

Spatially Heterogeneous Effects of a Public Works Program

The latest version can be found [here](#).

Joshua D Merfeld*

August 5, 2017

Abstract

Most research on labor market effects of the Mahatma Gandhi National Rural Employment Guarantee Scheme focuses on outcomes at the district level. This paper shows that such a focus masks substantial spatial heterogeneity: treated villages located near untreated areas see smaller increases in casual wages than treated villages located farther from untreated areas. I argue that worker mobility, rather than spatial differences in implementation or program leakages, drives this spatial heterogeneity. I also present evidence that the effects of the program on private-sector employment display similar intra-district heterogeneity. Finally, by exploiting the difference in wage changes over space, I show that a large portion of consumption increases are driven by wage increases, not program employment. Overall, these results suggest that a district-level focus underestimates the true effect of the program on wages and also support the argument that increasing rural wages is an effective poverty-fighting tool in developing countries.

Keywords: India, Public Works, Labor, Wages, Spillovers

JEL Codes: D50, H53, I38, J38, J46

*Evans School of Public Policy and Governance, University of Washington; merfeld@uw.edu. Conversations with and comments from Leigh Anderson, Justin Downs, James Long, Mark Long, Karthik Muralidharan, and Katya Roshchina were extremely helpful. I also thank participants at various seminars and conferences for helpful comments and suggestions. Any remaining errors are of course my own.

1 Introduction

India's Mahatma Gandhi National Rural Employment Guarantee Scheme (NREGS)¹, which was launched in 2006 and is now rolled out throughout the country, is the largest public works program in the world. In fiscal year 2015/16, the program generated more than two billion persondays of labor and amounted to more than two percent of the federal government budget.² Given the program's size, it is perhaps unsurprising that it has spawned a vast body of literature.

The program was phased in over three rounds, starting in 2006. However, this rollout was not randomized; the poorest districts were generally the first to receive the program. As such, while the phased rollout makes differences-in-differences attractive, the systematic differences in district characteristics make the identification assumption of parallel trends somewhat tenuous (Zimmermann, 2012; Sukhtankar, 2016). In addition, it is not clear what treatment effect estimates averaged over the entire country measure, as implementation of the program varies substantially from state to state and even from district to district (Liu and Barrett, 2012; Ravallion et al., 2013; Dutta et al., 2014; Banerjee et al., 2015; Imbert and Papp, 2015). In reality, these estimates are the net effect of the program, inclusive of implementation efficacy (Sukhtankar, 2016).

Despite these difficulties, a growing body of literature has analyzed the performance and far-reaching effects of the program. This includes education (Afridi et al., 2012; Li and Sekhri, 2013; Shah and Steinberg, 2015), nutrition and consumption/income (Liu and Deininger, 2010; Jha et al., 2011; Ravi and Engler, 2015), and corruption (Imbert and Papp, 2011; Niehaus and Sukhtankar, 2013; Banerjee et al., 2015; Muralidharan et al., 2016a).

¹The program was originally passed as the National Rural Employment Guarantee Act in 2005 and implemented as the scheme starting in 2006. The program was renamed in 2009, but I use the old acronym throughout this paper.

²Figures are from <http://nrega.nic.in/>, the official program website.

However, given the aims and scope of the program, much of the literature has focused on its labor market effects. While findings vary, there is broad agreement with one basic fact: the program led to a modest increase in prevailing casual wages in rural India (Azam, 2012; Berg et al., 2014; Imbert and Papp, 2015; Muralidharan et al., 2016b; Sukhtankar, 2016).³

In addition to a common finding, much of the literature examining the effects of the program's rollout on labor market outcomes shares another characteristic: it analyzes the effects of the program at the district level (Azam, 2012; Zimmermann, 2012; Berg et al., 2014; Imbert and Papp, 2015). This level of analysis is attractive for two reasons. First, the program was rolled out at the district level. Second, given low levels of rural-to-rural migration in India (Rosenzweig, 1988; Behrman, 1999; Anant et al., 2006; Munshi and Rosenzweig, 2009), districts are often considered to be relatively distinct labor markets. Empirical work in India has long focused on districts, especially for labor-market outcomes; for non-NREGS examples, see, for example, Rosenzweig (1978, 1984), Jayachandran (2006), Topalova (2010), Kaur (2014), and Shah and Steinberg (2017).

In this paper, I show that a district-level focus obscures substantial heterogeneity, with important policy implications. I use the Additional Rural Incomes Survey/Rural Economic Demographic Survey (ARIS/REDS), collected by the National Council of Applied Economics Research, to show that wage changes due to the program are spatially heterogeneous. The first wave of the survey was collected prior to the implementation of NREGS and the second wave was collected between the second and third phases on the rollout. Importantly, unlike other datasets, I am able to identify the location of each village in the survey using GPS data. For each district, I construct two geospatial variables: whether phase one or phase two districts ("treated" districts) share a border with phase three dis-

³In this context, "casual wages" refer to wages in daily markets for labor. The duration of this type of contract is generally only a day and the contract carries no expectation or promise of future work.

tricts (“untreated” districts), and vice versa; and how long that border is. For districts that border a district of the opposite treatment status (“border districts”), I construct a variable for surveyed villages equal to their distance to the nearest district of the opposite treatment status.

With these variables, I first document differences in households within districts. In line with previous NREGS studies, phase one and two districts had much lower wage rates prior to implementation of the program than phase three districts (Zimmermann, 2012; Imbert and Papp, 2015). However, this difference attenuates within 9 kilometers of the border between treated and untreated districts; in other words, districts on either side of the border are more similar to one another than the average treated and untreated village. I then estimate the effects of NREGS by slowly restricting estimation to villages closer and closer to the border. While there are large effects on casual wages when all households are included in the estimation, the estimated effect starts to attenuate around 15 kilometers from the border between a treated and untreated district and the estimated effect completely disappears when looking only at households located within six kilometers of the border. In addition, this attenuation begins earlier and is stronger for male wages than for female wages, consistent with a model in which women face higher travel costs—which may not be strictly monetary—than men.

I then further explore whether this heterogeneity is due to failure of the identification assumption of parallel pre-program trends, differences in NREGS implementation, or spillovers. While data issues prevent me from explicitly testing the identification assumption using the ARIS/REDS data, I present graphical evidence that pre-program trends, from 1982 to 1999, are unlikely to be responsible for the results. In addition, an analysis of aggregated district-level data using the National Sample Survey shows that treated districts that share a larger portion of their border with untreated districts see a smaller increase in the wage

than treated districts that share a smaller portion of their border with treated districts, a finding which supports results using the ARIS/REDS. Moreover, while this effect is seen using the pre- and post-program data, there is no such effect when using two pre-program waves, which supports the assumption of parallel trends in treated and untreated districts, at least with regards to spatial heterogeneity.

I then offer several other pieces of evidence to explore the source of this heterogeneity. First, NREGS implementation apparently does not differ by distance to border; a number of NREGS variables from 2012-2013—including number of days of employment created and total labor expenditures—are uncorrelated with the distance variable. Second, I show that this pattern is only seen in border districts. If districts in different phases were on different trends prior to the program or if households located in different locations within a district were trending differentially—for example due to differences in implementation due to geographic location—we should see an attenuation in the effect for non-border districts, as well. However, there is no attenuation in the estimated impact of the program in these districts. Then, I show that program leakages do not appear to explain the attenuation.⁴ Finally, I show that untreated households living close to the border commute longer distances to casual non-farm employment. I interpret this as evidence of a type of “wage arbitrage”—workers moving to higher-wage opportunities—with the eventual result being a smaller increase of the casual wage in border villages of treated districts. The spillover effect apparently operates in the direction of treated districts; laborers from untreated districts travel into treated districts in response to the program.

I find insignificant—though somewhat large and imprecisely estimated—overall effects of the program on private-sector employment. The finding is consistent with other studies that have found evidence NREGS crowds out some types of private employment (Zimmer-

⁴In this context, I define the possibility that untreated villages had access to NREGS as leakages.

mann, 2012; Imbert and Papp, 2015; Merfeld, 2017).⁵ However, this finding again masks substantial heterogeneity: the effects on private employment appear to mirror the effects on wages, with interior areas seeing decreases in private employment relative to border areas.

Finally, I examine the effects of the program on household per capita income. Interestingly, the effects again appear to be highly correlated with the wage increase: estimated increases in household income towards the interior of treated districts almost completely disappear at the border. While most previous research on the effects of NREGS on consumption or income (Liu and Deininger, 2010; Jha et al., 2011; Ravi and Engler, 2015) necessarily calculates the total effects of the program—that is, the net effect of both access to employment and an increase in the prevailing wage rate—I am able to disentangle these effects. Using the fact that wages are unchanged at the border, I assume that any effects on income at the border are due to access to program employment—not a wage increase—while effects in the interior of treated districts are due to both. While imprecision prevents me from calculating the exact percent of the consumption increase that is due to the wage increase, results suggest it is substantial. This estimate supports recent experimental estimates that the majority of program effects operate through wage increases (Muralidharan et al., 2016b).

This paper contributes to three separate bodies of literature. First, I provide further evidence of the broad labor-market impacts of NREGS (Azam, 2012; Zimmermann, 2012; Berg et al., 2014; Imbert and Papp, 2015; Muralidharan et al., 2016b) as well as public works programs, more generally (see, e.g., Subbarao et al. (2003) and Beegle et al. (2017)). However, unlike most previous studies, I am able to explore the spatially heterogeneous impacts of the rollout of the program. Consistent with the literature, I find overall increases in the wage rate. However, I also find that the estimated increase disappears at

⁵While this result is expected in competitive markets, the theoretical result is ambiguous under non-competitive markets (Basu et al., 2009).

the border with untreated districts. In addition, increases in the prevailing wage apparently explain a substantial portion of overall effects on household income. This finding suggests that increasing the wage rate in rural areas may be an effective tool in the fight against poverty. Finally, the importance of wage increases may help explain the results of recent studies that find few effects of public works programs (Beegle et al., 2017), since many of these programs may be too small to have appreciable effects on prevailing wages.

In addition, some previous research finds that NREGS had a larger impact on female wages than male wages (Azam, 2012). This study presents evidence that men are more mobile than women and that the effect on male wages attenuates sooner than the effect on female wages. It true, district-level difference-in-differences estimates of the effect on male wages may be underestimating the true impact of the program more than the estimates of the effect on female wages, explaining part of the gender effects.

Second, this study sheds some light on the functioning of labor markets in India (Jayachandran, 2006; Kaur, 2014). While spot markets for casual labor in India are generally assumed to work relatively well (Rosenzweig, 1980), it is unlikely that markets are complete (Rosenzweig, 1988; Behrman, 1999). These results suggest that labor is somewhat mobile, albeit within a relatively small radius of around 15 kilometers. In addition, while the minimum wage is unlikely to bind or be enforced in many developing country contexts (Behrman, 1999), NREGS appears to operate as a de facto minimum wage, increasing the bargaining power of casual laborers (Basu et al., 2009).

Finally, I contribute to the literature on treatment effects and spillovers of public policy interventions in developing countries (Duflo, 2000; Miguel and Kremer, 2004; Angelucci and De Giorgi, 2009; Cunha et al., 2011). In particular, this study reinforces the importance of capturing general equilibrium effects when evaluating large public programs and that failure to do so may result in biased estimates of the program's impact (Muralidharan

et al., 2016b). In addition, the focus on the district-level effects of NREGS theoretically underestimates the true effect of the program in two ways: the wage in untreated districts increases more in areas closer to treated districts than in areas farther from treated districts, while the wage in treated districts increases less in areas closer to untreated districts than in areas farther from untreated districts. The treatment effect we are interested in is the effect of the program on wages absent spillovers since the program will be expanded throughout the country. Yet, a district-level focus overestimates the counterfactual wage in untreated districts and underestimates the wage in treated districts, although the latter effect apparently dominates in the present study. These two facts lead to an underestimate of the true wage effect of the program.⁶ With this in mind and insofar as increasing rural wages is an explicit goal of policymakers, NREGS may have been more effective than previously estimated.

The closest paper to this one is Muralidharan et al. (2016b), who study the randomized rollout of an improvement in the implementation of NREGS. They find similar spatial effects and at similar distances. Their paper is different in two important respects. First, their randomized intervention provides a cleaner identification strategy and cleaner estimates of the intervention's impacts. However, they randomize improvements only in Andhra Pradesh. Since Andhra Pradesh consistently performs better than other states in implementation of NREGS (Banerjee et al., 2015; Imbert and Papp, 2015), it is not clear whether their findings would translate to other Indian states. Second, they do not study the implementation of the program itself. Rather, they estimate the effects of a technological reform that improved implementation of the program. This study, on the other hand, examines the effects of the program when it is first rolled out. Given these differences, one might expect

⁶While I refer to the effect “disappearing” at the border, it is important to keep in mind that the program has very real effects at the border. In particular, the wage in both treated *and* untreated districts increases. As such, it is not so much that the effect disappears; rather, the identification strategy results in an underestimate of the actual effect of the program.

findings to differ. However, the results presented in this paper also suffer from one major limitation: they are imprecisely estimated and much larger than previous estimates of the program's impacts. Nonetheless, the pattern is clear and the overall findings from both studies are largely similar, suggesting that the effects found in both studies are driven by mechanisms common to all of India and not specific to Andhra Pradesh. I argue that labor mobility is one such mechanism.

The rest of the paper is organized as follows. In the next section, I describe the operation of NREGS. Section 3 develops a simple model to explain possible heterogeneity in the effect of NREGS on wages. I discuss the methods and data in Section 4 before moving to results in Section 5. Section 6 concludes.

2 The Program

The Mahatma Gandhi National Rural Employment Guarantee Act was passed in 2005 after “more than a decade of sustained high growth in GDP... was perceived not to have made a sufficient dent in poverty in the rural India” (Azam, 2012, p. 1). The program guarantees up to 100 days of employment to any rural household that requests it. Employment is varied, but is generally focused on improving productivity in rural areas by building infrastructure, including irrigation, roads, and land improvements (Liu and Deininger, 2010). Unlike previous workfare programs, NREGS is explicitly designed as an employment *guarantee*; a household has the right to demand—and receive—employment. However, in practice, supply-side constraints are often binding, resulting in substantial unmet demand (Ravallion et al., 2013; Dutta et al., 2014; Mukhopadhyay et al., 2015).

In order to demand employment, households must first apply for a job card. The household applies for the card at the local Gram Panchayat (GP)⁷ and the card includes a picture

⁷The Gram Panchayat is the lowest level of administration in the Indian federal structure and generally con-

of every household member eligible for the program. With this card in hand, household members can request employment. According to the design of the program, individuals request work from their local GP, and the legislation stipulates that work be assigned within five kilometers of the individual's place of residence, otherwise a travel stipend is to be provided (Azam, 2012). If individuals do not receive employment within 15 days of requesting it, they are eligible for unemployment allowance. However, underreporting of unemployment is common and unemployment compensation is rare (Sharma, 2009).

The program's design begets a relatively complicated funding structure. The federal government is responsible for all unskilled wages and 75% of skilled wages and materials, while states are responsible for the remaining 25% (Azam, 2012). As such, funding flows from the federal government to the state before being dispersed through lower administrative levels and eventually to the GP's account, while funding requests from GPs to the state also flow through two separate administrative levels (Banerjee et al., 2015).⁸

Several specific features of the program are worth noting. First, the minimum wages set by the program⁹ are self-targeting and are relatively low (and the labor relatively difficult) so that only needy households will apply for employment. Nonetheless, program wages vary from state to state, and some wages were set above prevailing wage rates at the start of the program. In some cases, program wages were more than double the prevailing wage rate at the time, especially for women.¹⁰ Second, the program guarantees equal wages to both men and women and also mandates that a certain number of beneficiaries be women.¹¹ In

sists of several villages.

⁸GPs request funding from the block, which is one administrative level above the GP. Subsequently, the block forwards requests to the district, which then requests funding from the state pool. Banerjee et al. (2015) provide an excellent overview of this process.

⁹The original legislation allows the states to define the program wage, as long as it is above the legislated minimum.

¹⁰http://www.levyinstitute.org/pubs/EFFE/Mehrotra_Rio_May9_08.pdf

¹¹The legislation states that "priority shall be given to women in such a way that at least one-third of the beneficiaries shall be women who have registered and requested for work under this Act" (Ministry of Law and Justice, 2005, p. 14).

addition, childcare at worksites is required so that mothers with young children can participate in the program. However, many worksites do not adhere to this childcare requirement (Khawlneikim and Mital, 2016; Khera and Nayak, 2009; Reddy and Upendranadh, 2010; Sharma, 2009).

Given the demand-driven nature of the program, implementation is relatively decentralized, with employment and worksites started and administered by the GP. While the GP is the basis for employment generation, the program was rolled out at the district level, in three phases. Districts in the first phase received the program in 2006, districts in the second phase received the program in 2007, and districts in the third phase received the program in 2008. Given the phased rollout, the majority of research on the labor-market effects of NREGS utilize differences-in-differences. However, the rollout was not randomized. Rather, the Indian government assigned districts to phases based on a government-constructed rating, which resulted in the poorest districts generally receiving the program first (Zimmermann, 2012). While the rating used to assign districts to phases is not publicly available, underlying data used in the construction of this rating are available in the 2003 report from the Planning Commission (Planning Commission, 2003; cited in Zimmerman, 2012). These ratings show that earlier phase districts tended to be poorer and have more scheduled castes/tribes, lower agricultural productivity, and lower wages. As such, there are concerns regarding the identification assumption of differences-in-differences, that treated districts and untreated districts had identical labor-market trends prior to implementation of the program.

3 Model

Before discussing the data and methods, I present a simple model of the program's effects. Much of the model below is borrowed from Imbert and Papp (2015), but with the addition

of household types and travel costs. There are two types of households, denoted by superscripts $j = 1, 0$, where 1 denotes a household with access to NREGS (“treated” household) and 0 denotes a household without access to NREGS (“untreated” household). For treated and untreated households, the total labor endowment, \bar{T} , consists of leisure and total labor supplied. For treated households, total labor supplied includes labor supplied to the market, $L_i^{1,T}$; labor supplied to household production, $L_i^{1,f}$; NREGS labor, $\bar{L}^{1,N}$; and leisure, l_i^1 . Thus, the total labor endowment is given by $T = L_i^{1,T} + L_i^{1,f} + \bar{L}^{1,N} + l_i^1$, and labor allocation is allowed to vary by households i .

Treated and untreated households are assumed to reside in separate locations, and, as such, can face separate equilibrium wages, denoted w^j . However, in this model, I assume wages are the same in the pre-program period¹² and both treated and untreated households in the pre-program period resemble untreated households in the post-program period. Treated households have access to rationed NREGS labor, $\bar{L}^{1,N}$, and treated households allocate labor up to the rationed level. I model NREGS labor this way since actual labor is supplied well below the program maximum of 100 days due to rationing, which is common (Ravallion et al., 2013; Dutta et al., 2014; Mukhopadhyay et al., 2015). Treated households also engage in home production, with total labor demanded including both household labor, $L_i^{1,f}$ and hired labor, $L_i^{1,h}$. Households are assumed to maximize a utility function with respect to both consumption, x_i^1 , and leisure. The problem for treated households is thus

$$\max_{x_i^1, L_i^{1,f}, L_i^{1,T}, L_i^{1,h}} u(x_i^1, l_i^1) \quad (1)$$

¹²This is an assumption which greatly simplifies the model. Since I empirically estimate effects using differences-in-differences estimates, allowing the wage to differ in the pre-program period has no effect on the hypothesized *change* in wages.

subject to

$$\begin{aligned} x_i^1 &\leq A_i f(L_i^{1,f} + L_i^{1,h}) - w^T L_i^{1,h} + w^T L_i^{1,T} + w^1 \bar{L}^{1,N} \\ T &= L_i^{1,T} + L_i^{1,f} + \bar{L}^{1,N} + l_i^1, \end{aligned}$$

where w^1 is the NREGS wage rate, exogenously fixed by the state government. Finally, $f(\cdot)$ is the household production function, which satisfies the usual assumptions— $f'(\cdot) > 0$ and $f''(\cdot) < 0$ —household and hired labor are assumed to be perfectly substitutable, and A_i is a productivity factor, which differs by household. Without the addition of untreated households and as shown by Imbert and Papp (2015), households with high A_i will be net demanders of labor and households with low A_i will be net suppliers of labor. In addition, an increase of NREGS employment, $L^{1,N}$, increases the prevailing wage rate, w^T , and decreases private-sector employment.

Let us now consider the addition of untreated households, which do not have access to NREGS labor. However, untreated households are able to travel to the treated area to engage in casual work. Travel is costly, with costs $c_i(d_i)$, where d_i is distance. In addition, costs are increasing in distance: $c'_i(\cdot) > 0$. Without loss of generality, assume all untreated households reside in the same location and face the same travel costs, so that $c_i(d_i) = c(d)$. Thus, the total labor endowment for untreated households includes labor supplied to the market in untreated areas, $L_i^{0,U}$; labor supplied to the market in treated areas, $L_i^{0,T}$; labor supplied to household production, $L_i^{0,f}$; and leisure, l_i^0 . The problem for untreated households is

$$\max_{x_i^0, L_i^{0,f}, L_i^{0,U}, L_i^{0,T}, L_i^{0,h}} u(x_i^0, l_i^0) \quad (2)$$

subject to

$$x_i^0 = A_i f(L_i^{0,f} + L_i^{0,h}) - w^0 L_i^{0,h} + w^0 L_i^{0,U} + (w^1 - c(d)) L_i^{0,T}$$

$$T = L_i^{0,T} + L_i^{0,U} + L_i^{0,f} + l_i^0.$$

Note that the wage an untreated household receives for casual labor in the treated area is given by the net of actual wage minus travel costs, or $w^T - c(d)$.

To solve this problem, it is sufficient to examine the first-order conditions of untreated households. In the model above, the wage will be strictly higher in treated districts with no travel. As such, even if I allowed for the allocation of labor by treated households in untreated districts, no households would allocate positive amounts of labor. As such, any equilibrium will be driven only by movements of untreated households.¹³

Intuitively, untreated households will allocate labor to casual labor in treated areas such that the wage rate in treated areas is higher than the wage rate in untreated areas by exactly the travel cost, $c(d)$. In other words, with no labor market frictions other than travel costs, labor is allocated to the point that $w^1 = w^0 + c(d)$. To see this, first note that untreated households will equate their marginal utility of working in untreated and treated areas. For example, suppose the equality does not hold, and $w^1 > w^0 + c(d)$. Then, untreated households will reallocate labor from the untreated area to the treated area in order to equate marginal utilities across labor types. This reallocation increases labor supply in treated areas, decreasing the wage, and decreases labor supply in untreated areas, increasing the wage. This continues until $w^1 = w^0 + c(d)$. On the other hand, suppose that $w^1 < w^0 + c(d)$ and that untreated households are allocating at least some labor to casual work in the treated area. Then, untreated households can reallocate labor from the treated area to the untreated

¹³This is another simplifying assumption. Given that the empirical strategy relies on an examination of changes over time, this assumption does not affect the direction of the change; the hypothesized direction of the *change* is unaffected even if we assume there is travel in the pre-program equilibrium.

area. This reallocation will cease only when $w^1 = w^0 + c(d)$.

From this equality follow four predictions. First, $\frac{\partial w^1}{\partial d} > 0$. In other words, the closer the treated area is to an untreated area, the lower the wage in the treated area will be. Similarly, the farther the treated area is from an untreated area, the higher the wage will be. Second, $\frac{\partial w^0}{\partial d} < 0$: the wage in untreated areas will be higher the closer the untreated area is to the treated area, as more untreated labor flows out of the untreated area, raising wages. From these facts, it follows that treated villages close to the border with untreated districts will have lower wage increases than treated villages farther from the border. Similarly, the wage will increase relatively more in untreated areas closer to the border with treated districts than in untreated areas farther from the border. Third, note that this equality no longer holds in treated areas far enough from untreated areas. In particular, if $w^1 < w^0 + c(d)$ when $\int_i L_i^{0,T} di = 0$ —that is, when no untreated households are allocating labor to work in treated households—then no travel takes place and w^1 is permanently higher than w^0 . This implies that $\frac{\partial w^1}{\partial d} = 0$ for d sufficiently large. Finally, it is easy to see that decreasing (increasing) travel costs will close (widen) the wage gap between treated and untreated areas. While the model assumes travel costs are monetary, this does not need to be the case. For example, women are assumed to be less mobile than men in India (Khera and Nayak, 2009). One possible reason is that women are responsible for childcare. As such, travel may impose non-monetary costs on women, leading to an increase in the wage gap between treated and untreated areas for women relative to men, holding distance constant.

4 Data and Methods

To analyze the effects of NREGS, I mainly use the Additional Rural Incomes Survey/Rural Economic Demographic Survey (ARIS/REDS), collected by the National Council of Applied Economic Research. The ARIS/REDS is a panel survey that has been conducted peri-

odically since 1969.¹⁴ When first collected, the data was nationally representative, although it oversampled high-income households and areas suitable to green-revolution crops. However, with recent changes in demographics and administration, the data are no longer representative of the entire country (Foster and Rosenzweig, 2010). As such, I interpret the results as representative of overall spatial patterns—including ratios of magnitudes—but not as representative of the actual magnitude of effects. I use the fifth and sixth rounds of the survey—collected in 1999 and 2008, respectively—for my main analyses. I also use the 1982 round to explore pre-NREGS trends in the spatial heterogeneity of wages.

Since the second wave of data I use (2008) was collected after implementation of the first two NREGS phases but before implementation of the third phase, I refer to phase one and phase two districts as “treated” districts and phase three districts as “untreated” districts. Figure 1 provides an overview of the phased rollout of NREGS along with the location of ARIS/REDS villages. There is a lot of geographic variation in the location of surveyed villages, with all of the major Indian states covered except Jammu and Kashmir. Since the main strategy of this paper is to compare villages near borders of treated and untreated districts, it is instructive to look at the geographic variation in this specific characteristic. In particular, there appears to be a relatively dense concentration of villages near treated and untreated districts and the southern and central regions. On the other hand, treated villages far from untreated districts tend to be located in the eastern half of the country, while untreated villages far from treated districts tend to be in the northwestern region of the country.

¹⁴In all analyses in this paper, I do not use the panel nature of the data. Rather, I treat both waves as repeated cross-sections.

4.1 Methods

I estimate regressions of the form

$$y_{idt} = \alpha_0 + \alpha_1 NREGS_d + \alpha_2 Post_t + \alpha_3 NREGS_d \times Post_t + \delta X_{idt} + \phi H_{idt} + \gamma D_d + \beta Z_{idt} + \varepsilon_{idt}, \quad (3)$$

where y_{idt} is the outcome of interest for individual (or household)¹⁵ i in district d in wave t ; $NREGS_d$ is an indicator variable equal to one if the district is a treated (phase-one or phase-two) district; $Post_t$ is an indicator variable for observations in the second wave, 2008; X_{idt} is a vector of time-variant individual controls; H_{idt} is a vector of time-variant household controls; D_d is a vector of district controls; Z_{idt} is rainfall in each year, measured in standard deviations from the mean (z-score), calculated separately for each household; and ε_{idt} is a conditional mean-zero error term. Much of the previous literature has employed district-level fixed effects in estimation. Unfortunately, due to sample sizes—especially after restricting estimation to those living only a certain distance from borders—using district fixed effects results in overly sensitive estimates and even more imprecisely estimated effects than without their inclusion. As such, I am unable to implement the above strategy with district fixed effects. This may explain part of the difference in magnitudes between the present results and previous literature.

An alternative strategy would be to estimate the effects of the program using different windows of distances from the border. For example, comparing households between 40 and 50 kilometers of the border, then households between 35 and 45 kilometers of the border, etc., all the way down to within six kilometers of the border. However, this strategy suffers from the same limitation as district fixed effects: the sample size is not large enough

¹⁵The main analyses around wages and private-sector employment use the individual as the unit of analysis but the regressions focusing on household consumption use the household as the relevant unit of analysis.

and the resulting estimates are imprecise and sensitive.

The coefficient of interest is α_3 , which gives the difference-in-differences estimate of the effect of NREGS on outcomes. This coefficient captures the difference, across treatment status, in changes of each outcome at the district level. This implies a very specific assumption required for unbiased estimates of the effect of NREGS: the trends, across time, of outcomes in treated and untreated districts are identical in the absence of the program. In other words, if treated districts had not received the program, then treated and untreated districts would have had the same change, from 1999 to 2008, in each of the outcomes. I discuss possible failure of this assumption of difference-in-differences in the results section. Finally, given that outcomes are likely serially correlated over time, I cluster standard errors at the district level to allow for arbitrary correlation across both waves of the survey (Bertrand et al., 2004).

The second prediction from the model in Section 3 is that wages in untreated areas near the border with treated areas will increase more than wages in untreated areas farther from border. Unfortunately, there is no clear comparison group for estimating differences-in-differences. As such, I also estimate an alternative specification. I drop the interaction between post and NREGS and include instead an interaction between post and distance to border as well as between post and an indicator variable for whether the observation is more than six kilometers to the border. I estimate this separately for NREGS and non-NREGS districts. In essence, this specification is comparing the change in wages within NREGS districts and within non-NREGS districts (last two columns) separately. This specification is

$$\begin{aligned}
y_{idt} = & \alpha_0 + \alpha_1 dist_i + \alpha_2 Post_t + \alpha_3 dist_i \times Post_t \\
& + \delta X_{idt} + \phi H_{idt} + \gamma D_d + \beta Z_{idt} + \varepsilon_{idt},
\end{aligned} \tag{4}$$

where $dist_i$ is the distance variable, either distance to border or an indicator variable equal to one if the observation resides more than six kilometers from the border, and the other variables are defined as in (3).

Unfortunately, I am not able to empirically estimate pre-program trends with the ARIS/REDS data. In order to explore pre-program trends, I turn instead to data from the National Sample Survey (NSS), which is a nationally representative survey conducted by the Ministry of Statistics and Programme Implementation.¹⁶ Since these data are only aggregated to the district level, I am unable to measure distances to border for each village. Instead, I construct a variable defined as the percentage of each district's border shared with a district of the opposite treatment status. In other words, for treated districts, I measure the total border shared with untreated districts, while for untreated districts, I measure the total border shared with treated districts. I then implement a triple-difference specification:

$$\begin{aligned}
y_{idt} = & \alpha_0 + \alpha_1 NREGS_i + \alpha_2 Post_t + \alpha_3 Percent_d + \alpha_4 NREGS_i \times Post_t \\
& + \alpha_5 NREGS_i \times Percent_d + \alpha_6 Post_t \times Percent_d \\
& + \alpha_7 Post_t \times NREGS_i \times Percent_d + Controls + \varepsilon_{idt},
\end{aligned} \tag{5}$$

where $Percent_d$ is the border percent variable and $Controls$ are all of the district, household-, and individual-level controls used above. The coefficient of interest is now α_7 , which essentially represents whether the effect of NREGS ($Post_t \times NREGS_d$) differs

¹⁶See <http://mospi.nic.in/national-sample-survey-office-nsso>.

by percentage of border shared with a district of the opposite treatment status. I estimate this specification using 2004/05 and 2007/08 data (with the latter being the post wave), which analyzes the effect of NREGS at the same point in time as the ARIS/REDS data. I also estimate this specification using 1999/2000 and 2004/05 data (with the latter being the post wave) in order to determine whether pre-program trends appear to be correlated with one type of geographic variable (percent of border).

In another set of results, I examine the effect of distance to border on NREGS implementation. I use the Indian government's NREGS website¹⁷ to acquire data related to quality of NREGS implementation. These variables include number of rolls filled, number of individuals worked, number of person-days created, number of households that reached the limit of 100 days, total labor expenditure, percent labor expenditure (of total labor and materials expenditure), and total number of works. I include the value of these variables from 2013. In these regressions I include only NREGS districts in the ARIS/REDS. I then estimate regressions of the form

$$y_d = \alpha_0 + \alpha_1 dist_d + \phi D_d + \varepsilon_d, \quad (6)$$

where D_d is a vector of district-level controls and ε_d is a conditional mean-zero error term. In these regressions, α_1 is the coefficient of interest and represents differences in the implementation of NREGS by distance to the border.

I construct several geovariables for use in this analysis. For each district in the country, I first assign it a value of treated (if in the first or second phase) or untreated (if in the third phase). I then construct a variable, which I call *border district*, that equals one if a treated district borders an untreated district or if an untreated district borders a treated district and zero otherwise. In other words, *border district* identifies borders across which there is a

¹⁷<http://www.nrega.nic.in/netnrega/home.aspx>

difference in treatment status in 2008. I then calculate the total distance of such borders for each district in the country.

In order to estimate treatment heterogeneity, I estimate a series of differences-in-differences by restricting estimation based on distance to the nearest border. I begin by estimating regressions using all individuals. Then, I restrict estimation to only individuals living within a certain distance of the border (e.g. 30, 15, 9, etc.).

I define this distance variable as follow: For all villages in treated border districts, I construct another variable equal to the distance from each village to the closest border with an untreated district. Similarly, for villages in untreated border districts, I construct this variable as the distance from each village to the closest border with a treated district. To analyze spatial heterogeneity in effects of the program, I will essentially be comparing differences-in-differences estimates of all villages compared to villages close to borders. Figure 2 presents a histogram of this variable across treated (NREGS) and untreated (non-NREGS) districts. Mass is highest between zero and ten kilometers to the border, although there are still thousands of observations outside of 30 kilometers. The clustering of observations at certain distances suggests that there could be large variations in estimated effect sizes as the distance is restricted. For this reason, the closest distance to the border I use is six kilometers, as a relatively higher percentage of observations is dropped as estimation is restricted to distances closer to the border.

In one set of estimates below, I use non-border districts. These are defined as treated districts that do not border an untreated district and untreated districts that do not border a treated district. For villages in these districts, I calculate another distance variable, equal to the distance from each village to the nearest border shared with a district of the same treatment status. In other words, I calculate the distance from treated villages to the nearest border with a treated district and the distance from untreated villages to the nearest border

with an untreated district.

Figure 3 presents a graphical example of the creation of these two distance variables using actual districts from the ARIS/REDS. The triangles represent surveyed villages while the dark-gray districts—Chikballapura and Kolar—are untreated districts and the light-gray district—Chittoor—is a treated district. The blue lines represent the construction of the main distance variable used in the analyses below: distance to the closest district of the opposite treatment status. On the other hand, the red lines show construction of a distance variable that represents the distance from each surveyed village to the nearest district of the same treatment status. Throughout this paper, “distance to border” refers to the blue line distances unless otherwise noted.

In individual-level regressions, I include age, age squared, education, and gender. In both individual- and household-level regressions, I also include household characteristics, including its size and demographic make-up. I also include three village-level characteristics: (log of) population density, (log of) distance to the nearest town, and (log of) distance to the district headquarters. Additionally, I include a number of district-level variables. Using the 2001 census,¹⁸ I include a number of district-level characteristics: population (log), percent rural population, percent scheduled castes, percent scheduled tribes, literacy rate, labor force participation rate, the makeup of the rural labor force, and district area (log of square kilometers), the last of which was calculated using GIS software. Since the differences-in-differences estimator is measuring change and these variables could arguably affect the rating used to assign districts to phases—and thus a district’s treatment status—I allow the effect of each of these variables to affect the trend across waves, by including these variables as an interaction between each variable and $Post_{dt}$ (Imbert and Papp, 2015). Finally, I compute the district-level wage using the 1999 ARIS/REDS, which

¹⁸These variables technically come from after the 1999 wave. However, I prefer to use the 2001 census over the previous census, conducted ten years prior.

I again only include as an interaction with $Post_{dt}$.

In all regressions, I also include village population (log), distance to the nearest town (log of kilometers), and distance to the district headquarters (log of kilometers). In addition, I calculate a rainfall variable for each village, equal to the yearly deviation from the mean measured in standard deviations. Finally, I allow the effects rainfall to vary by NREGS status, with the former included to control for possible differences in the cyclicalities of wages driven by rainfall (Jayachandran, 2006; Imbert and Papp, 2015).

I use two main variables to measure labor-market outcomes. The first is the wage. I construct this variable by including both casual agricultural and casual non-agricultural wages at the individual level. For individuals that worked more than one type of job, I sum hours worked and total remuneration (cash and in-kind) to create the wage variable. The other main labor-market outcome I use is private-sector employment. I construct this variable by including all days worked in casual labor, own-account agricultural labor, own-account non-farm labor, own-account livestock labor, construction/maintenance labor, household (non-leisure) labor, and other economic activities, which includes the collection of firewood, water, etc.

4.2 Summary Statistics

I present summary statistics from 1999, prior to the program, for the sample in Table 1. In all columns, the level of observation is the individual, although they are referred to as households in the table due to the construction of the location variable. The first group of households, “All Households”, includes all households in border districts, regardless of distance to the border. This sample is similar to samples used by previous studies to estimate the impact of NREGS. As previously noted, others have shown that the differences in labor-market outcomes between treated (NREGS) individuals and untreated (non-

NREGS) individuals are quite large. To explore this in the ARIS/REDS data, the third column (“p-value”) presents the results from a test of means across treated and untreated districts. These results suggest large differences across treatment status. For example, the casual wage in treated districts—recall that the first districts to receive the program tended to be worse off (according to a government index) than districts that received the program in later phases—is substantially lower than in untreated districts. The same is true if we disaggregate the overall wage into male and female wages. In addition, in both treated and untreated districts, male wages are significantly—40-50 percent—higher than female wages. Another important outcome, household (food) consumption, is likewise much lower in treated districts.¹⁹

However, this difference in outcomes is much smaller when we focus only on individuals that live within 9 kilometers of borders between treated and untreated districts. For example, the difference in wages falls from almost 25 cents for all households to just 4 cents within 9 kilometers of the border. Interestingly, this difference is driven by competing forces: the male wage is still slightly higher in treated districts (around 10 cents) but the female wage is lower in treated districts by around 14 cents. The difference in monthly food expenditure at the household level likewise falls by almost half. Insofar as trends in labor-market outcomes are correlated with the starting values of these outcomes, the differences documented in Table 1 suggest that heterogeneous effects may be driven by differences in households, not the program itself.

¹⁹The wage and consumption variables are represented as levels in US dollars in the table for ease of interpretation. Throughout the analyses below, I take the log of rupees for wage and consumption variables.

5 Results

I begin with a graphical representation of wage changes from 1999 to 2008 in Figure 4. Only individuals located in border districts are included in the figure (and only these individuals are included in regressions in this section, unless otherwise noted). The x-axis denotes the distance to the border between treated and untreated districts, with negative values indicating movements towards the interior of untreated districts and positive values indicating movements towards the interior of treated districts. A value of zero indicates the location of the border. As expected after an examination of the summary statistics in Table 1, the wage in 1999 is monotonically decreasing as we move from the interior of untreated districts—presumably the location of better-off households—towards the border and then towards the interior of treated districts. In 2008, we see a similar pattern in untreated districts; the wage is decreasing as we move towards the border. However, there is a clear difference in the pattern as we move towards the interior of treated districts; the wage seems to be increasing relative to 1999.

Table 2 presents regression estimates of this effect. The first column estimates the effect of NREGS when using all individuals with a wage observation. The estimate indicates that NREGS increased the wage in treated districts by around 22 percent relative to untreated districts, an effect larger than previous estimates using other datasets, which may be due to the exclusion of fixed effects in estimation. Each subsequent column restricts the sample to individuals located closer and closer to the border. The second column restricts estimation to only individuals located within 30 kilometers of the border. The effect is just as large as the effect when using all households. The effect is relatively stable until restricting estimation to households within 9-12 kilometers, at which point the estimated effect of NREGS begins to quickly attenuate until it almost completely disappears within six kilometers of the border.

Figure 5 presents this change graphically. I compute these effects as follows: First, I estimate a sequence of separate regressions, restricting the sample to individuals living within some distance, k , of the border. I estimate regressions for all values of k from 9 to 40 and then plot the effect of each separate regression in Figure 5. The x-axis represents these values of k . For example, at $x = 25$ in the graph, the point estimate is for the regression restricting the sample to only individuals within 25 kilometers of the border. Figure 5 confirms the effects found in Table 2; the effect is relatively flat for all values of x greater than 20 but then drops precipitously. Apparently, the attenuation starts at around 14 kilometers. The flat section of the graph—as well as the empirical findings in Table 2—supports the prediction of the theoretical model that the change in wage rate will be constant for sufficiently large distances from the border. This result suggests that labor is mobile in a radius of around 10-15 kilometers. Since the distance is defined using the border, the true radius may be slightly larger than 10-15 kilometers unless untreated households are located right at the border. A slightly larger radius comports with the findings in Muralidharan et al. (2016b), who found spillover effects in a radius of around 20 kilometers.

The second prediction from the model in Section 3 is that wages in untreated areas near the border with treated areas will increase more than wages in untreated areas farther from border. To test this, I present an alternative specification in Table 3. The first two columns explore wage changes only within treated districts while the last two columns explore these changes within untreated districts. While the coefficients in the first and second columns are strongly significant ($p < 0.01$ and $p < 0.05$, respectively), neither of the coefficients for untreated districts is significant at conventional levels. In treated districts, households located more than six kilometers from the border see an average increase in the wage that is approximately 10.2 percent larger than the effect for the households within six kilometers of the border. Overall, these findings support the theoretical model, although the effect is only significant within treated districts.

Finally, the model predicts that increases in travel costs will result in a larger wage gap between treated and untreated districts. Empirically, this will show up as a larger effect of NREGS on wages and this effect will not begin to attenuate until closer to the border. The difference in wages by gender is an ideal test for this prediction. For a number of reasons, women in India tend to be less mobile than men (Khera and Nayak, 2009). As such, we might assume women face higher travel costs than men and thus that we will see less of a decrease in the female wage rate closer to the border. Table 4 presents estimates of the effects of NREGS by gender. There are relatively few observations for women. Since I am interested in comparing the same distance to border across gender, I therefore only look at all individuals within 9 kilometers of the border for both genders. The first three columns in Table 4 present results for male wages. The change in wages within 9 kilometers of the border (9.8 percent) is less than half as large as the change in all wages (20.2 percent). For female wages, on the other hand, there is no difference in the change in wages within 9 kilometers relative to the change in all wages. In addition, the difference between the effect on male wages and female wages within 9 kilometers of the border is significant ($p=0.018$). Insofar as women face higher travel costs than men, the results in Table 4 again support the theoretical model. However, caution must be taken in interpreting the magnitude of these effects; while the overall pattern is clear, a lack of power precludes accurate estimates of the true magnitude, especially as the sample size gets smaller closer to the border.

5.1 Causes of Heterogeneity

Given how different households near the border are from households towards the interior of districts, there is a concern that the identifying assumption of parallel trends does not hold. In particular, maybe poorer households—those that are more likely to live towards the interior of treated districts—had wages trending upwards relative to other households.

In this case, the estimates presented above may be driven by underlying trends, not the program itself.²⁰ Although the gender effects in Table 4 are suggestive evidence that this is not the case, I present further evidence that this is not true in Figure 6.

The previous wave of the ARIS/REDS was collected in 1982. While 17 years is a long time, trends from 1982 to 1999 may nonetheless provide some evidence that the assumption of (spatially) parallel trends is plausible. In Figure 6, I present a graph of nominal wages in 1982 and 1999 by distance to border. Not surprisingly, there is a large overall increase in wages over the time period. However, the spatial pattern is almost identical in both years; if anything, it appears that wages in the interior of *untreated* districts may have been trending slightly upwards. It does not appear that wages were trending in the same direction as the results presented above. While this is not definitive evidence, it nonetheless suggests the parallel trends assumption is plausible.²¹

I present additional evidence of spatially similar pre-program trends using NSS data. The NSS data do not allow any disaggregation below the district level. However, if the spatial heterogeneity identified above is a true representation of the effect of NREGS, we might expect to see similar heterogeneity based on a district's neighbors. In other words, treated districts that share a longer border with untreated districts presumably will see a smaller wage increase than treated districts that share a shorter border with untreated districts. The first column in Table 5 presents these results. The data come from 2004/05 and 2007/08, and thus identify the effect at the same time as the ARIS/REDS data. In line with expectations, we see that a longer border does indeed appear to decrease the effects of NREGS. While this finding supports the previous results, the real benefit of using the NSS data is

²⁰There is some evidence that this is unlikely. For example, Azam (2012) argues that NREGS was implemented due to the perception that decades of sustained economic growth had not had a sufficient impact on poverty. This would seem to suggest that wages were not trending upwards for the poorest households.

²¹Due to data limitations from the 1982 wave, all data is collapsed to the household level. This results in substantially fewer overall observations and precludes estimating effects empirically; the resulting estimates are too imprecise to be informative and the smaller sample size, especially in regressions restricted by distance, results in very sensitive estimates.

that it allows for an empirical exploration of pre-program trends. Using the 1999/2000 and 2004/05 waves of the NSS, column two presents results of a placebo test. If wages were trending differentially based on this border percent variable, then a differences-in-differences estimate using two pre-program trends will find similar results. However, this is not what we find; the results using the 1999/2000 and 2004/05 waves of the NSS present no evidence of such a trend. Overall, these results suggest the (conditional) parallel trends assumption is plausible in this case.

If household wages were trending differently due to specific characteristics—like wage levels prior to the program—we would also expect phase three districts to have different trends from phase one and two districts. In the previous subsection, I showed that the effects of NREGS were spatially heterogeneous within treated (untreated) districts that bordered untreated (treated) districts. Instead of labor mobility, it may be that NREGS districts implement the program differently near borders compared to the interior. For example, if district capitals or towns tend to be located in the interior of treated districts, then the program may be relatively better implemented in the interior. In Table 6, I explore this possibility by estimating the effect of distance to border on several characteristics related to implementation of the program. Importantly, the coefficients on distance are very small and never significant. The percent of agricultural labor (out of total labor force) in the district prior to implementation of NREGS has a consistently positive effect on NREGS implementation, with more individuals receiving work and more money spent on works.²²

If households near the border were trending differentially from households towards the interior, then we will see a similar attenuation if we estimate the effect at the border of two treated districts relative to the border of two untreated districts. Table 7 presents additional

²²The coefficient on agricultural labor is very large for a number of estimates. However, it is important to note that this is not the true effect seen in the data. For example, almost 60 percent of districts in the regressions have a value of agricultural labor between 0.3 and 0.4. As such, while the coefficients are very large, the largest possible effect seen in the data is much smaller. Standardizing the variables results in more intuitive values. Results available upon request.

evidence that this is not the case. The first column presents results for all households. The change in wages in treated districts is approximately 19.9 log-points larger than the change in untreated districts. Importantly, as we restrict the sample to households closer to the border with a district of the same treatment status, the effect stays stable; in the last column, the effect of NREGS on wages at the border between two treated districts relative to the border between two untreated districts is slightly larger than the effect if we use all households in treated districts and untreated districts and is still significant. In conjunction with Table 6, this is suggestive evidence that intra-district differences in program implementation are not driving the effects nor that differences in trends for border households relative to interior households are responsible.

Another possible explanation for these effects is program leakage. In other words, we may see no effect at the border because untreated households are somehow accessing NREGS even though their district does not yet have the program. Figure 7 presents suggestive evidence that this is not the case. In 2008, relatively few households in non-NREGS districts had access to the program. On the other hand, relatively many households in NREGS districts had access. Moreover, the difference manifests itself right where we would expect: at the border.

Finally, I present one piece of evidence that labor market spillovers—through movement of labor—are the mechanism. Ideally, the survey would ask how far each individual travels to work on a daily basis for each type of labor. While the ARIS/REDS does not have this specific question for all labor, it does ask a similar question for non-farm casual labor. Unfortunately, this question is only asked in 2008, which precludes an estimate of relative changes in travel time. However, if labor mobility is causing the heterogeneity, we should see workers near the border in untreated districts traveling longer, while we would expect relatively consistent travel times through treated districts. Figure 8 presents graphical evi-

dence that this is indeed the case. In particular, we see that travel time to non-farm casual employment is longest right near the border in untreated districts and then decreases as we move towards the interior.²³ In addition, travel time is almost constant throughout treated districts. Overall, Figure 8—in addition to the differential effects by gender—supports the hypothesis that individuals in untreated districts are traveling to treated districts to engage in casual labor.

5.2 Private-Sector Employment and Income

Much of the previous NREGS literature has examined the program’s effect on rural wages. However, two other outcomes are also commonly studied: private-sector employment and income. I first look at private-sector employment. Many economic models predict decreases in private-sector employment after the implementation of a large public works program like NREGS. Intuitively, if NREGS drives up wages, we would expect a lower level of private-sector employment in equilibrium (Imbert and Papp, 2015). Figure 9 presents a graph of private sector employment (days - log) across both survey waves. Private employment is defined as total time spent in casual labor, salary labor, agricultural self-employment, and non-agricultural self-employment. This includes “productive” household activities—including the collection of firewood, water, etc.—but does not include leisure. Consistent with the theory, there appears to be a large decrease in private-sector employment when comparing NREGS districts from 1999 to 2008. In addition, the graph appears relatively flat at the border in treated districts, which is consistent with a negligible wage change in that region.

The dependent variable is now (log of days of) total employment, so I no longer restrict

²³We would not expect longer travel times in the interior of untreated districts because these households are located far enough from the border that traveling for casual labor is not undertaken, as suggested by the model.

estimation to only individuals with observed wages. As such, this greatly increases the number of observations. Table 8 presents these results. In the first column, the estimation includes all households in border districts. The overall estimate is quite imprecise and insignificant. However, this again conceals a more complex picture. The results in the second column include only households within 9 kilometers of the border, while the results in the third column include only households outside 30 kilometers from the border. The estimates are very imprecise so it is difficult to pinpoint the true effect; nonetheless, a formal test for equality of the coefficients in columns (2) and (3) rejects equality ($p=0.023$). These results are as we would expect given the wage patterns documented above: in the ARIS/REDS data, private-sector employment decreases statistically significantly more in areas farther from the border—where the wage increase was highest—relative to areas near the border.

Finally, I explore the effect of NREGS on household per capita income. While an increase in wages might be assumed to be a de facto increase in the welfare of poor rural households, it is nonetheless an imperfect proxy. As such, studying the effects of NREGS on income offers a more direct measure of the welfare effects of the program. I begin with a graphical representation of per capita income in the data in Figure 10. The pattern in Figure 10 is similar to the wage and private employment changes shown in Figure 4 and Figure 9, as there again appears to be a relative increase in the outcome as we move towards the interior of NREGS districts. Table 9 presents these results in regression form. These results confirm the basic pattern in the figure: the effect of NREGS on income is smaller near the border. Despite the imprecision in the estimates, the effect outside of 30 kilometers is significantly larger than the effect within 9 kilometers ($p=0.054$). While the point estimates suggest the entirety of the consumption increase is due to the wage increase, the large standard errors suggest caution in interpreting the coefficients. Nonetheless, the overall pattern clearly indicates that wage changes are an important drive of consumption

changes in rural India.

6 Conclusion

In this paper, I document the effects of the Mahatma Gandhi National Rural Employment Guarantee Scheme (NREGS) in India. In particular, I show that the effects of the program are spatially heterogeneous: treated areas located near untreated areas see smaller wages increases than treated areas located farther from untreated areas. In addition This heterogeneity does not appear to be driven by pre-program trends or program leakage. Rather, evidence indicates that labor mobility is driving these effects; the effect is more concentrated among male wages and workers living in untreated areas located near treated areas are more likely to travel longer distances to work than other untreated workers. However, while the pattern is clear across all of the results presents, a small sample size and demanding empirical strategy suggest caution in interpreting the magnitude of the estimated effects.

The results in this paper have several implications. First, these results suggest that previous studies of the effect of the program on wages underestimate the true effect of the program. This bias theoretically operates through two separate channels: an overestimate of the counterfactual wage in untreated districts and an underestimate of the counterfactual wage in treated districts, although the present results suggest the former is more important. In addition, when the program is implemented in phase three districts, wages in phase one and two districts may increase. Second, this paper presents suggestive evidence that the wage increases due to the program had an appreciable effect on household incomes, which supports recent experimental evidence (Muralidharan et al., 2016b). This result suggests raising rural wages is an effective poverty-fighting tool. This may help explain studies which find smaller welfare effects of public works programs in other countries (e.g. Beegle

et al., 2017), especially if the program is not at a large enough scale to have appreciable impacts on prevailing wages.

References

- Afridi, F., Mukhopadhyay, A., and Sahoo, S. (2012). Female labour force participation and child education in India: The effect of the National Rural Employment Guarantee Scheme. *Working Paper*.
- Anant, T., Hasan, R., Mohapatra, P., Nagaraj, R., and Sasikumar, S. (2006). Labor markets in India: issues and perspectives. In Felipe, J. and Hasan, R., editors, *Labor Markets in Asia*, pages 205–300. Palgrave Macmillan.
- Angelucci, M. and De Giorgi, G. (2009). Indirect effects of an aid program: How do cash transfers affect ineligibles' consumption? *The American Economic Review*, 99(1):486–508.
- Azam, M. (2012). The impact of Indian job guarantee scheme on labor market outcomes: Evidence from a natural experiment. *IZA Discussion Paper*.
- Banerjee, A., Duflo, E., Imbert, C., Mathew, S., and Pande, R. (2015). Can e-governance reduce capture of public programs? Experimental evidence from a financial reform of India's employment guarantee. *Working Paper*.
- Basu, A. K., Chau, N. H., and Kanbur, R. (2009). A theory of employment guarantees: Contestability, credibility and distributional concerns. *Journal of Public Economics*, 93(3):482–497.
- Beegle, K., Galasso, E., and Goldberg, J. (2017). Direct and indirect effects of malawi's public works program on food security. *Journal of Development Economics*.
- Behrman, J. R. (1999). Labor markets in developing countries. *Handbook of Labor Economics*, 3:2859–2939.
- Berg, E., Bhattacharyya, S., Rajasekhar, D., and Manjula, R. (2014). Can Public Employment Schemes Increase Equilibrium Wages? Evidence from a natural experiment in India. *The Centre for Market and Public Organisation*.
- Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics*, 119(1):249–275.
- Cunha, J. M., De Giorgi, G., and Jayachandran, S. (2011). The price effects of cash versus in-kind transfers. *National Bureau of Economic Research Working Paper*.
- Duflo, E. (2000). Schooling and labor market consequences of school construction in Indonesia: Evidence from an unusual policy experiment. *National Bureau of Economic Research Working Paper*.

- Dutta, P., Murgai, R., Ravallion, M., and Van de Walle, D. (2014). Does India's employment guarantee scheme guarantee employment? *World Bank Policy Research Working Paper*.
- Foster, A. D. and Rosenzweig, M. R. (2010). Is there surplus labor in rural India? *Working Paper*.
- Imbert, C. and Papp, J. (2011). Estimating leakages in India's employment guarantee. In Khera, R., editor, *Battle for Employment Guarantee*, pages 269–78. Oxford University Press.
- Imbert, C. and Papp, J. (2015). Labor market effects of social programs: Evidence from India's employment guarantee. *American Economic Journal: Applied Economics*, 7(2):233–263.
- Jayachandran, S. (2006). Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of Political Economy*, 114(3):538–575.
- Jha, R., Bhattacharyya, S., and Gaiha, R. (2011). Social safety nets and nutrient deprivation: An analysis of the National Rural Employment Guarantee Program and the Public Distribution System in India. *Journal of Asian Economics*, 22(2):189–201.
- Kaur, S. (2014). Nominal wage rigidity in village labor markets. *National Bureau of Economic Research Working Paper*.
- Khawlneikim, S. and Mital, M. (2016). Socio-economic impact of Rural Employment Guarantee Scheme in Churachandpur District, Manipur: A beneficiary perspective. *International Journal of Scientific Research*, 4(12).
- Khera, R. and Nayak, N. (2009). Women workers and perceptions of the national rural employment guarantee act. *Economic and Political Weekly*, pages 49–57.
- Li, T. and Sekhri, S. (2013). The unintended consequences of employment based safety net programs. *Working Paper*.
- Liu, Y. and Barrett, C. B. (2012). Heterogeneous pro-poor targeting in India's Mahatma Gandhi National Rural Employment Guarantee Scheme. *Working Paper*.
- Liu, Y. and Deininger, K. (2010). Poverty impacts of India's National Rural Employment Guarantee Scheme: Evidence from Andhra Pradesh. *Selected paper prepared for presentation at the Agricultural & Applied Economics Association*.
- Merfeld, J. (2017). Public Works, Crowd Out, and Non-Farm Self-Employment: Evidence from India. *Working Paper*.
- Miguel, E. and Kremer, M. (2004). Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, 72(1):159–217.

- Ministry of Law and Justice (2005). The National Rural Employment Guarantee Act. <http://nrega.nic.in/rajaswa.pdf>.
- Mukhopadhyay, A., Sharan, M., et al. (2015). The National Rural Employment Guarantee Scheme in Rajasthan: Rationed funds and their allocation across villages. *Working Paper*.
- Munshi, K. and Rosenzweig, M. (2009). Why is mobility in India so low? Social insurance, inequality, and growth. *National Bureau of Economic Research Working Paper*.
- Muralidharan, K., Niehaus, P., and Sukhtankar, S. (2016a). Building state capacity: Evidence from biometric smartcards in India. *The American Economic Review*, 106(10):2895–2929.
- Muralidharan, K., Niehaus, P., and Sukhtankar, S. (2016b). General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence from India. *Department of Economics, University of California, San Diego, working paper*.
- Niehaus, P. and Sukhtankar, S. (2013). The marginal rate of corruption in public programs: Evidence from India. *Journal of Public Economics*, 104:52–64.
- Planning Commission (2003). Report of the Task Force: Identification of Districts for Wage and Self Employment Programmes.
- Ravallion, M., Van de Walle, D. P., Dutta, P., and Murgai, R. (2013). Testing information constraints on India's largest antipoverty program. *Working Paper*.
- Ravi, S. and Engler, M. (2015). Workfare as an effective way to fight poverty: The case of India's NREGS. *World Development*, 67:57–71.
- Reddy, D. N. and Upendranadh, C. (2010). National rural employment guarantee: Issues, concerns and prospects. *Oxfam India Working Papers Series*.
- Rosenzweig, M. R. (1978). Rural wages, labor supply, and land reform: A theoretical and empirical analysis. *The American Economic Review*, 68(5):847–861.
- Rosenzweig, M. R. (1980). Neoclassical theory and the optimizing peasant: An econometric analysis of market family labor supply in a developing country. *The Quarterly Journal of Economics*, 94(1):31–55.
- Rosenzweig, M. R. (1984). Determinants of wage rates and labor supply behavior in the rural sector of a developing country. In Binswanger, H. and Rosenzweig, M., editors, *Contractual Arrangements, Employment and Wages in Rural Labour Markets in Asia*, pages 1–41. Yale University Press, New Haven.
- Rosenzweig, M. R. (1988). Labor markets in low-income countries. *Handbook of Development Economics*, 1:713–762.

- Shah, M. and Steinberg, B. M. (2015). Workfare and human capital investment: Evidence from India. *National Bureau of Economic Research Working Paper*.
- Shah, M. and Steinberg, B. M. (2017). Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital. *Journal of Political Economy*, 125(2):000–000.
- Sharma, A. (2009). *NCAER-PIF Study on Evaluating performance of National Rural Employment Guarantee Act*. National Council of Applied Economic Research.
- Subbarao, K. et al. (2003). *Systemic shocks and social protection: Role and effectiveness of public works programs*. Social Protection, World Bank.
- Sukhtankar, S. (2016). India’s National Rural Employment Guarantee Scheme: What Do We Really Know about the World’s Largest Workfare Program? *Working Paper*.
- Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics*, 2(4):1–41.
- Zimmermann, L. (2012). Labor market impacts of a large-scale public works program: Evidence from the Indian Employment Guarantee Scheme. *IZA Discussion Paper*.

Figure 1: NREGS Districts and ARIS/REDS Villages

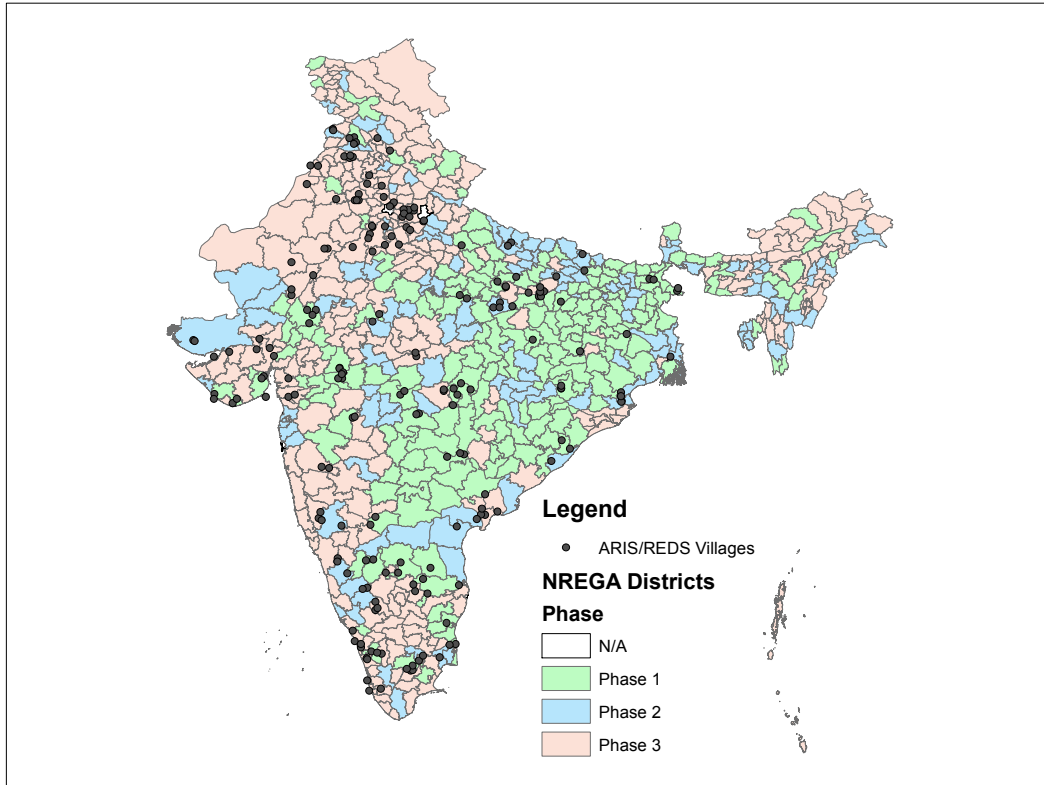

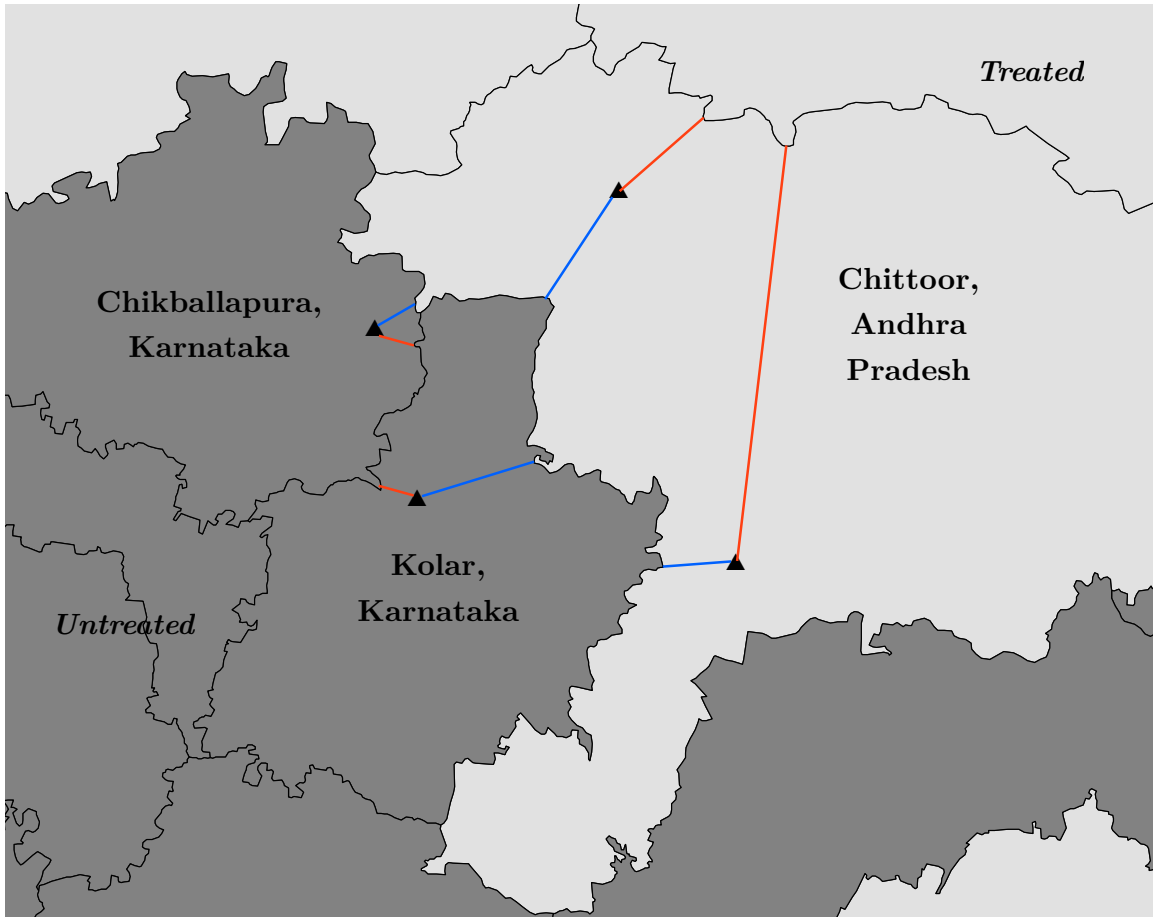


Figure 2: Village Distance to Border with Other Phase District



The figure shows the histogram of distance to border for NREGS and non-NREGS villages. For villages in NREGS districts (phase one and two districts), the distance is calculated as the shortest distance to a non-NREGS district (phase three districts). For villages in non-NREGS districts, the distance is calculated as the shortest distance to a NREGS district. Distance is top-coded at 100 in the figure, but not in analyses in the paper. Only districts that border a district from the other type of district are included.

Figure 3: Creation of Distance Variables




The dark gray districts—Chikballapura and Kolar—are untreated (phase three) districts. The light gray district—Chittoor—is a treated (phase one or two) district. The triangles represent actual villages in the dataset. Blue lines indicate the distance to the nearest district of the opposite treatment status while red lines indicate the distance to the nearest district of the same treatment type.

Table 1: Summary Statistics (1999)

	All Households			Within 9km of border		
	Non-NREGS	NREGS	p-value	Non-NREGS	NREGS	p-value
Wage (USD)	1.242 (0.507)	0.992 (0.456)	0.000	1.201 (0.537)	1.161 (0.464)	0.183
Wage - male (USD)	1.389 (0.513)	1.093 (0.486)	0.000	1.365 (0.547)	1.260 (0.470)	0.004
Wage - female (USD)	0.919 (0.305)	0.790 (0.300)	0.000	0.817 (0.227)	0.959 (0.380)	0.000
Food Consumption (USD)	70.434 (25.035)	61.018 (22.706)	0.000	70.930 (27.285)	65.737 (28.525)	0.002
Male	0.687 (0.464)	0.666 (0.472)	0.177	0.700 (0.458)	0.673 (0.470)	0.330
Age	34.681 (13.044)	34.385 (12.842)	0.480	34.995 (13.532)	33.133 (12.413)	0.017
Education	0.653 (0.833)	0.532 (0.853)	0.000	0.615 (0.783)	0.603 (0.921)	0.815
Household size	6.362 (3.445)	6.258 (3.192)	0.330	6.025 (3.032)	6.171 (2.628)	0.389
Head education	0.466 (0.724)	0.371 (0.671)	0.000	0.423 (0.646)	0.362 (0.696)	0.136
Head male	0.942 (0.233)	0.935 (0.246)	0.358	0.932 (0.252)	0.937 (0.243)	0.747
Head age	48.833 (13.461)	47.744 (13.098)	0.011	49.100 (13.377)	46.203 (13.569)	0.000
Observations	1995	1837		751	428	

Standard deviations are in parentheses. Observations are individuals. The first three columns include all individuals that reside in border districts. The last three columns include only individuals that live within 9 kilometers of the border with a district of the opposite treatment status.

Figure 4: Distance to Border and Wages



The figure shows a lowess curve of the relationship between distance to border and the wage rate for NREGS and non-NREGS households before (1999) and after (2008) implementation of the program. For households in NREGS districts (phase one and two districts), the distance is calculated as the shortest distance to a non-NREGS district (phase three districts). For households in non-NREGS districts, the distance is calculated as the shortest distance to a NREGS district. The figure is restricted to households within 100 km of the nearest district in the figure, but not in analyses in the paper.


Table 2: Effects of NREGS on Wages - Distance to Border

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	All Villages	Within 30 km	Within 20 km	Within 15 km	Within 12 km	Within 9 km	Within 6 km
Post times NREGS	0.219*** (0.060)	0.196*** (0.066)	0.199*** (0.071)	0.204*** (0.079)	0.197* (0.105)	0.155 (0.099)	-0.018 (0.078)
District Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Rainfall	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7714	5504	4357	3464	2731	2329	1449

Standard errors are in parentheses. Standard errors are clustered at the district level. Wage (log) is the dependent variable in all columns. The first column includes all households in "border" districts, which consist of NREGS districts that border non-NREGS districts and non-NREGS districts that border NREGS districts. The second through seventh columns restrict estimation to households within a certain distance (defined in each column label) of the border. The level of analysis in all columns is the individual.

* p<0.1 ** p<0.05 *** p<0.01

Figure 5: Effect by Distance to Border



To create the figure, the sample is varied by using different cut-offs at different distances to the border, from 9 to 40. For example, for a distance of 25, the plotted point represents the coefficient estimate from a difference-in-differences specification including only households within 25 kilometers of a NREGS/non-NREGS border.

Table 3: Effects by District Type

	NREGS Districts		Non-NREGS Districts	
	(1)	(2)	(3)	(4)
	All Villages	All Villages	All Villages	All Villages
Post times (log) Distance to Border	0.053*** (0.017)		0.002 (0.018)	
Post times Distance > 6		0.102** (0.048)		-0.047 (0.062)
District Controls	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes	Yes
Rainfall	Yes	Yes	Yes	Yes
Observations	4109	4109	3605	3605

Standard errors are in parentheses. Standard errors are clustered at the district level. Wage (log) is the dependent variable in all columns. The first two columns include only individuals in NREGS districts. The last two columns include only individuals in non-NREGS districts. The coefficient in the first row is an interaction between Post and Distance to Border; as such, the coefficient represents the change in wages within treated districts (first two columns) or within untreated districts (last two columns) based on the distance to the nearest border with an untreated or treated district. The coefficient in the second row is an interaction between Post and an indicator variable indicating whether the observation is located outside of 6 kilometers of the border.

* p<0.1 ** p<0.05 *** p<0.01

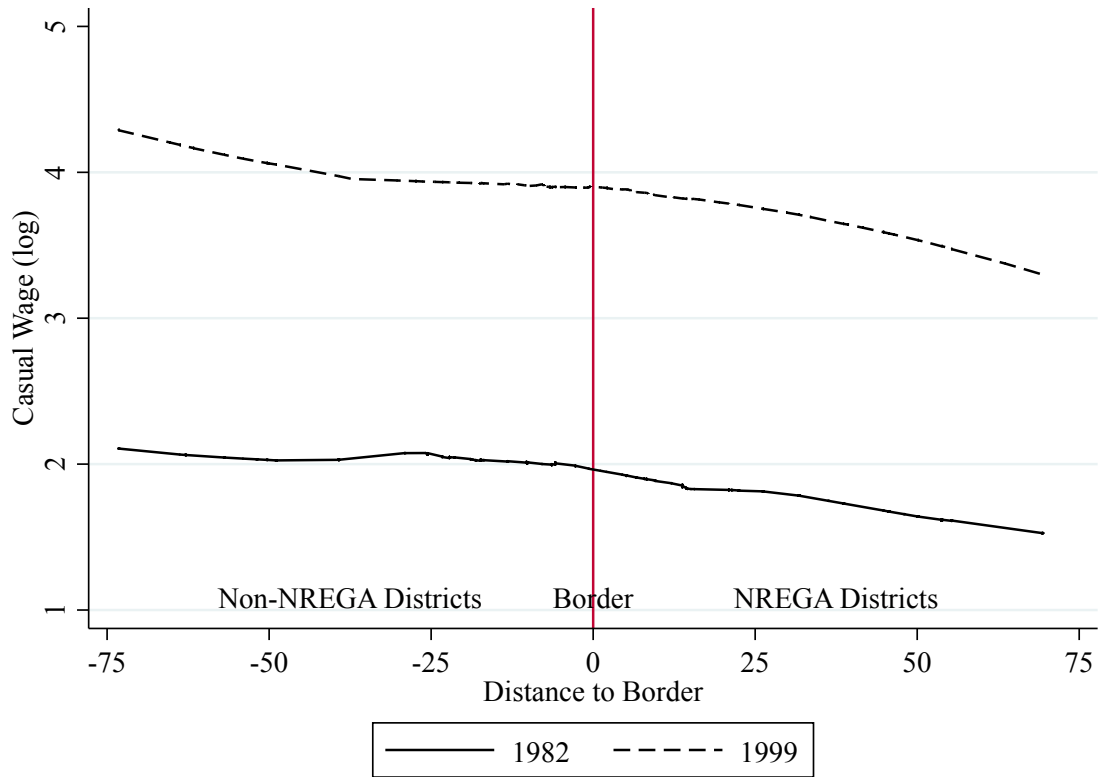
Table 4: Effects of NREGS on Wages by Gender

	Male wages (log)			Female wages (log)		
	(1)	(2)	(3)	(4)	(5)	(6)
	All	15 km	9 km	All	15 km	9 km
Post times NREGS	0.202*** (0.057)	0.188** (0.078)	0.098 (0.088)	0.278*** (0.078)	0.227** (0.108)	0.363** (0.136)
District Controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes	Yes	Yes	Yes
Rainfall	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5395	2382	1641	2319	1082	688

Standard errors are in parentheses. Standard errors are clustered at the district level. Wage (log) is the dependent variable in all columns. All columns include only households in "border" districts, which consist of NREGS districts that border non-NREGS districts and non-NREGS districts that border NREGS districts. The first column for each gender includes all observations in these border districts. The second column includes only individuals that reside within 15 kilometers of the border. The third column includes only individuals that reside within 9 kilometers of the border between these districts. The level of analysis in all columns is the individual.

* p<0.1 ** p<0.05 *** p<0.01

Figure 6: Distance to Border and Wages - Pre-Program Trends



Due to data limitations, the unit of analysis is the household. The figure shows a lowess curve of the relationship between distance to border and the wage rate for NREGS and non-NREGS households before implementation of the program. For households in NREGS districts (phase one and two districts), the distance is calculated as the shortest distance to a non-NREGS district (phase three districts). For households in non-NREGS districts, the distance is calculated as the shortest distance to a NREGS district. The figure is restricted to households within 75 km of the nearest district in the figure.

Table 5: Effects and District Border Percentage

	(1)	(2)
	Wage (log) - Effect	Wage (log) - Placebo
Post times NREGS times Border Percent	-0.281*** (0.090)	
Post (placebo) times NREGS times Border Percent		-0.033 (0.155)
Observations	49686	44407

Standard errors are in parentheses. Standard errors are clustered at the district level. Wage (log) is the dependent variable in both columns. Data comes from the National Sample Survey (NSS). In the first column, the coefficient is a triple interaction between post and NREGS (the difference-in-differences estimator at the district level) and a variable equal to the percentage of each district's borders that is shared with a district of the opposite treatment status. In other words, larger values for "Border Percent" indicates the district shares a longer border (as a percentage of own border) with a district of the opposite treatment status. The first column uses 2004/05 and 2007/08 waves, the same years in the ARIS/REDS data. The second column uses data from 1999/2000 and 2004/05 to estimate a placebo test for pre-program trends at the district level.

* p<0.1 ** p<0.05 *** p<0.01

Table 6: Distance to Border and NREGS Implementation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Rolls Filled	Ind. Worked	Person-Days	HH Limit	Labor Exp. (R's lakhs)	Percent Labor Expenditure	Total Works
Distance to Non-NREGS (km - log)	0.021 (0.076)	0.009 (0.083)	0.035 (0.086)	0.028 (0.116)	0.021 (0.097)	0.003 (0.012)	-0.014 (0.093)
Village Population Density (log)	0.130* (0.075)	0.128* (0.074)	0.129 (0.078)	0.063 (0.081)	0.141 (0.092)	0.012 (0.007)	0.104 (0.077)
Village Population (log)	0.128 (0.113)	0.187 (0.165)	0.197 (0.126)	0.108 (0.213)	0.142 (0.134)	0.058*** (0.016)	-0.243* (0.125)
Distance to District HQ (km - log)	-0.033 (0.154)	0.014 (0.217)	-0.067 (0.201)	-0.210 (0.312)	-0.005 (0.193)	0.023 (0.020)	0.220 (0.193)
Distance to Nearest Town (km - log)	0.081 (0.145)	0.174 (0.159)	0.116 (0.170)	0.097 (0.227)	0.162 (0.175)	0.008 (0.021)	-0.039 (0.138)
District Population (log)	0.388 (0.286)	-0.479 (0.444)	-0.132 (0.412)	0.810 (0.503)	-0.208 (0.425)	-0.031 (0.042)	0.263 (0.345)
District Percent Rural	4.639** (1.788)	0.329 (1.914)	1.735 (1.901)	4.486 (3.320)	1.042 (2.103)	0.184 (0.216)	1.704 (1.927)
District Percent Ag. Labor	3.828*** (1.235)	2.004 (1.644)	3.102** (1.344)	5.057** (2.080)	2.713* (1.404)	0.483*** (0.160)	1.172 (1.549)
District Wage (R - log)	-0.309 (0.442)	-0.356 (0.417)	-0.003 (0.453)	0.193 (0.859)	-0.007 (0.528)	0.001 (0.066)	-0.155 (0.461)
Observations	71	71	71	69	71	71	71

Standard errors are in parentheses. The dependent variable in each column is indicated by the column header. For columns (1), (2), (3), (4), (5), and (7), the dependent variable is logged. NREGS implementation data comes from 2012-2013. Only treated districts in the ARIS/REDS data are included. Also included are literacy rate, percent SC, percent ST, and labor force participation rate.

* p<0.1 ** p<0.05 *** p<0.01

Table 7: Wages and Distance to Same Border

	(1)	(2)	(3)	(4)
	All Villages	Within 15 km	Within 12 km	Within 9 km
Post times NREGS	0.199*** (0.065)	0.226*** (0.074)	0.230*** (0.082)	0.229** (0.097)
District Controls	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes	Yes
Rainfall	Yes	Yes	Yes	Yes
Observations	7286	4286	3688	3310

Standard errors are in parentheses. Standard errors are clustered at the district level. Wage (log) is the dependent variable in all columns. In columns (2) through (4), distance is defined as the distance from treated areas to the nearest border with a treated district and as the distance from untreated areas to the nearest border with untreated districts. In other words, the relevant comparison is the change in wages closer to the border between two treated districts relative to the change in wages closer to the border between two untreated districts.

* p<0.1 ** p<0.05 *** p<0.01

Figure 7: Public Works in 2008



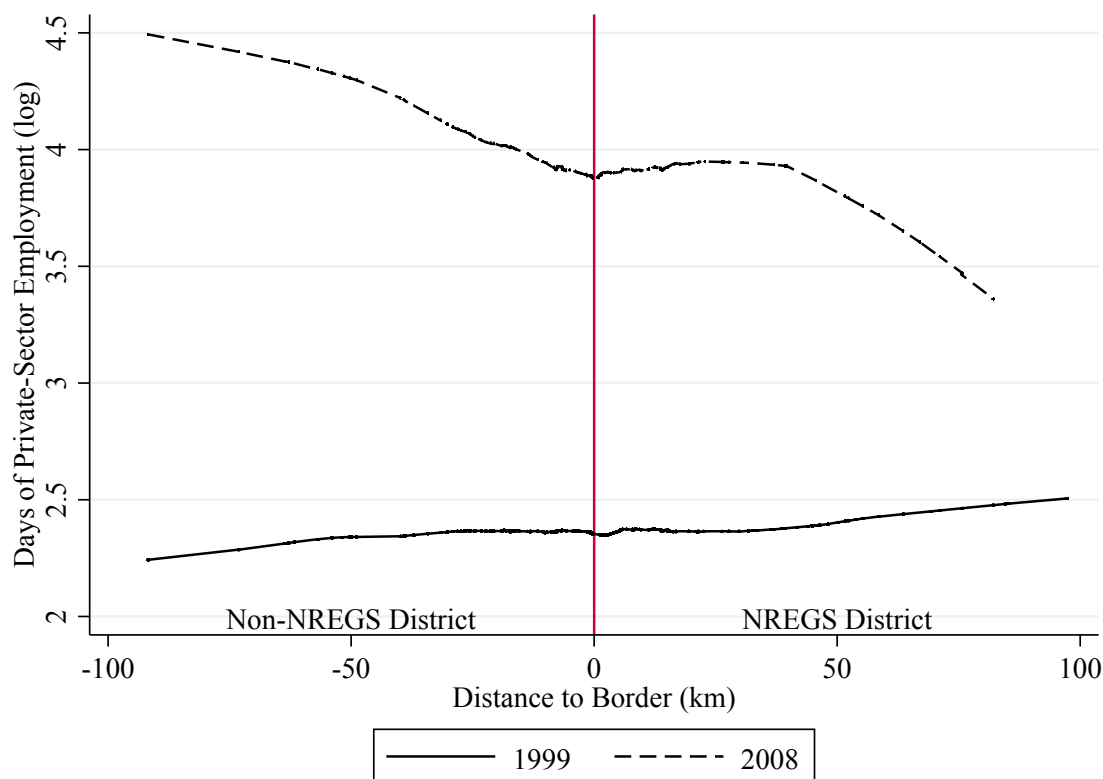


Figure 8: Travel Time to Non-Farm Employment (2008)



The y-axis is travel time to non-farm employment (minutes) in 2008. The x-axis is distance to border between NREGS and non-NREGS districts.

Figure 9: Private-Sector Labor by Wave



The y-axis is (log of) days of private-sector employment. The x-axis is distance to border between NREGS and non-NREGS districts.


Table 8: Effects of NREGS on Private-Sector Employment

	(1) All Villages	(2) Within 9 km	(3) Outside 30 km
Post times NREGS	-0.120 (0.118)	0.173 (0.175)	-0.341** (0.165)
District Controls	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes
Rainfall	Yes	Yes	Yes
Observations	52289	16369	14948

Standard errors are in parentheses. Standard errors are clustered at the district level. Days of private-sector employment (log) is the dependent variable in all columns. All columns include only households in “border” districts, which consist of NREGS districts that border non-NREGS districts and non-NREGS districts that border NREGS districts. The first column includes all households in these border districts. The second column includes only households within 9 kilometers of the border between these districts, while the third column includes only households outside 30 kilometers of the border. The level of analysis in all columns is the individual.

* p<0.1 ** p<0.05 *** p<0.01

Figure 10: Household Income (2006)



The y-axis is (log of) per capita income at the household level. The x-axis is distance to border between NREGS and non-NREGS districts.

Table 9: Effects of NREGS on Household Income

	(1) Outside 30 km	(2) Within 9 km
Post times NREGA	0.439** (0.189)	0.002 (0.132)
District Controls	Yes	Yes
Household Controls	Yes	Yes
Rainfall	Yes	Yes
Observations	2649	3099

Standard errors are in parentheses. Standard errors are clustered at the district level. Household per capita income (log) is the dependent variable in all columns. Both columns include only border districts. The first column includes all households more than 30 kilometers from the border, while the second column includes only households within 9 kilometers of the border. The level of analysis in all columns is the household.

* p<0.1 ** p<0.05 *** p<0.01