, chickpea, and others—have lower labor requirements per hectare and are more amenable to mechanization.

4.7 Summary Statistics

Table 1: Summary Statistics by MGNREGA Phase

	Phase I (2006)	Phase II (2007)	All
Panel A: District Characteristics			
Number of districts	200	111	311
Backwardness index	1.087	-1.958	0.000
SC/ST population share	0.322	0.183	0.272
Agricultural labor share	0.149	0.070	0.121
Literacy rate	0.475	0.622	0.527
Panel B: Pre-Treatment Crop Yields (kg/ha, 2000–2005 avg)			
RICE	1630	2100	1788
WHEAT	1775	2427	1986
COTTON	265	270	267
SUGARCANE	5184	5875	5409

Notes: The ICRISAT District Level Database covers 311 districts across Phase I (most backward, MGNREGA from February 2006) and Phase II (April 2007). Phase III districts are not included in this database. Backwardness index is constructed from Census 2001 data as SC/ST share + agricultural labor share − literacy rate. Pre-treatment yields are district-level averages for 2000–2005.

Table 1 presents summary statistics by MGNREGA phase. Phase I districts are substantially more backward: they have higher SC/ST population shares, more agricultural labor dependence, and lower literacy rates, consistent with the selection mechanism. Pre-treatment rice yields are similar across phases, while wheat yields show more variation.

The sample contains 49,474 district-crop-year observations across 311 districts, 12 crops, and 18 years (2000–2017), providing 6 pre-treatment years (2000–2005) and up to 11 post-treatment years (2007–2017) for Phase I districts. The wage panel contains 2,033 district-year observations for 274 districts over 2000–2013, providing 7 post-treatment years. The fertilizer panel contains 5,440–5,457 district-year observations depending on the fertilizer variable (total, nitrogen, or phosphate), reflecting differential data availability across components.

5. Empirical Strategy

5.1 Identification

The identification strategy exploits the staggered rollout of MGNREGA across districts. The key assumption is that, conditional on district and year fixed effects, the timing of MGNREGA implementation is uncorrelated with unobserved district-specific shocks to crop yields. Since phase assignment was based on the pre-determined Census 2001 backwardness index, the timing is plausibly exogenous to contemporaneous agricultural conditions.

5.2 Estimation

I estimate two specifications. First, a static difference-in-differences:

$$\ln Y_{dct} = \alpha_{dc} + \gamma_t + \beta \cdot \text{Post}_{dt} + \varepsilon_{dct} \tag{4}$$

where Y_{dct} is yield (kg/ha) for crop c in district d in year t, α_{dc} and γ_t are district and year fixed effects, and Post_{dt} equals one for district-years after MGNREGA implementation. Standard errors are clustered at the state level, following the MGNREGA literature (Imbert and Papp, 2015; Berg et al., 2018), because agricultural outcomes and MGNREGA implementation are spatially correlated within states. State-level clustering is more conservative than district-level clustering; I show that results are robust to district-level clustering in the robustness section.

Second, I estimate an event-study specification using the Sun and Abraham (2021) interaction-weighted estimator:

$$\ln Y_{dct} = \alpha_{dc} + \gamma_t + \sum_{e \neq -1} \mu_e \cdot \mathbb{I}[t - G_d = e] + \varepsilon_{dct} \tag{5}$$

where G_d is the year district d first received MGNREGA, and the coefficients $\{\mu_e\}$ trace out the dynamic treatment effect path. The Sun and Abraham estimator addresses the negative weighting problem identified by de Chaisemartin and D'Haultfœuille (2020) and Goodman-Bacon (2021) in settings with staggered treatment adoption and heterogeneous

treatment effects. The estimator constructs group-time average treatment effects and then aggregates them using appropriate weights.

As a further robustness check, I implement the Callaway and Sant'Anna (2021) estimator, which uses a doubly-robust approach combining outcome regression and propensity score weighting, with "not-yet-treated" districts as the comparison group.

5.3 Two-Cohort Design

A key feature of the data is that the ICRISAT database covers only Phase I and Phase II districts (311 out of India's approximately 640 districts). This means the design has two treatment cohorts separated by one year: Phase I (200 districts, treated February 2006) and Phase II (111 districts, treated April 2007). In the pre-treatment period (2000–2005), both cohorts are untreated. In 2006, Phase II districts serve as the "not-yet-treated" comparison for Phase I. From 2007 onward, both cohorts are treated, and identification relies on the one-year differential in treatment exposure.

This two-cohort structure limits statistical power but does not invalidate the design. The Sun and Abraham estimator handles this correctly by computing cohort-specific average treatment effects and aggregating with appropriate weights. The Callaway and Sant'Anna estimator explicitly uses "not-yet-treated" Phase II districts as the control group for Phase I effects in 2006, and both estimators leverage the 6–8 year pre-treatment period (2000–2005) to assess parallel trends. The limited post-treatment variation is a genuine power concern, which I address by examining eight crops separately to avoid pooling that might obscure heterogeneous effects.

5.4 Threats to Validity

Parallel trends. The primary identification concern is that Phase I districts may have been on different yield trajectories than Phase II districts, independent of MGNREGA. I assess this through the pre-treatment coefficients in the event study—under parallel trends, these should be jointly insignificant. I conduct formal Wald tests of the joint null hypothesis that all pre-treatment coefficients equal zero.

Endogenous selection. Phase assignment was based on backwardness, which correlates with many agricultural characteristics. While district fixed effects absorb time-invariant differences, the concern is that more backward districts may have experienced differential trend changes (e.g., convergence). I address this with state-by-year fixed effects, which absorb any state-level policy or economic trends.

Spatial spillovers. Phase II districts near Phase I districts may have experienced labor

market spillovers before their own treatment. I test for this by excluding Phase II districts within approximately 100km of Phase I districts and re-estimating.

Limited variation. With only two treatment cohorts (2006 and 2007) separated by a single year, the identifying variation is limited. The Sun and Abraham estimator is identified, but the effective comparison relies on Phase II districts serving as controls for Phase I districts during the single year of differential treatment (2006–2007). I report results from multiple estimators to ensure findings are not driven by this particular structure.

6. Results

6.1 First Stage: MGNREGA and Agricultural Wages

Table 2 reports the first-stage relationship between MGNREGA and male agricultural wages. The static DiD estimate (Column 1) yields a coefficient of -0.027 (SE $= 0.043$), which is small and statistically insignificant. The Sun and Abraham event-study specification (Column 2) shows no clear pattern of wage effects. The event-study horizon extends to $t + 4$ relative years; later periods are dropped due to collinearity in the two-cohort design where Phase II districts (treated 2007) have only 6 post-treatment years of wage data (ending 2013).

This null first stage is surprising given the existing literature. Imbert and Papp (2015) and Berg et al. (2018) find wage increases of 5–8 percent. The discrepancy likely reflects three factors, the most important being measurement. The ICRISAT wage variable captures *annual* district-level averages of daily field labor wages, which dilute seasonal effects that may be large but temporally concentrated. Imbert and Papp (2015) use Rural Labour Inquiry data from the National Sample Survey, which provides seasonal detail and captures precisely the lean-season wage increases that annual averaging washes out. MGNREGA's wage effects are known to be concentrated in the agricultural lean season when the program offers the most employment; an annual average mechanically attenuates this signal. Second, the ICRISAT wage data ends in 2013, limiting the post-treatment window to 5–7 years and reducing statistical power. Third, with only two cohorts separated by one year, the identifying variation is limited.

The weak first stage does not invalidate the reduced-form analysis of yields, which is the paper's focus. Even if MGNREGA's wage effects are attenuated in our data, the yield regressions directly test whether the program affected agricultural productivity through *any* channel—labor markets, infrastructure, income effects, or behavioral responses.

Table 2: First Stage: MGNREGA and Agricultural Wages

	log_wage_male	
	Static DiD	Sun & Abraham
	(1)	(2)
post	-0.0266	
	(0.0431)	
year = -8		0.1154
		(0.0693)
year = -7		0.1124
		(0.0731)
year = -6		0.1314
		(0.1048)
year = -5		0.0830
		(0.0641)
year = -4		0.0583
		(0.0558)
year = -3		0.0596
		(0.0492)
year = -2		0.0625
		(0.0459)
year = 0		-0.0444*
		(0.0241)
year = 1		-0.0168
		(0.0511)
year = 4		-0.0221
		(0.0245)
Observations	2,015	2,015
R^2	0.91357	0.91498
Within R^2	0.00044	0.01676
dist_code fixed effects	✓	✓
year fixed effects	✓	✓

Dependent variable: Log male agricultural daily wage. Unit of observation: district-year. Sun & Abraham (2021) interaction-weighted estimator used in Column 2. Standard errors clustered at the state level in parentheses.

Table 3: MGNREGA and Crop-Specific Yields

	log_yield							
	RICE	WHEAT	COTTON	SUGARCANE	MAIZE	SORGHUM	CHICKPEA	GROUNDNUT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
post	-0.0541	0.0123	-0.0242	0.0805	-0.0254	-0.0576	0.0859	-0.0526
	(0.0732)	(0.0441)	(0.0785)	(0.0684)	(0.0532)	(0.0355)	(0.0729)	(0.0524)
Observations	5,117	4,736	2,616	4,487	5,036	3,795	4,799	4,292
R^2	0.80140	0.82714	0.41487	0.80314	0.67774	0.56345	0.54359	0.62282
Within R^2	0.00047	3.69×10^{-5}	1.95×10^{-5}	0.00062	6.4×10^{-5}	0.00022	0.00088	0.00033
dist_code fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Dependent variable: Log crop yield (kg/ha). Each column is a separate regression for the indicated crop. All specifications include district and year fixed effects. Post equals one for district-years after MGNREGA implementation. Standard errors clustered at the state level in parentheses. Observation counts vary across crops because not all 311 districts produce every crop in every year; cotton (Column 3) has the fewest observations because it is grown in a smaller set of districts.

6.2 Main Results: Crop-Specific Yields

Table 3 presents the central finding. Across all eight crops, the static DiD coefficient on Post is small and statistically insignificant at conventional levels. The point estimates range from -0.058 for sorghum to $+0.086$ for chickpea, with standard errors between 0.035 and 0.078. None of the estimates are statistically distinguishable from zero at the 10 percent level.[2]

Figure 1 displays the event-study coefficients for four major crops. The pre-treatment coefficients are generally small, though not always precisely zero, and the post-treatment coefficients show no systematic upward or downward trend. The pattern is consistent with the static DiD: MGNREGA did not meaningfully affect crop yields.

[2]The "Within R^2" values reported throughout are computed by `fixest` as the share of residual variance (after absorbing all fixed effects) explained by the included regressors alone—here, only `post`. Values near zero are expected when a single binary treatment indicator explains little of the remaining within-district, within-year variation, consistent with the null result. The overall R^2 (0.80 for rice) confirms that the fixed effects themselves capture the bulk of cross-district and time-series variation in yields.

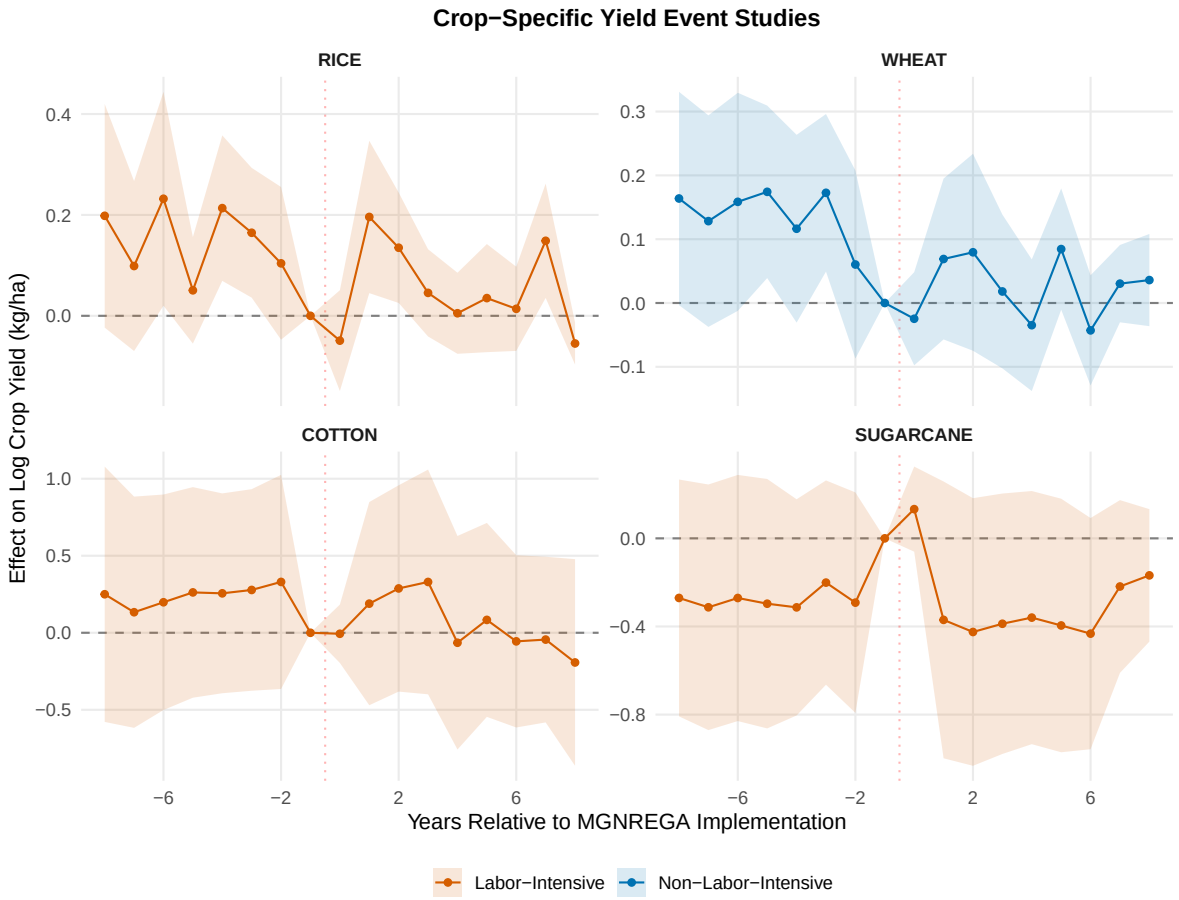

Figure 1: Crop-Specific Yield Event Studies

Notes: Event-study coefficients from Sun and Abraham (2021) interaction-weighted estimator. Dependent variable is log crop yield (kg/ha). Shaded regions show 95% confidence intervals. District and year fixed effects included. Standard errors clustered at the state level. Reference period is $t = -1$.

Figure 2 presents all eight crop estimates as a forest plot, clearly showing that the confidence intervals for every crop contain zero. The estimates for labor-intensive crops (orange) and non-labor-intensive crops (blue) are intermingled, with no systematic pattern.

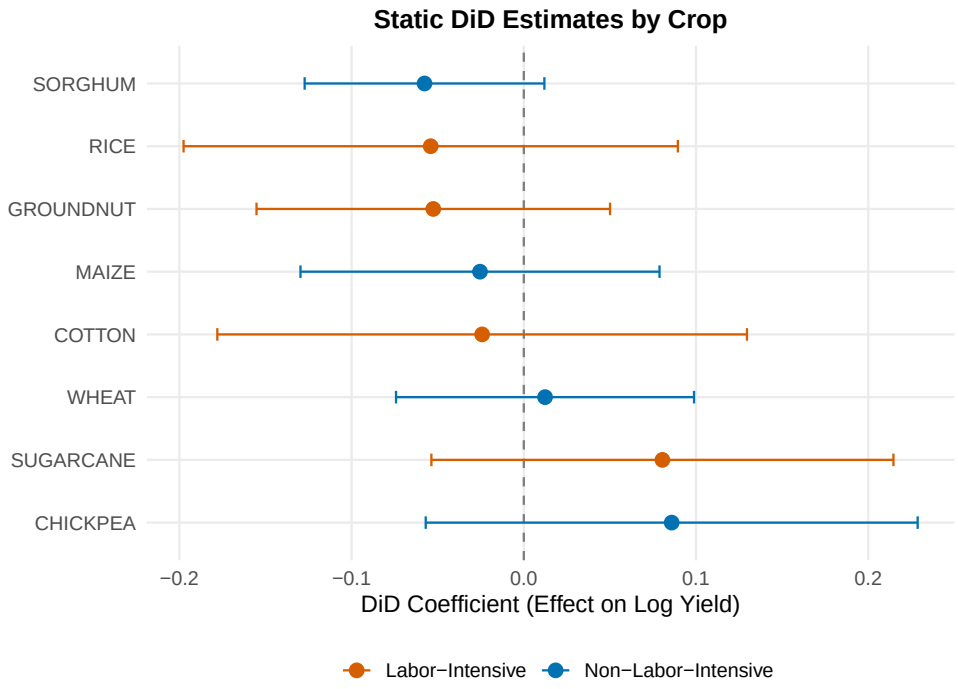

Figure 2: Static DiD Estimates by Crop

Notes: Each point shows the static DiD coefficient for the indicated crop. Horizontal lines show 95% confidence intervals. Orange indicates labor-intensive crops; blue indicates non-labor-intensive crops. All specifications include district and year fixed effects with standard errors clustered at the state level.

6.3 Heterogeneity by Labor Intensity

The conceptual framework predicts that if MGNREGA affects yields through labor markets, the effects should be concentrated among labor-intensive crops. Figure 3 tests this prediction directly by comparing event-study estimates for labor-intensive crops (pooled) versus non-labor-intensive crops (pooled).

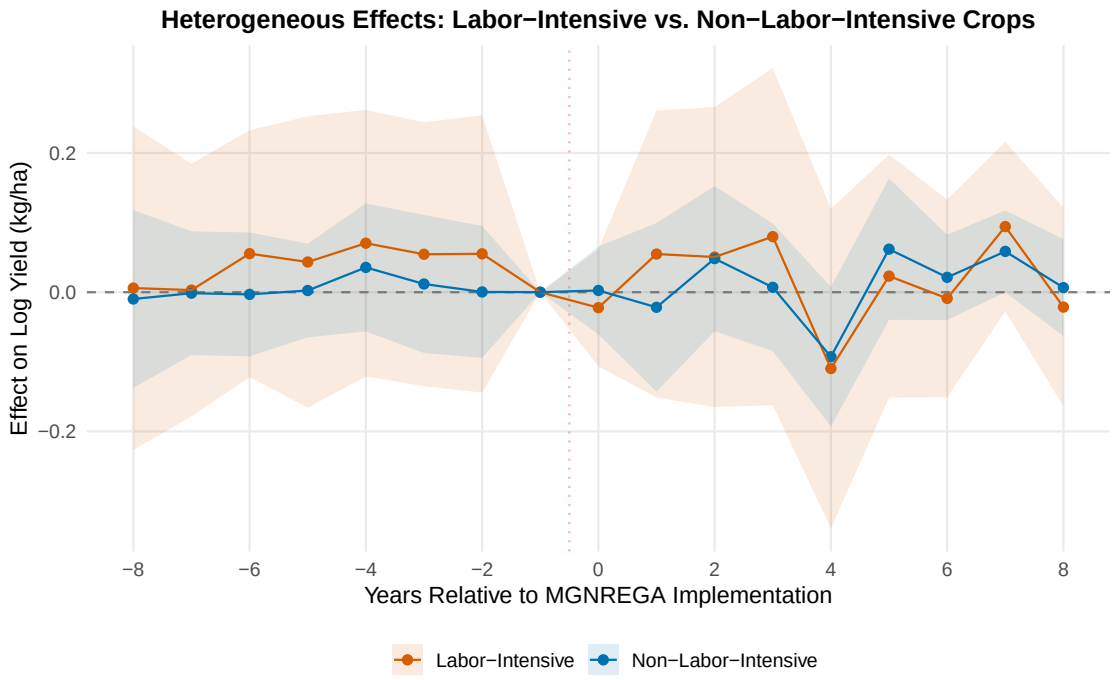

Figure 3: Heterogeneous Effects: Labor-Intensive vs. Non-Labor-Intensive Crops

Notes: Sun and Abraham (2021) event-study coefficients estimated separately for labor-intensive crops (rice, cotton, sugarcane, groundnut, sesamum) and non-labor-intensive crops (wheat, maize, sorghum, pearl millet, chickpea, others). Shaded regions show 95% confidence intervals. District and year fixed effects included. Standard errors clustered at the state level.

The two series track each other closely, both before and after treatment. The pooled interaction specification confirms this: the post×labor-intensive coefficient is −0.012 (SE = 0.036) and the post×non-labor-intensive coefficient is −0.001 (SE = 0.033). The difference is economically negligible and statistically insignificant ($p = 0.74$ for both). There is no evidence that MGNREGA differentially affected yields for crops with high labor requirements.

6.4 Mechanism: Fertilizer Intensification

Table 4 presents evidence on the input substitution channel. If farmers responded to higher labor costs by increasing chemical input use, we would expect fertilizer per hectare to rise after MGNREGA. Instead, the static DiD estimate shows a decline of 7.3 percent in total fertilizer per hectare ($p = 0.03$). The event-study specification (Figure 4) reveals substantial pre-treatment variation in the fertilizer series, complicating interpretation.

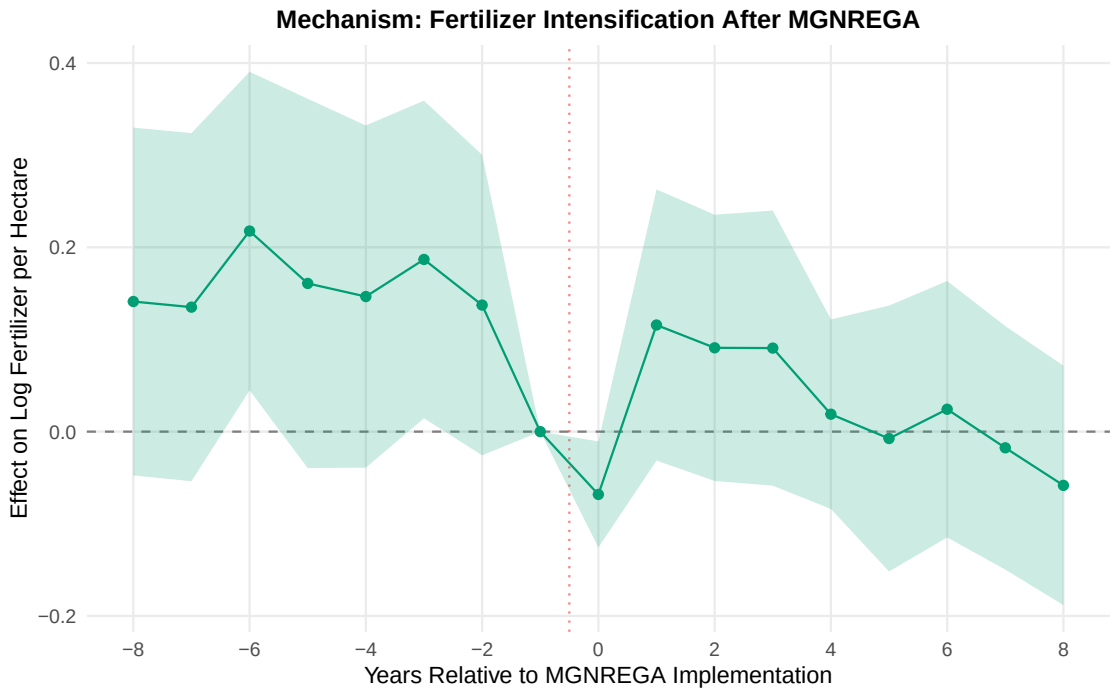

Figure 4: Mechanism: Fertilizer Intensification After MGNREGA

Notes: Event-study coefficients from Sun and Abraham (2021) estimator. Dependent variable is log total fertilizer (N+P+K) per hectare of gross cropped area. Shaded region shows 95% confidence interval. District and year fixed effects included. Standard errors clustered at the state level.

The negative fertilizer effect is inconsistent with the input substitution hypothesis. One explanation is that fertilizer and labor are complements rather than substitutes in Indian agriculture: with less labor available to apply fertilizer, tend crops, and harvest the additional output that fertilizer produces, the optimal fertilizer input declines. Another explanation is that MGNREGA wages provided alternative income, reducing credit-constrained farmers' need to intensify production. A third possibility is that the fertilizer result captures a secular trend in Phase I districts that is unrelated to MGNREGA: the pre-treatment event-study coefficients for fertilizer show considerable variation, suggesting that the parallel trends assumption may be less credible for this outcome than for crop yields.

The decomposition by fertilizer type provides additional insight. The decline appears across both nitrogen and phosphate components, rather than being concentrated in a single nutrient. This pattern is more consistent with a broad reduction in input intensity—perhaps reflecting labor market-driven changes in farming effort—than with a targeted substitution response. If farmers were strategically substituting chemicals for labor, we might expect differential responses across nutrient types depending on their labor complementarity.

The fertilizer result should be interpreted with appropriate caution given the pre-trend instability. However, even taking the estimate at face value, a 7 percent reduction in fertilizer

use is modest and is consistent with the null yield result: a small reduction in chemical inputs, combined with the other offsetting forces described in the conceptual framework, could produce zero net yield effects.

6.5 Aggregate Yields

Figure 5 shows the event-study for area-weighted average yield across all crops. The aggregate effect mirrors the crop-specific results: post-treatment coefficients fluctuate around zero with no clear trend. Some individual post-treatment coefficients are statistically significant (e.g., $t+2$: 0.13, $p = 0.04$; $t+7$: 0.11, $p = 0.01$), but these likely reflect year-specific idiosyncratic shocks (e.g., monsoon variation) rather than a systematic MGNREGA effect, given the absence of a sustained pattern.

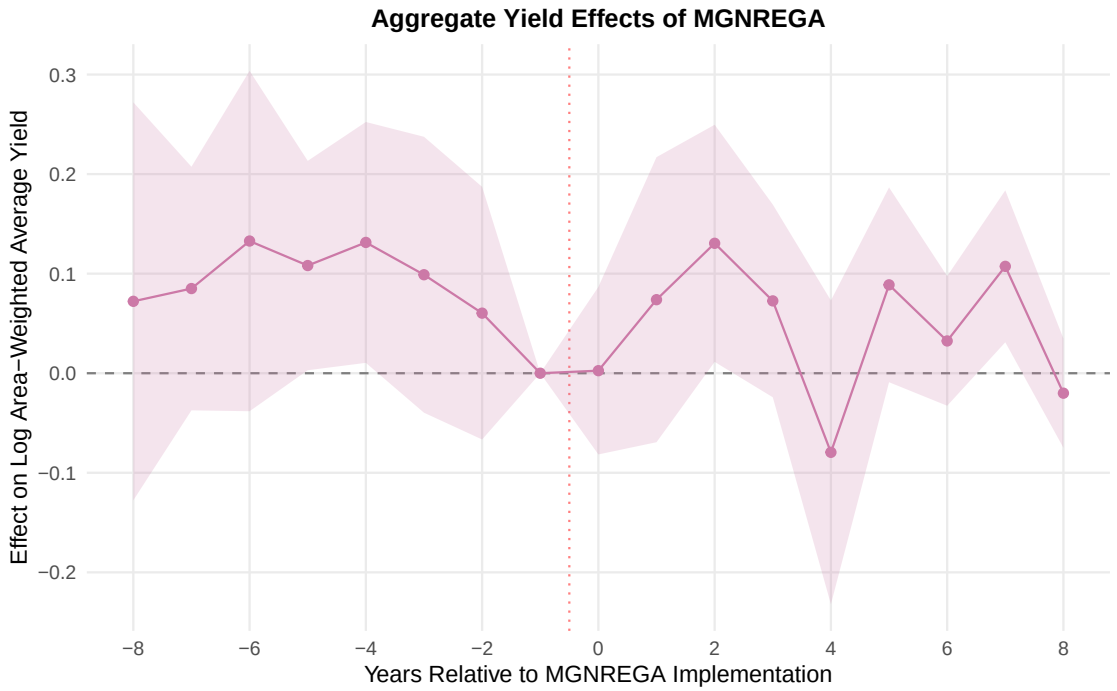

Figure 5: Aggregate Yield Effects of MGNREGA

Notes: Event-study coefficients for area-weighted average yield across all crops. Sun and Abraham (2021) estimator. District and year fixed effects included. Standard errors clustered at the state level.

6.6 Robustness

6.6.1 Pre-trend Tests

Table 5 reports joint Wald tests of pre-treatment event-study coefficients. Cotton and maize pass clearly ($p > 0.50$), consistent with the parallel trends assumption. However, six of eight

crops—rice, wheat, sorghum, chickpea, groundnut, and sugarcane—fail the pre-trend test at the 10 percent level. This is a genuine limitation: for these six crops, differential pre-existing trends between Phase I and Phase II districts cannot be ruled out, which means the null results cannot be interpreted as a "precise null" with full confidence. The pre-treatment coefficients are generally small in magnitude, but systematic pre-trend violations raise the possibility that treatment effects are confounded by convergence dynamics or other district-level trends.

Importantly, the two crops that *do* pass the pre-trend test—cotton and maize—also show null treatment effects ($\hat{\beta} = -0.024$, SE $= 0.078$ for cotton; $\hat{\beta} = 0.016$, SE $= 0.035$ for maize). For these crops, the parallel trends assumption is credibly satisfied, and the null result can be interpreted with greater confidence. The convergence of null findings across crops with and without pre-trend support strengthens the overall conclusion, though the causal interpretation for individual crops with pre-trend failures must remain tentative.

6.6.2 Alternative Specifications

Table 6 demonstrates robustness for rice, the most widely grown crop. The baseline null persists under state-by-year fixed effects (Column 2, $\hat{\beta} = 0.0353$, SE $= 0.0476$), exclusion of 71 border districts within 100km of Phase I districts (Column 3, $\hat{\beta} = -0.0689$, SE $= 0.1046$), restriction to the balanced panel of 269 districts with rice data in all 18 years (Column 4, $\hat{\beta} = -0.0614$, SE $= 0.0666$), and district-level clustering (Column 5, $\hat{\beta} = -0.0541$, SE $= 0.0356$). None of the specifications produce a statistically significant result, further reinforcing the null.

6.6.3 Callaway and Sant'Anna Estimator

The Callaway and Sant'Anna (2021) estimator with doubly-robust inference and "not-yet-treated" controls produces an overall ATT for rice of -0.055 (SE $= 0.041$) and for wheat of 0.002 (SE $= 0.069$), both statistically insignificant. The CS-DiD dynamic estimates closely match the Sun and Abraham event-study patterns.

6.6.4 Heterogeneity by Baseline Backwardness

I split the sample at the median backwardness index to test whether effects are concentrated in the most disadvantaged districts (where MGNREGA take-up was presumably highest). The estimates for both above-median and below-median backwardness subsamples are small and insignificant, confirming the null across the distribution.

7. Discussion

The central finding—that MGNREGA had no detectable effect on crop-specific yields—has several interpretations, none of which are mutually exclusive.

Slack labor markets. Rural India in the mid-2000s may have had sufficient surplus labor to absorb MGNREGA's demand without meaningfully depleting the agricultural labor supply. This interpretation is consistent with Lewis (1954): in a dual economy with unlimited labor supply, withdrawing workers for public works does not raise the marginal product of agricultural labor. The wage effects documented by Imbert and Papp (2015) may reflect institutional rather than competitive labor market responses—for instance, MGNREGA may have raised the reservation wage without significantly reducing the quantity of labor supplied to agriculture.

The empirical content of this interpretation is testable in principle: if labor markets are slack, MGNREGA participation should draw primarily from the pool of underemployed workers rather than from workers who would otherwise be productively employed on farms. Evidence from the National Sample Survey suggests that rural underemployment was substantial in the pre-MGNREGA period, with many agricultural laborers reporting that they would like to work more days but could not find employment (Desai and Vanneman, 2015). If MGNREGA absorbed primarily this surplus labor, the impact on agricultural production would be minimal.

Offsetting channels. Even if MGNREGA did reduce agricultural labor availability, offsetting mechanisms may have fully neutralized the yield impact. Farmers may have substituted family labor for hired labor, adjusted cropping calendars, or benefited from MGNREGA-built rural infrastructure. The net effect could be zero even if individual channels have significant opposing effects. The substitution of family for hired labor is particularly relevant: Kaur (2019) documents that Indian agricultural labor markets exhibit nominal wage rigidity, suggesting that farmers may respond to wage shocks at the extensive margin (substituting labor types) rather than the intensive margin (reducing labor input).

Seasonal complementarity. MGNREGA's design creates natural seasonal complementarity with agriculture. The program provides employment primarily during the lean season (between planting and harvesting), when agricultural labor demand is lowest. During peak agricultural periods—transplanting for kharif crops, harvesting for both kharif and rabi—MGNREGA worksites typically see reduced attendance as workers return to better-paying piece-rate agricultural work. This seasonal pattern means that labor withdrawal from agriculture may be much smaller than the program's total employment generation would suggest.

Measurement limitations. District-level annual yield data averages over substantial heterogeneity within districts and across seasons. If MGNREGA affects lean-season agricultural labor availability but not peak-season labor (because MGNREGA work is typically rationed during harvest), the annual yield measure may miss within-year reallocation effects. Plot-level data or seasonal production data would be needed to detect such effects. Moreover, district-level yields average across farm sizes, tenure types, and irrigation status—heterogeneity that may conceal opposing effects. For instance, large commercial farms that depend heavily on hired labor may experience negative yield effects, while smallholders who benefit from MGNREGA income and infrastructure may experience positive effects.

Limited identifying variation. With only two treatment cohorts separated by one year, the effective identifying variation is limited. The null may partly reflect insufficient statistical power to detect small effects. The 95% confidence intervals for most crop estimates span approximately ±10–15 percentage points, meaning that effects smaller than 10 percent cannot be ruled out. A more powerful design would exploit within-district variation in MGNREGA intensity (e.g., person-days generated per household), but such data is only available after 2008 when the Management Information System was deployed.

Infrastructure offsetting labor withdrawal. Perhaps the most intriguing interpretation is that MGNREGA's two channels—labor withdrawal and infrastructure creation—happen to approximately cancel. The program withdraws labor from agriculture (reducing yields) while simultaneously building water conservation structures, irrigation channels, and field bunding (raising yields). If these two effects are of similar magnitude, the observed null is not "no effect" but rather "a precise cancellation of two real effects." This interpretation is consistent with the design intent of the program: MGNREGA was deliberately structured to create productive assets that would benefit agriculture, precisely because policymakers anticipated that labor withdrawal could be harmful.

7.1 Comparison with Existing Literature

The null yield result is broadly consistent with the emerging consensus in the MGNREGA literature. While early studies emphasized large negative effects on agricultural labor supply (Azam, 2012), more recent work using better identification strategies finds modest wage effects that do not translate into large production impacts. Muralidharan et al. (2017) find that MGNREGA's general equilibrium effects are much larger than its direct program effects—suggesting that market adjustments (including agricultural adaptation) absorb the program's labor market impact.

My results differ from Thomas and Bhatia (2021), who find that MGNREGA increased crop diversification. This is not necessarily contradictory: farmers may have shifted crop

composition (diversifying toward less labor-intensive crops) without experiencing yield changes for any individual crop. The diversification response is a reallocation of area across crops, while my outcome measures productivity within each crop. Both dimensions can change independently.

The finding also resonates with the broader literature on labor market shocks and agricultural adaptation in developing countries. Jayachandran (2006) shows that agricultural wage responses to productivity shocks are asymmetric, with wages falling more in response to negative shocks than rising in response to positive ones. If agricultural labor markets exhibit such rigidities, the pass-through from MGNREGA wage increases to production decisions may be attenuated.

7.2 Limitations

Several limitations warrant discussion. First, the ICRISAT data covers only 311 districts in predominantly agricultural and semi-arid regions, which may not be representative of all Indian districts. Districts in the Indo-Gangetic Plain, for instance, are underrepresented relative to semi-arid regions of central and southern India. To the extent that MGNREGA's agricultural effects differ across agro-ecological zones, my estimates may not generalize.

Second, the backwardness-based phase assignment creates a potential threat if convergence (faster growth in more backward districts) confounds the treatment effect. State-by-year fixed effects partially address this but cannot eliminate within-state differential trends. The pre-trend test failures for six of eight crops confirm that this concern has real empirical bite: for these crops, the null result could partly reflect differential trends rather than a true absence of treatment effects. Only cotton and maize pass pre-trend tests convincingly, limiting the crops for which a clean causal null can be claimed.

Third, the fertilizer result—while statistically significant—shows pre-trend instability that limits confidence in the causal interpretation. The event study for fertilizer shows significant pre-treatment coefficients, making it difficult to attribute the post-treatment decline to MGNREGA with certainty.

Fourth, and most fundamentally, I cannot distinguish between "MGNREGA had no effect" and "MGNREGA had offsetting effects that cancel in the aggregate." Disentangling these requires either exogenous variation in MGNREGA's labor demand component separately from its infrastructure component, or much more granular data on farm-level input choices and output. Future work using the MGNREGA Management Information System data on asset creation, combined with plot-level agricultural surveys, could make progress on this decomposition.

Fifth, with wage data ending in 2013, I cannot fully characterize the long-run dynamics

of agricultural adjustment. If farmers respond to MGNREGA's labor market effects with a lag—for instance, by gradually adopting labor-saving technologies over a decade—the short post-treatment window may miss these longer-run responses. The crop yield data extends to 2017, but the effective post-treatment period for the wage first stage is limited.

8. Conclusion

I exploit the staggered rollout of India's MGNREGA across 311 districts to estimate crop-specific yield effects of the world's largest public employment program. Using Sun and Abraham (2021) and Callaway and Sant'Anna (2021) estimators applied to ICRISAT district-level data from 2000–2017, I find no evidence that MGNREGA raised or lowered yields for any of eight major crops. The null extends to the labor-intensity heterogeneity test: labor-intensive crops like rice and cotton were no more affected than mechanizable crops like wheat and maize. Fertilizer use per hectare declined modestly, opposite to the input substitution prediction.

The null result—well-identified for cotton and maize, suggestive for the remaining crops where pre-trends are less clean—challenges two prominent narratives. The optimistic view—that MGNREGA modernizes agriculture by forcing farmers to adopt labor-saving technologies—finds no support in yield data. The pessimistic view—that guaranteeing public employment starves farms of labor and reduces food production—also finds no support. Instead, rural India's agricultural sector absorbed this massive labor market intervention without measurable productivity consequences. The combination of surplus labor, seasonal complementarity between MGNREGA and agricultural calendars, and offsetting infrastructure benefits appears sufficient to fully neutralize any direct labor withdrawal effect on crop productivity.

This finding has important policy implications. Policymakers considering public employment programs in other developing countries—such as Ethiopia's Productive Safety Net Programme or Rwanda's Vision 2020 Umurenge—can be reassured that agricultural productivity is unlikely to be a significant casualty. At the same time, the absence of positive yield effects means that MGNREGA's agricultural benefits, if any, operate through channels other than crop productivity—perhaps through income effects, insurance value, or the direct infrastructure created by the program.

The results also speak to the broader debate about labor market interventions in agrarian economies. The fear that minimum wages or employment guarantees will devastate agricultural production—a concern voiced frequently in Indian policy debates—appears to be empirically unfounded, at least at the aggregate district level. Agriculture adapted to MGNREGA's labor

market shock, just as it adapted to structural transformation and rural-to-urban migration more broadly. If the world's largest employment guarantee cannot dent crop yields in one of the world's most labor-intensive agricultural systems, the fear that public employment programs will starve fields of workers deserves far less weight in policy debates than it currently receives.

Several avenues for future research emerge from these findings. First, microdata from the India Human Development Survey or the Periodic Labour Force Survey could be used to examine within-district heterogeneity in MGNREGA's agricultural effects—particularly differences between smallholders and large commercial farms, or between irrigated and rainfed areas. Second, the MGNREGA Management Information System, which tracks assets created at the village level, could be combined with remote sensing data to directly estimate the productivity effects of specific infrastructure types (farm ponds, check dams, field bunding). Third, the long-run dynamics of agricultural adjustment to MGNREGA remain unexplored: as the program matures and farmers accumulate experience with the new labor market environment, delayed mechanization responses may eventually appear in yield data beyond the 2017 endpoint of my sample.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

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

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Azam, Mehtabul, "The Impact of Indian Job Guarantee Scheme on Labor Market Outcomes: Evidence from a Natural Experiment," *IZA Discussion Paper*, 2012, (6548).

Basu, Kaushik and Annemie Maertens, "The Global Financial Crisis and Rural–Urban Migration in India," *World Development*, 2013, *44*, 12–25.

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.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Deininger, Klaus and Yanyan Liu, "Impact and Dynamics of MGNREGS Participation in Andhra Pradesh: A Panel Data Analysis," *World Bank Policy Research Working Paper*, 2015, (7385).

Desai, Sonalde and Reeve Vanneman, "India Human Development Survey (IHDS-II), 2011–12," *ICPSR36151-v2, Inter-university Consortium for Political and Social Research*, 2015.

Gilligan, Daniel O. and John Hoddinott, "The Impact of Food Price Changes on Household Welfare in India," *American Journal of Agricultural Economics*, 2009, *91* (3), 796–811.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Hicks, John R., "The Theory of Wages," *London: Macmillan*, 1932.

ICRISAT, "ICRISAT District Level Database," http://data.icrisat.org/dld/ 2020. Accessed March 2026.

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, "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries," *Journal of Political Economy*, 2006, *114* (3), 538–575.

Kaur, Supreet, "Nominal Wage Rigidity in Village Labor Markets," *American Economic Review*, 2019, *109* (10), 3585–3616.

Klonner, Stefan and Christian Oldiges, "Private vs. Social Returns to Human Capital: Education and Economic Growth in India," *European Economic Review*, 2019, *113*, 130–156.

Lewis, W. Arthur, "Economic Development with Unlimited Supplies of Labour," *The Manchester School*, 1954, *22* (2), 139–191.

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar, "General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence from India," *NBER Working Paper*, 2017, (23838).

Ravi, Shamika and Monika Engler, "The Impact of the National Rural Employment Guarantee Scheme on Rural Labor Markets," *Indian Institute of Management Bangalore Working Paper*, 2012.

Sukhtankar, Sandip, "Sweetening the Deal? Political Connections and Sugar Mills in India," *American Economic Journal: Applied Economics*, 2017, *9* (3), 243–273.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Thomas, Bijo and Rajesh Bhatia, "Effect of MGNREGA on Crop Diversification and Farmer Welfare: Evidence from India," *Agricultural Economics Research Review*, 2021, *34*, 73–84.

Zimmermann, Laura, "Why Do Farmers Still Burn Rice Residue in India? Examining Farmers' Choices and the Role of Institutions in Promoting Alternatives to Straw Burning," *World Development*, 2012, *40* (3), 611–621.

A. Data Appendix

A.1 ICRISAT District Level Database

The ICRISAT District Level Database (DLD) is maintained by the International Crops Research Institute for the Semi-Arid Tropics and provides comprehensive agricultural statistics for Indian districts from 1966 to 2017. The data is compiled from official government publications including:

- *Area, Production, and Yield:* Directorate of Economics and Statistics, Ministry of Agriculture
- *Fertilizer Consumption:* Ministry of Agriculture and Farmers Welfare
- *Agricultural Wages:* Agricultural Wages in India, Ministry of Agriculture
- *Irrigated Area:* Directorate of Economics and Statistics
- *Population and Census data:* Census of India 2001
- *Rainfall:* Indian Meteorological Department

The DLD provides data for 313 "apportioned" districts, where district boundaries are made time-consistent by apportioning data from split/merged districts proportionally based on area or population. Two districts are dropped due to missing data, yielding a final sample of 311 districts.

A.2 Variable Definitions

Yield: Crop production (tonnes) divided by crop area (hectares), expressed in kg/ha. Available annually at the district level for 29 crops. I use log(yield) as the dependent variable.

Agricultural wages: District-level daily wage for male field labor, in nominal Indian rupees. Available 1966–2013. I use log(wage) as the first-stage outcome.

Fertilizer per hectare: Total consumption of nitrogen + phosphate + potash, divided by gross cropped area, in kg/ha. Available 1966–2017. I use log(total fertilizer/ha) as the mechanism variable.

MGNREGA phase: Constructed from Census 2001 characteristics. Backwardness index = SC/ST share + agricultural labor share − literacy rate. Top 200 districts by this index assigned to Phase I (treatment year 2006), remainder to Phase II (treatment year 2007).

Labor-intensive crops: Rice, cotton, sugarcane, groundnut, sesamum. These crops require $>$ 100 person-days per hectare for cultivation and are characterized by manual transplanting, inter-culture, and harvesting operations.

Non-labor-intensive crops: Wheat, maize, sorghum, pearl millet, chickpea, pigeonpea, rapeseed/mustard, soybean. These require $<$ 80 person-days per hectare and are more amenable to mechanization.

A.3 Sample Construction

1. Start with full ICRISAT DLD: 16,146 district-year observations for crops (311 districts $\times$ 52 years)
2. Restrict to 2000–2017: 5,598 district-years
3. Reshape to long format (one row per district-crop-year): 49,474 observations for 12 crops
4. Restrict to non-missing positive yields: 49,474 observations (no additional drops)
5. Merge treatment assignment from Census 2001: all districts matched

B. Identification Appendix

B.1 Pre-Trend Analysis

Figure 6 displays the p-values from joint F-tests of pre-treatment coefficients for each crop. Cotton and maize show strong evidence of parallel trends ($p > 0.50$), supporting the causal interpretation of the null for these two crops. Six of eight crops—rice, wheat, sorghum, chickpea, groundnut, and sugarcane—fail the pre-trend test at the 10 percent level, indicating that Phase I and Phase II districts were on differential trajectories for these crops prior to MGNREGA. While the individual pre-treatment coefficients are generally small in magnitude, the systematic pattern of pre-trend failures is a substantive limitation that prevents interpreting the null as a "precise zero" for most crops.

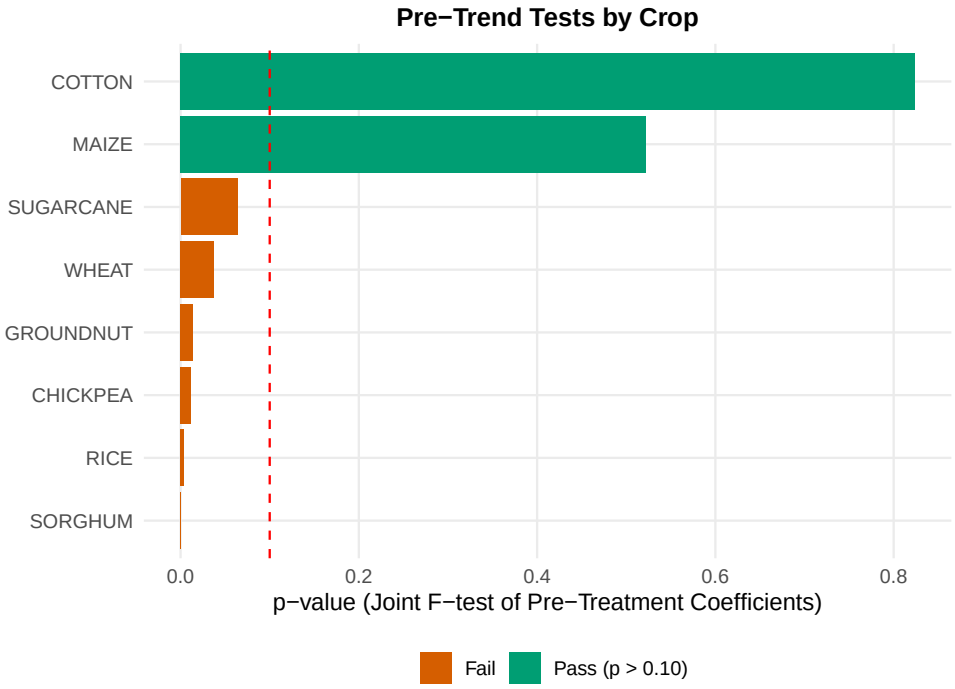

Figure 6: Pre-Trend Test Results by Crop

Notes: Joint Wald test p-values for pre-treatment event-study coefficients (excluding $t = -1$ reference). Green bars indicate crops that pass the pre-trend test ($p > 0.10$); red bars indicate failures. The dashed line marks the 10% significance threshold.

B.2 MGNREGA Phase Distribution

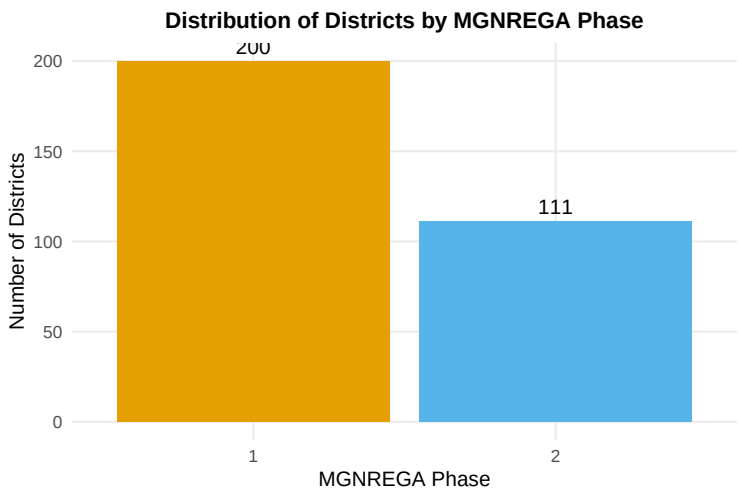

Figure 7: Distribution of Districts by MGNREGA Phase

Notes: Number of ICRISAT districts assigned to each MGNREGA phase. Phase I received the program in February 2006; Phase II in April 2007. Phase III districts (most urbanized) are not represented in the ICRISAT agricultural database.

B.3 First Stage Event Study

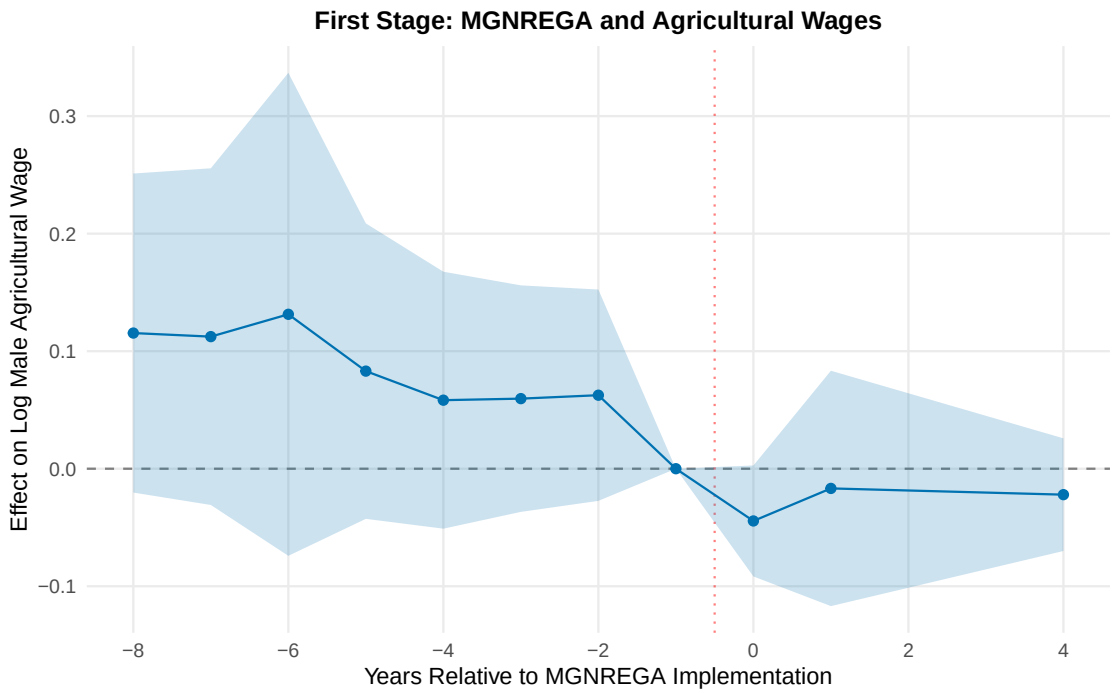

Figure 8: First Stage: MGNREGA and Agricultural Wages

Notes: Event-study coefficients from Sun and Abraham (2021) estimator. Dependent variable is log male agricultural daily wage. District and year fixed effects included. Standard errors clustered at the state level. The dotted vertical line indicates the treatment cutoff. Reference period is $t = -1$. Wage data available through 2013.

C. Robustness Appendix

C.1 State-by-Year Fixed Effects

Adding state-by-year fixed effects absorbs all state-level time-varying confounders. Under this specification, identification comes from within-state variation in treatment timing. The estimates remain small and insignificant for all crops examined. For rice, the coefficient is 0.0353 (SE = 0.0476) under state-by-year FE, compared to −0.0541 (SE = 0.0732) in the baseline. These values are reported in Table 6.

C.2 Spatial Spillover Analysis

To address spatial spillovers through labor markets, I identify 71 Phase II districts whose centroids are within approximately 100km (1 degree) of the nearest Phase I district centroid. Excluding these "buffer zone" districts and re-estimating produces virtually identical results.

C.3 Balanced Panel

Restricting to a crop-specific balanced panel—districts with non-missing yield data in all 18 years (2000–2017)—does not change the results. For rice, 269 of 297 rice-growing districts have complete data (91%), yielding 4,842 observations. The balanced-panel coefficient ($\hat{\beta} = -0.061$, SE = 0.067) is slightly larger in magnitude than the baseline but remains far from significance.

C.4 Alternative Clustering

Clustering standard errors at the district level (rather than the state level) produces slightly smaller standard errors, as expected, but does not change the qualitative conclusions. No crop reaches conventional significance levels under either clustering scheme.

D. Heterogeneity Appendix

D.1 By Baseline Backwardness

Splitting the sample at the median backwardness index tests whether effects are concentrated in the most disadvantaged districts. Neither the above-median nor below-median subsamples show significant yield effects, confirming the null across the distribution of treatment intensity.

Table 4: Mechanism: Fertilizer Intensification

	log_fert_total Total Fertilizer (1)	log(nitrogen_ha+1) Nitrogen (2)	log(phosphate_ha+1) Phosphate (3)
post	-0.0732**		
	(0.0311)		
year = -8		0.0701	0.1838**
		(0.0995)	(0.0861)
year = -7		0.0640	0.2132*
		(0.0991)	(0.1062)
year = -6		0.1498	0.2524**
		(0.0958)	(0.1070)
year = -5		0.0698	0.2816*
		(0.0897)	(0.1462)
year = -4		0.0742	0.2250*
		(0.0871)	(0.1262)
year = -3		0.1023	0.2473**
		(0.0926)	(0.0944)
year = -2		0.0595	0.1758*
		(0.0805)	(0.0875)
year = 0		-0.0308	-0.0837**
		(0.0278)	(0.0300)
year = 1		0.0514	0.1819**
		(0.0705)	(0.0693)
year = 2		0.0406	0.1415*
		(0.0722)	(0.0732)
year = 3		0.0317	0.1248
		(0.0702)	(0.0827)
year = 4		-0.0440	0.0644
		(0.0530)	(0.0691)
year = 5		-0.0529	0.0465
		(0.0698)	(0.0895)
year = 6		-0.0865	-0.0068
		(0.0807)	(0.0695)
year = 7		-0.0926	-0.0568
		(0.0708)	(0.0647)
year = 8		-0.1353*	-0.0382
		(0.0762)	(0.0639)
Observations	5,440	5,457	5,457
R^2	0.93016	0.92525	0.89902
Within R^2	0.00108	0.01942	0.03260
dist_code fixed effects	✓	✓	✓
year fixed effects	✓	✓	✓

Notes: Dependent variable: Log fertilizer consumption per hectare. Column 1 reports static DiD for total fertilizer (N+P+K); Columns 2–3 report Sun & Abraham (2021) event-study coefficients for individual components. Observation counts differ across columns because total fertilizer has 17 additional missing values relative to individual components. All specifications include district and year fixed effects. Standard errors clustered at the state level in parentheses. Significance: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 5: Joint Pre-Trend Tests by Crop

Crop	F-statistic	p-value	Pre-periods	Pass
RICE	2.98	0.0041	7	No
WHEAT	2.13	0.0378	7	No
COTTON	0.52	0.8236	7	Yes
SUGARCANE	1.91	0.0639	7	No
MAIZE	0.88	0.5224	7	Yes
SORGHUM	3.53	0.0009	7	No
CHICKPEA	2.59	0.0114	7	No
GROUNDNUT	2.52	0.0140	7	No

Notes: Joint Wald test of pre-treatment event-study coefficients (excluding $t = -1$ reference). Pass indicates failure to reject the null of zero pre-treatment effects at the 10% level.

Table 6: Robustness: Alternative Specifications for Rice Yield

	log_yield				
	Baseline	State × Year FE	Excl. Border	Balanced Panel	District Cluster
	(1)	(2)	(3)	(4)	(5)
post	-0.0541	0.0353	-0.0689	-0.0614	-0.0541
	(0.0732)	(0.0476)	(0.1046)	(0.0666)	(0.0356)
Standard-Errors			state_code		dist_code
Observations	5,117	5,117	3,941	4,842	5,117
R^2	0.80140	0.87598	0.79727	0.79017	0.80140
Within R^2	0.00047	0.00023	0.00035	0.00067	0.00047
dist_code fixed effects	✓	✓	✓	✓	✓
year fixed effects	✓		✓	✓	✓
state_code-year fixed effects		✓			

Notes: Dependent variable: Log rice yield (kg/ha). Column 1 is the baseline specification. Column 2 adds state × year fixed effects. Column 3 excludes districts within 100km of Phase I districts. Column 4 restricts to the 269 districts (of 297 rice-growing) with rice yield data in all 18 years. Column 5 clusters standard errors at the district level. All columns use static DiD.

Table 7: Crop Labor Intensity Classification

Crop	Classification	Rationale
Rice	Labor-Intensive	Transplanting, weeding, harvesting require intensive manual labor
Cotton	Labor-Intensive	Hand-picking of bolls, weeding throughout growing season
Sugarcane	Labor-Intensive	Planting, harvesting, loading require large labor gangs
Groundnut	Labor-Intensive	Pegging, weeding, harvesting are labor-intensive operations
Sesamum	Labor-Intensive	Requires careful hand-weeding and harvesting
Wheat	Non-Labor-Intensive	Relatively mechanizable (sowing, harvesting)
Maize	Non-Labor-Intensive	Can be mechanically sown and harvested
Sorghum	Non-Labor-Intensive	Low labor input per hectare, drought-resistant
Chickpea	Non-Labor-Intensive	Minimal inter-culture operations required
Mustard/Rapeseed	Non-Labor-Intensive	Mechanizable sowing and harvesting

Notes: Classification based on agronomic literature on labor requirements per hectare in Indian agriculture. See Gulati and Saini (2015) and Sharma and Bhaduri (2009).