Three Essays on Labor and Government Policy in Developing Countries

Joshua D Merfeld

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2018

Reading Committee:

C. Leigh Anderson, Chair

Brian Dillon

James Long

Mark Long

Victor Menaldo

Program Authorized to Offer Degree:

Evans School of Public Policy and Governance

University of Washington

Abstract

Three Essays on Labor and Government Policy in Developing Countries

Joshua D Merfeld

Doctor of Philosophy in Public Policy and Management

Chair of the Supervisory Committee:

C. Leigh Anderson

Evans School of Public Policy and Governance

In this dissertation, I examine household labor allocation in two developing countries: India and Malawi. I focus on the interaction between government policies and household labor allocation decisions. The first and second papers look at labor allocation and an Indian government program, while the first and third papers explore labor allocation to non-farm self-employment. In the first paper, I analyze the effects of an increase in the rural wage in India – induced by the government's Mahatma Gandhi National Rural Employment Guarantee Scheme (NREGS) – on non-farm self-employment. Households significantly decrease labor allocation to non-farm self-employment following implementation of the program. Moreover, this effect is higher in areas with higher rainfall variation, which I use as a proxy for agricultural production risk. Taken together, these findings support the view that many non-farm enterprises are subsistence enterprises, started due to a lack of remunerative labor opportunities and to diversify production risk. In the second paper, I continue analyzing the effects of NREGS. Using a unique dataset with GPS location data, I show that the wage effects of the program are spatially heterogeneous. The wage effects are largest on the interior of treated districts, far from untreated areas. In contrast, areas on the border between treated and untreated areas see no wage increase at all. Estimates suggest labor is mobile on a daily basis in a radius of around 15 to 20 kilometers. In the third paper, coauthor Peter Brummund and I study labor allocation across agricultural and non-farm production in Malawian households. Economic theory suggests that households should equate the marginal revenue product of labor (MRPL) across productive activities within the household. We test this assumption by estimating production functions for agricultural and non-farm production. We show that MRPL in agricultural production is significantly higher than MRPL in non-farm production. Moreover, consistent with a large body of literature on the sectoral productivity gap, we show that the average product of labor is higher in non-farm production. These results suggest a rethinking of how we measure the sectoral productivity gap may be warranted.

Acknowledgements

First and foremost, I want to thank my wonderful wife, Young Eun, for her love, support, and encouragement during the last five years. Having someone by my side who understood the difficulties of being a graduate student was invaluable. Young Eun is also responsible for bringing our beautiful daughter, Elin, into this world. The two of them have made the last six months of hard work and stress much easier to bear. I also thank my parents, Mary Anne and Michal, and my brother, Sean, for their love and support over the last few years.

I additionally want to thank all of my committee members – Leigh Anderson, Brian Dillon, James Long, Mark Long, and Victor Menaldo – for their guidance and advice. This dissertation would not have been possible without them. I especially thank Peter Brummund for his help with the final chapter of this dissertation. The published version of the paper will be coauthored by the two of us.

Finally, I want to thank my entire cohort – Young Eun, Kate, and Amy – for their support. Completing this process would have been much more arduous if not for being surrounded by such intelligent, caring friends.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Rural households in developing countries have access to very few productive resources. Farmers lack easy access to many modern inputs, including fertilizers, draught animals, and machinery. In this context, labor is one of these households' most productive assets. As such, understanding how and why these households allocate their labor the way they do is imperative to developing new and improving existing development policy. In this dissertation, I examine household labor allocation in two developing countries: India and Malawi. I focus on the interaction between government policies and household labor allocation; the first two essays explicitly evaluate the effects of an Indian government active labor market program while the third essay informs a current policy debate surrounding structural transformation and development policy. Additionally, the first and third essays build on this structural transformation theme, focusing on labor allocation to the non-farm sector by way of non-farm self-employment.

In the first essay of this dissertation, I explore how and why households allocate labor to non-farm self-employment. The classical economic development literature argues that growth is accompanied by workers moving out of the agricultural sector and into the more "modern" industrial sector (Lewis, 1954; Kuznets, 1957; Ranis and Fei, 1961). While large proportions of the labor force remain tied to their lands, recent evidence suggests the non-farm sector is growing in many developing countries (Lanjouw and Shariff, 2004; Lanjouw and Murgai, 2009). Does this recent growth spurt in the rural non-farm sector – and non-farm self-employment, in particular – represent a turning point, a growing non-agricultural sector that can drive further transformations in the rural economy (Matsuyama, 1992; Tiffin and Irz, 2006)? Or, instead, is this growth evidence of increasing stress in rural labor markets, as individuals flock to the non-farm sector as a means of subsistence (Barrett et al., 2001; Lanjouw and Lanjouw, 2001; Kijima and Lanjouw, 2005; Lanjouw and Murgai, 2009; Binswanger, 2012)?

I first make a simple observation regarding the reactions of different types of enterprises: households and individuals that enter into non-farm self-employment for subsistence reasons are more likely to exit the sector when wages increase or when less risky employment opportunities become available. In other words, the labor elasticity of non-farm self-employment with respect to the market wage should be much higher for subsistence enterprises than other firms. Households that enter non-farm self-employment for profit-making reasons, on the other hand, are less likely to exit the sector with small wage increases. To help tease out these two competing narratives, I explore the effects of India's Mahatma Gandhi National Rural Employment Guarantee Scheme (NREGS). NREGS was the world's largest workfare program at the time of its implementation in 2006 and increased rural wages significantly, by around five percent (Azam, 2011; Imbert and Papp, 2015). In addition, the program has helped households cope with shocks and diversify into riskier activities (Gehrke, 2014; Zimmermann, 2015). These findings suggest the program may plausibly lead to the exit of many subsistence enterprises.

Using a differences-in-differences strategy, I first show that the program has a significant effect on non-farm self-employment for prime-age adults. NREGS decreases the probability an individual engages in any non-farm self-employment during the last week by 1.5 percentage points, a more than 10 percent decrease relative to the pre-program mean. In addition, the program decreases days in non-farm self-employment in the last week by 3.0 percent. I then show that the program's effects are concentrated in the dry season, consistent with prior evidence of the program's wage impacts (Imbert and Papp, 2015) and with the de facto seasonal implementation of the program (Lakha and Taneja, 2009; Sharma, 2009a; Sukhtankar, 2016). Consistent with implementation being concentrated in the dry season, there is an insignificant change in days of non-farm self-employment in the rainy season but a large and statistically significant effect in the dry season: NREGS decreases days of non-farm self-employment by 4.4 percent. A back-of-the-envelope calculation suggests a labor supply elasticity of 0.94, which is three times as large as estimates of overall labor elasticity in India (Evenson and Binswanger, 1980; Muralidharan et al., 2017), including estimates calculated using the same data and program as this paper (Imbert and Papp, 2015). This evidence is consistent with non-farm self-employment playing a significant subsistence role during the dry season. Further analyses suggest households in riskier agricultural areas are more likely to reallocate labor in response to the program, evidence again consistent with non-farm self-employment playing a subsistence and risk-coping role.

In the second, paper I continue exploring the effects of NREGS. Much of the previous literature examining the effects of NREGS focused on labor market outcomes at the district level (Azam, 2011; Zimmermann, 2012; Berg et al., 2014; Imbert and Papp, 2015). This level of focus is attractive for two reasons. First, the program was rolled out at the district level, making analysis at this level only natural. 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 and have been the focus of previous non-NREGS studies, as well Rosenzweig (1978, 1984); Jayachandran (2006); Topalova (2010); Shah and Steinberg (2017). In the second essay, 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 districts ("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.

I estimate the effects of NREGS by slowly restricting estimation to villages closer and closer to the border. While there are large estimated 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. It is difficult to explicitly test the parallel trends assumption given the lack of data. Nonetheless, 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 estimated 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 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 (Zimmermann, 2012; Imbert and Papp, 2015). 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. Similarly, the effects of the program on household per capita income again appear to be highly spatially heterogeneous: estimated increases in household income towards the interior of treated districts almost completely disappear at the border, consistent with the argument that labor market spillovers are biasing the estimated effects of the program. 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 estimated effects on 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).

The third paper returns to the theme of non-farm self-employment, this time focusing on household labor allocation across production activities. It is commonly assumed that agriculture's share of GDP decreases as a country develops (Lewis, 1954; Ranis and Fei, 1961). This relationship holds both in the cross-section, with relatively more developed countries deriving a smaller percentage of GDP from agricultural sources (Chenery et al., 1975; Gollin et al., 2014), and within countries, where the non-farm sector tends to be more productive – as measured by the average product of labor – than the agricultural sector (Gollin et al., 2014; McCullough, 2017; Young, 2013). However, these findings do not suggest a reallocation of labor is *necessarily* warranted, since a reallocation of labor and its effect on income relate to the marginal product, not the average product. There are many reasons to suggest conclusions using the average product may differ. A large body of research shows that agricultural households diversify into non-farm self-employment for a number of reasons, including production shocks, household shocks, agricultural seasonality, and missing markets (Barrett et al., 2001; Haggblade et al., 2010; Lanjouw and Lanjouw, 2001; Nagler and Naudé, 2014). In addition, diversification seems to be the norm, not the exception (Davis et al., 2017). Under these scenarios, households may be moving into the non-farm sector due to a lack of more remunerative options and a desire to mitigate risk, which may actually lead to a lower marginal product of labor in non-farm production. In this paper, we show that using the marginal product as opposed to the average product leads to substantially different conclusions.

To test the assumption of equality of the marginal revenue products of labor across household activities, we use three waves of the Malawi Integrated Household Survey (IHS). We examine whether the marginal product of labor is equal across agricultural and non-agricultural production within a household. To the best of our knowledge, this is the first paper to implement such a test. We control for unobservable household characteristics using household fixed effects. We estimate a number of specifications, including: separate production functions for plots and non-farm enterprises; a pooled production function; and a collapsed production function. Importantly, we first show that the average product of labor is consistently higher in non-farm production than agricultural production. That is, there is nothing unusual about our Malawian subsample relative to the previous literature. However, in contrast to the conclusions drawn by prior work, we consistently find agricultural MRPL to be higher than non-farm MRPL; in our preferred specification, agricultural MRPL is four times higher than non-farm MRPL. We show that this difference does not appear to be driven by temporal variation in labor allocation; the estimated difference is consistent when we split the sample into high non-farm labor months and low non-farm labor months.

We then show that focusing on a single statistic obscures substantial heterogeneity. Focusing on the MRPL estimates at the household/wave level, we show that a number of production characteristics significantly predict deviations from MRPL equality. We find evidence that price risk plays an important role in household labor allocation. In particular, deviations from equality are much higher for plots planted with tobacco and cotton, two pure cash crops, than for plots planted with maize, a common subsistence crop in Malawi. Interestingly, households that hire for agricultural production have much higher deviations from equality (with agricultural MRPL being higher than non-farm MRPL) than households that do not hire for agricultural productions. However, this relationship is reversed for households that hire for non-farm production. These differences by household type are driven completely by those households that reported hiring, suggesting that households buying labor on the market make substantially different labor allocation decisions than subsistence households.

We also split the sample by output, acreage, and crop sales. For all three variables, households above the median value show substantially larger MRPL deviations than households below the median: for output, we are unable to reject the null hypothesis of MRPL equality for households with below median total output, while the MRPL difference is around twice as large for the upper half of the distribution (relative to the lower half) of both acreage and crop sales. Insofar as these variables are proxies for market access, this evidence is consistent with price risk (Barrett, 1996).

Consistent with production risk, we also find evidence that rainfall variability is associated with higher agricultural MRPL and suggestive evidence that it is also associated with deviations from equality. We only have rainfall variables for the first and second waves; higher rainfall variation (as proxied by the rainfall coefficient of variation, CV) is positively correlated with deviations from equality in these two waves. Moreover, when we assign wave two CV values to wave three households, we find similar deviations, though substantial measurement error hinders our ability to make any inferences. Overall, these results reinforce the commonly held belief that agriculture is risky and decreasing that risk might allow some farmers to increase expected incomes.

This dissertation contributes to three broad bodies of literature. First, this dissertation informs the literature on risk and household production in developing countries. A large body of literature examines the effects of risk on household decisions in developing countries, including two types of decisions particularly germane to the current paper: production and labor allocation (Morduch, 1995; Karlan et al., 2014; Mobarak and Rosenzweig, 2014; Heltberg et al., 2015; Emerick et al., 2016; Cole et al., 2017). Risk often induces households to diversify into less risky activities, which both the first and third papers in this dissertation confirm. The first paper shows households in rural India make a similar choice with respect to non-farm self-employment and that a significant proportion of households completely substitute out of non-farm self-employment when other options for risk management arise. The third paper leads to a similar conclusion; households in rural Malawi also appear to use non-farm production as a risk-diversification strategy, as the median household overallocates labor to non-farm production. The use of non-farm self-employment as a risk-mitigation and coping strategy is a finding echoed in previous research (Reardon, 1997; Barrett et al., 2001; Lanjouw and Lanjouw, 2001; Kijima and Lanjouw, 2005). This result suggests that market failures in developing countries – in this case, the lack of credit and/or insurance markets – tie the hands of many households, inducing them to make decisions that reduce risk but also decrease incomes.

Second, this dissertation contributes to the large literature on the labor-market impacts of NREGS, as well as government labor policies more broadly. Much of the NREGS literature finds that the program increased prevailing wages in rural areas (Azam, 2011; Imbert and Papp, 2015; Deininger et al., 2016). The first essay in this dissertation adds to the NREGS literature by documenting effects on one type of labor – non-farm self-employment – and by showing that these effects appear to be driven by the general equilibrium wage increase and risk reduction effects of the program, not by direct program employment. This finding also supports other research that finds a substantial reallocation of labor following implementation of the program (Islam and Sivasankaran, 2014; Shah and Steinberg, 2015). In addition, the current results suggest the program plays an important role in helping households manage and cope with production risks, which complements similar findings in the literature (Gehrke, 2014; Zimmermann, 2015). The second paper, on the other hand, speaks to the labor market spillovers induced by the program, adding to the literature on spillovers of public policy 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). Overall, the first and second papers explore important general equilibrium effects of NREGS. Insofar as general equilibrium effects are responsible for much of the household-level responses to the program, this may be a partial explanation for the lack of effects found in smaller evaluations of similar programs (Gilligan et al., 2009; Subbarao et al., 2013; Beegle et al., 2017).

Finally, the conclusions of the third essay, especially, are salient to the literature examining the "productivity gap" between the non-farm and agricultural sectors in developing countries (Gollin et al., 2014; McCullough, 2017; Young, 2013). Unlike those studies, which focus on all forms of non-farm and agricultural production, we focus only on agricultural and non-farm household production in households that operate both types of enterprise simultaneously. Nonetheless, we show that this restriction does not affect the productivity gap as commonly defined: the average product of labor is higher in non-farm production than in agricultural production in our sample, as well. However, given that we find the marginal product of labor to be *lower* in non-farm production than agricultural production, our results suggest we revisit how we measure the sectoral productivity gap in developing countries.

Chapter 2

Moving up or Just Surviving? Non-Farm Self-Employment in India

The classical economic development literature argues that growth is accompanied by a reduction in agriculture's share and an increase in non-agriculture's share of employment. Yet, growth of the non-farm sector does not necessarily signal increasing levels of development, as the sector may serve as subsistence employment for many individuals. This ambiguity is heightened by a surprising lack of microevidence regarding sectoral and occupational choice and, especially, how government policies impact these decisions. In this paper, I make a simple observation regarding how non-farm self-employment reacts to market conditions: households and individuals that enter into non-farm self-employment for subsistence reasons are more likely to exit the sector when wages increase or when more stable employment becomes available. With this assumption as a starting point, I examine the effects of the Mahatma Gandhi National Rural Employment Guarantee Act, which increased prevailing wages in rural India. I find that the program significantly decreases days spent in non-farm self-employment. In addition, the implied labor elasticity is almost three time higher than economy-wide estimates, suggesting rural non-farm self-employment is a sector of last resort for many individuals. Additional analyses suggest this impact is driven primarily by two mechanisms: higher wages and alternative options for risk-management.

2.1 Introduction

The design of effective development policy depends on a more comprehensive understanding of how, and why, individuals allocate their labor the way they do. The classical economic development literature, for example, argues that growth is accompanied by workers moving out of the agricultural sector and into the more "modern" industrial sector (Lewis, 1954; Kuznets, 1957; Ranis and Fei, 1961). In line with this view, it is a relatively established empirical finding that the average product of labor in agriculture tends to be much lower than the average product of labor in non-farm production (Young, 2013; Gollin et al., 2014). While large proportions of the labor force remain tied to their lands, recent evidence suggests the non-farm sector is growing in many developing countries (Lanjouw and Shariff, 2004; Lanjouw and Murgai, 2009). Does this recent growth spurt in the rural non-farm sector – and non-farm self-employment, in particular – represent a turning point, a growing non-agricultural sector that can drive further transformations in the rural economy (Matsuyama, 1992; Tiffin and Irz, 2006)? Or, instead, is this growth evidence of increasing stress in rural labor markets, as individuals flock to the non-farm sector as a means of subsistence (Barrett et al., 2001; Lanjouw and Lanjouw, 2001; Kijima and Lanjouw, 2005; Lanjouw and Murgai, 2009; Binswanger, 2012)? The answers to these questions are fundamental to the formulation of efficacious public policy, especially policies that target the rural non-farm sector.

The view of a low-productivity rural non-farm sector has some empirical support. The vast majority of the self-employed in developing countries are poor (Cho et al., 2016). Moreover, microenterprises make up the majority of employment in developing countries and a much larger share of employment than in developed countries (Gindling and Newhouse, 2014). For example, according to the World Bank's 2013 World Development Report, more than 80 percent of employment in services and 60 percent of employment in manufacturing in India is in microenterprises; this share is even higher in other developing countries, like Ethiopia, where upwards of 90 percent of employment in both services and manufacturing is in microenterprises. Moreover, one type of microenterprise is apparently becoming more common: in India, the prevalence of non-farm self-employment among prime-aged adults increased by 40 percent in the earlier 2000s, from an average of around 0.57 days per week in 1999/2000 to 0.80 days per week in 2004/05.[1]

Given the pervasive nature of these enterprises, it is no surprise that increasing their productivity is a common development strategy. In fact, growth of microenterprises, including non-farm self-employment, is one of the many outcomes targeted by the recent microfinance revolution (Armendáriz and Morduch, 2010). However, given that microfinance does not seem to work for all microenterprises (De Mel et al., 2008, 2009; Banerjee et al., 2015a), a more thorough understanding of the drivers of self-employment growth can improve policymaking and targeting. In an attempt to improve our understanding of these drivers, this paper asks whether the non-farm self-employment sector in India is characterized by subsistence entrepreneurship – which, due to low entry barriers, particularly suits poor households coping with shocks, seasonality, and missing markets (Barrett et al., 2001; Lanjouw and Lanjouw, 2001; Haggblade et al., 2010; Nagler and Naudé, 2014) but may offer few opportunities for growth – or by opportunity entrepreneurship, which may offer longer-term prospects for poverty alleviation and further economic growth. While the classical view of economic development is that movements out of agriculture often accompany development and reductions in poverty (Dercon, 2009), it is not clear that this is indeed the case if much of this movement is driven by subsistence concerns.

In this paper, I first make a simple observation regarding the reactions of different types of enterprises: households and individuals that enter into non-farm self-employment for subsistence reasons are more likely to exit the sector when wages increase or when less

[1] Numbers are from the author's calculations, using rounds 55 and 61 of the Indian National Sample Survey.

risky employment opportunities become available. In other words, the labor elasticity of non-farm self-employment with respect to the market wage should be much higher for subsistence enterprises than other firms. Households that enter non-farm self-employment for profit-making reasons,[2] on the other hand, are less likely to exit the sector with small wage increases. With this assumption as a starting point, I examine the effects of India's Mahatma Gandhi National Rural Employment Guarantee Act (NREGS)[3] on the labor allocation decisions of individuals in rural India. NREGS is the largest public works – or "workfare" – program in the world, generating more than two billion person-days of employment in 2015/16.[4] Previous research has found that the program significantly increased prevailing wages in rural areas of India (Azam, 2011; Imbert and Papp, 2015; Deininger et al., 2016) and has helped households cope with shocks and diversify into riskier activities (Gehrke, 2014; Zimmermann, 2015). These findings suggest the program may plausibly lead to the exit of many subsistence enterprises.

Using a differences-in-differences strategy, I first show that the program has a significant effect on non-farm self-employment for prime-age adults. NREGS decreases the probability an individual engages in any non-farm self-employment during the last week by 1.5 percentage points, a more than 10 percent decrease relative to the pre-program mean. In addition, the program decreases days in non-farm self-employment in the last week by 3.0 percent. I then show that the program's effects are concentrated in the dry season, consistent with prior evidence of the program's wage impacts (Imbert and Papp, 2015) and with the de facto seasonal implementation of the program (Lakha and Taneja, 2009; Sharma, 2009a; Sukhtankar, 2016). Consistent with implementation being concentrated in the dry season, there is an insignificant change in days of non-farm self-employment in the rainy

[2]We might think of these households as more reminiscent of *entrepreneurs*.

[3]NREGS was originally passed into law as the National Rural Employment Guarantee Act (NREGS). Its name was later amended to include the name of the Indian independence movement leader. Throughout this paper, I refer to the program as NREGS.

[4]Figures from http://NREGS.nic.in/

season but a large and statistically significant effect in the dry season: NREGS decreases days of non-farm self-employment by 4.4 percent. A back-of-the-envelope calculation suggests a labor supply elasticity of 0.94, which is three times as large as estimates of overall labor elasticity in India (Evenson and Binswanger, 1980; Muralidharan et al., 2017), including estimates calculated using the same data and program as this paper (Imbert and Papp, 2015). This evidence is consistent with non-farm self-employment playing a significant subsistence role during the dry season. I also present results that show pre-program trends are unlikely to be responsible for this finding.

This paper then shows that more prevalent enterprises – by industry – also tend to have lower wages and that enterprises in these industries also show substantial seasonality relative to higher-wage industries. This is consistent with self-employment having low capital requirements – and thus low entry barriers – but also low returns (Haggblade et al., 2010). Also in line with the subsistence enterprise argument, results show that pre-program industry wages predict the effects of the program on industry prevalence in the dry season but not in the rainy season; industries with the lowest wages prior to implementation of NREGS see the largest decrease in days in non-farm self-employment in the dry season, when labor demand is lowest. However, two final sets of results show that non-farm self-employment nonetheless plays an important role for many households. First, the inclusion of days of NREGS employment as a covariate in the main results specification has a very minimal effect on the estimated treatment effect of the program, suggesting non-wage benefits play an important role in the allocation of labor to non-farm self-employment. Second, and more importantly, the effect of NREGS on days in non-farm self-employment is significantly stronger in areas with higher variation in rainfall, which I use as a proxy for agricultural production risk. This evidence supports the argument that individuals in rural India consider non-farm self-employment to be an integral part of risk management and coping strategies.

This paper contributes to several bodies of literature. First, this paper informs the literature on risk and household production in developing countries. A large body of literature examines the effects of risk on household decisions in developing countries, including two types of decisions particularly germane to the current paper: production and labor allocation (Morduch, 1995; Karlan et al., 2014; Mobarak and Rosenzweig, 2014; Heltberg et al., 2015; Emerick et al., 2016; Cole et al., 2017). Risk often induces households to diversify into less risky activities. However, the primary downside of this decision is that low-risk activities also tend to be low-return. As such, households under risk often make the deliberate choice to reduce their expected income. This paper offers empirical evidence that households in rural India make a similar choice with respect to non-farm self-employment and that a significant proportion of households completely substitute out of non-farm self-employment when other options for risk management arise. The use of non-farm self-employment as a risk-mitigation and coping strategy is a finding echoed in previous research (Reardon, 1997; Barrett et al., 2001; Lanjouw and Lanjouw, 2001; Kijima and Lanjouw, 2005; Brummund and Merfeld, 2017). This result suggests that market failures in developing countries – in this case, the lack of credit and/or insurance markets – tie the hands of many households, inducing them to make decisions that reduce risk but also decrease incomes.

Second, this paper contributes to the growing literature on general equilibrium and indirect effects of government policies in developing countries, including the implications for program evaluation (Duflo, 2000; Miguel and Kremer, 2004; Angelucci and De Giorgi, 2009; Cunha et al., 2011; Muralidharan et al., 2016b; Merfeld, 2017b). Much of this literature points to the importance of accounting for indirect effects of public policies. The wage increase induced by NREGS has important general equilibrium effects that, if not accounted for, will bias traditional treatment effect estimates (Ravallion, 2007, 2009). In addition, the fact that NREGS apparently helps households manage and cope with production risk suggests a simple comparison of short-term outcomes – like consumption – may miss important benefits that accrue to households over time as they substitute into alternative forms of employment and even invest in alternative production technologies (Emerick et al., 2016). In this sense, it may be difficult to form a complete picture of these effects with respect to NREGS, as most studies – including this one – necessarily examine relatively short-term impacts using quasi-experimental methods. Nonetheless, even after accounting for these indirect benefits, it may very well be the case that alternative government policies – basic-income schemes, for example – still dominate public works programs in certain contexts (Murgai et al., 2015). In this regard, it is clear that NREGS is displacing some alternative forms of employment and that the opportunity cost of participation in NREGS is not zero.

Finally, this paper contributes to the literature examining the effects and implementation of NREGS and public works programs more broadly. A significant portion of this literature examines the labor market impacts of the program, specifically the effect on wages. Much of this literature finds that the program increased prevailing wages in rural areas (Azam, 2011; Imbert and Papp, 2015; Deininger et al., 2016), although the concentration of effects differs, with some authors finding increases mainly in the rainy season (Imbert and Papp, 2015) and others finding an increase in female wages but not male wages (Azam, 2011). On the other hand, Zimmermann (2015) provides evidence the program serves more of a safety net role, and finds few effects on private-sector wages. I add to the NREGS literature by documenting effects on one type of labor – non-farm self-employment – as well as showing that these effects appear to be driven by the general equilibrium wage increase and risk reduction effects of the program, not by direct program employment. This finding also supports other research that finds a substantial reallocation of labor following implementation of the program (Islam and Sivasankaran, 2014; Shah and Steinberg, 2015). In addition, the current results suggest the program plays an important role in helping households manage and cope with production risks, which complements similar findings in the literature (Gehrke, 2014; Zimmermann, 2015). This finding also speaks to the effects of public works programs more generally, such as programs in Malawi (Beegle et al., 2017), Ethiopia (Gilligan et al., 2009), and elsewhere (see, for instance, Subbarao et al. (2013) for an overview of existing evidence).

2.2 Institutional Setting

Creating employment opportunities and raising wages are key instruments in the fight against poverty (Devereux and Solomon, 2006). While governments use a number of different policies in an attempt to increase employment, perhaps none is as ambitious as India's NREGS. Unlike other workfare programs, NREGS treats employment as a fundamental right, with "workers' rights to demand and get work... enshrined in a legislative framework" (**?**, p. 6).

NREGS was originally passed into law as the National Rural Employment Guarantee Act in 2005 (**?**).[5] It is designed as a legal guarantee of employment to households in rural areas. Every rural household is entitled to up to 100 days of manual labor at the statutory minimum wage, which was, on average, approximately 2 USD per day in 2011 (Azam, 2011), although wages vary by state. According to a presentation from one senior adviser to the government of India,[6] at the start of the program, the official NREGS wage rate varied between 50 rupees (in Gujarat) and 125 rupees (in Kerala). Notably, these official wages often more than doubled the prevailing wage rate at the time for men. For women, the difference was even greater, with the same presentation noting that the prevailing wage rate for women in Rajasthan was only 10 rupees a day prior to implementation of the program,

[5]All information in this section comes from the original legislation, unless otherwise cited.
[6]http://www.levyinstitute.org/pubs/EFFE/Mehrotra_Rio_May9_08.pdf

while official NREGS wages were 73 rupees a day. Although the wages were higher than the prevailing wage rates in many states, they are nonetheless designed as a self-targeting mechanism; since the wages are still relatively low, only the poorest households should be interested in applying for jobs through the program (Dutta et al., 2012).

While the program is open to all households in rural areas, interested households must apply for a job card. The job cards are administered by the Gram Panchayat, which is the lowest level of administration in the Indian administrative structure. After receiving a job card, households can apply to the local government for work at any time during the year and must be assigned to a job within 15 days of requesting employment. If applicants do not receive a job within 15 days, they are eligible to receive unemployment compensation – which is equal to between one-fourth and one-half of the official wage rate – until a job is secured. The costs of the program are shared between the federal and state governments, although the federal government bears the majority of costs. The federal government is responsible for all wages of unskilled laborers and three-quarters of all costs for materials as well as skilled and semi-skilled workers. The state government, on the other hand, is responsible for the remaining one-quarter of costs for materials and is completely responsible for unemployment compensation.

A key issue in implementation of the program concerns late payment of wages, which by law must be paid within 15 days. In one study, more than a third of all respondents reported having to look for work elsewhere due to delays in being paid (Khawlneikim and Mital, 2016), while in another study, every single woman interviewed at worksites in Uttar Pradesh reported wages had not been paid on time (Khera and Nayak, 2009). Insofar as women are more likely to require prompt payment – perhaps because they are the sole wage-earners in the household or they rely on the wages to feed their children – these delays may be leading women to look for opportunities eslewhere (Khera and Nayak, 2009).

Actual implementation of the law took place in three phases. In the first phase, 200 districts received the program in February of 2006. An additional 130 districts received the program between April and May of 2007, while the program was implemented in the remaining districts in April of 2008, bringing the total number of districts with the program to 644 (**?**). Figure 2.1 shows the distribution of districts and phases across India. There is some spatial clustering of phase-one districts – represented by green – in the eastern half of country. This is not surprising, as the selection of districts was done based on a number of characteristics, including poverty. According to government documents, the five main variables used to select districts were agricultural productivity per worker, agricultural wage rates, population of scheduled castes and scheduled tribes, agricultural productivity per hectare, and the poverty rate (Planning Commission, 2003a). As such, many of the poorest districts were the first to receive the program, and many of these districts are located in the eastern half of the country. It is not clear how this would be correlated with non-farm self-employment, but, given this purposive selection, differences-in-differences estimates – comparing differences in trends – are more likely to provide valid causal inference than a simple comparison of early- and late-phase districts. Nonetheless, given the purposive selection of districts, the parallel trends assumption may be violated. As such, I present evidence in support of this assumption below.

2.3 The Model

In this section, I develop a simple model and derive three predictions regarding the effects of NREGS on non-farm self-employment. Suppose a single-person household maximizes its utility of consumption and leisure and, for simplicity, assume the household has only three income sources: farm self-employment, non-farm self-employment, and wage em-

Figure 2.1: Districts and NREGS Phase

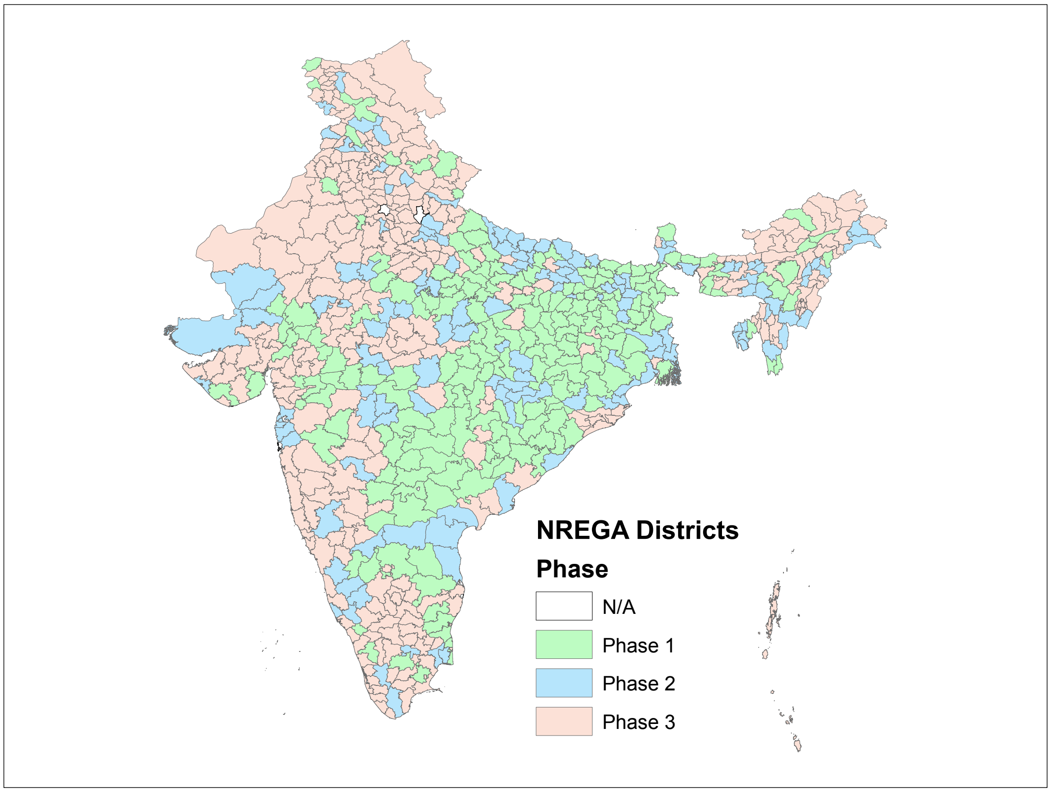

Phase one districts received the program in 2006. Phase two districts received the program in 2007. Phase three districts received the program in 2008.

ployment.[7] In a static framework, the household's problem is:

$$\max_{L_F, L_N, L_w, c, l} u(c, l)$$

$$\text{subject to} \qquad c \leq p_F f(L_F; A) + p_N g(L_N; Z) + wL_w$$

$$\bar{L} = L_F + L_N + L_w + l$$

where c is consumption; p_F is the price of agricultural output; L_F is agricultural labor and $f(L_F; A)$ is the agricultural production function, satisfying $f'(L_F) > 0$ and $f''(L_F) < 0$; A is a measure of capital (e.g. land), assumed to be fixed; p_N is the price of the good produced through non-farm self-employment; L_N is non-farm enterprise labor and $g(L_N; Z)$ is the non-farm production function, satisfying the same assumptions as the agricultural production function; Z is capital used in non-farm production, also assumed to be fixed; w is the wage rate; L_w is amount of wage labor; and $\bar{L}$ is the household's total labor endowment. For simplicity, I assume there is no cap on wage labor.[8]

Assuming an interior solution, the optimal decision is given by:

$$w = p_N g'(L_N) = p_F f'(L_F; A) \tag{2.1}$$

In other words, the household equates returns to non-farm production – the marginal revenue product of labor, or MRPL – with the wage rate. Given the assumptions on the production functions, an increase in the prevailing wage rate will result in a decrease in the amount of labor the household applies to non-farm self-employment, or $\frac{dL_N}{dw} < 0$. The NREGS literature has consistently found that the program increased prevailing wage rates in rural areas. This leads to the first result of the model:

[7] I also assume separability. However, assuming non-separability does not change the comparative statics.
[8] All of the results in this section hold with the inclusion of a cap on wage labor.

Result 1 *Households will allocate less labor to their non-farm enterprises following implementation of NREGS.*

Two corner solutions are possible with respect to non-farm labor. First, it is possible that some households will prefer to allocate all of their labor to non-farm self-employment. In particular, this will be the case if $p_N g'(\bar{L}) > w$ and $p_N g'(\bar{L}) > p_F f'(0)$.[9] If this is true after implementation of NREGS, then we will see no effect of NREGS on these households. Second, some households may prefer to allocate no labor to non-farm employment. For these households, it must be true that $p_N g'(0) < w$ or $p_N g'(0) < p_F f'(\bar{L})$. This leads to the second prediction of the model:

Result 2 *Fewer households will operate non-farm enterprises following implementation of NREGS.*

Households for which $w_{pre} < p_N f'(0)$ and $\bar{w} > p_N f'(0)$ – where w_{pre} is the prevailing wage rate prior to implementation of NREGS and $\bar{w}$ is the wage rate after implementation – will allocate some positive amount of labor to self-employment prior to NREGS but no labor to self-employment after implementation. This change will appear as a closure of a non-farm enterprise in the data.

The first result also leads to another prediction. If households allocate less labor to non-farm production and the production function is concave, the marginal product of labor will increase, or $\frac{dMRPL}{dL_N} < 0$. Since $\frac{dL_N}{dw} < 0$, then, at the new optimum: $\frac{dMRPL}{dw} = \frac{dMRPL}{dL_N}\frac{dL_N}{dw} > 0$. This leads to the final prediction of the model:

Result 3 *The marginal (and average) revenue product of labor will be higher for surviving enterprises following implementation of NREGS.*

[9] We can also allow for an optimum in which the household allocates labor to both non-farm and farm production, but not wage employment. This will be true if $p_N g'(L_N^*) > w$ and $p_F f'(\bar{L} - L_N^*) > w$.

It is important to note that Result 3 is really due to two separate effects. First, households will allocate less labor to non-farm production. This suggests that MRPL will be higher in these enterprises following implementation of NREGS, and thus the overall average productivity will be higher. Call this the “labor substitution” effect. Second, Result 2 suggests low productivity enterprises will close completely. This effect will also increase average productivity in NREGS districts. Call this the “composition” effect.

Figure 2.2 presents graphical examples of the model. In this graphical depiction of the model, I ignore farm labor for simplicity.[10] Given the assumptions on the production function, the MRPL curves will be downward sloping with the shape of the curves in Figure 2.2. Assuming fixed capital, each row depicts a different possible marginal revenue product of labor curve for a household enterprise. The left column is the household’s labor allocation to its non-farm enterprise prior to implementation of NREGS, while the right column is the household’s allocation decision following implementation.

Figure 2.2 presents three separate households. Row A shows a household that is completely unaffected by NREGS. The household’s MRPL in non-farm self-employment is high enough that the household will not reallocate any of its labor following the increase in the wage rate. This household’s estimated MRPL at the optimum is thus unaffected. Row B shows an example of the labor substitution effect. The household applies labor up to $\bar{L}$ to non-farm self-employment prior to implementation of NREGS, but reduces its allocation of labor to a following implementation. Assuming a concave production function, this implies that MRPL at the optimum for household B will increase following implementation.[11] Finally, row C shows a household that is completely crowded out by the program. While the household applies labor up to a to non-farm self-employment prior to the program, the wage change is such that the household applies no labor to non-farm self-employment

[10] As above, the results hold if we redefine the labor endowment in Figure 2.2 to be non-farm labor.
[11] This increase is caused by a shift in labor allocation, not a change in the production function itself.

Figure 2.2: NREGS and Marginal Revenue Product of Labor

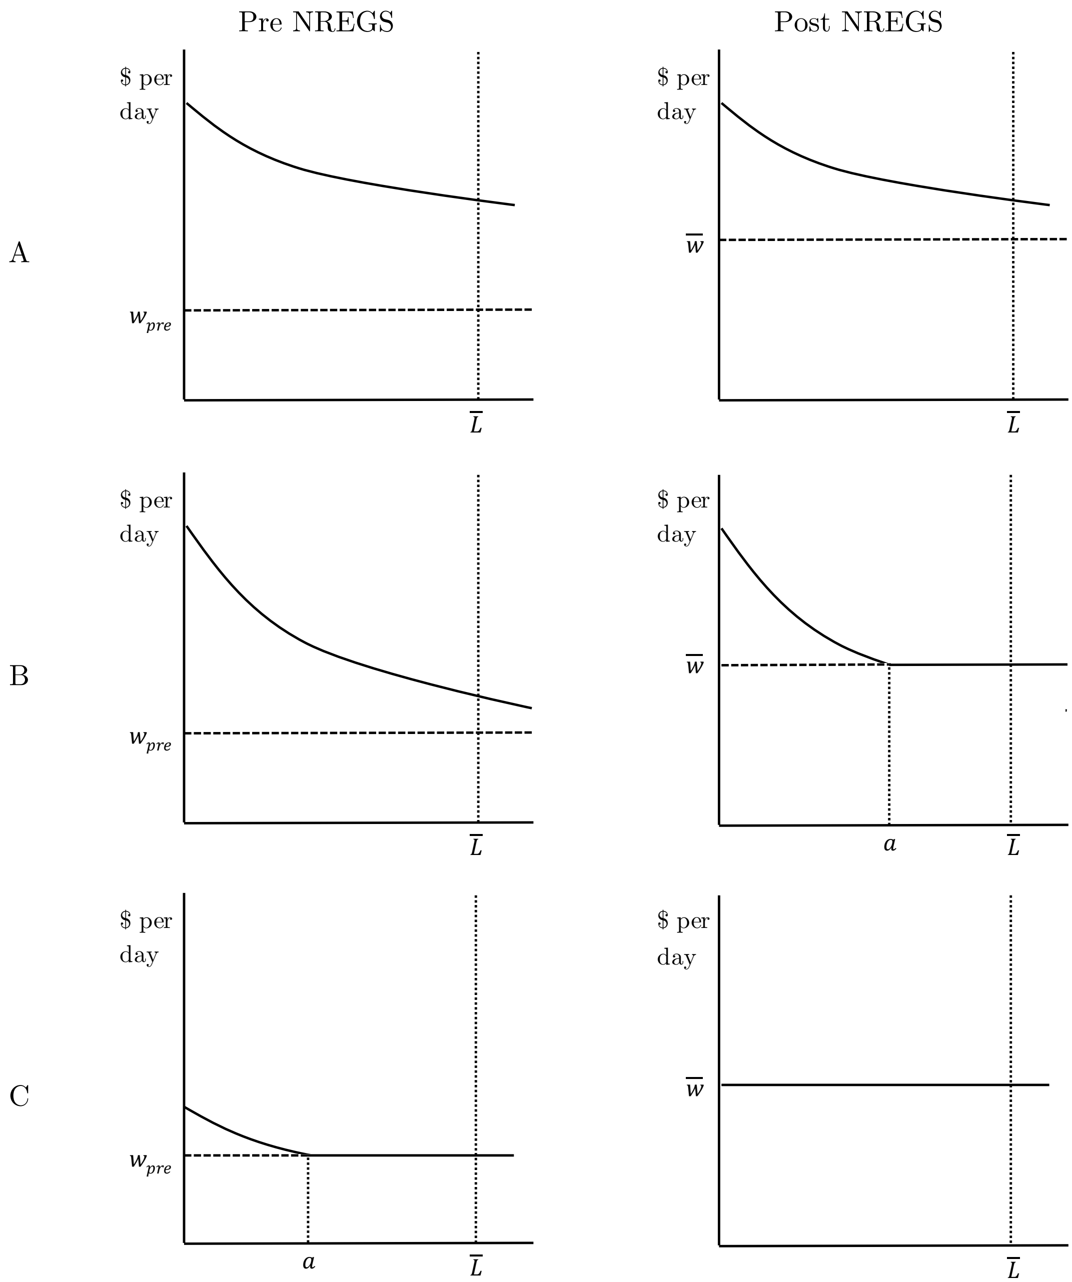

The left column represents non-farm self-employment for three separate example households prior to NREGS implementation. The right column shows the same three households after implementation of NREGS. Row A shows a household that is unaffected by NREGS implementation; MRPL is higher than the post-NREGS wage rate and, as such, the household does not reallocate any labor. Row B shows a household that applies less labor to non-farm self-employment following implementation but continues to operate the non-farm enterprise. Row C shows a household that is completely crowded out of the non-farm self-employment sector by NREGS; the post-NREGS wage rate is high enough that the household allocates no labor to non-farm self-employment following implementation.

following implementation. Since the household was less productive – as measured by its MRPL – than household B prior to the program, overall population MRPL will increase when household C reallocates labor. This is the composition effect.

2.4 Data and Estimation Strategy

In this paper, I predominately use the Indian National Sample Survey (NSS). The survey is collected on an annual basis by the Ministry of Statistics and Program Implementation. However, not all of the survey rounds are "thick" rounds, with large sample sizes. I use rounds 61 and 64, collected in 2004/05 and 2007/08, respectively, which are the same rounds used by Azam (2011) and Imbert and Papp (2015). The surveys are repeated cross-sections and are representative of the entire population of India. In order to make the results nationally representative, I utilize the NSS survey weights in all of the analyses that follow.

The first round I use comes prior to implementation of the program and the second round comes after the first two phases have been implemented, but prior to implementation of the third phase.[12] Given the timing of the two rounds, I employ a differences-in-differences estimator. I estimate regressions of the form:

$$\begin{aligned} y_{hdt} = \alpha_d + \beta_1 Post_t + \beta_2 NREGS_d \times Post_t + \phi_{hdt} + \\ \gamma_d \times Post_t + \eta_{st} + \psi_{mt} + \varepsilon_{hdt}, \end{aligned}$$

where y_{hdt} is the outcome of interest for household h in district d at time t, $Post_t$ is a dummy variable indicating that the observation is from the 2007/08 round, ϕ_{hdt} is a vector of household-level controls, η_{st} is state/wave fixed effects, ψ_{mt} is a vector of dummy variables indicating the month and year of the interview, and ε_{hdt} is a household-specific

[12] Unfortunately, there is no thick round between the first and second phase.

error term. In addition, all estimates include district-level fixed effects (α_d). I also include district-level variables measured prior to the program; these are listed in Table 2.A1. I allow these variables to affect the outcome trend by including them as an interaction with the post dummy ($\gamma_d \times Post_t$). I cluster all standard errors at the district level, allowing for arbitrary correlation across time (Bertrand et al., 2002).

I also use the Additional Rural Incomes Survey/Rural Economic & Demographic Survey, known as ARIS/REDS, which has been collected periodically since 1969.[13] The surveys were originally nationally representative of the rural population of India and covered 4,527 households from 259 villages. The two most recent surveys are from 1999 and 2008 but are unfortunately no longer nationally representative. The 2008 wave resurveyed all households from 1999 plus approximately eight new households in each village, with a total sample of size of 9,500. Unlike the NSS, the ARIS/REDS contains information on remuneration of non-farm self-employment. I use the ARIS/REDS to examine productivity of non-farm self-employment and the differences-in-differences estimator using the ARIS/REDS identifies the same effect as the NSS as it compares phases one and two to phase three. The ARIS/REDS also has a village-level module that provides information on distance to the nearest bank, which I use to explore effect heterogeneity of NREGS. In addition, geospatial information at the household level facilitates accurate measures of rainfall, which I use to explore heterogeneity in NREGS' impacts.

As with all differences-in-differences estimators, identification relies on the assumption of parallel trends: that the trends in outcomes between treatment districts and comparison districts would have been the same in the absence of the program. In practice, this assumption could be violated if, for example, other policies were implemented in only phase one districts between the two survey rounds or if treated and untreated districts were trending in

[13]The survey is organized by the National Council of Applied Economic Research (NCAER). NCAER has collected four rounds between 1971 and 2008. See http://www.ncaer.org/data_details.php?dID=9.

different directions prior to implementation of the program. I present supporting evidence for the assumption of parallel trends in subsection 2.5.1.

I also estimate robustness checks comparing the effects of NREGS across seasons and states. First, lower-return non-farm self-employment may be more common in the dry season when there are fewer employment options (Barrett et al., 2001; Haggblade et al., 2010; Lanjouw and Lanjouw, 2001; Nagler and Naudé, 2014). In addition, previous NREGS research has found differential effects of the program on wages depending on the season (Imbert and Papp, 2015; Zimmermann, 2012). As such, I analyze the effects of NREGS across the rainy and dry seasons, by estimating separate regressions for each season. In all regressions, the dry season is defined as January through June and the rainy season is defined as July through December.

As an additional check, and again following Imbert and Papp (2015), I create a variable for the "star states" that identifies the best-performing states. These states are Andhra Pradesh, Chhattisgarh, Himachal Pradesh, Madhya Pradesh, Tamil Nadu, Rajasthan, and Uttarakhand. If the effects I find are indeed due to the program, then effects may be concentrated in star states.

2.4.1 Main Outcome Variables

NREGS may impact several outcomes related to non-farm self-employment. Most importantly, the program may impact the number of days in non-farm self-employment. The NSS employment module inquires about labor allocation in the last week. Each respondent can report labor allocation in half-day increments. I define self-employment as individuals that responded they either worked in a household enterprise as an own account worker, as an employer, or as an unpaid family worker. In addition, I define non-farm as any industry other than agriculture, livestock, fishing, or mining. Days spent in non-farm self-employment ranges from zero to seven. For the analyses with days below, I use the log of days plus one as the dependent variable. I also present results using levels in the appendix. I create a second non-farm self-employment variable ("non-farm worker") that equals one if the individual reported spending any time in non-farm self-employment and zero otherwise. Finally, in the first set of results below, I collapse responses to the household level, by summing total days and taking the count and max of non-farm worker.

ARIS/REDS allows an examination of the effects of NREGS on non-farm enterprise productivity. Rather than estimating production functions, I define productivity using (log of) average revenue per worker-day.[14] In these analyses, I aggregate these variables to the household level.

2.4.2 Summary Statistics

Table 2.1 presents (unweighted) district summary statistics by NREGS status. Table 2.A1 in the appendix lists the sources for each of the variables in Table 2.1. Treated (NREGS) and untreated (non-NREGS) districts appear geographically similar. For example, both types of districts are approximately of the same size in population as well as area. However, other statistics make clear that districts receiving the program first are, on average, poorer and worse off than districts that receive it later. For example, the average wage in NREGS districts was almost 25 percent lower than the wage in non-NREGS districts. This may be partially explained by the composition of the workforce: workers in NREGS districts are more likely to be engaged in agriculture (cultivators and agricultural workers) than workers in non-NREGS districts.

Three other statistics also display the difference in wealth across districts. First, NREGS

[14] If production follows a Cobb-Douglas, then the marginal and average products of labor are proportional: $MPL = (1-\alpha)APL$, where MPL is the marginal product of labor and APL is the average product of labor.

Table 2.1: District Summary Statistics

	Pre-Program	
	(1)	(2)
	NREGS	Non-NREGS
Population (thousands)	1749.684	1765.740
	(1354.585)	(1255.065)
Percent rural	0.840	0.716
	(0.112)	(0.168)
Area (km^2)	5958.587	5717.904
	(4659.635)	(5422.598)
Average wage (R's)	49.406	65.297
	(22.912)	(39.717)
Percent scheduled castes	0.145	0.145
	(0.084)	(0.083)
Percent scheduled tribes	0.218	0.114
	(0.271)	(0.244)
Sex ratio	944.562	928.673
	(46.365)	(63.584)
Literacy rate	0.495	0.582
	(0.118)	(0.103)
Labor force participation rate	0.414	0.405
	(0.068)	(0.072)
Cultivators (percent)	0.166	0.150
	(0.084)	(0.095)
Agricultural workers (percent)	0.124	0.076
	(0.062)	(0.051)
Household industry workers (percent)	0.015	0.016
	(0.013)	(0.014)
Other workers (percent)	0.110	0.166
	(0.053)	(0.062)
Observations	288	228

Standard deviations are in parentheses. All statistics are unweighted and are from the 2004/05 National Sample Survey, prior to implementation of NREGS. Table 2.A1 of the appendix details the source for each variable. The first column includes all NREGS districts (phase one and two districts) and the second column includes all non-NREGS districts (phase three districts). Wage is in rupees and is defined as the median wage in 2004/05. The sex ratio is defined as number of females per 1000 males.

Table 2.2: Individual Summary Statistics

	Pre-Program (2004/05)	
	(1)	(2)
	NREGS	Non-NREGS
Wage (R's)	48.467	63.649
	(25.764)	(43.841)
Non-farm worker (last week - yes=1)	0.141	0.123
	(0.348)	(0.329)
Days non-farm self-employment (last week)	0.762	0.688
	(2.081)	(2.009)
Male	0.503	0.505
	(0.500)	(0.500)
Age	32.482	32.463
	(11.998)	(12.102)
Education	2.275	2.508
	(1.491)	(1.526)
Household head male (yes=1)	0.929	0.920
	(0.257)	(0.271)
Household head age	47.004	48.185
	(12.160)	(12.375)
Household head education	2.148	2.244
	(1.503)	(1.509)
Household size	6.084	6.098
	(3.125)	(2.942)
Percent children	0.281	0.261
	(0.213)	(0.211)
Percent prime male	0.342	0.351
	(0.174)	(0.179)
Percent prime female	0.324	0.329
	(0.145)	(0.148)
Percent elderly female	0.025	0.028
	(0.064)	(0.067)
Observations	120618	91881

Standard deviations are in parentheses. All statistics are unweighted and are from the 2004/05 National Sample Survey, prior to implementation of NREGS. The first column includes all NREGS districts (phase one and two districts) and the second column includes all non-NREGS districts (phase three districts). Wage is in rupees. There are 16379 wage observations from NREGS districts and 11514 wage observations from non-NREGS districts. Education is coded from 1-6: 1 indicates less than a primary education; 2 indicates a primary education; 3 indicates a middle-school education; 4 indicates a secondary education; 5 indicates a higher-secondary education; and 6 indicates a college degree. The percent variables following household size are the composition of the individual's household. Children are defined as individuals under 15; prime males are defined as males between 15 and 60 years of age; prime females are defined as females of the same age; and elderly females are defined as females 60 or older.

districts are, on average, much more rural than non-NREGS districts. Second, a higher percentage of the population of NREGS districts is made up of scheduled tribes and castes – some of the most disadvantaged castes in India – than non-NREGS districts. Finally, the literacy rate is much higher in non-NREGS districts: while more than 58 percent of adults in non-NREGS districts are literate, less than half are literate in NREGS districts. On the whole, these statistics suggest that we must be careful interpreting the results from differences-in-differences, as pre-program trends may be quite different.

Table 2.2 presents (unweighted) summary statistics at the individual level from the NSS, for individuals between 15 and 60 years of age. Again, we see that the wage was much lower in NREGS districts prior to the program. Regarding the two main outcomes of interest, individuals in NREGS districts are slightly more likely to have worked in non-farm self-employment during the last week. Similarly, they worked, on average, slightly more days in non-farm self-employment than individuals in non-NREGS districts. Many of the other variables do not suggest large differences by district type, except education, which is slightly higher in non-NREGS districts and consistent with the story above .[15]

2.5 Results

Table 2.3 presents the estimated effect of NREGS on non-farm self-employment at the household level. All three coefficients are in the hypothesized direction and are significant or marginally significant. In the data, households decrease their total days of non-farm self-employment by approximately 5.9 percent. In addition, approximately seven percent fewer households (0.02 divided by the pre-program average of 0.288) operate a non-farm

[15]Table 2.A3 in the appendix presents unweighted summary statistics for all individuals over the age of ten. The statistics suggest that individuals younger than 15 and older than 60 are much less likely to engage in non-farm self-employment, with the average days in non-farm self-employment almost 20 percent lower than in Table 2.2. This motivates the restriction of the sample to only individuals between 15 and 60 years of age.

Table 2.3: Effect of NREGS on Household Non-Farm Self-Employment

	Household level		
	(1)	(2)	(3)
	Self-emp (days - log)	Self-emp (count)	NFE HH
Post times NREGS	−0.059*	−0.044	−0.020*
	(0.033)	(0.029)	(0.012)
State/Wave	Yes	Yes	Yes
Interview Month/Wave	Yes	Yes	Yes
Observations	141049	141049	141049

Standard errors are in parentheses and are clustered at the district level. All regressions are defined at the household level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. Column one includes total number of days (log + 1) of non-farm self-employment all household members reported engaging in during the previous week. Column two includes total number of household members that reported engaging in any non-farm self-employment during the previous week. The dependent variable in column three equals one if any household member reported engaging in non-farm self-employment during the previous week.

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

enterprise, relative to the pre-program mean, following NREGS implementation.

Table 2.4 presents results at the individual level. The first column is log of days in self-employment over the last week. We see a decrease in days of non-farm self-employment of approximately 3.0 percent among individuals in NREGS districts relative to individuals in non-NREGS districts. Similarly, treated individuals are 1.5 percentage points less likely to report having worked in non-farm self-employment during the last week. While the coefficient is relatively small, it is nonetheless qualitatively meaningful; it represents an effect more than ten percent as large as the dependent variable mean prior to implementation of the program. Both of the effects in Table 2.4 are statistically significant.[16]

We might expect the effects of NREGS to differ by season for two reasons. First, non-farm self-employment may show seasonality. Second, the wage effect of NREGS shows seasonal variability (Berg et al., 2014; Imbert and Papp, 2015; Deininger et al., 2016), perhaps due to elite pressures to only implement the program during times of low labor demand (Lakha and Taneja, 2009; Sharma, 2009a; Sukhtankar, 2016). If NREGS is well

[16]Table 2.A4 in the appendix also presents results for days of non-farm self-employment in levels. The main results are qualitatively unchanged and, if anything, the effect of NREGS appears to be larger.

Table 2.4: Effect of NREGS on Individual Non-Farm Self-Employment

	Individual level	
	(1)	(2)
	Self-emp (log days)	Self-emp (Yes=1)
Post times NREGS	−0.030**	−0.015*
	(0.015)	(0.008)
State/Wave	Yes	Yes
Interview Month/Wave	Yes	Yes
Observations	366116	366116

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. Column one includes total number of days (log + 1) of non-farm self-employment the individual reported engaging in during the previous week. The dependent variable in column two equals one if the individual reported engaging in any non-farm self-employment during the previous week.

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

implemented only in the dry season and if non-farm self-employment is more likely to be subsistence during the dry season, we should see relatively small effects of the program during the rainy season but large effects during the dry season. The first column of Table 2.5 presents the results disaggregated by season. Column one of Panel A confirms expectations: NREGS has no effect on non-farm self-employment in the rainy season. However, column two shows a large and highly significant decrease during the dry season. The point estimate indicates that non-farm self-employment decreased by 4.4 percent in NREGS districts relative to non-NREGS districts during the dry season. When estimated in levels, the effect is much larger; Appendix Table 2.A4 suggests a drop of around one quarter relative to the pre-program mean.

This finding is consistent with non-farm self-employment playing different roles during the rainy and dry seasons. In particular, the results are indicative of non-farm self-employment being less productive in the dry season than the rainy season. However, if wages increase more in the dry season than in the rainy season (Imbert and Papp, 2015), sector of last result need not be the explanation, a point to which I return below. Imbert and Papp (2015), using the same data, find an increase in the wage rate of approximately 4.7

Table 2.5: Effect of NREGS by Season and State

Non-Farm Self-Employment (days - log)	(1) Dry/Rainy	(2) Star State
Post times NREGS (Rainy)	-0.015	
	(0.016)	
Post times NREGS (Dry)	-0.044**	
	(0.018)	
Post times NREGS (Non-Star)		-0.016
		(0.015)
Post times NREGS (Star)		-0.032
		(0.021)
State/Wave	Yes	No
Interview Month/Wave	No	Yes
Observations	366933	366876

The dependent variable in all columns is log of days of non-farm self-employment plus one. Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. In all regressions, the dependent variable is days (log + 1) of non-farm self-employment during the previous week. The top panel reports effects by season. The dry season is defined as January through June while the rainy season is defined as July through December. The bottom panel reports effects by Star state status. Star states are defined as Andhra Pradesh, Chattisgarh, Himachal Pradesh, Madhya Pradesh, Rajasthan, Uttarkhand and Tamil Nadu (Imbert and Papp, 2015).

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

percent and an implied labor elasticity of 0.31 in the dry season. Their estimate of labor elasticity is consistent with both earlier (Evenson and Binswanger, 1980) and later (Muralidharan et al., 2017) estimates for rural India. Using their estimated wage increase of 4.7 percent and the estimated decrease in non-farm self-employment of 4.4 percent in Table 2.5, the implied labor elasticity of non-farm self-employment is 0.94, much higher than the overall labor elasticity in the economy, even using estimates from the same program and with the same data (Imbert and Papp, 2015). Non-farm self-employment is apparently much more sensitive to wage changes in the dry season than is overall employment, which is consistent with at least part of the sector being subsistence entrepreneurship, or a sector of last resort.

The second column of Table 2.5 disaggregates the effect into non-star states and star states (Imbert and Papp, 2015). Consistent with previous evidence, the effect of NREGS on non-farm self-employment is greatest in the states that best implemented the program, the "star states." However, none of the results nor the difference across states is significant.[17]

Table 2.A5 and Table 2.A6 in the appendix present the individual-level results for all individuals over 10 years of age. The overall effect is smaller in the larger sample, which is not surprising given that the newly included individuals were much less likely to engage in non-farm self-employment prior to implementation of the program, as shown in the summary statistics in Table 2.A3. However, effect seasonality is still clear in the more inclusive sample: the largest decrease in non-farm self-employment comes in the dry season. As an additional check, I also re-estimate the overall effect and effects by season using ordered probit, since the labor dependent variable can take on only 15 ordered values at the individual level (zero to seven days in half-day increments). The results in Appendix Table 2.A7 are qualitatively similar to OLS estimates.

2.5.1 Pre-Program Trends

While the previous section presented evidence that NREGS decreased the prevalence of non-farm self-employment, the causal interpretation requires an assumption of common trends. To examine the plausibility of this assumption, Figure 2.3 and Figure 2.4 graphically present pre-program trends using the 1999/2000 NSS (round 55) and 2004/05 NSS (round 61). Figure 2.3 presents trends in days of non-farm self-employment.[18] Non-farm self-employment is increasing slightly in both NREGS and non-NREGS districts between 1999/2000 and 2004/05, which is consistent with the non-farm sector helping to absorb an increasing rural labor force (Binswanger, 2012). However, days in non-farm self-employment in treated districts are not rising as quickly as in untreated districts, a finding which might invalidate the causal interpretation above. Figure 2.4 presents the same

[17] In results not shown but available upon request, I test for significance by completely interacting a star state dummy with all other variables in the model.

[18] As in the analyses above, the dependent variable is log of days plus one.

Figure 2.3: Pre-Program Trends in Days of Non-Farm Self-Employment

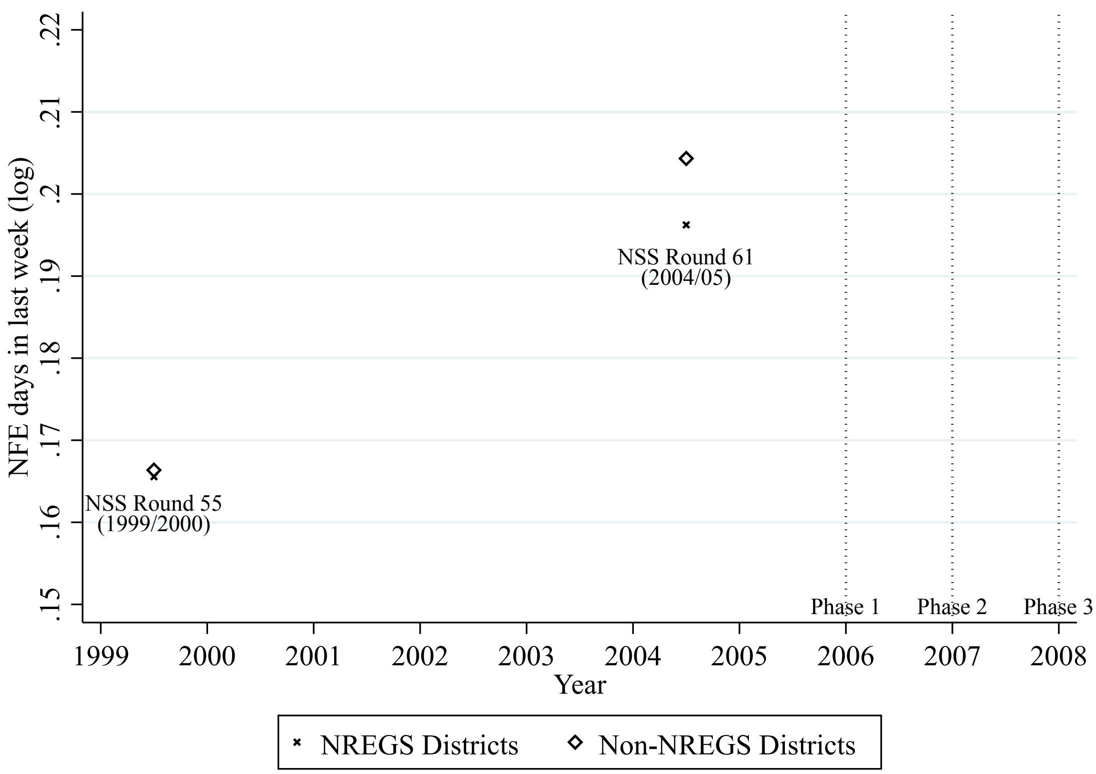

The y-axis measures average non-farm self-employment days over the last week across individuals in NREGS and non-NREGS districts. All statistics are weighted to be representative of the population.

trend but disaggregates the trend into the rainy and dry season. Recall that we expect to see effects of NREGS only during the dry season. In fact, Figure 2.4 suggests that treated districts are actually trending slightly upwards relative to untreated districts during the dry season, which is in the opposite direction of the expected and empirical result. Most of the negative trend in Figure 2.3 is due to the rainy season. If this is true, pre-program trends are unlikely to explain the effects of NREGS on non-farm self-employment found in the previous section.

To empirically analyze the magnitude of these trends, I re-estimate the regressions from Table 2.5 but use 2004/05 as the "post" wave and 1999/2000 as the pre-program wave. If

Figure 2.4: Pre-Program Trends in Days of Non-Farm Self-Employment by Season

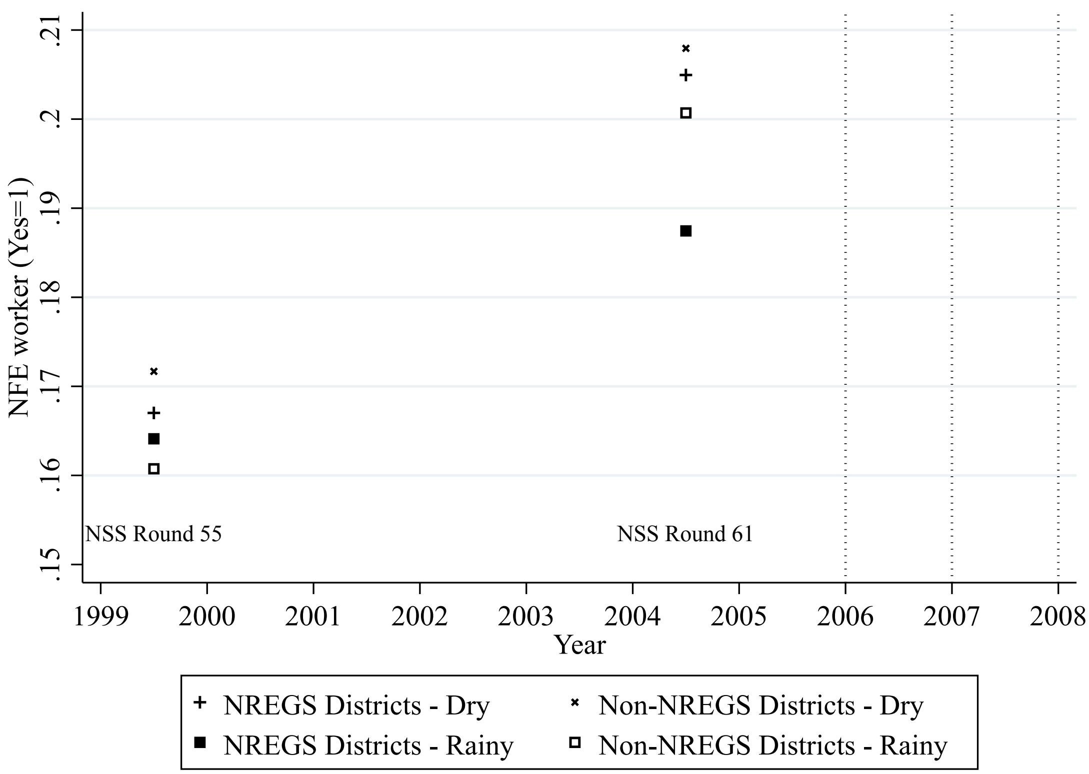

The y-axis measures average non-farm self-employment days over the last week across individuals in NREGS and non-NREGS districts. All statistics are weighted to be representative of the population.

Table 2.6: Effect of NREGS on Non-Farm Self-Employment - Placebo

	Individual level	
	(1)	(2)
Non-Farm Self-Employment (days - log)	Self-emp (log days)	Self-emp (Yes=1)
Post times NREGS	-0.011	-0.005
	(0.011)	(0.006)
State/Wave	Yes	Yes
Subround/Wave	Yes	Yes
Observations	325332	325332

The dependent variable in all columns is log of days of non-farm self-employment plus one. Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. The estimates include the 55th wave of the NSS–collected in 1999/2000–and the 61st wave–collected in 2004/05. Post is now defined as 2004/05, before the program was implemented, in order to assess whether pre-program outcomes are trending differentially. Column one includes total number of days (log + 1) of non-farm self-employment the individual reported engaging in during the previous week. The dependent variable in column two equals one if the individual reported engaging in any non-farm self-employment during the previous week.

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

Table 2.7: Effect of NREGS by Season and State - Placebo

	(1)	(2)
Non-Farm Self-Employment (days - log)	Dry/Rainy	Star State
Post times NREGS (Rainy)	-0.015	
	(0.013)	
Post times NREGS (Dry)	-0.007	
	(0.015)	
Post times NREGS (Non-Star)		-0.017
		(0.011)
Post times NREGS (Star)		0.008
		(0.017)
State/Wave	Yes	No
Interview Subround/Wave	No	Yes
Observations	325332	325332

The dependent variable in all columns is log of days of non-farm self-employment plus one. Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. The estimates include the 55th wave of the NSS–collected in 1999/2000–and the 61st wave–collected in 2004/05. Post is now defined as 2004/05, before the program was implemented, in order to assess whether pre-program outcomes are trending differentially. In all regressions, the dependent variable is days (log + 1) of non-farm self-employment during the previous week. The top panel reports effects by season. The dry season is defined as January through June while the rainy season is defined as July through December. The bottom panel reports effects by star state status. Star states are defined as Andhra Pradesh, Chattisgarh, Himachal Pradesh, Madhya Pradesh, Rajasthan, Uttarkhand and Tamil Nadu (Imbert and Papp, 2015).

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

differential trends are driving the results, then we should see similar results in this placebo analysis. Table 2.7 presents overall results. In the rainy season, the pre-program trend is the same magnitude as the estimated program effect. Although the effect during the rainy season in the previous section was insignificant, the pre-program trends nonetheless suggest the actual effect of the program during the rainy season is near zero, which would be consistent with NREGS being implemented mostly during the dry season. The estimated pre-program trend during the dry season, where we expect a negative coefficient, is slightly negative. However, the coefficient is only 15 percent as large as the estimated program effect in the previous section. As such, it appears unlikely that pre-program trends are responsible for the large estimated effect of NREGS in the dry season.

2.5.2 Enterprise Revenue and Productivity

While the NSS allows an examination of overall employment, it does not allow an examination of the productivity effects of NREGS. As such, I now turn to the ARIS/REDS. Table 2.8 uses the ARIS/REDS data, at the household level, to look at four separate dependent variables: 1) whether the household engages in non-farm self-employment; 2) the total number of workers employed; 3) total family workers employed; and 4) revenue per worker-day. Columns one through four confirm the NSS estimates: households are less likely to operate a non-farm enterprise and employ fewer overall workers when they do.

The final column is log of revenue per worker-day for only those households that operated a non-farm enterprise. Consistent with the model, the average product of labor (proxied by revenue per worker) apparently increases in NREGS districts relative to non-NREGS districts.[19] However, the increase is statistically insignificant. As such, while the coefficient

[19]The model predicts an increase in the marginal product of labor. Assuming Cobb-Douglas production technology, the marginal product of labor is proportional to the average product of labor, and thus changes in the average product will mirror changes in the marginal product.

Table 2.8: Effect of NREGS on Non-Farm Self-Employment Income

	All		NFE Households		
	(1)	(2)	(3)	(4)	(5)
	NFE HH	Workers (log)	Workers (log)	Days (log)	Rev. per Day
Post times NREGS	−0.176***	−0.178**	−0.225*	−0.774	0.092
	(0.066)	(0.072)	(0.121)	(0.495)	(0.368)
State/Wave	Yes	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes	Yes
Observations	11674	11674	1698	1698	1626

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the household level. All regressions use the ARIS/REDS. District fixed effects are included in all regressions. The first two columns include all households. The dependent variable in the first column is a dummy variable equal to one if the household operated a non-farm enterprise of any type and the second column is the log of number of household members that worked in non-farm self-employment plus one. The third through fifth columns include only households that operated a non-farm enterprise. The dependent variable in the third column is (log of) number of total workers plus one. The dependent variable in the fourth column is (log of) days of workers. The dependent variable in the last column is (log of) revenue per worker-day.

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

is in the hypothesized direction, there is not enough power to infer anything about the effects of NREGS on enterprise productivity.

2.5.3 NREGS Employment, Wage Increase, and Market Failures

The effect of NREGS on non-farm self-employment could theoretically be caused or mitigated by several mechanisms. First, households may be reallocating labor from non-farm self-employment to NREGS itself. If NREGS wages are higher than returns to non-farm self-employment, then NREGS employment may be directly responsible for the decrease in days of non-farm self-employment. Second, the rural wage increase could induce households to change its reallocation away from non-farm self-employment. If more remunerative off-farm (non-NREGS) opportunities arise in the dry season due to the implementation of NREGS, then a household with low-return non-farm self-employment may exit the sector and instead reallocate its labor towards more remunerative opportunities.

A third possible mechanism is risk. The rural non-farm sector plays an important risk management and coping role in many developing countries (Barrett et al., 2001; Haggblade et al., 2010; Lanjouw and Lanjouw, 2001; Nagler and Naudé, 2014). The common strategy of diversification can serve to decrease household risk exposure, but also decrease the expected value of returns. As such, decreasing the risk a household faces can improve its welfare, even without increasing its income. There is some evidence that NREGS is an important insurance tool for many households (Gehrke, 2014; Zimmermann, 2015). In the current context, the introduction of de facto insurance through NREGS employment might induce individuals to decrease their reliance on low-risk, low-return non-farm self-employment.

Finally, the model above makes an important assumption that non-farm capital is fixed. If this is not the case, then the model predictions may not hold if NREGS affects a household's ability to overcome credit constraints. I explore this possibility briefly in the last section.

NREGS Employment

The decrease in non-farm self-employment could be driven by NREGS employment directly. It is, unfortunately, difficult to tease out this direct effect using the NSS data since the decision to allocate one's labor to NREGS or non-farm self-employment is endogenous. One possibility is to include days in public works as a covariate, which is what I do in Table 2.9. With the caveat that days in public works is unlikely to be exogenous with respect to days in non-farm self-employment, Table 2.9 presents the main results but including public works as a covariate. Not surprisingly, the coefficient on public works is negative and highly significant; those individuals that engage in public works employment significantly decrease their time in non-farm self-employment. However, the estimated effects of the program are almost completely unchanged. This suggests that other factors – such as the increasing wage rate – are responsible for the significant shift out of non-

Table 2.9: Effect of NREGS - Public Works Days as Covariate

	(1)
	aysinNFSelf-Emp.
Days of Public Works (log)	−0.100***
	(0.006)
Post times NREGS (Rainy)	-0.015
	(0.016)
Post times NREGS (Dry)	-0.043**
	(0.018)
State/Wave	Yes
Interview Month/Wave	No
Observations	366173

The dependent variable in all columns is log of days of non-farm self-employment plus one. Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. Individuals between 15 and 60 years of age are included in estimation. Column one includes total number of days (log + 1) of non-farm self-employment the individual reported engaging in during the previous week. The dependent variable in column two equals one if the individual reported engaging in any non-farm self-employment during the previous week. In both regressions, log of days of public works is included as a covariate.

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

farm self-employment, not direct access to NREGS employment. However, this finding is also consistent with diversification into non-farm self-employment serving as a risk-management strategy. If individuals believe they will have access to remunerative forms of employment in the case of an adverse shock – crop failure caused by drought, for example – then they may choose to no longer engage in non-farm self-employment even if they are not participating in NREGS at the time the decision to exit the sector is made. The next two subsections focus on these possibilities.

Wage Increase

The above results suggest that non-farm self-employment has a relatively large proportion of subsistence enterprises, since workers are more likely to substitute out of these enterprises than other employment options following an increase in the wage rate. However, the empirical effect documented above is not necessarily due to the wage increase alone. To further investigate this possibility, I next present evidence that lower-wage industries see larger effects after NREGS implementation.

Figure 2.5 presents the industry wage in 2004/05 on the x-axis and the industry prevalence, measured as percent of individuals that worked in that industry, in 2004/05. The NSS unfortunately does not include data on self-employment income. Thus, it is impossible to construct a measure of self-employment income by industry using the NSS. As such, the wages are constructed using casual workers, not non-farm self-employment workers. The assumption I make is that patterns of the overall wages in that industry reflect patterns of self-employment income in that industry, as well, which seems plausible. Table 2.A2 in the appendix presents these numbers in table form. There appears to be a relationship between wage and prevalence, with more prevalent industries (e.g. retail trade, clothing manufacturing, and other) also being those with lower wages. This is consistent with households entering less productive industries, which may require less capital and, thus, have lower barriers to entry (Haggblade et al., 2010).

Given that the dry season may make low-return entrepreneurship more likely, we might expect to see differences in industry prevalence across seasons. Figure 2.6 explores this possibility. Manufacture (clothing) – on the far left of the figure – has the lowest average wage in 2004/05, while manufacture (wood) – on the far right – has the highest wage. The wage is increasing from left to right. Interestingly, clothing manufacturing and retail trade, two industries with the highest prevalence and lowest wages, also show the largest absolute differences in prevalence between the rainy and dry seasons. This is again suggestive of households opening low return enterprises during the dry season.

Given these findings and the fact that NREGS increased wages significantly in the dry season but had no effect in the rainy season, if the wage increase is partly responsible for

Figure 2.5: Industry Prevalence and Wages - 2004/05

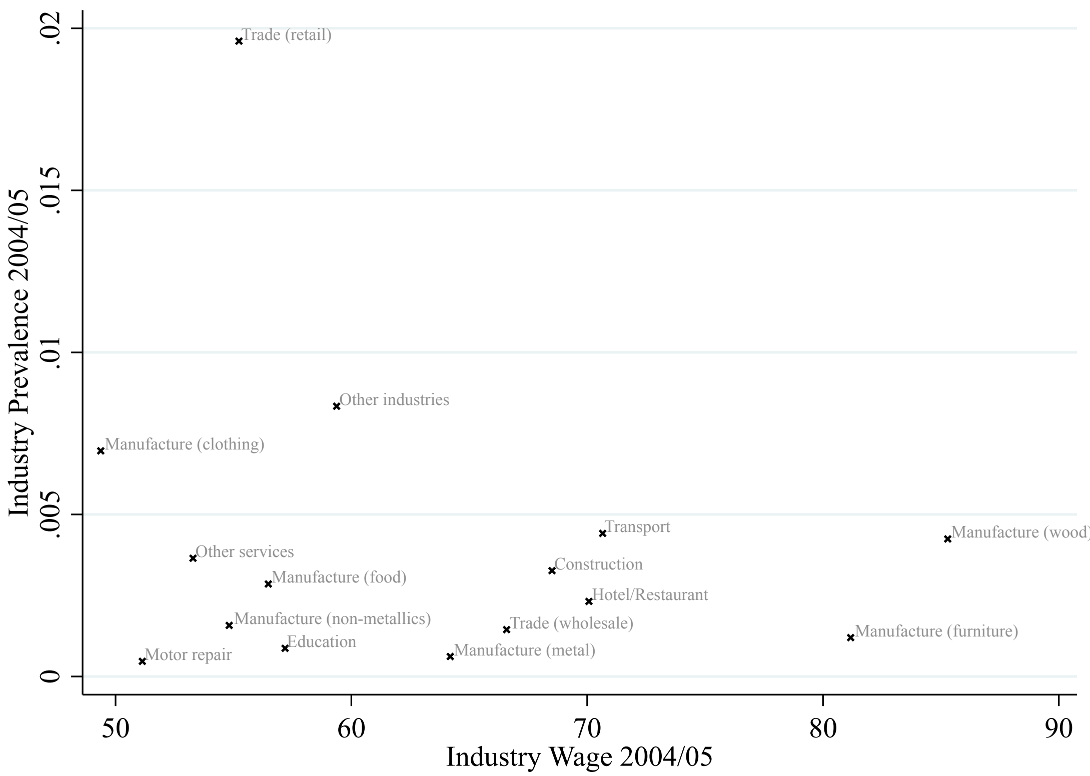

The x-axis represents the industry-level wage in 2004/05. The wage is constructed using casual wages–not self-employment income–since individuals engaged in non-farm self-employment do not report earnings in the NSS. The y-axis represents the prevalence of each industry in 2004/05. The prevalence is defined as the percentage of individuals that reported any self-employment in the industry. Both wages and prevalence are weighted to be representative of the population.

Figure 2.6: Industry Prevalence by Season - 2004/05

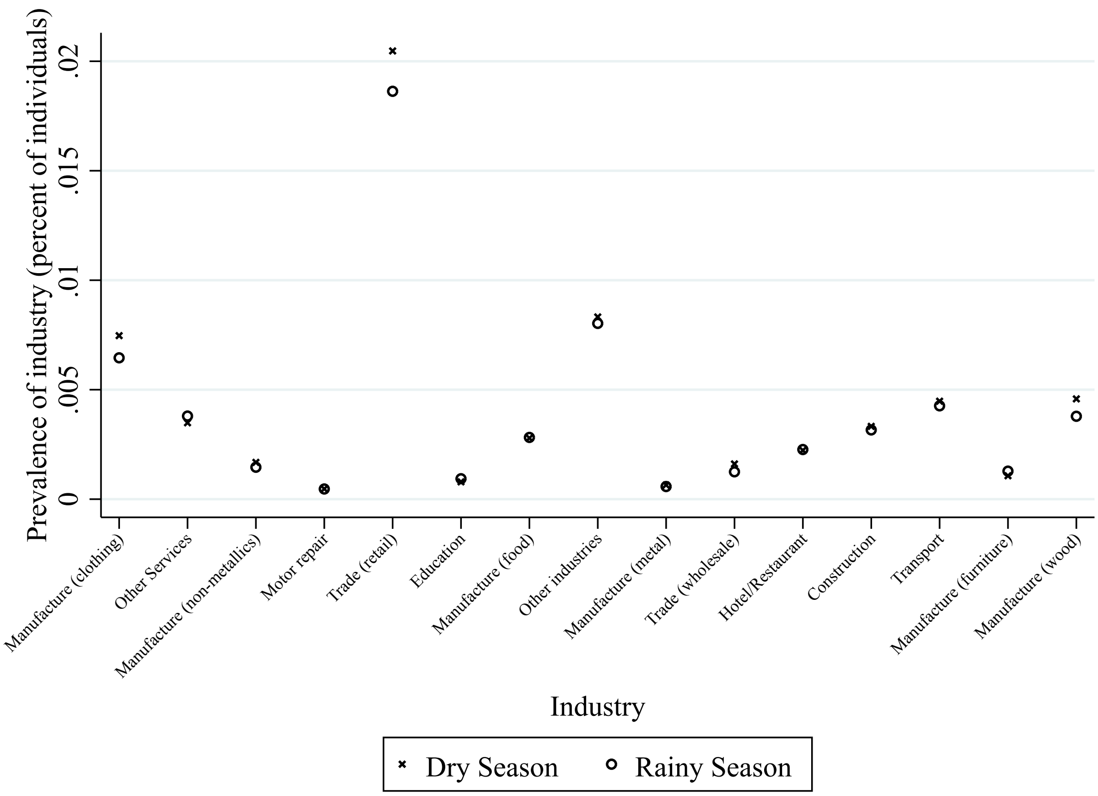

The y-axis represents the prevalence of each industry in 2004/05. The dry season is defined as January through June and the rainy season is defined as July through December. The prevalence is defined as the percentage of individuals that reported any self-employment in the industry. Statistics are weighted to be representative of the population.

the decrease observed in the previous section, we would expect the effect of NREGS on industry prevalence to be driven by the wage rate prior to its implementation in the dry season but not the rainy season. Figure 2.7 presents results from a set of regressions. For each industry, a single regression of the effect of NREGS was estimated for both the dry and rainy season. Panel A presents the results for the dry season in graphical form. The individual points indicate the effect for an individual industry, while the line is an OLS line of best fit for the relationship, weighted by initial industry prevalence. In line with expectations, the effect of NREGS on industry prevalence is largest (most negative) for industries that had the lowest wages prior to implementation of the program. Even with only 15 industry categories, the effect is close to significant (p=0.105).[20] In fact, of the six lowest wage industries, three of them saw the largest effects of the program. Similarly, only three industries saw an increase in prevalence due to the program, two of which had the highest wages prior to implementation. Panel B explores the same relationship during the rainy season. The relationship between initial wage and the effect of NREGS is qualitatively small and in the opposite direction. Overall, Figure 2.7 presents additional evidence that the effects of NREGS are concentrated in lower return enterprises operated in the dry season.

Previous research suggests that local demand is an important aspect in the growth of non-farm enterprises. In particular, this research has found that local agricultural productivity partly determines the fortunes of local non-farm enterprises (Hornbeck and Keskin, 2012; Santangelo, 2016). This is especially true for enterprises that sell their goods only locally, as opposed to enterprises that produce goods for export to other regions of the country or even internationally. It seems plausible that lower wage industries are more likely to represent enterprises that sell their goods locally (or, perhaps more accurately, higher wage industries are more likely to produce tradables). If true, this reinforces the results above

[20]The p-value is constructed by bootstrapping 1,000 replications of the multi-step estimator.

Figure 2.7: Effect of NREGS by Industry Wage

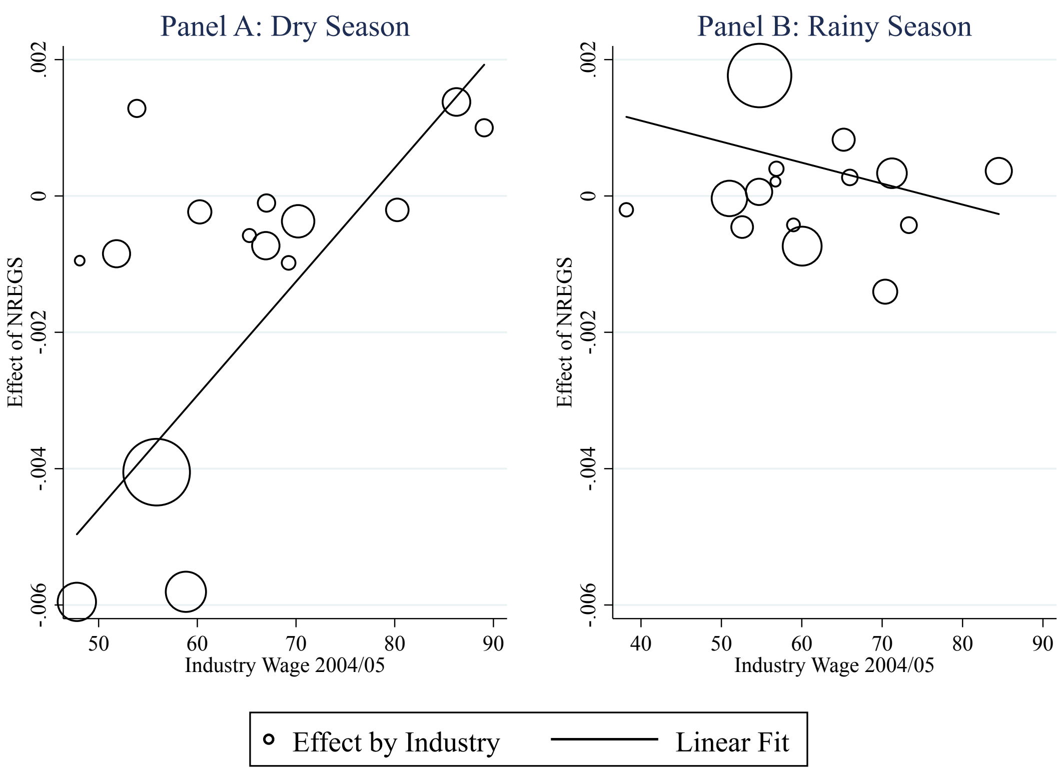

In both panels, the y-axis represents the effect of NREGS on industry prevalence. The x-axis represents the overall industry wage in 2004/05, for the relevant season, prior to NREGS implementation. The dry season is defined as January through June and the rainy season is defined as July through December. The prevalence is defined as the percentage of individuals that reported any self-employment in the industry. The estimates are weighted to be representative of the population. In addition, the linear fit and the figure are weighted by pre-NREGS industry prevalence.

and suggests they might even be underestimates of the true (direct) impact of the program; even with increases in household income and consumption due to NREGS, the prevalence of low-wage industries declined markedly in treated districts relative to untreated districts.

NREGS as Insurance?

Another possible mechanism is risk. Income diversification is a common risk-management strategy in developing countries (Barrett et al., 2001). This is especially true in areas that rely on agriculture, as agricultural production is generally considered to be a risky undertaking in developing countries due to its reliance on rainfall. Individuals can use non-farm self-employment to either manage risk ex ante or cope with risk ex post, following an unexpected shock (Reardon, 1997). If NREGS serves as an implicit form of insurance – guaranteeing individuals employment and a minimum income following a negative shock – then individuals may close any non-farm enterprises that operated principally as a form of risk management.

The most obvious "risky" activity for most households in rural India is agriculture. Rainfall can be highly variable from year to year, which also causes variability in income. This suggests that insofar as non-farm self-employment serves as a risk-management activity, individuals in areas with higher rainfall variability are more likely to diversify into non-farm self-employment. However, households and individuals may operate a non-farm enterprise not just in expectation of possible shocks, but also explicitly in response to a shock. As such, we might also expect individuals to increase their labor allocation to non-farm self-employment following negative rainfall shocks.

Before examining whether these mechanisms might be responsible for the demonstrated effects of NREGS, I first present evidence regarding whether the above assumptions about the operation of non-farm enterprises are indeed true. In Table 2.10, I present estimates using the ARIS/REDS data, which are at the household level, of whether the effects of NREGS are correlated with rainfall variability. In all columns, I continue to include current rainfall, current rainfall squared, previous year's rainfall, and previous year's rainfall squared. In addition, I allow the effects of rainfall to vary by NREGS status. The inclusion of these rainfall variables should help ensure that any effects attributed to rainfall variability are not confounded with current rainfall increasing (or decreasing) rainfall variability.[21]

In column one, I explore whether rainfall variability affects the decision to diversify into non-farm self-employment. The coefficient on the standard deviation of rainfall is positive and marginally significant (p=0.168), which I interpret as suggestive evidence that households are more likely to operate non-farm enterprises in areas with higher rainfall variability, at least in the ARIS/REDS sample. In columns two through four, I present triple-difference estimates of whether the effect of NREGS varies by risk, proxied by the standard deviation of rainfall. The corresponding coefficient in Table 2.10 is *Post times NREGS times Rain SD*. In all three columns, the triple interaction is negative and highly significant ($p<0.05$).

These results show that the effect of NREGS is larger (more negative) in areas with more variable rainfall. In other words, individuals are more likely to substitute out of non-farm self-employment in areas of riskier agricultural production. This finding is consistent with non-farm self-employment playing an important role as a risk-management strategy in rural India and with NREGS offering an alternative option for risk management (or risk coping).[22]

[21] I choose the ARIS/REDS data due to the fact that it contains latitude and longitude information for villages. The NSS data are aggregated to the district level, so any rainfall variables are necessarily constructed with a single value for the entire district.

[22] Moreover, in results available upon request, I show that the effect of current rainfall does not mediate the effects of NREGS on non-farm self-employment. In other words, the effect appears to be driven by rainfall variability rather than levels.

Table 2.10: Rainfall, Risk, Non-Farm Self-Employment, and the Effects of NREGS

	(1) NFE HH	(2) NFE HH	(3) Workers (log)	(4) Days (log)
Rainfall SD	0.003	0.008***	0.008***	0.056***
	(0.002)	(0.003)	(0.003)	(0.018)
Post times NREGS times Rain SD		−0.002**	−0.002**	−0.012**
		(0.001)	(0.001)	(0.006)
Post times NREGS		0.336*	0.362	2.057
		(0.199)	(0.227)	(1.404)
State/Wave	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
Observations	11674	11674	11674	11674

Standard errors are in parentheses and clustered at the district level. All regressions are at the household level. All regressions use the ARIS/REDS. District fixed effects are included in all regressions. The ARIS/REDS data has spatial information at the household level, so coefficients are identified by variation within districts in rainfall variability.

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

Credit Market Failures

The preceding model makes an important assumption: non-farm capital is fixed. If we relax this assumption, comparative statistics can change, especially if we introduce non-convexities. However, a change requires other market failures; if credit and savings markets function well, households will employ an efficient level of capital (Blattman et al., 2013). Without credit and savings markets, households may not be able to overcome entry barriers. Much evidence suggests that credit markets do not function particularly well in developing countries (Blattman et al., 2013; Haushofer and Shapiro, 2016; Bandiera et al., 2017). In addition, without access to savings accounts, many households may fail to save, making short-term decisions that do not align with their long-term interests (Frederick et al., 2002; Blattman et al., 2013). These common market failures result in non-convexities in production that have a number of implications, possibly including poverty traps (Banerjee and Newman, 1993; Aghion and Bolton, 1997; McKenzie and Woodruff, 2006).

The combination of credit/savings market failures and non-convexities – e.g. high start-up costs – may prevent poorer households from engaging in more remunerative non-farm employment. McKenzie and Woodruff (2006) point out three important implications, two of which are especially germane to the present discussion. First, the number of firms in the market is not efficient. In particular, there are not enough firms because some would-be entrepreneurs are not able to overcome the start-up costs and thus do not enter the market. Second, entrepreneurs that are able to overcome entry barriers may still not produce at an efficient level if they are not able to invest sufficiently in capital.

These two possibilities present difficulties for the simple model above. We can incorporate non-convexities in production in a simple way, by allowing for two different types of enterprises: a less productive type that requires little to no capital and a more productive type with substantial barriers to entry. This is not unlike some dual-sector models, with a more informal – but less productive – sector acting as a queue for employment in a formal, more productive sector (Harris and Todaro, 1970; Fields, 1975).

The model in the previous section unambiguously predicts a decrease in non-farm self-employment due to an increase in the prevailing wage. However, if the wage change increases household income sufficiently to overcome entry barriers, some households may actually *enter* into non-farm self-employment. In fact, some previous studies have found that the program increased asset accumulation and savings (Deininger and Liu, 2013; Ravi and Engler, 2015; Deininger et al., 2016). If this mechanism is operating, then the previous estimates of the effects of NREGS on days in non-farm self-employment might be underestimating the "true" wage effect of the program.

Credit market failures may prevent households from entering remunerative forms of non-farm self-employment if there are substantial startup costs. If this is the case, we should see heterogeneity in the effect of NREGS on non-farm self-employment based on access to credit. The ARIS/REDS includes a village-level module with information on distance to the nearest bank. If credit market failures are preventing some households from entering

Table 2.11: Effect of NREGS and Access to Financial Institutions

	(1) Bank	(2) Rural Bank	(3) Cooperative
Post times NREGS times Financial	−0.014	−0.013	0.000
Institution Dummy	(0.048)	(0.051)	(0.045)
Post times NREGS	−0.153**	−0.165**	−0.165**
	(0.071)	(0.066)	(0.071)
State/Wave	Yes	Yes	Yes
Month	Yes	Yes	Yes
Observations	11674	11674	11674

Standard errors are in parentheses and clustered at the district level. All regressions are at the household level. All regressions use the ARIS/REDS. District fixed effects are included in all regressions. The financial institution in the first through fourth columns are a bank, rural bank, cooperative, and money lender, respectively. Distances are defined from the enumerated village.

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

non-farm entrepreneurship, we should see a larger effect of NREGS where these market failures are more likely. In Table 2.11, I assume being located closer to a financial institution improves access to credit or savings. If credit markets are preventing households from entering non-farm self-employment, then (ceteris paribus) households closer to credit should see a more negative effect of NREGS, since the program would be alleviating less of a market failure than for households farther from credit.

In all three columns of Table 2.11, the dependent variable is an indicator variable for whether the household operates a non-farm enterprise. If the alleviation of credit market failures is mitigating the negative effects of NREGS on non-farm self-employment, then we expect to see positive coefficients on the triple interaction terms. In all three cases, the coefficient is small and insignificant, suggesting households located farther from a financial institution react similarly to the implementation of NREGS as households located closer to a financial institution. As such, it does not appear that NREGS has provided any alleviation of credit constraints.

However, distance to a financial institution may be correlated with other omitted variables that may mediate the effects of NREGS. In particular, distance to a financial institution may be correlated with local wages or distance to other locations, such as large towns or markets. If this is true, then the results in Table 2.11 may be biased. To alleviate some concerns regarding this possibility, I present further results in appendix tables Table 2.A8, Table 2.A9, and Table 2.A10. In each table, the first column replicates the results for the first column in Table 2.11. The second column adds a control for the local village-level wage in 1999 (prior to NREGS). The third column adds an additional control for village population. The fourth column adds additional controls for distance to the district headquarters as well as distance to the nearest large town. The addition of these controls has little effect on the estimated effect of NREGS by distance to the nearest financial institution. Finally, Table 2.A11 of the appendix presents correlations for the distance variables and shows that the lack of significance is not likely to be driven by multicollinearity.

2.6 Conclusion

In this paper, I present evidence that implementation of the Mahatma Gandhi National Rural Employment Guarantee Act induced significant changes in the rural non-farm self-employment sector. The program led to a decrease in the prevalence of non-farm self-employment at both the household and individual level. Overall, the evidence suggests that the program reduced individuals' reliance on less remunerative forms of non-farm self-employment and offered these households an alternative risk-management option.

It is important not to read the results as suggesting that non-farm self-employment serves little purpose in rural India. While individuals substitute out of non-farm self-employment following the rural wage increase, non-farm self-employment was still the utility-maximizing form of employment available to these individuals prior to the introduction of the program. In reality, without access to this form of employment, the incomes of the poorest households are likely to be even lower (Lanjouw, 1999). With this in mind, it appears that

NREGS is working as intended, even if not directly through the program; introduction of NREGS has allowed at least some individuals to move out of lower return non-farm self-employment and into higher return alternatives.

Similarly, the results presented in this paper do not suggest no enterprises within the sector are worth being promoted; the rural non-farm sector is heterogeneous, and so is non-farm self-employment (Haggblade et al., 2002). There is unlikely to be a single policy appropriate for the entire sector (Lanjouw, 1999). Rather, the results suggest we need to improve our understanding of which kinds of self-employment have growth prospects and which kinds of self-employment serve as an economic refuge for the poor. The forces driving the results in this paper are likely related to the forces driving heterogeneity in the effects of microfinance, more generally (Morduch, 1999; Karlan and Morduch, 2009), and the effects of microfinance on non-farm entrepreneurship, specifically (Haggblade et al., 2002). Recognizing this heterogeneity and focusing on entrepreneurial support as a livelihood strategy rather than a driver of growth may be more effective as a poverty-reduction strategy (Cho et al., 2016).

A final note regarding the benefits of the Mahatma Gandhi National Rural Employment Guarantee Act is warranted. While the current results indicate some individuals are leaving behind less remunerative forms of employment, this says nothing of the overall benefits and costs of the program. In fact, these results also suggest caution when interpreting the overall benefits. Some individuals appear to be substituting time in non-farm self-employment for time in NREGS. For these individuals, the relevant "benefit" is not the overall wage received from the program, but rather the difference between the program wage and the self-employment wage. This finding has important implications for benefit-cost analysis of the program, of which previous attempts have been somewhat pessimistic (Murgai et al., 2015). However, the fact that NREGS also appears to offer households an alternative risk-management strategy means that a simple comparison of monetary returns is not the relevant comparison these households make. As such, the results in this paper and others (Gehrke, 2014; Zimmermann, 2015) also suggest that a more through accounting of the benefits of NREGS is required.

2.7 Paper 1 Appendix

Table 2.A1: District-Level Controls

Variable	Source
Population	2001 Census
Percent rural	2001 Census
Area	2001 Census
Casual wage	2004/05 NSS
Percent scheduled castes	2001 Census
Percent scheduled tribes	2001 Census
Sex ratio	2001 Census
Literacy rate	2001 Census
Labor force participation rate	2001 Census
Labor force composition	2001 Census
Agricultural productivity	2000-2001 Agricultural Census
Rainfall for previous year	University of East Anglia[1]
Rainfall for current year	University of East Anglia

[1] http://www.cru.uea.ac.uk/data

Table 2.A2: Summary Statistics - Non-Farm Self-Employment Industries

	Pre-Program (2004/05)	
	Operation (Yes=1)	Wage (Rupees)
Manufacturing (wood)	0.004	85.289
Manufacturing (furniture)	0.001	81.175
Transport	0.004	70.657
Hotel and restaurant	0.002	70.075
Construction	0.003	68.511
Trade (wholesale)	0.001	66.583
Manufacturing (metal)	0.001	64.194
Other industries	0.008	59.375
Education	0.001	57.191
Manufacturing (food)	0.003	56.479
Trade (retail)	0.020	55.233
Manufacturing (non-metallic)	0.002	54.819
Other services	0.004	53.286
Motor vehicle repair	0.000	51.136
Manufacturing (clothing)	0.007	49.374
All industries	0.069	65.125
Observations	398006	

All statistics are weighted and are from the 2004/05 NSS, which was collected prior to implementation of NREGS. The first column shows the prevalence of each industry. For example, 0.4 percent of individuals reported engaging in manufacturing (wood) as non-farm self-employment. The second column is mean wage for each industry. However, wages are defined for all workers in the industry, since those individuals reporting non-farm self-employment do not report earnings in the NSS.

Table 2.A3: Summary Statistics - All Individuals over Ten

	Pre-Program (2004/05)	
	(1)	(2)
	NREGS	Non-NREGS
Wage (R's)	47.907	63.136
	(25.675)	(43.546)
Non-farm worker (last week - yes=1)	0.115	0.101
	(0.319)	(0.301)
Days non-farm self-employment (last week)	0.622	0.561
	(1.904)	(1.833)
Male	0.507	0.509
	(0.500)	(0.500)
Age	32.474	32.989
	(16.977)	(17.403)
Education	2.106	2.294
	(1.391)	(1.438)
Household head male (yes=1)	0.925	0.917
	(0.263)	(0.276)
Household head age	47.812	49.040
	(12.702)	(13.030)
Household head education	2.121	2.208
	(1.491)	(1.497)
Household size	6.150	6.154
	(3.154)	(2.966)
Percent children	0.304	0.285
	(0.216)	(0.214)
Percent prime male	0.314	0.320
	(0.173)	(0.178)
Percent prime female	0.304	0.307
	(0.145)	(0.148)
Percent elderly female	0.039	0.044
	(0.098)	(0.102)
Observations	159710	122403

Standard deviations are in parentheses. All statistics are unweighted and are from the 2004/05 NSS, which was collected prior to implementation of NREGS. The first column includes all individuals in NREGS districts. The second column includes all individuals in non-NREGS districts. Wage is in rupees. Education is coded from 1-6: 1 indicates less than a primary education; 2 indicates a primary education; 3 indicates a middle-school education; 4 indicates a secondary education; 5 indicates a higher-secondary education; and 6 indicates a college degree. The percent variables following household size are the composition of the individual's household. Children are defined as individuals under 15; prime males are defined as males between 15 and 60 years of age; prime females are defined as females of the same age; and elderly females are defined as females 60 or older. All individuals over ten years of age are included in the statistics.

Table 2.A4: Effect of NREGS on Days of Non-Farm Self-Employment - Levels

		By Season	
	(1)	(2)	(3)
	All	Dry	Rainy
Post times NREGS	−0.094*	−0.200***	0.012
	(0.049)	(0.068)	(0.058)
State/Wave	Yes	Yes	Yes
Interview Month/Wave	Yes	No	No
Observations	366876	182889	184044

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. Individuals between 15 and 60 years of age are included in estimation. The dependent variable in all columns is the number of days of non-farm self-employment (in levels).

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

Table 2.A5: Effect of NREGS on Non-Farm Self-Employment - All Individuals over Ten

	Individual level	
	(1)	(2)
	Self-emp (log days)	Self-emp (Yes=1)
Post times NREGS	−0.020*	−0.010
	(0.011)	(0.006)
State/Wave	Yes	Yes
Interview Month/Wave	Yes	Yes
Observations	539813	539813

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. All individuals over ten years of age are included in estimation. Column one includes total number of days (log + 1) of non-farm self-employment the individual reported engaging in during the previous week. The dependent variable in column two equals one if the individual reported engaging in any non-farm self-employment during the previous week.

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

Table 2.A6: Effect of NREGS by Season - All Individuals over Ten

	(1) Dry Season	(2) Rainy Season
Post times NREGS	-0.043***	0.002
	(0.016)	(0.014)
State/Wave	Yes	Yes
Observations	268672	271217

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. All individuals over ten years of age are included in estimation. Column one includes total number of days (log + 1) of non-farm self-employment the individual reported engaging in during the previous week. The dependent variable in column two equals one if the individual reported engaging in any non-farm self-employment during the previous week.

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

Table 2.A7: Effect of NREGS on Non-Farm Self-Employment - Ordered Probit

		By Season	
	(1) All	(2) Dry	(3) Rainy
Post times NREGS	-0.119**	-0.228***	0.013
	(0.052)	(0.073)	(0.063)
State/Wave	Yes	Yes	Yes
Interview Month/Wave	Yes	No	No
Observations	366876	182889	184044

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the individual level. All regressions use the NSS and are weighted to be representative of the population. District fixed effects are included in all regressions. Individuals between 15 and 60 years of age are included in estimation. The dependent variable in all columns is the number of days of non-farm self-employment and takes on 15 possible values from zero to seven days in half-day increments. All results are estimated using ordered probit.

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

Table 2.A8: NREGS and Rural Banks

	(1) NFE HH	(2) NFE HH	(3) NFE HH	(4) NFE HH
Post times NREGS times Rural Bank Dummy	-0.013	-0.021	-0.025	-0.012
	(0.051)	(0.050)	(0.050)	(0.051)
Post times NREGS	-0.165**	-0.146**	-0.107	-0.135*
	(0.066)	(0.067)	(0.076)	(0.076)
State/Wave	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
Village Wage (1999)	No	Yes	Yes	Yes
Village Population	No	No	Yes	Yes
Village Distance Variables	No	No	No	Yes
Observations	11674	11575	11238	11238

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the household level. All regressions use the ARIS/REDS. Due to sample size, district fixed effects are not included in the estimation. The dependent variable in all columns is an indicator for whether the household operated a non-farm enterprise or not.

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

Table 2.A9: NREGS and Cooperatives

	(1) NFE HH	(2) NFE HH	(3) NFE HH	(4) NFE HH
1.post1.nrega1.coop_dum	0.000	0.015	0.030	0.037
	(0.045)	(0.044)	(0.044)	(0.046)
Post times NREGS	-0.165**	-0.151**	-0.127	-0.158*
	(0.071)	(0.072)	(0.080)	(0.081)
State/Wave	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
Village Wage (1999)	No	Yes	Yes	Yes
Village Population	No	No	Yes	Yes
Village Distance Variables	No	No	No	Yes
Observations	11674	11575	11238	11238

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the household level. All regressions use the ARIS/REDS. Due to sample size, district fixed effects are not included in the estimation. The dependent variable in all columns is an indicator for whether the household operated a non-farm enterprise or not.

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

Table 2.A10: Effect of NREGS and Rural Banking

	(1) NFE HH	(2) NFE HH	(3) NFE HH	(4) NFE HH
1.post1.nrega1.bank_dum	−0.014	−0.025	0.008	0.006
	(0.048)	(0.051)	(0.052)	(0.050)
Post times NREGS	−0.153**	−0.130*	−0.120	−0.137*
	(0.071)	(0.070)	(0.075)	(0.076)
State/Wave	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes
Village Wage (1999)	No	Yes	Yes	Yes
Village Population	No	No	Yes	Yes
Village Distance Variables	No	No	No	Yes
Observations	11674	11575	11238	11238

Standard errors are in parentheses and clustered at the district level. All regressions are defined at the household level. All regressions use the ARIS/REDS. Due to sample size, district fixed effects are not included in the estimation. The dependent variable in all columns is an indicator for whether the household operated a non-farm enterprise or not. Distances are defined from the enumerated village. "Village Distance Variables" include distance to nearest town and distance to the district headquarters.

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

Table 2.A11: ARIS Distance Variables - Sums and Correlations

	(1) Distance HQ	(2) Distance Town	(3) Road (yes=1)	(4) Lender (yes=1)	(5) Bank (yes=1)	(6) Rural Bank (yes=1)	(7) Cooperative (yes=1)
Panel A: Summary Statistics							
Mean	52.464	13.256	0.249	0.109	0.661	0.335	0.529
Standard Deviation	63.320	14.211	0.433	0.312	0.475	0.473	0.500
Observations	221	221	221	221	221	221	221
Panel B: Correlations							
	Distance HQ	Distance Town	Road (yes=1)	Lender (yes=1)	Bank (yes=1)	Rural Bank (yes=1)	Cooperative (yes=1)
Distance HQ	1.000						
Distance Town	0.240	1.000					
Road	0.246	0.142	1.000				
Lender	−0.020	0.050	0.102	1.000			
Bank	−0.123	0.003	0.103	0.220	1.000		
Rural Bank	−0.014	0.015	0.190	0.122	0.164	1.000	
Cooperative	−0.009	−0.073	0.019	0.125	0.377	0.189	1.000

Statistics are at the village level in 1999. Each observation is a single village. All dummy variables indicate whether the specified infrastructure is located within the enumerated village.

Chapter 3

Spatially Heterogeneous Effects of a Public Works Program

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

3.1 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., 2015b; 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., 2015b; Muralidharan et al., 2016a).

[1]The 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.
[2]Figures are from http://NREGS.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 agreement with one basic fact: the program led to a modest increase in prevailing casual wages in rural India (Azam, 2011; 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, 2011; 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-

[3] In 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-

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

mann, 2012; Imbert and Papp, 2015; Merfeld, 2017a).[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, 2011; 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

[5] While 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, 2011). 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., 2015b; 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

[6]While 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.

3.2 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, 2011, 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

[7]The 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, 2011). If individuals do no 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, 2009b).

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, 2011). 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., 2015b).[8]

Several specific features of the program are worth noting. First, the minimum wages set by the program[9] 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 addition, childcare at worksites is required so that mothers with young children can partic-

sists of several villages.

[8]GPs 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. (2015b) provide an excellent overview of this process.

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

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

[11]The legislation state 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" (**?**, p. 14).

ipate 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, 2009b).

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.

3.3 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 period[12] 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

$$\max_{x_i^1, L_i^{1,f}, L_i^{1,T}, L_i^{1,h}} u\left(x_i^1, l_i^1\right) \tag{3.1}$$

[12]This 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

$$
\max_{x_i^0, L_i^{0,f}, L_i^{0,U}, L_i^{0,T}, L_i^{0,h}} u\left(x_i^0, l_i^0\right) \tag{3.2}
$$

subject to

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

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 $w^T 1 c(d)$.

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

[13] This 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.

3.4 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 periodically 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 3.1 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.

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

Figure 3.1: NREGS Districts and ARIS/REDS Villages

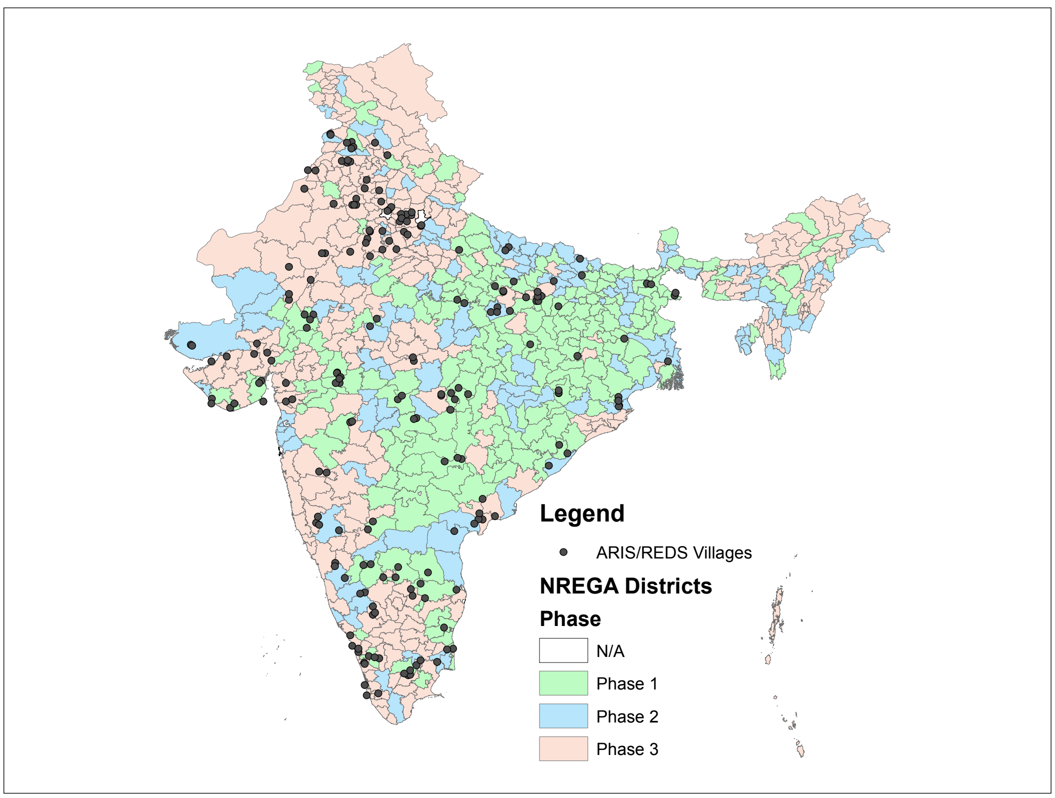

3.4.1 Methods

I estimate regressions of the form

$$
\begin{aligned}
y_{idt} = \alpha_0 + \alpha_1 NREGS_d + \alpha_2 Post_t + \alpha_3 NREGS_d \times Post_t \\
+ \delta X_{idt} + \phi H_{idt} + \gamma D_d + \beta Z_{idt} + \varepsilon_{idt},
\end{aligned}
\tag{3.3}
$$

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_{dt} 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

[15]The 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

$$\begin{aligned} y_{idt} = \alpha_0 + \alpha_1 dist_i + \alpha_2 Post_t + \alpha_3 dist_i \times Post_t \\ + \delta X_{idt} + \phi H_{idt} + \gamma D_d + \beta Z_{idt} + \varepsilon_{idt}, \end{aligned} \tag{3.4}$$

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:

$$\begin{aligned} y_{idt} = \alpha_0 + \alpha_1 NREGS_i + \alpha_2 Post_t + \alpha_3 Percent_d + \alpha_4 NREGS_i \times Post_t \\ + \alpha_5 NREGS_i \times Percent_d + \alpha_6 Post_t \times Percent_d \\ + \alpha_7 Post_t \times NREGS_i \times Percent_d + Controls + \varepsilon_{idt}, \end{aligned} \tag{3.5}$$

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

[16]See 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 website[17] 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

$$y_d = \alpha_0 + \alpha_1 dist_d + \phi D_d + \varepsilon_d, \tag{3.6}$$

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.NREGS.nic.in/netNREGS/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 3.2 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

Figure 3.2: Village Distance to Border with Other Phase District

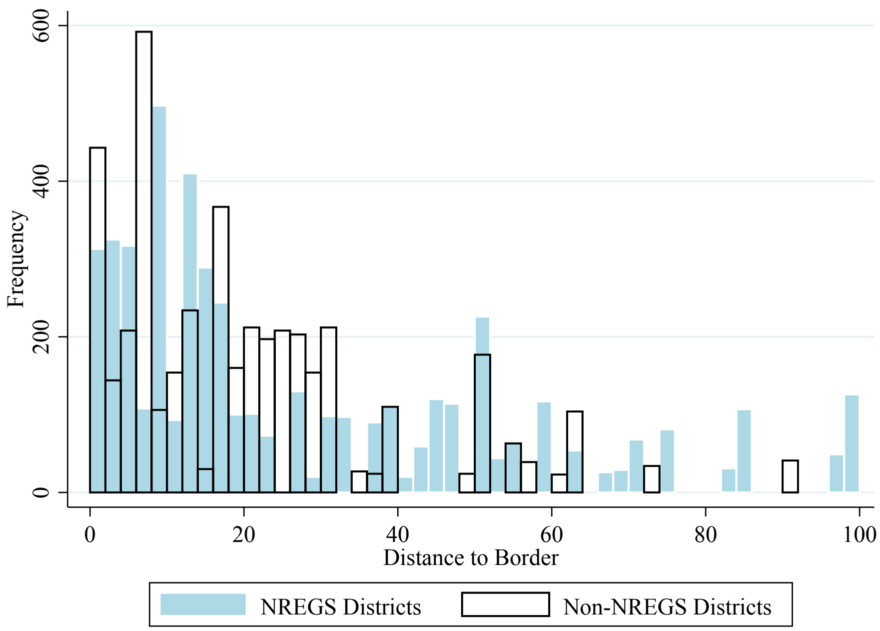

The figure shows the histogram of distance to border for NREGS and non-NREGS villages. For villages in NREGS districts (phase one and two districts), the distance is calculated as the shortest distance to a non-NREGS district (phase three districts). For villages in non-NREGS districts, the distance 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.

with an untreated district.

Figure 3.3 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

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

Figure 3.3: Creation of Distance Variables

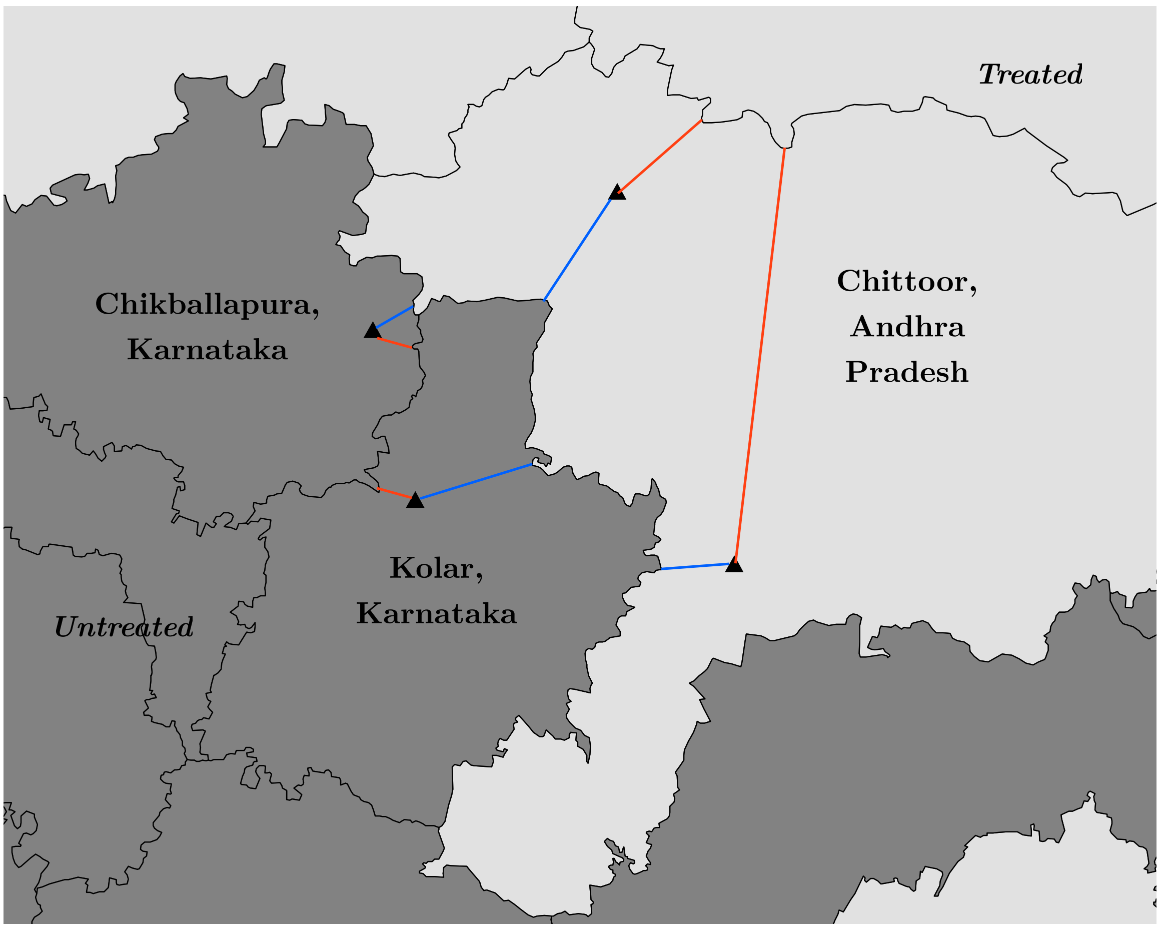

The dark gray districts – Chikballapura and Kolar – are untreated (phase three) districts. The light gray district – Chittoor – is a treated (phase one or two) district. The triangles represent actual villages in the dataset. Blue lines indicate the distance to the nearest district of the opposite treatment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

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 cyclicality 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.

3.4.2 Summary Statistics

I present summary statistics from 1999, prior to the program, for the sample in Table 4.1. 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 4.1 suggest that heterogeneous effects may be driven by differences in households, not the program itself.

[19]The 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.

Table 3.1: Summary Statistics (1999)

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

3.5 Results

I begin with a graphical representation of wage changes from 1999 to 2008 in Figure 3.4. 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 4.1, 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 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 3.2 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 3.4: Distance to Border and Wages

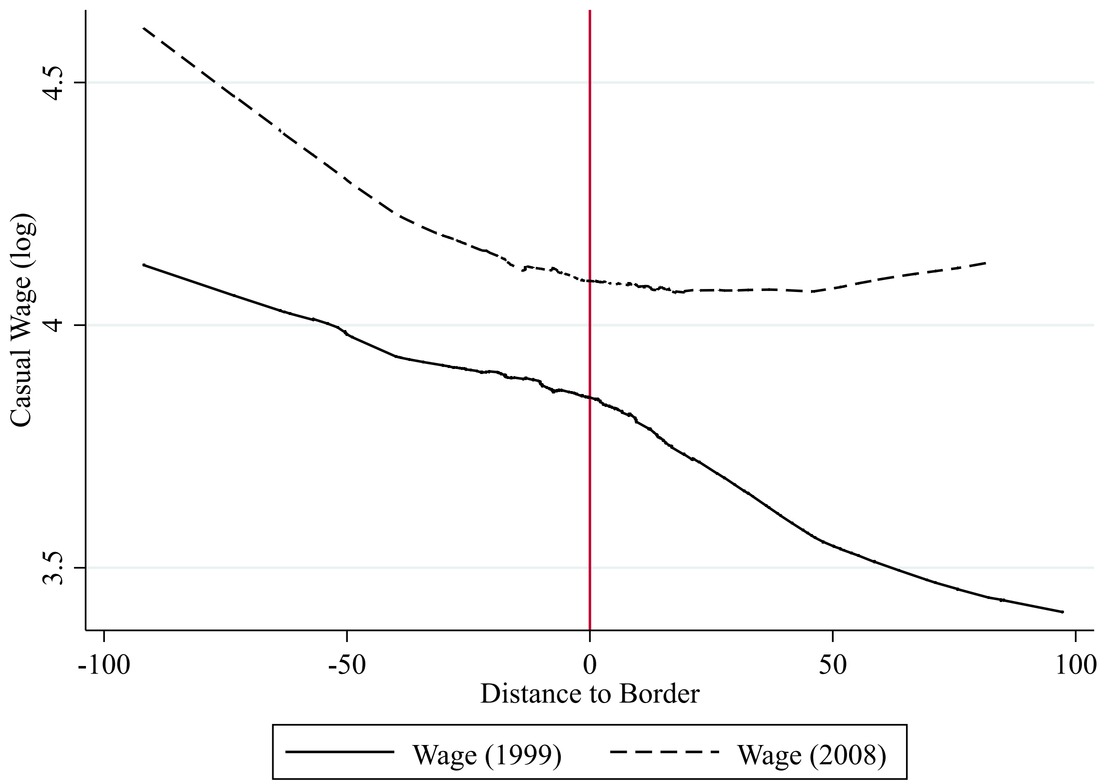

The figure shows a lowess curve of the relationship between distance to border and the wage rate for NREGS and non-NREGS households before (1999) and after (2008) implementation of the program. For households in NREGS districts (phase one and two districts), the distance is calculated as the shortest 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 3.2: Effects of NREGS on Wages - Distance to Border

	(1) All Villages	(2) Within 30 km	(3) Within 20 km	(4) Within 15 km	(5) Within 12 km	(6) Within 9 km	(7) Within 6 km
Post times	0.219***	0.196***	0.199***	0.204**	0.197*	0.155	−0.018
NREGS	(0.060)	(0.066)	(0.071)	(0.079)	(0.105)	(0.099)	(0.078)
Controls:							
District	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household	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 3.5 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 3.5. 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 3.5 confirms the effects found in Table 3.2; 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 3.2 – 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 3.3. 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.

Figure 3.5: Effect by Distance to Border

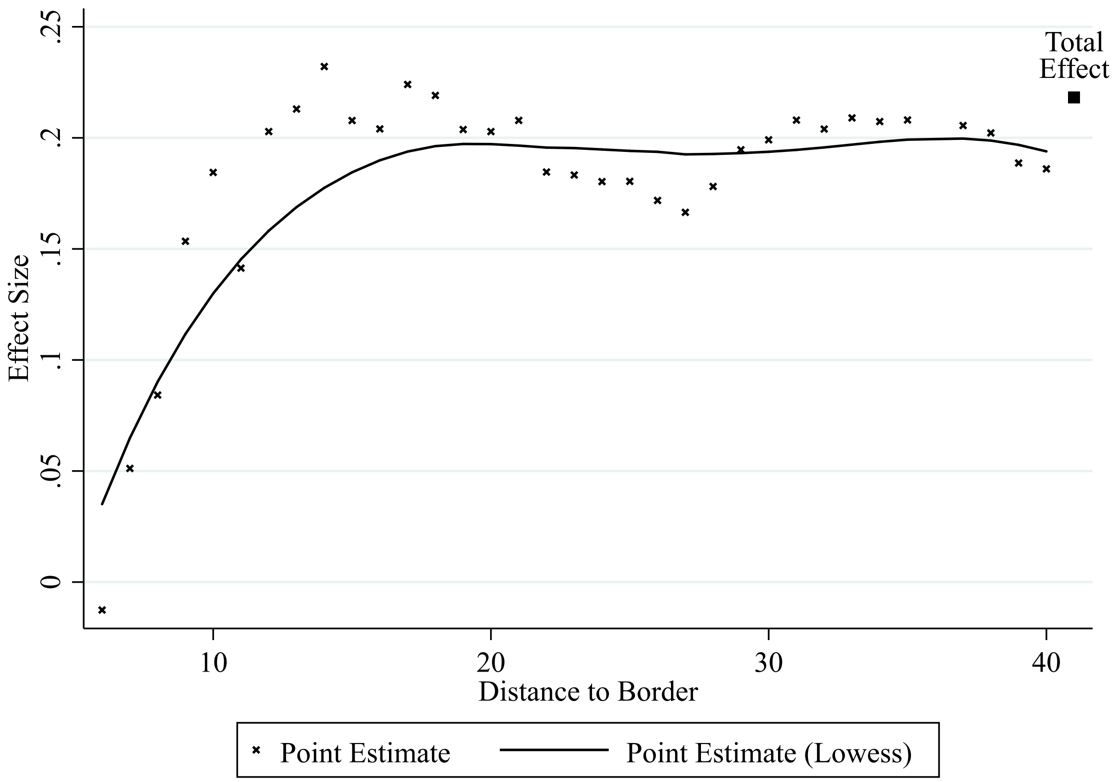

To create the figure, the sample is varied by using different cut-offs at different distances to the border, from 9 to 40. For example, for a distance of 25, the plotted point represents the ceofficient estimate from a difference-in-differences specification including only households within 25 kilometers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

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 travels costs than men and thus that we will see less of a decrease in the female wage rate closer to the border. Table 3.4 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 3.4 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 3.4 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.

3.5.1 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

Table 3.3: 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***		0.002	
	(0.017)		(0.018)	
Post times Distance> 6		0.102**		−0.047
		(0.048)		(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 3.4: 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***	0.188**	0.098	0.278***	0.227**	0.363**
	(0.057)	(0.078)	(0.088)	(0.078)	(0.108)	(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 3.6: Distance to Border and Wages - Pre-Program Trends

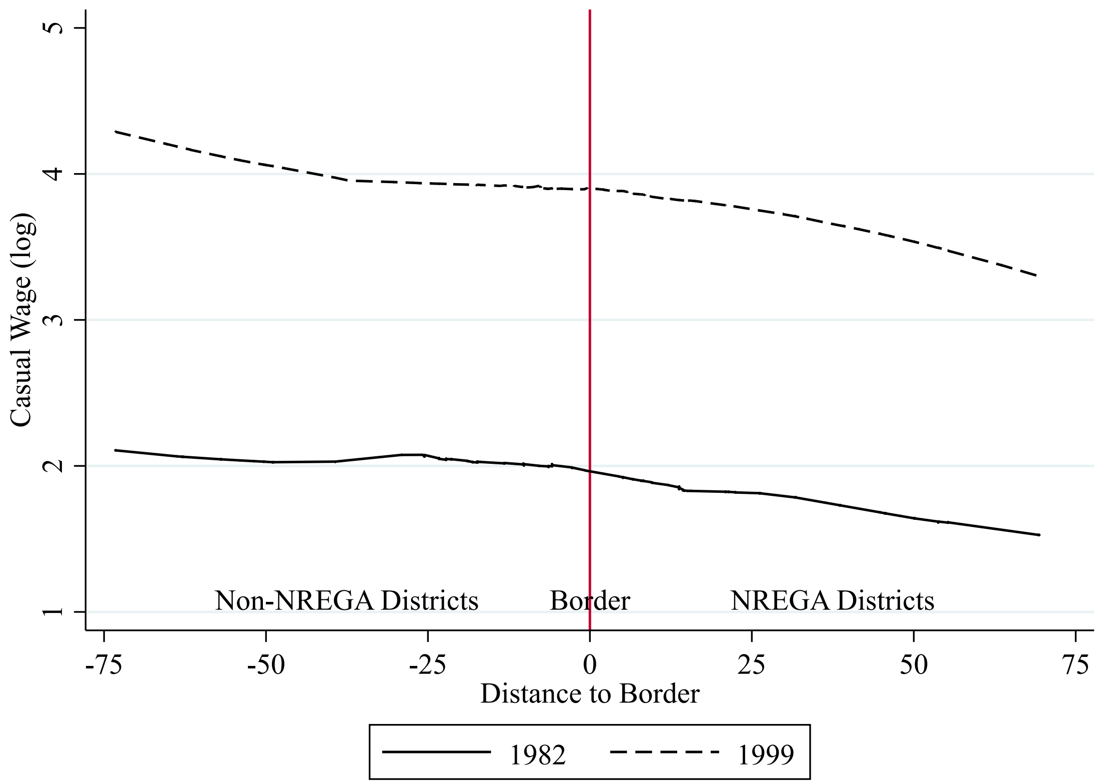

Due to data limitations, the unit of analysis is the household. The figure shows a lowess curve of the relationship between distance to border and the wage rate for NREGS and non-NREGS households before implementation of the program. For households in NREGS districts (phase one and two districts), the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

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 3.4 are suggestive evidence that this is not the case, I present further evidence that this is not true in Figure 3.6.

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

[20]There 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.

of (spatially) parallel trends is plausible. In Figure 3.6, 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 3.5 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 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

[21] Due 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.

Table 3.5: Effects and District Border Percentage

	(1) Wage (log) - Effect	(2) 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

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 3.6, 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]

[22] The 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

Table 3.6: Distance to Border and NREGS Implementation

	(1) Rolls Filled	(2) Ind. Worked	(3) Person-Days	(4) HH Limit	(5) Labor Exp. (R's lakhs)	(6) Percent Labor Expenditure	(7) Total Works
Distance to Non-NREGS (km - log)	0.021	0.009	0.035	0.028	0.021	0.003	−0.014
	(0.076)	(0.083)	(0.086)	(0.116)	(0.097)	(0.012)	(0.093)
Village Population Density (log)	0.130*	0.128*	0.129	0.063	0.141	0.012	0.104
	(0.075)	(0.074)	(0.078)	(0.081)	(0.092)	(0.007)	(0.077)
Village Population (log)	0.128	0.187	0.197	0.108	0.142	0.058***	−0.243*
	(0.113)	(0.165)	(0.126)	(0.213)	(0.134)	(0.016)	(0.125)
Distance to District HQ (km - log)	−0.033	0.014	−0.067	−0.210	−0.005	0.023	0.220
	(0.154)	(0.217)	(0.201)	(0.312)	(0.193)	(0.020)	(0.193)
Distance to Nearest Town (km - log)	0.081	0.174	0.116	0.097	0.162	0.008	−0.039
	(0.145)	(0.159)	(0.170)	(0.227)	(0.175)	(0.021)	(0.138)
District Population (log)	0.388	−0.479	−0.132	0.810	−0.208	−0.031	0.263
	(0.286)	(0.444)	(0.412)	(0.503)	(0.425)	(0.042)	(0.345)
District Percent Rural	4.639**	0.329	1.735	4.486	1.042	0.184	1.704
	(1.788)	(1.914)	(1.901)	(3.320)	(2.103)	(0.216)	(1.927)
District Percent Ag. Labor	3.828***	2.004	3.102**	5.057**	2.713*	0.483***	1.172
	(1.235)	(1.644)	(1.344)	(2.080)	(1.404)	(0.160)	(1.549)
District Wage (R - log)	−0.309	−0.356	−0.003	0.193	−0.007	0.001	−0.155
	(0.442)	(0.417)	(0.453)	(0.859)	(0.528)	(0.066)	(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 3.7: Wages and Distance to Same Border

	(1) All Villages	(2) Within 15 km	(3) Within 12 km	(4) Within 9 km
Post times NREGS	0.199***	0.226***	0.230***	0.229**
	(0.065)	(0.074)	(0.082)	(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

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 3.7 presents additional 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 3.6, 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.

largest possible effect seen in the data is much smaller. Standardizing the variables results in more intuitive values. Results available upon request.

Figure 3.7: Public Works in 2008

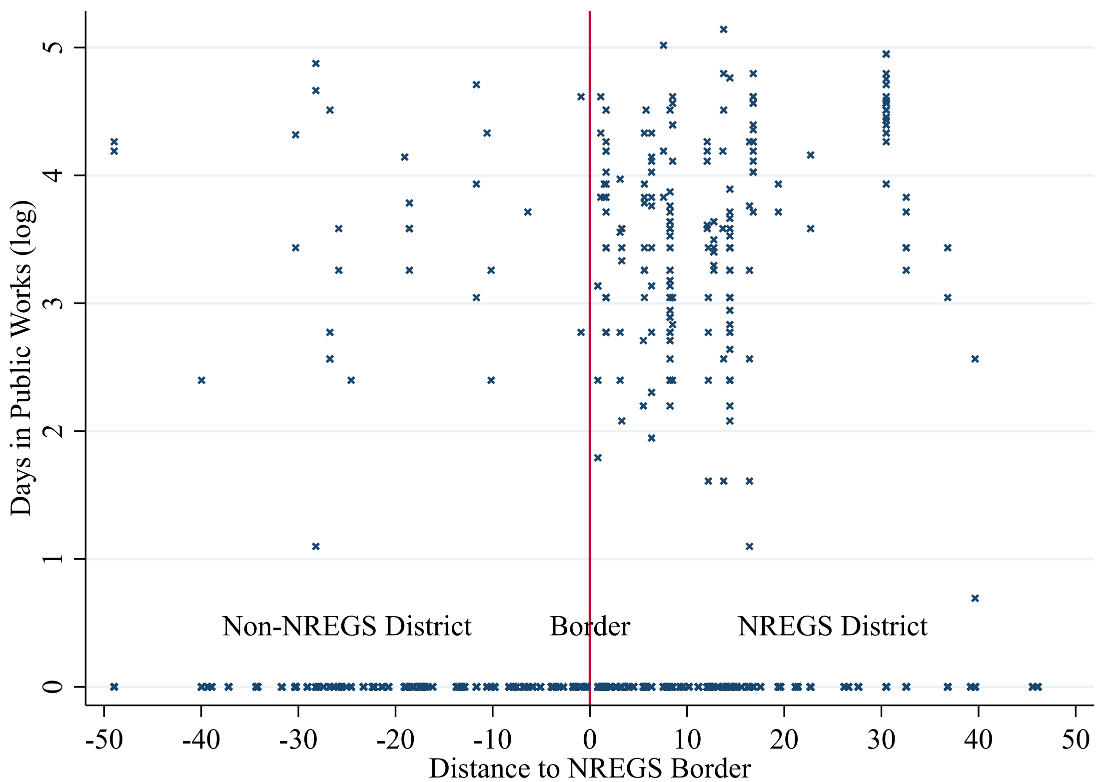

Another possible explanation for these effects is program leakage. In other words, we may see no effect at the border because untreated households are somehow accessing NREGS even though their district does not yet have the program. Figure 3.7 presents suggestive evidence that this is not the case. In 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 3.8 presents graphical evidence 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 3.8 – 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.

[23] We 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.

Figure 3.8: Travel Time to Non-Farm Employment (2008)

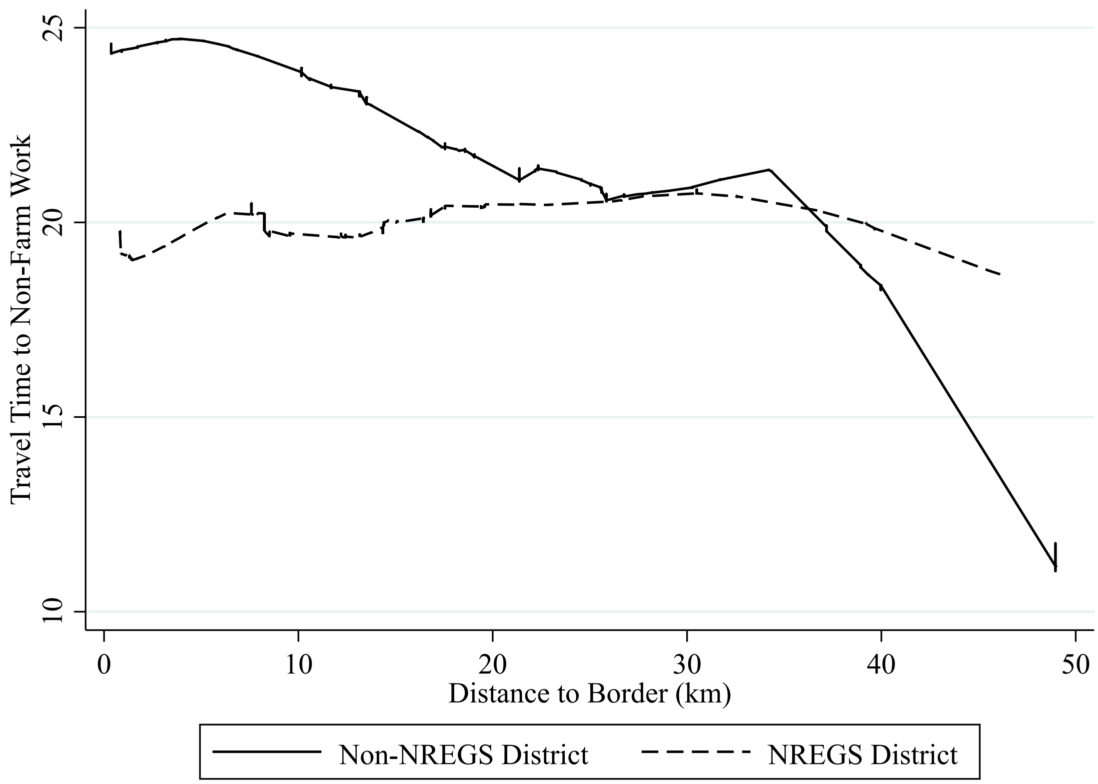

The y-axis is travel time to non-farm employment (minutes) in 2008. The x-axis is distance to border between NREGS and non-NREGS districts.

3.5.2 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 3.9 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 estimation to only individuals with observed wages. As such, this greatly increases the number of observations. Table 3.8 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

Figure 3.9: Private-Sector Labor by Wave

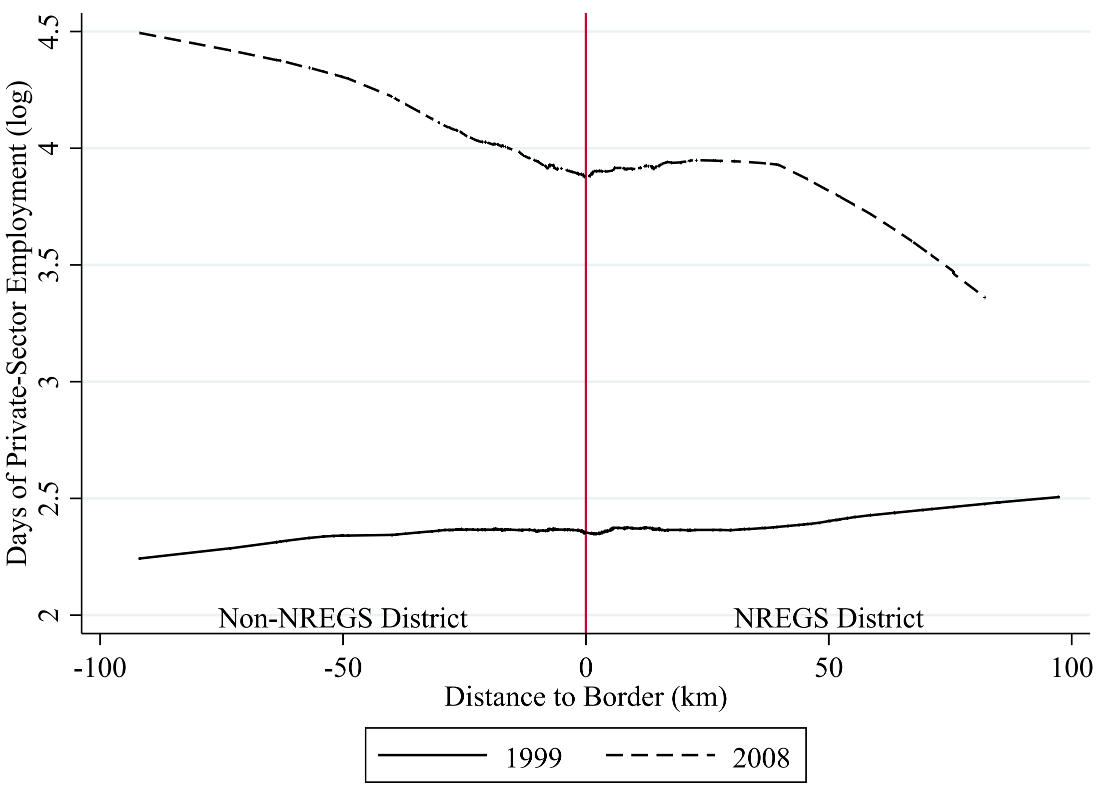

The y-axis is (log of) days of private-sector employment. The x-axis is distance to border between NREGS and non-NREGS districts.

Table 3.8: Effects of NREGS on Private-Sector Employment

	(1) All Villages	(2) Within 9 km	(3) Outside 30 km
Post times NREGS	−0.120	0.173	−0.341**
	(0.118)	(0.175)	(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

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 3.10. The pattern in Figure 3.10 is similar to the wage and private employment changes shown in Figure 3.4 and Figure 3.9, as there again appears to be a relative increase in the outcome as we move towards the interior of NREGS districts. Table 3.9 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

Figure 3.10: Household Income (2006)

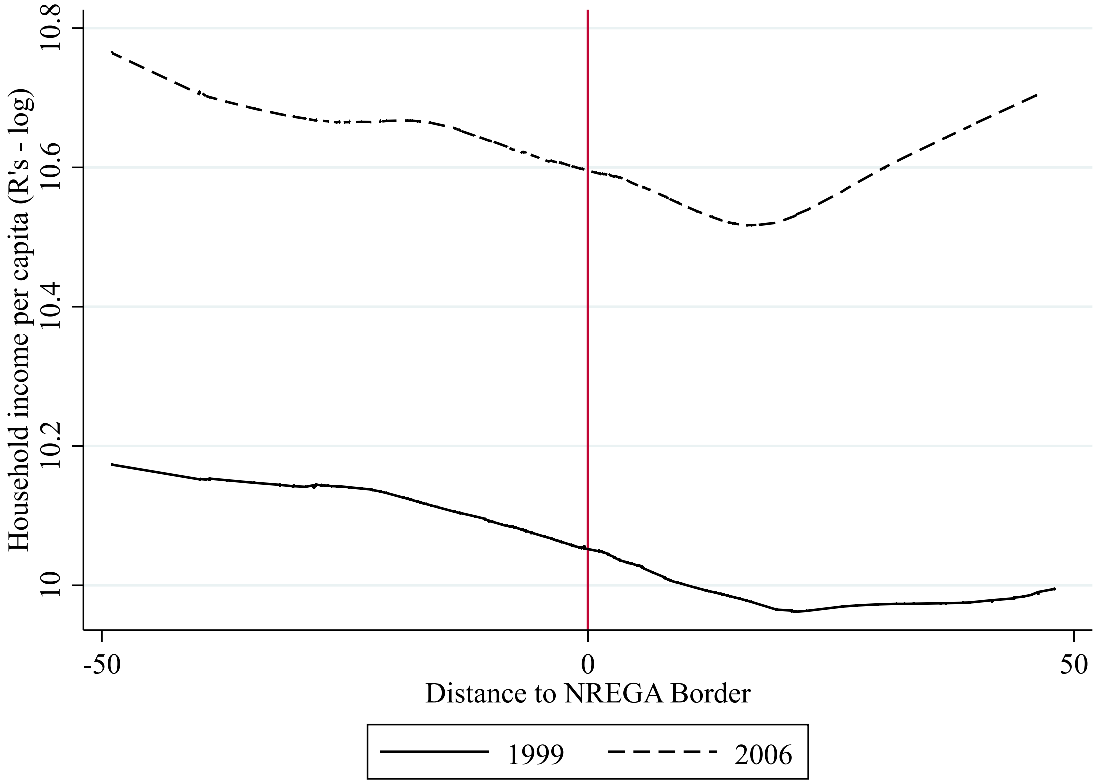

The y-axis is (log of) per capita income at the household level. The x-axis is distance to border between NREGS and non-NREGS districts.

overall pattern clearly indicates that wage changes are an important drive of consumption changes in rural India.

3.6 Conclusion

Table 3.9: Effects of NREGS on Household Income

	(1) Outside 30 km	(2) Within 9 km
Post times NREGA	0.439**	0.002
	(0.189)	(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

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.

Chapter 4

Should Farmers Farm More? Evidence from Malawi

According to standard economic theory, households will equate the marginal, though not necessarily the average, revenue product of an input across different activities within the household. However, this result relies on restrictive assumptions, several of which may be violated in developing countries. Using data on agricultural plots and non-farm enterprises in Malawi, we test whether the marginal revenue product of labor (MRPL) is equal across agricultural and non-agricultural production within a household. We control for many household characteristics using household fixed effects and find agricultural MRPL to be consistently higher than non-farm MRPL. We show both production risk and price risk affect farmer decision-making, with the latter partially offsetting the former in smaller, subsistence households. These results stand in contrast to a large body of literature that finds the average product of labor to be higher in the non-farm sector than the agricultural sector. Our findings are not a function of our data or sample selection, as we are able to recreate previous findings when using the average product. This suggests a rethinking of how we measure the sectoral productivity gap may be warranted, as our results point to substantially different policy prescriptions compared to the prior literature.

Keywords: Labor productivity, agriculture, non-farm production, productivity gap

JEL Codes: J24,J43,O13,R23

4.1 Introduction

A stylized fact of the development process is that agriculture's share of GDP decreases as a country develops (Lewis, 1954; Ranis and Fei, 1961). This relationship holds in the cross-section, with relatively more developed countries deriving a smaller percentage of GDP from agricultural sources (Chenery et al., 1975; Gollin et al., 2014). Moreover, even within countries, the non-farm sector tends to be more productive – as measured by the average product of labor – than the agricultural sector (Gollin et al., 2014; McCullough, 2017; Young, 2013). Given these persistent empirical patterns, it is perhaps no surprise that many development policies focus on non-farm growth; the recent microfinance revolution is but one example of this (Armendáriz and Morduch, 2010).

However, do the policy prescriptions necessarily follow from the empirical evidence? In other words, does the empirical evidence really suggest that reallocating labor towards the non-farm sector would increase incomes? Since a reallocation of labor and its effect on income relate to the marginal product, and not the average product, a higher average product in the non-farm sector does not necessarily imply a reallocation of labor is warranted. Whether the marginal product suggests the same likely relates to the reasons that households operate non-farm and agricultural enterprises simultaneously. A large body of research shows that agricultural households diversify into non-farm self-employment for a number of reasons, including production shocks, household shocks, agricultural seasonality, and missing markets (Barrett et al., 2001; Haggblade et al., 2010; Lanjouw and Lanjouw, 2001; Merfeld, 2018; Nagler and Naudé, 2014). In addition, diversification seems to be the norm, not the exception (Davis et al., 2017). Under these scenarios, households may not be moving into the non-farm sector chasing profits – as aggregate statistics suggest they might – but may instead be pushed into the sector due to a lack of more remunerative options and a desire to mitigate risk, which may actually lead to a lower marginal product of labor in non-farm production. We revisit this agricultural/non-farm productivity question and focus on the marginal product, as opposed to the average product. We show that the empirical choice between the two measures is an important one.

Regardless of the exact motivation for diversification, standard economic theory of profit maximization predicts the equality of the marginal revenue product of labor (MRPL) across productive activities – in the current context, agricultural and non-farm production – as well as the equality of MRPL with the market wage; if labor were allocated in any other way, it would be possible to increase profits by reallocating labor. However, this result relies on restrictive assumptions, several of which may be violated in developing countries, especially in agricultural households in rural areas. This is especially true for assumptions such as complete markets, which refer not just to labor markets, but also credit and insurance markets, which are unlikely to be complete in many contexts. As such, whether households equate marginal products across productive activities is an empirical question. Moreover, the answer to this question is not only interesting in its own right, but is integral to labor supply estimation (Abdulai and Regmi, 2000; Barrett et al., 2008; Jacoby, 1993; Seshan, 2014; Skoufias, 1994) and can even shed light on some of the underlying market conditions which characterize production in rural areas of developing countries. This, in turn, may help us better understand why households diversify into non-farm production and develop more appropriate development interventions. In this paper, we test this assumption of MRPL equality using household survey data from Malawi. In the first test of MRPL equality across agricultural and non-farm production within households, we show that this common assumption fails for the median household. Moreover, the patterns of failure allow us to characterize some important aspects of the environment in which these households operate.

The assumption of market completeness is challenged by the literature examining the agricultural household model (Singh et al., 1986). Under complete markets, agricultural households are able to treat the production and consumption decisions as recursive; households first maximize expected profit in production before making consumption decisions. This result, known as separation, suggests simple tests for market completeness. While early research was unable to reject this hypothesis Benjamin (1992), more recent research suggests otherwise (Dillon and Barrett, 2017; Dillon et al., 2017; LaFave and Thomas, 2016). This finding casts doubt on the assumption of complete markets that drives separation in the agricultural household model. As such, deviations from MRPL equality need not suggest irrational behavior on the part of households. In particular, differing risk profiles of production can lead to deviations from MRPL equality (Barrett et al., 2008; Stiglitz, 1974). If production risk differs across activities, households optimize by equalizing their expected marginal utilities across activities. Additionally, price risk – uncertainty over the market price for a good – can also affect MRPL equality. Barrett (1996) shows that price risk can affect households differently since it is likely to be correlated with production risk. In particular, households are predicted to behave differently depending on whether they are net buyers or net sellers of crops, with net sellers more likely to exhibit what we traditionally associate with risk: an underallocation of labor to the risky activity. The interplay of these two types of risk thus predicts substantial heterogeneity, a prediction confirmed in this paper.

To test this assumption of equality of the marginal revenue products of labor across household activities, we use three waves of the Malawi Integrated Household Survey (IHS). Without making further assumptions regarding market completeness, equality is only predicted for households that operate both agricultural and non-farm enterprises simultaneously. As such, we begin with summary statistics comparing households that operate both types of enterprises in the same wave to all households. There is little evidence that the relevant subsample is better or worse off than other Malawian households, as we find contrasting differences by enterprise type.

We then examine whether the marginal product of labor is equal across agricultural and non-agricultural production within a household. To the best of our knowledge, this is the first paper to implement such a test. We control for unobservable household characteristics using household fixed effects. We estimate a number of specifications, including: separate production functions for plots and non-farm enterprises; a pooled production function; and a collapsed production function.[1] Importantly, we first show that the average product of labor is consistently higher in non-farm production than agricultural production. That is, there is nothing unusual about our Malawian subsample relative to the previous literature. However, in contrast to the conclusions drawn by prior work, we consistently find agricultural MRPL to be higher than non-farm MRPL; in our preferred specification, agricultural MRPL is four times higher than non-farm MRPL. We show that this difference does not appear to be driven by temporal variation in labor allocation; the estimated difference is consistent when we split the sample into high non-farm labor months and low non-farm labor months.

We then show that focusing on a single statistic obscures substantial heterogeneity. Focusing on the MRPL estimates at the household/wave level, we show that a number of production characteristics significantly predict deviations from MRPL equality. We find evidence that price risk plays an important role in household labor allocation. In particular, deviations from equality are much higher for plots planted with tobacco and cotton, two pure cash crops, than for plots planted with maize, a common subsistence crop in Malawi. Interestingly, households that hire for agricultural production have much higher deviations from equality (with agricultural MRPL being higher than non-farm MRPL) than households that do not hire for agricultural productions. However, this relationship is reversed for households that hire for non-farm production: MRPL differences are much higher in

[1]We discuss these specifications at length in section 4.3.

households that do not hire. These differences by household type are driven completely by those households that reported hiring; households that do not report hiring for agricultural nor non-farm production are, unsurprisingly, very similar. These results suggest that households buying labor on the market make substantially different labor allocation decisions than subsistence households.

We also split the sample by output, acreage, and crop sales. For all three variables, households above the median value show substantially larger MRPL deviations than households below the median: for output, we are unable to reject the null hypothesis of MRPL equality for households with below median total output, while the MRPL difference is around twice as large for the upper half of the distribution (relative to the lower half) of both acreage and crop sales. Insofar as these variables are proxies for market access, this evidence is consistent with price risk (Barrett, 1996).

Consistent with production risk, we also find evidence that rainfall variability is associated with higher agricultural MRPL and suggestive evidence that it is also associated with deviations from equality. We only have rainfall variables for the first and second waves; higher rainfall variation (as proxied by the rainfall coefficient of variation, CV) is positively correlated with deviations from equality in these two waves. Moreover, when we assign wave two CV values to wave three households, we find similar deviations, though substantial measurement error hinders our ability to make any inferences. Overall, these results reinforce the commonly held belief that agriculture is risky and decreasing that risk might allow some farmers to increase expected incomes.

Our results are especially salient to the literature examining the "productivity gap" between the non-farm and agricultural sectors in developing countries (Gollin et al., 2014; McCullough, 2017; Young, 2013). Unlike those studies, which focus on all forms of non-farm and agricultural production, we focus only on agricultural and non-farm household production in households that operate both types of enterprise simultaneously. Nonetheless, we show that this restriction does not affect the productivity gap as commonly defined: the average product of labor is higher in non-farm production than in agricultural production in our sample, as well. However, given that we find the marginal product of labor to be *lower* in non-farm production than agricultural production, our results suggest we revisit the commonly held assumption that non-farm production is more "productive" at the margin than agricultural production in developing countries.

The rest of this paper is organized as follows. In section 4.2, we present a simple model that informs our study. We discuss methodology and summary statistics in section 4.3. We present results of our analyses in section 4.4 before concluding in section 4.5.

4.2 Theory

In the model that follows, we omit hired labor for simplicity.[2] We begin by elaborating a simple model, assuming no risk. We borrow much of the annotation used below from Barrett (1996). The household's problem is:

$$\max_{N,S,h,L_f,L_n,L_M,z_f,z_n} \quad u(N,S,h) \tag{4.1}$$

subject to

$$\begin{aligned} N + p_S S \leq & Y^* \\ Y^* = & p_S f(L_f, z_f; k_f) - p_{z_f} z_f + p_n n(L_n, z_n; k_n) - p_{z_n} z_n + w L_m \\ \bar{T} \geq & h + L_f + L_n + L_m, \end{aligned}$$

[2] In our data, only about 13 percent of agricultural plots and only 11 percent of non-farm enterprises reported hiring any labor. In the empirical specifications, we pool hired and family labor.

where Y^* is endogenous income; N is a non-staple good, available only through market purchase, with price normalized to one; p_S is the price of a staple good, S, available either through the market or through home production, with production function $f(\cdot)$, which in satisfies the common assumptions of differentiability and concavity; L_f is days of agricultural labor; h is leisure; z_f is other agricultural variable inputs; k_f is a vector of exogenous characteristics affecting agricultural output, including soil quality and weather; p_n is the price of the non-farm enterprise output, whose production function is given by $n(\cdot)$, which also satisfies the common assumptions; L_n is non-farm enterprise labor; k_n is exogenous characteristics affecting non-farm output; w is the market wage; L_m is days worked on the market; and $\bar{T}$ is the total time endowment.

At an interior solution – that is, for households that operate both types of enterprises and work on the market – both constraints will bind. The first order conditions are:

$$p_S \frac{\partial f}{\partial L_f} = p_n \frac{\partial n}{\partial L_n} = w \tag{4.2}$$

Focusing on labor, this is the standard treatment in the literature: at the optimum, the marginal revenue product of labor is equated across all activities and is equal to the wage rate, assuming the individual is active on the market.

However, if an individual does not work on the market or is not able to work as many days as preferred – perhaps due to labor market frictions – then the above equalities will no longer include the wage rate, but rather the shadow wage,

$$w^* = w + \frac{\lambda_B}{\lambda_T}, \tag{4.3}$$

where λ_B is the Lagrange multiplier on the budget constraint and λ_T is the multiplier on the time constraint (Barrett et al., 2008). This again assumes the household operates at least one non-farm and one agricultural enterprise. In this scenario, the marginal revenue products of labor across activities within the household are still equated under the most common assumptions. That is,

$$p_f \frac{\partial f}{\partial L_f} = p_n \frac{\partial n}{\partial L_n}. \tag{4.4}$$

It is this prediction of equality of marginal revenue products of labor across activities within the household that we test in this paper.

This result need not hold in the presence of risk. Here we discuss two types of risk that affect input allocation: output risk and price risk. Output risk is the risk associated with uncertainty in production. While farmers, for example, make agricultural decisions – planting, applying fertilizer, allocating labor, etc. – based on what they expect to happen, idiosyncratic shocks can cause actual outcomes to deviate from expected outcomes. Price risk, on the other hand, relates to uncertainty regarding the price of agricultural outputs following harvest. Under these circumstances, the above model needs to be amended; the household will maximize expected utility and the expectation will be taken over output.

To see this, we borrow the below model, with slight changes, from Barrett (1996). In a two-period model, households make labor allocation decisions before post-harvest prices are revealed. The household's problem is

$$\max_{h, L_f, L_n, L_M, z_f, z_n} \quad E\left(\max_{N,S} u(N, S, h)\right)$$

subject to

$$\begin{aligned} N + p_S S \leq & Y^* \\ Y^* = & p_S f(L_f, z_f; k_f) - p_{z_f} z_f + p_n n(L_n, z_n; k_n) - p_{z_n} z_n + w L_m \\ \bar{T} \geq & h + L_f + L_n + L_m, \end{aligned}$$

where all variables are defined as in Equation 4.1. In this model, households must make all labor allocation decisions prior to the realization of prices. By duality, we can instead work with the indirect utility function, $V(h, Y^*, P)$ – where P is a vector of all prices – which is homogeneous of degree zero in income and prices. The household's problem is now,

$$\max_{h, L_f, L_n, L_M, z_f, z_n} E\left(V(h, Y^*, P)\right) \tag{4.5}$$

subject to

$$Y^* = p_S f(L_f, z_f; k_f) - p_{z_f} z_f + p_n n(L_n, z_n; k_n) - p_{z_n} z_n + w(\bar{T} - L_f - L_n - h)$$

Here, we present only the main result. Interested readers are directed to Barrett (1996) for a full exposition.

Focusing on farm labor, Barrett (1996) shows that $cov(V_Y, P_S)$ is crucial to predicting labor allocation. In pure producer theory, Arrow-Pratt income risk aversion implies $cov(V_Y, P_S) < 0$ (Sandmo, 1971) and, thus, that a household under uncertainty produces less – and thus applies less labor and has a higher MRPL – than would be the case without uncertainty (Barrett, 1996). However, this result focuses only on pure producers and does not take into account a household model of production, in which households can sell or *consume* the agricultural goods they produce. The insight of Barrett (1996) is that labor allocation hinges on the importance of the crop to households as a consumption good relative to the importance of the crop as a source of income (cash). In particular, he shows that if a household is a net *seller* of a crop, then it will underemploy labor: $MRPL_{ag} > w$. On the other hand, if a household is a net *buyer* of a good, then it will overemploy labor: $MRPL_{ag} < w$.

Note that, by assumption, households do not consume their non-farm good. Thus, non-farm production will follow the familiar result if non-farm production is risky: $MRPL_{nf} > w$. However, if non-farm production is a relatively risk-free activity, then $MRPL_{nf} = w$. Putting together these predictions, we can say, for net buyers of crops, that: $MRPL_{ag} < MRPL_{nf}$. However, for net sellers of crops, the direction is ambiguous and depends on the relative risk and importance of each productive activity.

4.3 Methods and Data

The basic steps involved in testing MRPL equality are as follows, and are similar to those in Linde-Rahr (2005), who studies allocative efficiency across plots planted with different crops. First, we estimate production functions for both agricultural and non-farm enterprises. Second, we compute the marginal revenue products of labor across different activities within the household. With these MRPL estimates, we then explicitly test whether marginal revenue products of labor – or shadow wages – are equal across activities within the household.

In order to estimate the production functions, we present production function results using both a Cobb-Douglas production function and a translog production function. However, given that we reject the nested Cobb-Douglas within the translog production function, we present MRPL results only for translog production function estimates. We estimate:

$$\begin{aligned}\ln Q_{iht} = \alpha_h + I(\textit{Agriculture} = 1) \times (\sum_j \beta_j \ln \gamma_{jiht} + \frac{1}{2}\sum_j\sum_k \beta_{jk} \ln \gamma_{jiht} \ln \gamma_{kiht} + \delta C_{iht} \\ + D_{dt} + \eta_m) + \sum_j \beta_j \ln \gamma_{jiht} + \frac{1}{2}\sum_j\sum_k \beta_{jk} \ln \gamma_{jiht} \ln \gamma_{kiht} + \delta C_{iht} + D_{dt} + \eta_m \\ + I(\textit{Agriculture} = 1) + \varepsilon_{iht},\end{aligned} \tag{4.6}$$

where α_h is an intercept that varies by household (household fixed effects); $I(\textit{Agriculture} = 1)$ is a dummy variable equal to one if the observation is an agricultural plot and equal to zero if the observation is a non-farm enterprise; γ_{jiht} and γ_{kiht} are inputs j and k for enterprise *i* in household *h* in wave *t*; C_{iht} is a vector of controls that may affect output and which differ depending on whether the enterprise is (non-)agricultural; D_{dt} is district/wave fixed effects; η_m is a set of dummy variables indicating the month of interview; and ε_{iht} is a conditional mean-zero error term. We include labor (log of days), acres (log), and fertilizer (log of kg) as productive inputs in the agricultural production functions and labor (log of days) and total costs (log of MWK) as productive inputs in the non-farm production functions. We set land and fertilizer to zero for all non-farm enterprises and non-farm costs to zero for all agricultural plots.

For agricultural plots, we restrict attention to plots planted with a select number of crops. Part of the process of constructing the marginal revenue products involves finding prices for the agricultural output. We construct these prices by taking medians at the lowest administrative level of aggregation with a sufficient number of observations. The crops we choose offer a sufficient number of observations with which to create median prices.[3] In contrast, we do not need to restrict estimation in a similar way for non-farm enterprises; entrepreneurs are directly asked about their total output. Finally, we trim the top one per-

[3] We elaborate on this restriction below.

cent of agricultural and non-farm output and labor before estimating production functions, though we report non-trimmed results in the appendix.

For both the Cobb-Douglas and translog specifications, we pool the data and estimate a single production function for both agricultural and non-farm enterprises, but we allow the effect of all variables – other than the fixed effect – to vary by type of enterprise. We employ household fixed effects to identify the production functions. Since the identifying assumptions for a pooled sample are somewhat restrictive, we also estimate separate agricultural and non-farm production functions.

After estimating the production functions, we then calculate the marginal revenue product of labor for each input. Using the translog specification,[4] we construct our MRPL estimates for agricultural plots as

$$\frac{\partial Q}{\partial L} = \frac{\hat{Q}_{iht}}{L_{iht}} \left[\beta_L + \beta_{LL} \log L_{iht} + \beta_{LA} \log A_{iht} + \beta_{LF} \log F_{iht} \right], \tag{4.7}$$

where $\hat{Q}_{iht}$ is predicted output, β_{LL} is the coefficient on the labor squared term, β_{LA} is the coefficient on the interaction between labor and acreage, and β_{LF} is the coefficient on the interaction between labor and fertilizer use. We use predicted output as it is the best estimate we have of the farmer's expected output. We calculate the MRPL for non-farm enterprises as

$$\frac{\partial Q}{\partial L} = \frac{\hat{Q}_{iht}}{L_{iht}} \left[\beta_L + \beta_{LL} \log L_{iht} + \beta_{LC} \log C_{iht} \right], \tag{4.8}$$

where β_{LC} is the coefficient on the interaction between labor and costs.

The production function in Equation 4.6 is at the plot/enterprise level. Thus, estimating MRPLs using Equation 4.7 and Equation 4.8 presents a problem when we attempt to

[4]The translog results strongly reject the nested Cobb-Douglas, especially for non-farm production. As such, we only present MRPL estimates from the translog. We discuss this below.

aggregate to the household. In particular, many households operate more than one plot in the same wave. Since we construct MRPL for each plot and enterprise separately, we then need to aggregate these to the household/wave. We do this in two ways. First, we take the simple median across plots within each household. If a household operates an even number of plots, we use the mean of the two middle plots to construct MRPL. We construct MRPL similarly for non-farm enterprises, though most households operate only one, such that aggregation is not as problematic for non-farm production. Second, we compute the household's weighted average MRPL, weighting by labor allocation.

Aggregating over plots and enterprises is an important part of MRPL estimation. However, it is not clear what the "correct" aggregation method is. To bypass this problem completely, and due to issues with recall bias raised by Arthi et al. (2018),[5] we also estimate production functions in which we aggregate (separately) all plots and all enterprises to the household level. We sum output, labor, acreage, fertilizer, and non-farm costs to the household level. For crop dummies,[6] we collapse the data and leave these as indicator variables for whether the household grew that crop in that wave. For the plot characteristics, we include continuous variables for the percent of total household land with each characteristic.

MRPL equality holds only for households that operate both types of enterprises.[7] Therefore, we only use households that operate both non-farm and agricultural enterprises in the same wave to construct MRPL estimates. However, we present estimates of production functions using these households as well as separate estimates using all households.

Finally, in order to conduct inference, we bootstrap the standard errors with 1,000 repli-

[5]Arthi et al. (2018) find that reported labor allocation is mismeasured at the plot level. However, when labor is aggregated to the household, this mismeasurement largely disappears.
[6]We discuss these controls more below.
[7]Economic theory makes no predictions regarding marginal products in different households unless markets are complete, nor does it make predictions regarding the equality of marginal products across time.

cations. Since we are interested in a multi-step estimator, bootstrapping offers a convenient method to estimate valid standard errors. We employ household fixed effects and, as such, we set up the bootstrap to draw households, including all non-farm and agricultural enterprises operated by that household, across all waves.

4.3.1 Data

We use the Malawi Integrated Household Survey (IHS) in this paper. The IHS is part of the World Bank's Living Standards Measurement Study (LSMS) program. The IHS data consists of three waves: the first wave was collected from 2010-2011, the second in 2013, and the third in 2016-2017. There is a common panel sample in all three waves, such that we are able to follow some households across three separate years. However, the first and third waves also had large cross-section components, which dwarf the panel component in size. After restricting our sample and looking only at households that operate both types of enterprises in the same wave, we are left with 3,786 unique households. Of these, 3,506 households appear in one wave, 482 households appear in two waves, and 117 households appear in all three waves. These numbers correspond to 5,873 plots and 3,750 enterprises in the one-wave households, 881 plots and 547 enterprises in the two-wave households, and 250 plots and 129 enterprises in the three-wave households. Thus, our total sample with which we estimate production functions consists of 7,004 plots and 4,426 enterprises.

Many of the variables used in estimation come from different modules in the IHS. As such, certain variables are at different levels of aggregation (i.e. household, household-plot, or household-plot-crop). For example, although crop output is reported at the plot-crop level, labor is only reported at the plot level. Thus, missing plot-crop observations require that we drop the entire plot from the sample. In the following sub-sections, we document some of these idiosyncracies and the resulting decisions.

4.3.2 Output

The key dependent variable in the production functions is output value. The Malawi IHS, like many household surveys, asks farmers for output in kilograms, not in value. Since our methodology is designed to estimate the marginal revenue product of labor, we must construct crop prices to value output. The survey also asks households about crops sold during the year, including amount sold and the value of crops sold. Using these questions, we construct a price variable at the household level.

However, most households do not report selling crops. As such, we must construct aggregate prices. We construct aggregate prices at four separate levels – the enumeration area, the traditional authority, the district, and the region – by taking the median crop price at each level of aggregation separately. We then assign prices to households using the lowest level of aggregation at which there are at least five valid price observations.[8] If any regions have less than five prices observations for any given price, we assign a missing price value for all observations of that crop in that district. In this way, we end up with prices for 13 crops: maize, tobacco, groundnut, rice, sweet potato, potato, beans, soya, pigeon peas, cotton, sunflower, nkhwani, and tomato.

4.3.3 Labor and Selection of Plots

We use days of labor as the independent variable of interest. This variable includes both family and hired labor, which we aggregate into the single variable. However, family labor generally predominates, as is clear in the summary statistics below. A major issue is that labor is reported at the plot level, while crop output is reported at the plot-crop level. This means that the labor inputs we observe on a given plot are applied to all crops on that

[8] In the appendix, we report results requiring at least ten valid price observations and without any price requirements. Our conclusions are unchanged.

plot, and we are not able to disaggregate that labor by crop. As such, after constructing price observations, we also drop any plots that are planted with at least one crop with a missing price value. In this way, the entirety of labor allocated to the plot is applicable to the entirety of the output value we construct. This restriction does not appreciably affect our sample size; we lose less than ten percent of plots in each wave.

4.3.4 Other Controls

In addition to the productive inputs, we also control for a number of agriculture variables, including plot quality (as reported by the farmer), plot type, plot erosion, plot slope, and whether the plot is located in swamp/wetlands. In agricultural regressions, in addition to days of labor, we include the size of the plot (log of acres) and amount of fertilizer applied to the plot (log of kilograms) as productive inputs. In non-farm regressions, we include one separate productive input in addition to labor: log of total costs. This variable is directly asked of all households with a non-farm enterprise. Additional industry controls include indicator variables for industry and a single indicator variable for whether the non-farm enterprise has electricity. The use of value for a productive input could bias the coefficients in our production functions if input prices vary by region and/or household (Jacoby, 1993), although this is also somewhat of an issue when constructing output values for agriculture, as well. In order to mitigate these concerns, we include district/wave fixed effects to help alleviate both regional and temporal differences in input prices.

We believe including costs – even with caveats – is preferable to estimating the production functions without them. We estimate Cobb-Douglas production functions both with and without non-farm costs included as a productive input to test this assumption. The estimates with costs included are presented below in the results section (Table 4.2). Restricting the sample to only those households that operate both types of enterprises in the same wave, the coefficient on non-farm labor is 0.200. When we remove costs and re-estimate the production function,[9] the coefficient on labor increases to 0.342, an increase of more than 70 percent. As such, we conclude that the benefits of including costs outweigh the possible risks.

Finally, the dependent variable for non-farm production is constructed using a single survey question, which asks respondents for total output in the 30 days prior to the interview. We include month of interview fixed effects in order to control for any seasonality in labor allocation across the year. However, we also estimate production functions on separate subsamples, depending on month of interview, and present these results below.

4.3.5 Summary Statistics and Sample Restrictions

Summary statistics at the household/wave level are shown in Table 4.1. To calculate these statistics, we collapse all agricultural and non-farm statistics to the household level before taking logs. The first column presents statistics for all households that are in the Malawi IHS data and meet our criteria for calculating plot output, but are not in our final sample because they do not operate both types of enterprises in the same wave. The second column includes all households that are in our final sample. The third column presents the p-value of tests for equality across the two samples.[10]

Of the households that meet the output criteria but are not in our sample, approximately 86 percent of them have plots but no non-farm enterprise. Households in our final sample had slightly more output (that we could calculate) than households that were not in our sample. Moreover, this difference is magnified when we look at the labor statistics; dropped households actually apply more labor – mostly more family labor – but have lower output,

[9]Results are not shown but are available upon request.

[10]We construct these p-values by regressing each variable on a dummy variable indicating whether a household/wave observation is in our final sample, clustering the standard errors at the household level.

Table 4.1: Summary Statistics

	(1) Dropped Households	(2) Final Sample	(3) Diff (p-value)
Agricultural Production Statistics			
Household has plot in sample	0.856	1.000	
	(0.351)		
Ag output (2010 MWK - log + 1)	10.005	10.186	0.000
	(1.261)	(1.233)	
Total labor (log)	4.348	4.268	0.000
	(0.831)	(0.900)	
Total family labor (log + 1)	4.254	4.086	0.000
	(0.972)	(1.137)	
Total hired labor (log + 1)	0.422	0.771	0.000
	(0.996)	(1.261)	
Household hired for ag production (yes = 1)	0.179	0.316	0.000
	(0.383)	(0.465)	
Fertilizer (kg - log + 1)	1.377	1.499	0.000
	(1.594)	(1.609)	
Acres in sample	1.738	2.178	0.310
	(18.724)	(26.109)	
Maize in sample	0.930	0.923	0.160
	(0.256)	(0.267)	
Tobacco in sample	0.101	0.077	0.000
	(0.301)	(0.266)	
Non-FarmProduction Statistics			
Household has non-farm ent. in sample	0.144	1.000	
	(0.351)		
NF output (2010 MWK - log + 1)	9.613	9.094	0.000
	(1.522)	(1.391)	
Total labor (log)	3.085	2.868	0.000
	(0.809)	(0.852)	
Total family labor (log + 1)	2.954	2.790	0.000
	(0.803)	(0.845)	
Total hired labor (log + 1)	0.400	0.243	0.000
	(1.114)	(0.856)	
Household hired for NF production (yes = 1)	0.124	0.083	0.000
	(0.330)	(0.275)	
Last monthly costs (2010 MWK - log + 1)	8.500	7.808	0.000
	(2.768)	(2.825)	
Household Characteristics			
Male household head (yes = 1)	0.735	0.803	0.000
	(0.441)	(0.398)	
Household size	4.667	5.144	0.000
	(2.141)	(2.127)	
Total acres owned	3.621	4.941	0.159
	(65.536)	(51.902)	
N	19154	4105	

Standard deviations are in parentheses. Statistics are at the household/wave level. Agricultural and non-farm statistics were collapsed to the household/wave level before taking logs. The "Dropp Households" columns include all households that meet our criteria for valuing agricultural output (on at least one plot) or which operate a non-farm enterprise but do not operate both in the same wave. The "Final Sample" column includes the sample of households we use to calculate MRPL: those that operate at least one plot *and* at least one enterprise in the same wave. To calculate the p-value for their difference, we regress each variable in the far left column on a single dummy variable for whether the household is in our final sample or not. We cluster the standard errors at the household level.

suggesting our sample may consist of slightly better off households. Consistent with this argument, these households were also more likely to hire at least one day of labor for agricultural production and hired more labor for agricultural production overall.

Although output is higher for the group of households in our final sample, they also achieved this output on more land (though the difference across the two groups is not statistically significant, mostly due to the large variance in both groups). As such, the overall productivity question is not as clear cut as the combination of output and labor statistics suggested. Moreover, in line with the fact that maize is a common crop and foodstuff, around 93 percent of households in both groups grew maize on at least one plot that meets our output criteria. Along with acreage, maize is the only other agricultural statistic that does not differ significantly across the two groups. Interestingly, although there is suggestive evidence that final sample households are better off than other households, other households are actually more likely to grow tobacco on at least one plot. This is somewhat surprising since tobacco is a pure cash crop, which we might associate with better-off households.

The remaining 14 percent of dropped households operate non-farm enterprises.[11] While final sample households appeared to be slightly better off when looking at the agricultural statistics, we actually see the opposite in the non-farm statistics. Final sample households have lower output and only slightly higher labor allocation. Two statistics suggest that all households do not have higher output due to, for example, a reliance on non-farm production in the face of crop failure. These households are both more likely to hire at least one day and hire more days in general. Moreover, these households also have higher non-farm costs, which is suggestive evidence that these households are not simply operating more

[11] Since non-farm enterprises are never dropped due to restrictions on output – unlike agricultural plots – a household only appearing in the non-farm statistics does not imply the household does not operate any plots. Rather, it is possible that the household operates plots that get dropped in our price-creation procedure explained above.

labor- and less capital-intensive enterprises than final sample households. Results from t-tests suggest all differences in non-farm statistics are significant.

The last three variables presented are household-level variables. Final sample households are five percentage points more likely to have a male household head relative to other households. Households in our final sample are also slightly larger than the other group of households. Overall, it is difficult to come to any clear conclusions regarding how our sample restrictions impact how representative our results may be. While there are clear differences across groups (only acres in sample and maize are *not* significantly different across the two groups), the differences are sometimes tell conflicting stories. Thus, while households that operate both types of enterprises are clearly different from households that do not, it is not clear from these statistics whether they are better off or worse off.

One concern with our estimation method is that labor is seasonal, and thus we may not be making the proper temporal comparison across productive activities. This concern is compounded by the fact that agricultural labor is spread throughout the season whereas non-farm labor is only reported for the last 30 days. To help examine whether this is affecting our results, we present a graph of median non-farm output and labor by month of interview in Figure 4.1. In Malawi, there appears to be clear seasonality in labor allocation to non-farm production. In particular, both labor allocation and output appear to be lowest from February to around July, and then increase between August and January. Given that there does appear to be some seasonality in labor allocation and output, we will explore this further below, though we control for month of interview in all production function estimates.

Figure 4.1: Month of Interview and Non-Farm Production Characteristics

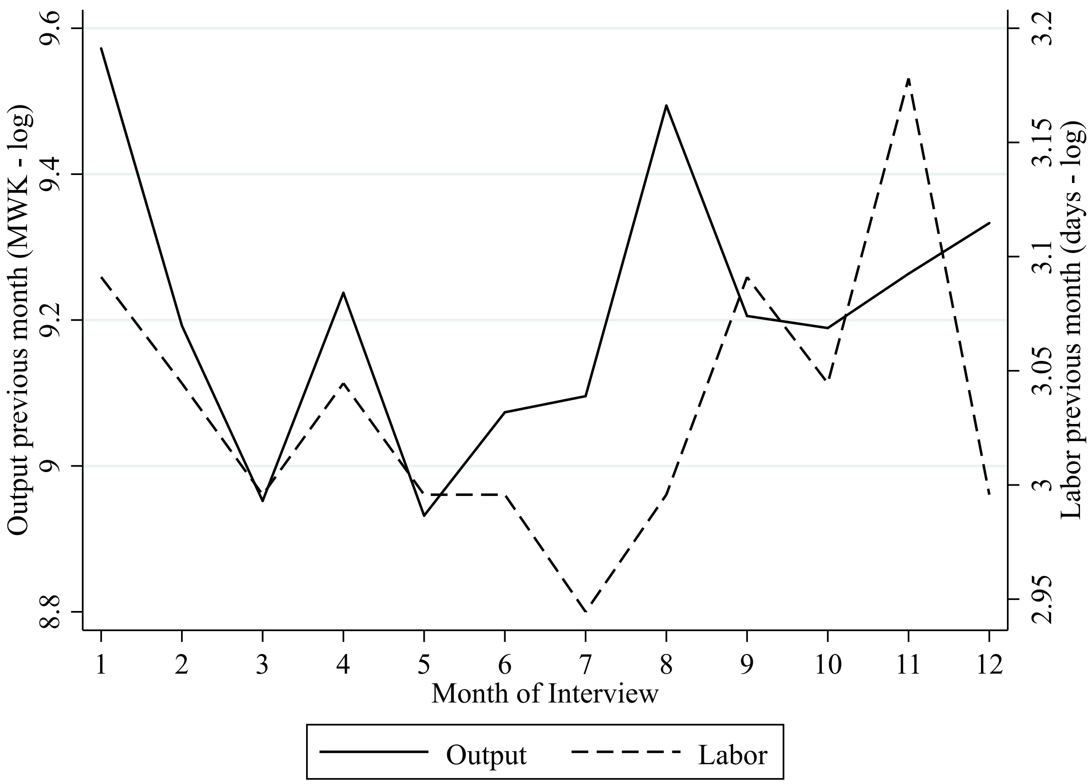

Output and labor are for non-farm production only.

4.4 Results

We begin our discussion of the results with the pooled production function estimates in Table 4.2. We present results for both the Cobb-Douglas and translog specifications as well as for all households (first two columns) and for households that operate both enterprises in the same wave, which is the subsample we use in MRPL construction (last two columns). Comparing the Cobb-Douglas results for all households (column 1) to final sample households (column 3) suggests the production technology is relatively similar. While labor appears slightly more productive – for both agricultural and non-farm production – for the select subsample of households, the opposite is true for land. The coefficients on fertilizer and non-farm costs, however, are relatively similar.

The translog results are harder to interpret given the large number of interaction terms. Moreover, it is clear moving from column two to column four that we lose a substantial amount of precision in our estimates. This is especially true for agriculture. However, the translog coefficients for non-farm production are remarkably similar for both groups of households, with the largest difference being just 0.006 (costs). Overall, it appears that production technology is relatively similar for both groups of households. A glance at the interaction terms suggests that the Cobb-Douglas specification might not be rejected for agriculture; there is only one significant interaction term (including the squared terms) in each column. The same cannot be said for non-farm production, however. In both columns two and four, the squared costs term and interaction term are significant, with the former having a t-statistic of more than 20 in each column. We formally test for the nested Cobb-Douglas in the translog specifications and present these results at the bottom of the table. Consistent with the above interpretation of the coefficients, we reject the Cobb-Douglas specification for agriculture for all households but fail to reject for our final sample. However, the F-test strongly rejects the Cobb-Douglas specification for non-farm

Table 4.2: Production Function Estimates

	All households		Final sample	
	(1)	(2)	(3)	(4)
	C-D	Translog	C-D	Translog
Ag Labor (L_a)	0.255***	0.149*	0.284***	0.027
	(0.017)	(0.088)	(0.027)	(0.125)
Acres (A)	0.365***	0.450***	0.336***	0.440***
	(0.016)	(0.088)	(0.028)	(0.145)
Fertilizer (F)	0.084***	0.250***	0.085***	0.216**
	(0.008)	(0.050)	(0.014)	(0.085)
$L_a \times L_a$		0.012		0.026
		(0.011)		(0.016)
$A \times A$		0.008		−0.005
		(0.015)		(0.022)
$F \times F$		−0.026***		−0.026**
		(0.008)		(0.013)
$L_a \times A$		−0.017		−0.020
		(0.017)		(0.029)
$L_a \times F$		−0.011		0.003
		(0.009)		(0.014)
$F \times A$		0.001		0.000
		(0.010)		(0.017)
NF Labor (L_n)	0.188***	0.303***	0.201***	0.305***
	(0.027)	(0.103)	(0.027)	(0.109)
Costs (C)	0.236***	−0.289***	0.232***	−0.295***
	(0.012)	(0.033)	(0.011)	(0.032)
$L_n \times L_n$		0.012		0.009
		(0.018)		(0.019)
$C \times C$		0.057***		0.057***
		(0.002)		(0.002)
$L_n \times C$		−0.037***		−0.036***
		(0.010)		(0.010)
Test for Nested Cobb-Douglas (p-value)				
Agriculture		0.007		0.192
Non-Farm		0.000		0.000

Standard errors clustered at the household level are in parentheses. Household fixed effects are included in all regressions. Also included are month of interview fixed effects and wave/district fixed effects. Month of interview and wave/district fixed effects are allowed to vary by type of production. In addition, we include crop dummies, plot quality variables, non-farm industry dummies, and a dummy indicating whether the non-farm industry has access to electricity. The F-tests present tests for a nested Cobb-Douglas production function in each translog; the p-value is constructed by testing whether all squared and interaction terms are simultaneously zero. The "All Households" columns include all households that operate at least one plot or at least one enterprise across all three waves of the Malawi LSMS. The "Final Sample" column includes the sample of households we use to calculate MRPL: those that operate at least one plot *and* at least one enterprise in the same wave. Output and non-farm costs are in (March) 2010 MWK.
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

production. Given these results, we use the translog specification in all results that follow.

Rather than estimating average revenue products (APL),[12] as has much of the literature, we instead estimate marginal products. Before presenting our MRPL estimates, we first verify that we find similar APL results in our final sample of households. Figure 4.2 presents kernel density estimates of APL using actual output (Panel A) and predicted output from the production functions (Panel B). In both cases, it appears that non-farm APL stochastically dominates agricultural APL. Table 4.3 presents the mean and median average product, again using both actual and predicted output. Consistent with previous literature, the average product in non-farm production is substantially higher than the average product in agricultural production. When using the mean, the APL is more than twice a large with the predicted output and more than 80 percent larger with the actual output. The difference is slightly smaller with the median: non-farm APL is around 52 percent larger when using actual output and 47 percent larger when using predicted output. All of our estimated differences are smaller than some of the previous literature, but similar to recent research using a similar identification strategy (Hicks et al., 2017). As such, the use of APL results in similar findings to those in the larger literature: it appears that non-farm production is more "productive" than agricultural production.

[12]We use the terms "average product" and "average revenue prouct" interchangeably.

Figure 4.2: Average Revenue Product of Labor

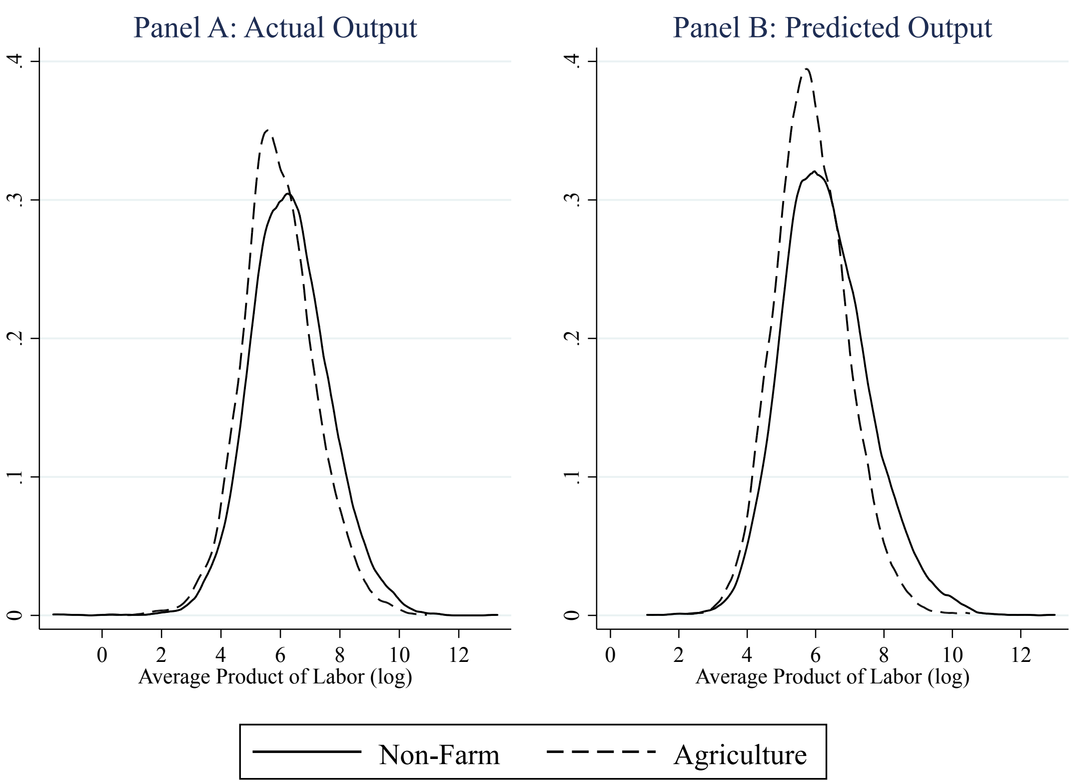

All statistics are calculated as output in MWK over number of days worked. Actual output uses reported output (non-farm) or constructed output using aggregate prices and reported harvest (agriculture). Predicted output uses output predicted from the production function estimates used to estimate the marginal product of labor in autorefpf1.

Higher non-farm average product does not imply that a reallocation of labor towards non-farm production would bring about welfare gains, however. A more appropriate statistic is the marginal revenue product of labor. We present these MRPL estimates in Table 4.4. Since MRPL equality is only theoretically predicted for households that operate both types of enterprise, we present MRPL results for only this subset of households. All MRPL estimates are in (March) 2010 Malawi Kwacha (MWK). We construct MRPL estimates using the production function in column four of Table 4.2 and present these in the first two columns of Table 4.4. Recall that we construct the MRPL estimates in two ways: taking the simple median across plots/enterprises and weighting MRPLs across plots/enterprises based on labor allocation. The simple median suggests an agricultural MRPL of about 171 MWK. In March of 2010, the exchange rate was approximately 150 MWK to USD, so this MRPL translates to slightly more than one US dollar per worker-day. The non-farm MRPL, on the other hand, is significantly smaller, at just 23 MWK. Finally, the difference, approximately 101 MWK, is highly significant.

The second column presents results using the same base production function as the first column, but aggregates the MRPL across plots/enterprises by a weighted average based on labor allocation. Two patterns emerge. First, all three MRPL estimates are substantially more precisely estimated than the simple medians. Second, the agricultural MRPL is smaller than in column one, possibly suggesting some plots with low labor allocation and high MRPL were inflating the estimated MRPL in the first column.[13] The end result is that the estimated MRPL difference is less than sixty percent as large in column two. Nonetheless, the increase in precision results in a highly significant difference of 56 MWK, or around 0.35 USD, per worker-day.

To get a sense of the magnitude of these differences, it is important to remember that these

[13]This is not unexpected. The MRPL formula requires dividing by total labor allocation on a plot. Very small reported labor allocation can thus result in very large estimated MRPL values.

Table 4.3: Mean and Median of Average Revenue Product of Labor

	(1) Actual Output	(2) Predicted Output
Panel A: Mean		
Agriculture	875	681
Non-Farm	1584	1567
Panel B: Median		
Agriculture	355	337
Non-Farm	541	495

All statistics are calculated as output in MWK over number of days worked. Actual output uses reported output (non-farm) or constructed output using aggregate prices and reported harvest (agriculture). Predicted output uses output predicted from the production function estimates used to estimate the marginal product of labor in autorefpf1.

are not purchasing power parity-adjusted exchange rates, so the relevant number used in the construction of poverty statistics would be significantly higher than 0.35 USD. For instance, the PPP-adjusted exchange rate in 2010 would result in an agricultural MRPL of around 1.25 dollars, a non-farm MRPL of around 0.4 dollars, and a difference of around 0.8 dollars. For a point of reference, Malawi's average per capita income was approximately one dollar (PPP adjusted) per day in 2010. These numbers point to possible harm from encouraging a reallocation of labor away from agricultural production and towards non-farm production. At the margin, reallocating a single day of labor in this direction is estimated to *decrease* income for the median household by around 55 Kwacha, which is a substantial percentage of average income in Malawi. This stands in stark contrast to the large body of literature on the intersectoral productivity gap.

Overall, the first two columns suggest that households could increase their income by reallocating some labor away from non-farm production and towards agricultural production. However, the estimates in the first two columns also rely on pooled production functions, in which we estimate agricultural and non-farm production functions in a single regression. This increases our statistical power, but at the cost of a more restrictive estimation strategy. We relax the restrictions on the fixed effects in columns three and four by estimating the

Table 4.4: MRPL Estimates

	Pooled Production Function		Separate Production Functions		Collapsed
	(1)	(2)	(3)	(4)	(5)
	Simple Median	Weighted by Labor	Simple Median	Weighted by Labor	Collapsed
Agriculture	127.013***	81.819***	170.996***	111.566***	99.882***
	(17.944)	(8.375)	(30.461)	(14.242)	(10.105)
Non-Farm	22.607	21.461**	19.839	18.418	30.522***
	(15.303)	(10.385)	(29.702)	(19.783)	(11.443)
Difference	100.606***	56.227***	134.491**	79.468***	68.121***
	(28.246)	(14.974)	(54.096)	(28.667)	(17.581)
N (Household/Wave)	4097	4097	4097	4097	4097

Standard errors are constructed through bootstrapping the MRPL construction 1,000 times. The bootstrap is set to draw households. MPRL estimates are in 2010 MWK. MRPL difference is constructed as agricultural MRPL minus non-farm MRPL. The MRPLs in the first and second columns are constructed by estimating a single production function but allowing the effects of all variables (except household fixed effects) to vary by enterprise type. MRPLs in the third and fourth columns are constructed by estimating separate production functions for agriculture and non-farm production. In the fifth column, MRPL is constructed by collapsing across plots and enterprises to the household/wave level and estimating a single production function, again allowing the effects of all variables (except household fixed effects) to vary by enterprise type. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

agricultural and non-farm production functions separately. Now, the fixed effect is allowed to affect agricultural and non-farm production differently, but at a cost: agricultural coefficients are now identified only by households in our sample that operate at least two plots – either in the same wave or across waves – while non-farm coefficients are identified by households that operate at least two non-farm enterprises.

This loss of power is clear in the resulting estimates: the standard errors on the agricultural and non-farm MRPLs almost double, regardless of the choice of aggregation. Reassuringly, however, the patterns and general conclusions are identical. Agricultural MRPL continues to be significantly higher than non-farm MRPL in both columns, though the estimated differences are slightly larger, 134 MWK in column three compared to 101 MWK in column one, and 79 MWK in column four compared to 56 MWK in column two.

Columns one through four all rely on somewhat ad hoc choices of MRPL aggregation. Moreover, there is some evidence that disaggregated labor statistics are more prone to bias than are aggregated labor statistics, as households tend to forget plots and distribute labor from those plots when reporting on other plots (Arthi et al., 2018).[14] As such, we present one final set of results in column five in which we aggregate all productive inputs to the household level *prior* to estimating production functions. We still pool these results and we present these MRPL estimates in column five. Again, the results are generally consistent with the previous four columns. Taking these results together, we interpret this as evidence that agricultural MRPL tends to be higher than non-farm MRPL within the household. This result stands in contrast to the agricultural productivity gap literature, which finds the average product of labor in non-farm production to be higher than the average product of labor in agricultural production. Importantly, *these differing results also suggest dramatically different policy responses.*

[14]We note, however, that it is not clear how this would affect our results, as we do not use all plots in estimation.

4.4.1 Robustness Checks

Given the relative consistency across specifications and the fact that weighting by labor (column two of Table 4.4) is not only the most precisely estimated but also the most conservative estimate of MRPL difference, we report all results using the specification in column two for the rest of the paper. In the appendix, we also explore the robustness of MRPL estimates to several variations in specifications. We estimate MRPL when using a minimum of 10 prices observations instead of five to construct aggregate crop prices (Table 4.A2) as well as no minimum number of price observations (Table 4.A3). In Table 4.A4, we present MRPL estimates when not trimming the top one percent of output and labor. Finally, in Table 4.A5, we estimate MRPL using actual output instead of predicted output. Our conclusions are unchanged and agricultural MRPL remains higher than non-farm MRPL in all models/specifications.

One concern already documented is seasonality. While we control for month of interview in all regressions, concerns nonetheless remain. Fortunately, the Malawi IHS allows us to examine whether seasonality may be affecting our results. Panel households were enumerated twice in every wave for the Malawi IHS. In general, approximately half of households responded to the non-farm module during the first visit – which tended to be during or just after the planting season – while approximately half of panel households responded to the non-farm module during the second visit, just after harvest.[15] Importantly, only panel households are visited twice. Moreover, panel households are the minority of our sample, as there are large cross-section components of waves one and three. As such, a large portion of households responded to the agricultural and non-farm modules at different visits. This is at least suggestive evidence that survey timing is unlikely to be driving our results.

As an additional – and more explicit – check, we can also estimate production functions

[15] All households should have reported agricultural output during the harvest visit.

Table 4.5: MRPL by Months of Interview

	Aug to Jan		Feb to July	
	(1)	(2)	(3)	(4)
	Simple Median	Weighted by Labor	Simple Median	Weighted by Labor
Agriculture	72.219***	71.955***	87.723***	89.321***
	(12.675)	(13.031)	(11.444)	(16.039)
Non-farm	18.883	18.772	34.801**	34.715**
	(12.974)	(12.906)	(14.193)	(15.956)
Difference	46.508**	45.952**	48.235***	49.265*
	(20.385)	(20.530)	(18.183)	(29.401)
N (Household/Wave)	1899	1899	2200	2200

Standard errors are constructed through bootstrapping the process 1,000 times. The bootstrap is set to draw households. The results re-estimate the results from columns one and two of Table 4.4, restricting estimation only to households surveyed during the "high" non-farm season of August to January (columns one and two) or the "low" non-farm season of February to July (columns three and four). * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

and construct MRPL estimates for each subsample, based on month of interview. If seasonality is affecting our results, then we are likely to see substantially different MRPL differences during planting and harvest seasons, especially since these two seasons apparently show large differences in non-farm labor allocation. Table 4.A1 in the appendix presents summary statistics for month of interview across all three waves.

We estimate production functions separately for the subsample of households that were interviewed during the "low non-farm output" season, which we define as February to July[16], and the subsample of households that were interviewed during times of higher non-farm labor allocation and output, from August to January. We present these results in Table 4.5. The first two columns restrict estimation to only non-farm enterprises that were interviewed from August to January, while the last two columns restrict estimation to only households interviewed from February to July. It appears that agricultural MRPL is slightly higher for those households interviewed from February to July. However, non-farm MRPL is also slightly higher, resulting in remarkably consistent MRPL differences across time

[16]This only refers to the non-farm module. The agricultural statistics still come from the post-harvest questionnaire.

Figure 4.3: Yearly Rainfall in Waves One and Two

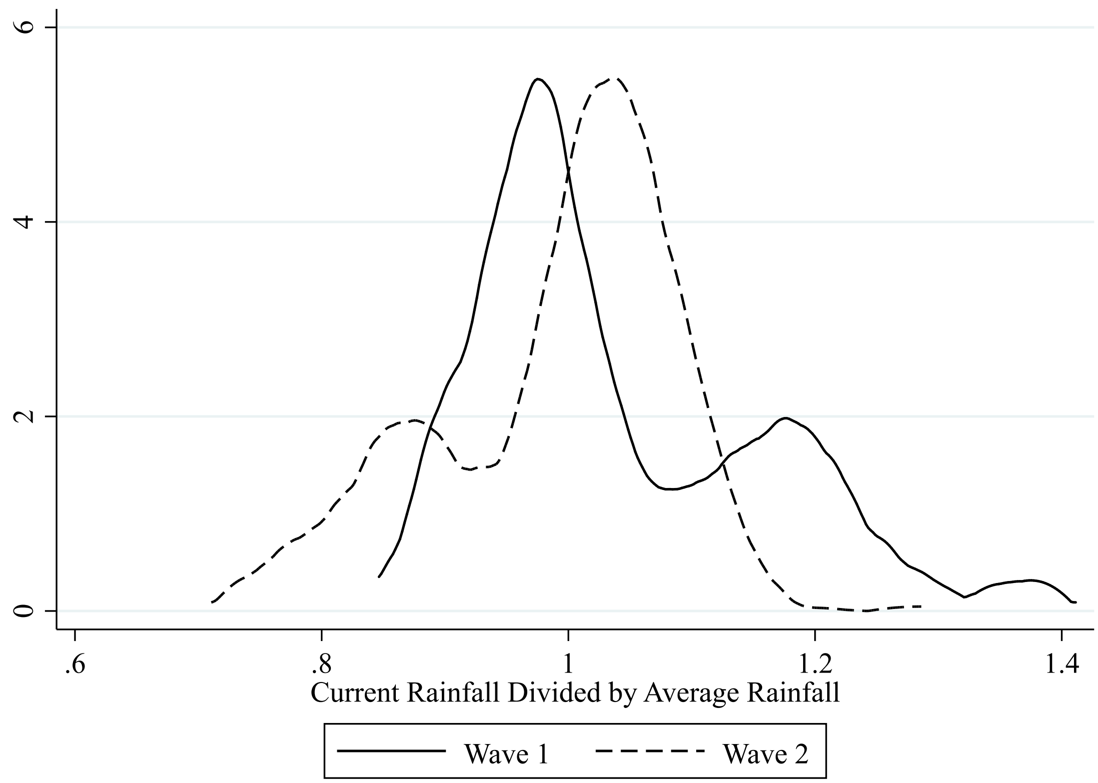

periods. We interpret this as suggestive evidence that temporal variations in labor allocation and survey timing are unlikely to be driving our results.

Another possible explanation is that rainfall happened to be very good in the three survey years, which might induce an increase in MRPL over what was expected by farmers.[17] Figure 4.3 presents kernel density estimates of rainfall in waves one and two (the only waves for which we have rainfall data). The figure shows that rainfall was relatively higher than normal in wave one but relatively lower than normal in wave two. While we do not have rainfall for wave three, we can look at other proxies for rainfall. Whenever farmers

[17]Recall that we use predicted output to construct our main MRPL estimates, which is likely a better predictor of households' ex ante expectations. However, results in the appendix also show that the choice of predict versus actual output does not affect our results.

report area harvested lower than area planted, they are asked for the cause. In wave three, around 60 percent of plot-crop observations report lower area harvest than area planted. Of these, between 80 and 90 percent (depending on whether we use the cross-section or panel subsample) blame drought or irregular rains. In other words, it appears that wave three was actually a relatively poor rainfall year in Malawi. Based on these facts, we conclude high rainfall is unlikely to be driving our high agricultural MRPL estimates.

4.4.2 Correlates of Labor Misallocation

A focus on the median obscures substantial heterogeneity in MRPLs. In this section, we explore this heterogeneity by exploring the correlates of labor (mis)allocation. We begin with a simple histogram of MRPL difference in Figure 4.4. For ease of presentation, we trim the top 5 percent of the sample. Consistent with our empirical results, the vast majority of the distribution lies to the right of zero, with the highest density around 50 MWK. While there is a wide distribution of MRPL differences, fully 86 percent of household/wave observations have a positive difference, which underscores the empirical results.

We next examine whether certain variables predict deviations from MRPL equality, with a particular focus on price and production risk. Table 4.6 and Table 4.7 examine price risk. In Table 4.6, we split the sample by crop choice and hiring. In Panel A, we look at households that grow maize (column one) and households that grow tobacco and/or cotton (column two). Maize is a common subsistence crop and tobacco and cotton are common cash crops. As shown by Barrett (1996), we expect larger deviations from equality for cash crops than subsistence crops when price risk affects decision-making. Consistent with this hypothesis, MRPL difference is more than twice as large for households that grow tobacco and/or cotton than for households that grow maize. However, given the small sample size for cash crops, there is substantial imprecision; the standard error is around 40 percent

Figure 4.4: MRPL Distribution - Pooled Production Function

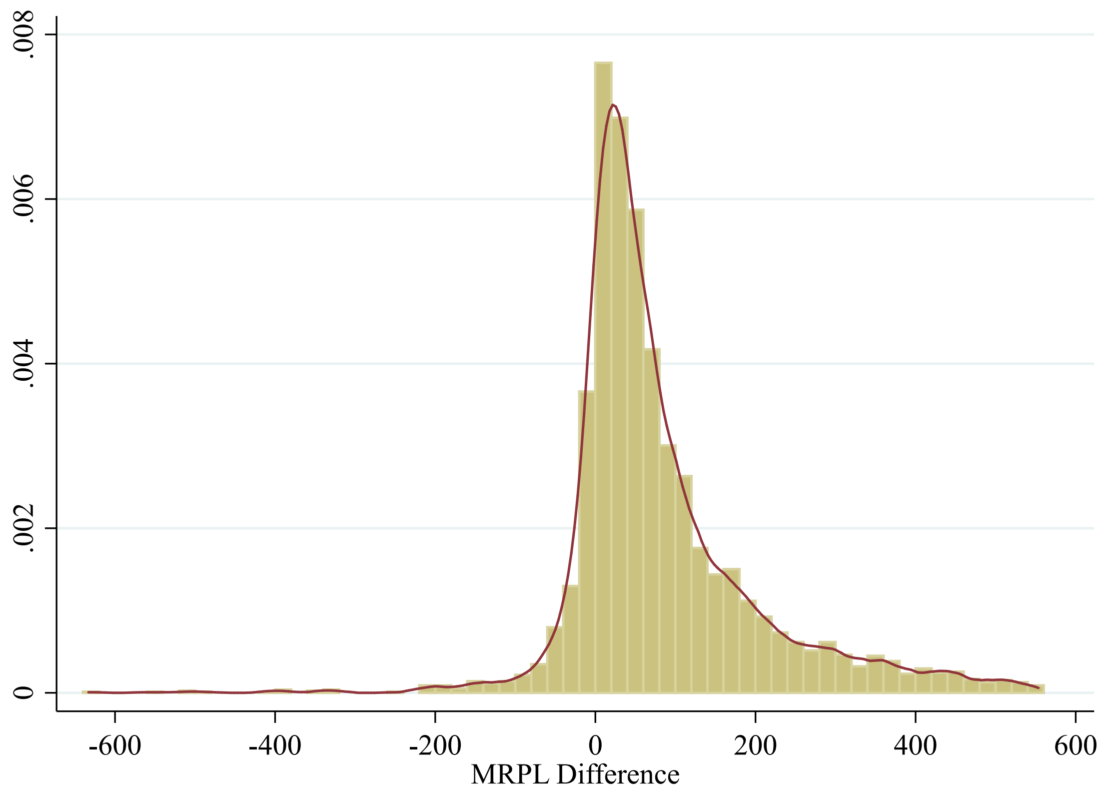

MRPL estimates are constructed as in column two of Table 4.4. MRPL difference is constructed as agricultural MRPL minus non-farm MRPL. The top five percent are trimmed from the figure for ease of presentation.

Table 4.6: MRPL, Crop Choice, and Hiring

	(1)	(2)
Panel A: Crop Choice	**Subsistence (Maize)**	**Cash (Tobacco/Cotton)**
Agriculture	73.552***	109.278**
	(11.398)	(52.374)
Non-Farm	13.775	9.198
	(13.688)	(29.418)
Difference	47.183**	113.605
	(19.350)	(70.341)
N (Household/Wave)	2422	405
Panel B: Hires for Ag	**Yes**	**No**
Agriculture	144.275***	71.485***
	(26.984)	(7.827)
Non-Farm	16.110	30.255***
	(23.307)	(30.255)
Difference	119.240***	40.046***
	(42.953)	(14.896)
N (Household/Wave)	1295	2800
Panel C: Hires for Non-Farm	**Yes**	**No**
Agriculture	104.212**	79.106***
	(52.327)	(8.330)
Non-Farm	110.335	13.268
	(73.866)	(8.874)
Difference	3.625	59.456***
	(88.822)	(14.462)
N (Household/Wave)	337	3758

A separate pooled production function is estimated for each column in each panel. Panel A splits the sample by crop choice, Panel B splits the sample by whether the household hires for agricultural production, and Panel C splits the sample by whether the household hires for non-farm production. Standard errors are constructed through bootstrapping the process 1,000 times. The bootstrap is set to draw households. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

larger than the *point estimate* for the maize subsample. Nonetheless, the pattern supports the hypothesis.

In Panel B and Panel C, we split the sample by households that hire for agriculture (Panel A) and for non-farm (Panel B) and estimate MRPL separately for each group. Households that hire for agricultural production have much higher deviations from equality (with agricultural MRPL being higher than non-farm MRPL) than households that do not hire for agricultural productions. However, this relationship is reversed for households that hire for non-farm production: MRPL differences are much higher in households that do not hire.

Table 4.7: MRPL and Production Characteristics

	(1) Above Median	(2) Below Median
Panel A: Output		
Ag Median	141.216***	40.476***
	(23.912)	(6.218)
NF Median	30.684	29.679**
	(19.583)	(11.689)
Diff. Median	122.875***	6.492
	(36.921)	(10.793)
N (Household/Wave)	2044	2053
Panel B: Acreage		
Ag Median	82.795***	72.194***
	(11.362)	(13.184)
NF Median	13.063	25.517*
	(12.999)	(14.783)
Diff. Median	67.382***	39.648**
	(22.566)	(19.481)
N (Household/Wave)	2045	2052
Panel C: Crop Sales		
Ag Median	85.513***	71.023***
	(13.528)	(10.785)
NF Median	7.403	35.247**
	(10.148)	(15.270)
Diff. Median	85.039***	31.222*
	(21.821)	(18.454)
N (Household/Wave)	1973	2124

A separate pooled production function is estimated for each column in each panel.Panel A splits the sample by median of total output, Panel B splits the sample by median total acreage, and Panel C splits the sample by median of crop sales (gross). Standard errors are constructed through bootstrapping the process 1,000 times. The bootstrap is set to draw households. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Interestingly, hiring appears to have the largest effect on MRPL in the activity for which the household hires. In other words, households that hire for agriculture actually have lower non-farm MRPL than households that do not hire for agriculture. There is a similar, though less pronounced, pattern for households that hire for non-farm production, as well. We might expect these findings in agriculture, where households that hire are more likely to sell crops on the market (Barrett, 1996). Since all non-farm production is (by assumption) sold on the market and not consumed, we would not expect such deviations based on hiring.

Table 4.7 presents additional results. We now split households based on the median of three separate variables: total output, total acreage, and total crop sales in Panel A, Panel B, and Panel C, respectively. Column one presents MRPL estimates for all households above the median of the respective variable while column two presents estimates for all households below the median. If these variables are correlated with marketable surplus, we expect larger MRPL differences for households above the median than below if price risk appreciably affects labor allocation decisions. In fact, this is exactly what we see. When splitting the sample by output, households above the median show a median MRPL difference of more than 120 MWK, while households below the median show a difference of just 6 MWK. We see a similar, though less pronounced, pattern for both acreage and crop sales. This is again evidence that price risk is an important component of agricultural decision-making.

Having established that price risk appears to be an important predictor of deviations from MRPL equality, we now move to rainfall. Figure 4.5 presents locally weighted regressions of MRPL on rainfall in waves one and two. Panel A presents results using rainfall in levels while Panel B presents results using rainfall CV. It is difficult to come to any conclusions regarding the direction of the correlation, other than that rainfall appears to have absolutely no effect on non-farm MRPL. To examine this relationship more formally, we estimate

Figure 4.5: MRPL and Rainfall

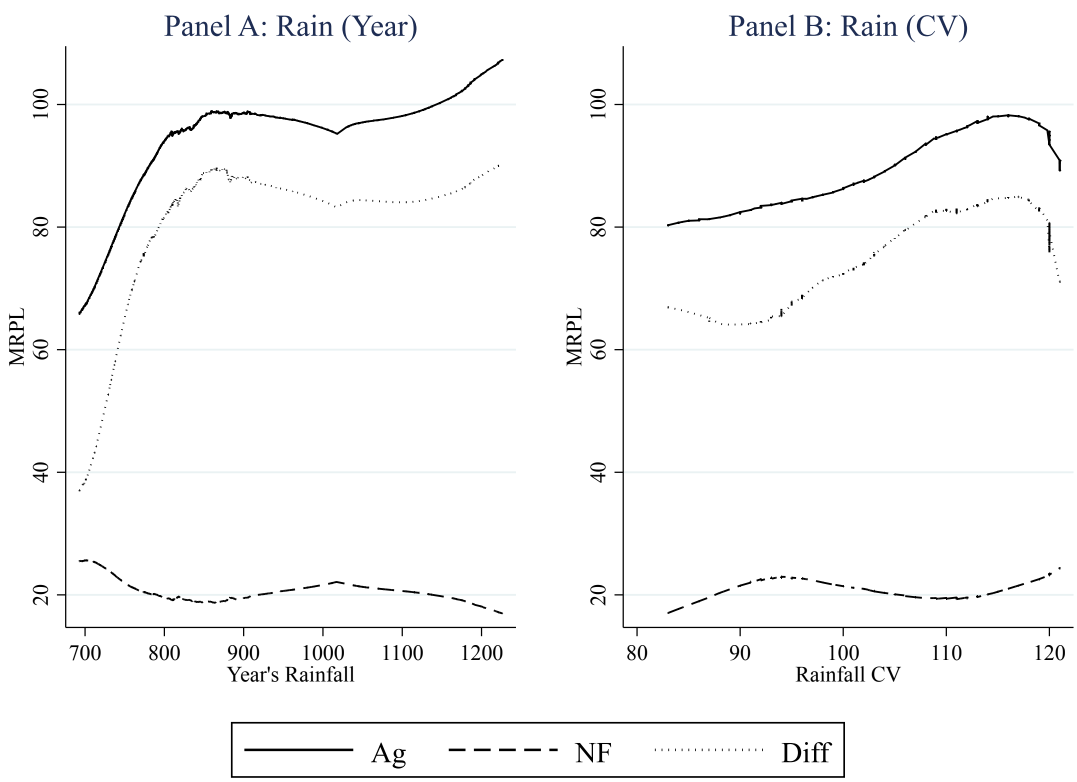

All figures are locally weighted regressions of MRPL on rainfall. In Panel A, the rainfall variable is total rainfall. In Panel B, the rainfall variable is the coefficient of variation of rainfall.

Table 4.8: Median Regression - MRPL and Rainfall

	(1)	(2)	(3)
MRPL:	Ag	NF	Diff
Panel A - Regressions one through three: Waves 1 and 2			
Current rainfall (mm)	0.023***	0.001	0.027*
	(0.007)	(0.008)	(0.015)
Rainfall CV	0.346**	0.099	0.354*
	(0.166)	(0.107)	(0.213)
Panel B - Regressions four through six: Wave 3 only			
Rainfall CV	1.787**	0.000	1.353
in Wave 2	(0.803)	(0.212)	(0.934)

The results are from six separate quantile (median) regressions using the different MRPL estimates as dependent variables. MRPL estimates come from column two of Table 4.4. In Panel A, we include waves one and two only. In Panel B, we use the panel component of the survey to match wave two rainfall CV values with wave three households. Standard errors are constructed through bootstrapping the process 1,000 times. The bootstrap is set to draw households. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

MRPL using our main results in Table 4.4. We then estimate quantile (median) regressions of MRPL on our rainfall values. We present these results in Table 4.8. In the table, there are six separate regressions. Panel A includes results for waves one and two. When including both rainfall variables in each regression, clear patterns emerge. Intuitively, higher rainfall is associated with higher agricultural MRPL and a higher difference. Rainfall CV – which we use as a proxy of production risk – is also positively correlated with both agricultural MRPL and MRPL difference. In Panel B, we merge wave two rainfall CV values with wave three households using the panel component of the IHS. There is substantial measurement error, as evidenced by the large standard errors and implausible coefficients. However, the pattern remains: rainfall CV is positively correlated with agricultural MRPL and MRPL difference, though the relationship is only significant for the former. Overall, it appears production risk affects labor allocation.

4.5 Conclusion

In this paper, we examine allocative efficiency across agriculture and non-farm production in rural households. To the best of our knowledge, this is the first paper to test allocative efficiency using the marginal product of labor across agricultural and non-farm activities within the household. We present evidence of labor misallocation in Malawi, though in a surprising direction; across a number of specifications, agricultural MRPL is consistently *higher* than non-farm MRPL. Moreover, we show that risk plays an important role in labor allocation. Our results suggest that removing risk – or, similarly, insuring farmers against risk – might result in higher incomes.

These results suggest that the median household could increase its mean income by reallocating labor out of non-farm production and into agricultural production. In other words, the results seemingly point to a misallocation of labor. However, failure of MRPL equality does not imply households can increase their expected *utility* by reallocating labor. The fact that risk appears to play an important role in this apparent misallocation suggests that households may indeed by making rational labor allocation decisions. These results show that risk leads households to protect themselves ex ante, but at the cost of a lower ex post income.

Moreover, our findings stand in contrast to a relatively large literature that finds average productivity is higher in the non-farm sector than the agricultural sector (Gollin et al., 2014; McCullough, 2017; Young, 2013). However, there are two large differences between most papers in that literature and this paper. First, we focus only on household production and restrict our sample to households that operate both non-farm and agricultural enterprises simultaneously. Second, we estimate the marginal product of labor – which has clear theoretical underpinnings – while the majority of the literature estimates the average product of labor. Importantly, the contrasting results point to different policy conclusions. While much of the literature leads one to believe that household welfare might increase if they reallocated labor out of agricultural production and into non-farm production, our results suggest the opposite. In fact, a back of the envelope suggests a reallocation of just one day of labor in this direction would decrease household income by around one-third of per-capita daily income.

Despite these differences, there are some important caveats to our findings. First, our sample is a relatively select group of households from just a single country. Previous literature has focused on a much broader set of countries. Second, our estimation strategy requires a specific set of identification assumptions. As with all such assumptions, additional studies that rely on different identification assumptions are required. Given the policy implications of our results, testing their robustness across different identification assumptions and country contexts is necessary.

4.6 Paper 3 Appendix

Table 4.A1: Month of Interview by Type of Enterprise

	(1) Agriculture	(2) Non-Farm
January	371	221
February	409	278
March	434	334
April	363	463
May	392	480
June	409	435
July	458	410
August	908	283
September	1442	593
October	1038	482
November	425	261
December	338	179

Counts are the number of observations in the restricted sample (households that operate both types of enterprises in the same wave) and when they responded to the agricultural module (first column) and non-farm module (second column). Cross-section households responded to both modules in the same sitting. Approximately half of panel households responded in the same sitting while the other half responded at different times.

Table 4.A2: MRPL Estimates with 10 Prices Observations

	Pooled Production Function		Separate Production Functions		Collapsed
	(1)	(2)	(3)	(4)	(5)
	Simple Median	Weighted by Labor	Simple Median	Weighted by Labor	Collapsed
Ag MRPL (Median)	121.732***	78.485***	166.973***	112.578***	94.436***
	(17.591)	(8.301)	(30.494)	(14.341)	(10.290)
NF MRPL (Median)	23.407	22.421**	21.691	20.539	31.940**
	(16.299)	(10.908)	(28.891)	(19.104)	(12.796)
Difference (Median)	90.160***	50.041***	129.881**	76.199***	59.107***
	(29.395)	(15.301)	(51.857)	(27.753)	(18.700)
N (Household/Wave)	3985	3985	3985	3985	3985

Standard errors are constructed through bootstrapping the MRPL construction 1,000 times. The bootstrap is set to draw households. MPRL estimates are in 2010 MWK. MRPL difference is constructed as agricultural MRPL minus non-farm MRPL. The MRPLs in the first and second columns are constructed by estimating a single production function but allowing the effects of all variables (except household fixed effects) to vary by enterprise type. MRPLs in the third and fourth columns are constructed by estimating separate production functions for agriculture and non-farm production. In the fifth column, MRPL is constructed by collapsing across plots and enterprises to the household/wave level and estimating a single production function, again allowing the effects of all variables (except household fixed effects) to vary by enterprise type. Unlike our main estimates, we require a minimum of ten price observations for the estimates in this table. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Table 4.A3: MRPL Estimates with no Minimum Price Observations

	Pooled Production Function		Separate Production Functions		Collapsed
	(1)	(2)	(3)	(4)	(5)
	Simple Median	Weighted by Labor	Simple Median	Weighted by Labor	Collapsed
Ag MRPL (Median)	144.143***	94.083***	201.662***	134.372***	101.685***
	(21.437)	(9.789)	(37.520)	(16.937)	(12.685)
NF MRPL (Median)	26.975	25.716**	27.823	26.133	32.661**
	(17.082)	(11.549)	(29.991)	(19.986)	(12.836)
Difference (Median)	114.588***	65.085***	158.979***	93.264***	69.563***
	(33.120)	(17.264)	(57.004)	(29.630)	(20.663)
N (Household/Wave)	4160	4160	4160	4160	4160

Standard errors are constructed through bootstrapping the MRPL construction 1,000 times. The bootstrap is set to draw households. MPRL estimates are in 2010 MWK. MRPL difference is constructed as agricultural MRPL minus non-farm MRPL. The MRPLs in the first and second columns are constructed by estimating a single production function but allowing the effects of all variables (except household fixed effects) to vary by enterprise type. MRPLs in the third and fourth columns are constructed by estimating separate production functions for agriculture and non-farm production. In the fifth column, MRPL is constructed by collapsing across plots and enterprises to the household/wave level and estimating a single production function, again allowing the effects of all variables (except household fixed effects) to vary by enterprise type. Unlike our main estimates, we do not require any minimum number of price observations for these estimates. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Table 4.A4: MRPL Estimates without Trimming

	Pooled Production Function		Separate Production Functions		Collapsed
	(1)	(2)	(3)	(4)	(5)
	Simple Median	Weighted by Labor	Simple Median	Weighted by Labor	Collapsed
Ag MRPL (Median)	150.591***	99.145***	197.038***	134.347***	112.322***
	(25.295)	(11.808)	(47.375)	(22.149)	(14.244)
NF MRPL (Median)	32.378	31.392**	40.383	37.789	35.553**
	(20.203)	(13.454)	(38.192)	(25.424)	(15.576)
Difference (Median)	114.029***	65.439***	137.176**	81.493**	75.230***
	(39.011)	(20.672)	(69.210)	(36.798)	(25.128)
N (Household/Wave)	4180	418	4180	4180	4180

Standard errors are constructed through bootstrapping the MRPL construction 1,000 times. The bootstrap is set to draw households. MPRL estimates are in 2010 MWK. MRPL difference is constructed as agricultural MRPL minus non-farm MRPL. The MRPLs in the first and second columns are constructed by estimating a single production function but allowing the effects of all variables (except household fixed effects) to vary by enterprise type. MRPLs in the third and fourth columns are constructed by estimating separate production functions for agriculture and non-farm production. In the fifth column, MRPL is constructed by collapsing across plots and enterprises to the household/wave level and estimating a single production function, again allowing the effects of all variables (except household fixed effects) to vary by enterprise type. Unlike our main estimates, we do not trim the top one percent of labor and output prior to estimating production functions. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Table 4.A5: MRPL Estimates Using Actual Output Instead of Predicted Output

	Pooled Production Function		Separate Production Functions		Collapsed
	(1)	(2)	(3)	(4)	(5)
	Simple Median	Weighted by Labor	Simple Median	Weighted by Labor	Collapsed
Ag MRPL (Median)	133.007***	86.289***	180.967***	116.759***	98.435***
	(19.472)	(8.699)	(32.658)	(14.675)	(9.877)
NF MRPL (Median)	21.931	20.840**	19.888	18.744	31.826***
	(15.547)	(10.586)	(29.825)	(19.841)	(12.064)
Difference (Median)	110.104***	63.762***	142.224**	83.901***	65.848***
	(31.569)	(16.234)	(56.740)	(29.172)	(18.024)
N (Household/Wave)	4097	4097	4097	4097	4097

Standard errors are constructed through bootstrapping the MRPL construction 1,000 times. The bootstrap is set to draw households. MPRL estimates are in 2010 MWK. MRPL difference is constructed as agricultural MRPL minus non-farm MRPL. The MRPLs in the first and second columns are constructed by estimating a single production function but allowing the effects of all variables (except household fixed effects) to vary by enterprise type. MRPLs in the third and fourth columns are constructed by estimating separate production functions for agriculture and non-farm production. In the fifth column, MRPL is constructed by collapsing across plots and enterprises to the household/wave level and estimating a single production function, again allowing the effects of all variables (except household fixed effects) to vary by enterprise type. Unlike our main estimates, we use actual output instead of predicted output for these estimates.
* $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$

Chapter 5

Conclusion

Bibliography

Abdulai, A. and Regmi, P. P. (2000). Estimating Labor Supply of Farm Households under Nonseparability: Empirical Evidence from Nepal. *Agricultural Economics*, 22(3):309–320.

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.

Aghion, P. and Bolton, P. (1997). A theory of trickle-down growth and development. *The Review of Economic Studies*, 64(2):151–172.

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.

Armendáriz, B. and Morduch, J. (2010). *The economics of microfinance*. MIT press.

Arthi, V., Beegle, K., De Weerdt, J., and Palacios-López, A. (2018). Not your average job: Measuring farm labor in Tanzania. *Journal of Development Economics*, 130:160–172.

Azam, M. (2011). The impact of Indian job guarantee scheme on labor market outcomes: Evidence from a natural experiment. *Available at SSRN 1941959*.

Bandiera, O., Burgess, R., Das, N., Gulesci, S., Rasul, I., and Sulaiman, M. (2017). Labor markets and poverty in village economies. *The Quarterly Journal of Economics*, 132(2):811–870.

Banerjee, A., Duflo, E., Glennerster, R., and Kinnan, C. (2015a). The miracle of microfinance? evidence from a randomized evaluation. *American Economic Journal: Applied Economics*, 7(1):22–53.

Banerjee, A., Duflo, E., Imbert, C., Mathew, S., and Pande, R. (2015b). Can e-governance reduce capture of public programs? Experimental evidence from a financial reform of India's employment guarantee. *Working Paper*.

Banerjee, A. V. and Newman, A. F. (1993). Occupational Choice and the Process of Development. *Journal of Political Economy*, 101(2):274–298.

Barrett, C. B. (1996). On Price Risk and the Inverse Farm Size-Productivity Relationship. *Journal of Development Economics*, 51(2):193–215.

Barrett, C. B., Reardon, T., and Webb, P. (2001). Nonfarm income diversification and household livelihood strategies in rural Africa: concepts, dynamics, and policy implications. *Food policy*, 26(4):315–331.

Barrett, C. B., Sherlund, S. M., and Adesina, A. A. (2008). Shadow Wages, Allocative Inefficiency, and Labor Supply in Smallholder Agriculture. *Agricultural Economics*, 38(1):21–34.

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*, 128:1–23.

Behrman, J. R. (1999). Labor markets in developing countries. *Handbook of Labor Economics*, 3:2859–2939.

Benjamin, D. (1992). Household composition, labor markets, and labor demand: testing for separation in agricultural household models. *Econometrica: Journal of the Econometric Society*, pages 287–322.

Berg, E., Bhattacharyya, S., Rajasekhar, D., and Manjula, R. (2014). Can Public Employment Schemes Increase Equilibrium Wages? Evidence from a Natural Experiment in India.

Bertrand, M., Duflo, E., and Mullainathan, S. (2002). How much should we trust differences-in-differences estimates? Technical report, National Bureau of Economic Research.

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.

Binswanger, H. (2012). India 1960-2010: Structural change, the rural non-farm sector, and the prospects for agriculture.

Blattman, C., Fiala, N., and Martinez, S. (2013). Generating skilled self-employment in developing countries: Experimental evidence from Uganda. *The Quarterly Journal of Economics*, 129(2):697–752.

Brummund, P. and Merfeld, J. D. (2017). Allocative Efficiency of Labor in Agricultural Households: Evidence from Malawi. *Working paper*.

Chenery, H. B., Syrquin, M., and Elkington, H. (1975). *Patterns of development, 1950-1970*, volume 75. Oxford University Press London.

Cho, Y., Robalino, D., and Watson, S. (2016). Supporting self-employment and small-scale entrepreneurship: potential programs to improve livelihoods for vulnerable workers. *IZA Journal of Labor Policy*, 5(1):7.

Cole, S., Giné, X., and Vickery, J. (2017). How does risk management influence production decisions? Evidence from a field experiment. *The Review of Financial Studies*, 30(6):1935–1970.

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*.

Davis, B., Di Giuseppe, S., and Zezza, A. (2017). Are African households (not) leaving agriculture? Patterns of households? income sources in rural Sub-Saharan Africa. *Food Policy*, 67:153–174.

De Mel, S., McKenzie, D., and Woodruff, C. (2008). Returns to capital in microenterprises: Evidence from a field experiment. *The Quarterly Journal of Economics*, pages 1329–1372.

De Mel, S., McKenzie, D., and Woodruff, C. (2009). Are women more credit constrained? experimental evidence on gender and microenterprise returns. *American Economic Journal: Applied Economics*, 1(3):1–32.

Deininger, K. and Liu, Y. (2013). Welfare and Poverty Impacts of India's National Rural Employment Guarantee Scheme: Evidence from Andhra Pradesh. *World Bank Policy Research Working Paper*.

Deininger, K., Nagarajan, H. K., and Singh, S. K. (2016). Short-term effects of India's employment guarantee program on labor markets and agricultural productivity. *World Bank Policy Research Working Paper*, (7665).

Dercon, S. (2009). Rural poverty: Old challenges in new contexts. *The World Bank Research Observer*, 24(1):1–28.

Devereux, S. and Solomon, C. (2006). *Employment creation programmes: The international experience*. ILO.

Dillon, B. and Barrett, C. B. (2017). Agricultural factor markets in Sub-Saharan Africa: an updated view with formal tests for market failure. *Food policy*, 67:64–77.

Dillon, B., Brummund, P., and Mwabu, G. (2017). Asymmetric Non-Separation and Rural Labor Markets. *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*.

Dutta, P., Murgai, R., Ravallion, M., and Van de Walle, D. P. (2012). Does India's employment guarantee scheme guarantee employment? *World Bank Policy Research Working Paper*, (6003).

Emerick, K., de Janvry, A., Sadoulet, E., Dar, M. H., et al. (2016). Technological innovations, downside risk, and the modernization of agriculture. *American Economic Review*, 106(6):1537–61.

Evenson, R. and Binswanger, H. (1980). Estimating Labor Demand Functions for Indian Agriculture. *Yale University Economic Growth Center Discussion Paper 356*.

Fields, G. S. (1975). Rural-Urban Migration, Urban Unemployment and Underemployment, and Job-Search Activity in LDCs. *Journal of Development Economics*, 2(2):165–187.

Foster, A. D. and Rosenzweig, M. R. (2010). Is there surplus labor in rural India? *Working Paper*.

Frederick, S., Loewenstein, G., and O'donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of economic literature*, 40(2):351–401.

Gehrke, E. (2014). An employment guarantee as risk insurance? assessing the effects of the nregs on agricultural production decisions. *The World Bank Economic Review*.

Gilligan, D. O., Hoddinott, J., and Taffesse, A. S. (2009). The impact of Ethiopia's Productive Safety Net Programme and its linkages. *The Journal of Development Studies*, 45(10):1684–1706.

Gindling, T. and Newhouse, D. (2014). Self-employment in the developing world. *World Development*, 56:313–331.

Gollin, D., Lagakos, D., and Waugh, M. E. (2014). The agricultural productivity gap. *The Quarterly Journal of Economics*, 129(2):939–993.

Haggblade, S., Hazell, P., and Reardon, T. (2002). *Strategies for stimulating poverty-alleviating growth in the rural nonfarm economy in developing countries*. International Food Policy Research Institute, Environmental and Production Technology Division.

Haggblade, S., Hazell, P., and Reardon, T. (2010). The rural non-farm economy: Prospects for growth and poverty reduction. *World Development*, 38(10):1429–1441.

Harris, J. R. and Todaro, M. P. (1970). Migration, Unemployment and Development: A Two-Sector Analysis. *The American Economic Review*, pages 126–142.

Haushofer, J. and Shapiro, J. (2016). The short-term impact of unconditional cash transfers to the poor: Experimental Evidence from Kenya. *The Quarterly Journal of Economics*, 131(4):1973–2042.

Heltberg, R., Oviedo, A. M., and Talukdar, F. (2015). What do household surveys really tell us about risk, shocks, and risk management in the developing world? *The Journal of Development Studies*, 51(3):209–225.

Hicks, J. H., Kleemans, M., Li, N. Y., and Miguel, E. (2017). Reevaluating agricultural productivity gaps with longitudinal microdata. Technical report, National Bureau of Economic Research Working Paper 23253.

Hornbeck, R. and Keskin, P. (2012). Does agriculture generate local economic spillovers? Short-run and long-run evidence from the Ogallala aquifer. Technical report, National Bureau of Economic Research.

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.

Islam, M. and Sivasankaran, A. (2014). How does child labor respond to changes in adult work opportunities? evidence from NREGA. In *International Conference on MGN-REGA. Mumbai: Impact Indira Gandhi institute of Development Research*.

Jacoby, H. G. (1993). Shadow Wages and Peasant Family Labour Supply: An Econometric Application to the Peruvian Sierra. *The Review of Economic Studies*, 60(4):903–921.

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.

Karlan, D. and Morduch, J. (2009). Access to finance. *Handbook of Development Economics*, 5:4704–4784.

Karlan, D., Osei, R., Osei-Akoto, I., and Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints. *The Quarterly Journal of Economics*, 129(2):597–652.

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.

Kijima, Y. and Lanjouw, P. (2005). Economic diversification and poverty in rural India. *Indian Journal of Labour Economics*, 48(2):349–374.

Kuznets, S. (1957). Quantitative aspects of the economic growth of nations: II. Industrial distribution of national product and labor force. *Economic Development and Cultural Change*, 5(S4):1–111.

LaFave, D. and Thomas, D. (2016). Farms, Families, and Markets: New Evidence on Completeness of Markets in Agricultural Settings. *Econometrica*, 84(5):1917–1960.

Lakha, S. and Taneja, P. (2009). Balancing democracy and globalisation: The role of the state in poverty alleviation in india. *South Asia: Journal of South Asian Studies*, 32(3):408–424.

Lanjouw, J. O. and Lanjouw, P. (2001). The rural non-farm sector: Issues and evidence from developing countries. *Agricultural Economics*, 26(1):1–23.

Lanjouw, P. (1999). The rural non-farm sector: A note on policy options. *Non-Farm Workshop Background paper, The World Bank*.

Lanjouw, P. and Murgai, R. (2009). Poverty decline, agricultural wages, and nonfarm employment in rural India: 1983–2004. *Agricultural Economics*, 40(2):243–263.

Lanjouw, P. and Shariff, A. (2004). Rural non-farm employment in India: Access, incomes and poverty impact. *Economic and Political Weekly*, pages 4429–4446.

Lewis, W. A. (1954). Economic development with unlimited supplies of labour. *The Manchester School*, 22(2):139–191.

Li, T. and Sekhri, S. (2013). The unintended consequences of employment based safety net programs. *Working Paper*.

Linde-Rahr, M. (2005). Differences in agricultural returns: an empirical test of efficiency in factor input allocation using vietnamese data. *Agricultural Economics*, 32(1):35–45.

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*, pages 25–27.

Matsuyama, K. (1992). Agricultural productivity, comparative advantage, and economic growth. *Journal of Economic Theory*, 58(2):317–334.

McCullough, E. B. (2017). Labor productivity and employment gaps in Sub-Saharan Africa. *Food Policy*, 67:133–152.

McKenzie, D. J. and Woodruff, C. (2006). Do entry costs provide an empirical basis for poverty traps? evidence from Mexican microenterprises. *Economic Development and Cultural Change*, 55(1):3–42.

Merfeld, J. (2017a). Public Works, Crowd Out, and Non-Farm Self-Employment: Evidence from India. *Working Paper*.

Merfeld, J. D. (2017b). Spatially Heterogeneous Effects of a Public Works Program. *Working Paper*.

Merfeld, J. D. (2018). Non-Farm Self-Employment, Risk, and Public Works: 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. shttp://nrega.nic.in/rajaswa.pdf.

Ministry of Rural Development (2013). Frequently asked questions on MGNREGA operational guidelines - 2013.

Ministry of Rural Development (2016). Mahatma Ghandi National Rural Employment Guarantee Act 2005 - Performance, initiatives & strategies (FY 15-16 & FY 17-16).

Mobarak, A. M. and Rosenzweig, M. (2014). Risk, insurance and wages in general equilibrium. Technical report, National Bureau of Economic Research.

Morduch, J. (1995). Income smoothing and consumption smoothing. *The Journal of Economic Perspectives*, 9(3):103–114.

Morduch, J. (1999). The microfinance promise. *Journal of economic literature*, 37(4):1569–1614.

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*.

Muralidharan, K., Niehaus, P., and Sukhtankar, S. (2017). General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence from India. *Department of Economics, University of California, San Diego, working paper*.

Murgai, R., Ravallion, M., and van de Walle, D. (2015). Is workfare cost-effective against poverty in a poor labor-surplus economy? *The World Bank Economic Review*, 30(3):413–445.

Nagler, P. and Naudé, W. (2014). Non-farm enterprises in rural Africa: New empirical evidence. *World Bank Policy Research Working Paper*, (7066).

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 (2003a). Identification of districts for wage and self employment programmes. *Report of the Task Force*.

Planning Commission (2003b). Report of the Task Force: Identification of Districts for Wage and Self Employment Programmes.

Ranis, G. and Fei, J. C. (1961). A theory of economic development. *The American economic review*, pages 533–565.

Ravallion, M. (2007). Evaluating anti-poverty programs. *Handbook of Development Economics*, 4:3787–3846.

Ravallion, M. (2009). Evaluation in the Practice of Development. *The World Bank Research Observer*, 24(1):29–53.

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.

Reardon, T. (1997). Using evidence of household income diversification to inform study of the rural nonfarm labor market in Africa. *World development*, 25(5):735–747.

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.

Sandmo, A. (1971). On the theory of the competitive firm under price uncertainty. *The American Economic Review*, 61(1):65–73.

Santangelo, G. (2016). Firms and farms: The impact of agricultural productivity on the local Indian economy. *Working paper*.

Seshan, G. K. (2014). The Impact of Trade Liberalisation on Household Welfare in a Developing Country With Imperfect Labour Markets. *Journal of Development Studies*, 50(2):226–243.

Shah, M. and Steinberg, B. M. (2015). Workfare and human capital investment: Evidence from India. Technical report, National Bureau of Economic Research.

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. (2009a). Evaluating the performance of National Rural Employment Guarantee Act. *NCAER-PIF*.

Sharma, A. (2009b). *NCAER-PIF Study on Evaluating performance of National Rural Employment Guarantee Act*. National Council of Applied Economic Research.

Singh, I., Squire, L., and Strauss, J. (1986). *Agricultural Household Models: Extensions, Applications, and Policy*. The World Bank.

Skoufias, E. (1994). Using Shadow Wages to Estimate Labor Supply of Agricultural Households. *American Journal of Agricultural Economics*, 76(2):215–227.

Stiglitz, J. E. (1974). Incentives and Risk Sharing in Sharecropping. *The Review of Economic Studies*, 41(2):219–255.

Subbarao, K., Del Ninno, C., Andrews, C., and Rodríguez-Alas, C. (2013). *Public Works as a Safety Net: Design, Evidence, and Implementation*. World Bank Publications.

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? *India Policy Forum*, pages 1–39.

Tiffin, R. and Irz, X. (2006). Is agriculture the engine of growth? *Agricultural Economics*, 35(1):79–89.

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.

Young, A. (2013). Inequality, the urban-rural gap, and migration. *The Quarterly Journal of Economics*, 128(4):1727–1785.

Zimmermann, L. (2012). Labor market impacts of a large-scale public works program: evidence from the Indian employment guarantee scheme.

Zimmermann, L. (2015). Why guarantee employment? Evidence from a large Indian public-works program. *Working paper*.