

Kwacha Gonna Do? Experimental Evidence about Labor Supply in Rural Malawi

Jessica Goldberg*

September 2011

Abstract

I use panel data from a unique field experiment to estimate the wage elasticity of working in the day labor market in rural Malawi. Once a week for 12 consecutive weeks, I make job offers to a pre-defined sample of 529 adults. The wage varies each week, ranging from MK 30 (\$US 0.21) to MK 140 (\$US 1.00) for a day's work. This approach provides exogenous variation in wages, allows me to observe the full distribution of wage offers rather than the censored distribution of accepted wages, and permits the inclusion of time and village or time and individual fixed effects. I estimate that the elasticity of employment is between 0.15 and 0.17, with no significant differences between men and women. I use auxiliary data from a nationally representative survey to confirm that equal elasticities for men and women is typical of the Malawian labor market during the agricultural off-season, the time of year when my experiment takes place. I collapse my data into a censored cross section that mimics data used in the previous literature to demonstrate that my low point estimates of the elasticity employment are due to my improved identification strategy, rather than reflecting inherent differences between Malawi and other developing countries. My results reject backward-bending labor supply curves for men or women, and suggest that labor supply is highly inflexible along the relevant margin in poor, rural markets. JEL Codes: O12, J22, J43.

*3115 G Tydings Hall, Department of Economics, University of Maryland, College Park MD 20742. E-mail goldberg@econ.umd.edu. This project was supported with research grants from the Center for International and Comparative Studies and Rackham Graduate School and an African Initiative Grant from the Center for Afroamerican and African Studies, all at the University of Michigan, I thank Brian Jacob, David Lam, Jeff Smith, and Dean Yang for their extremely helpful comments. All errors and omissions are my own.

1 Introduction

It is widely appreciated that labor is the most abundant resource of the poor. In agricultural economies the poor may work on their own land to produce goods for home consumption or market sale, and work for other people for wages. Paid employment often takes the form of casual day labor rather than longer-term arrangements governed by contracts, and can be an important source of cash as well as a mechanism for coping with negative shocks that reduce non-labor income. The importance of this type of labor is highlighted by public sector employment programs with dual goals of infrastructure development and income support. Malawi, which has invested \$40 million in its Community Livelihoods Support Fund, is one of 29 countries in sub-Saharan Africa with a public sector works program[18]. As another example, almost 45 million households were employed to do day labor through the National Rural Employment Guarantee Act in India in 2008-2009. Despite the importance of day labor to rural households and the large scale investments in programs to employ day laborers by governments in developing countries, little is known about the elasticity of employment for day laborers.

In fact, there is scant evidence on the elasticity of labor supply in any type of labor market in developing countries. The usual challenge in estimating the elasticity of labor supply is that wages are endogenous; in developing countries, there is the additional challenge of obtaining high quality data. I overcome the identification problem by randomizing wages for community agricultural development projects in 10 villages in rural Malawi. I estimate the probability of accepting employment in the day labor market, the relevant market for millions of individuals in poor, rural communities. My sample includes 529 adults from households that have supplied “ganyu,” or day labor, in the previous year. These individuals are offered employment one day per week for 12 consecutive weeks. Wages vary by village-week, ranging from MK 30 (\$US 0.21) to MK 140 (\$US 1.00) per day, and wages for each workday are announced one week in advance. I estimate the elasticity of working on a given day using administrative attendance records, and use surveys to study changes in labor supply in response to household shocks. I find that a ten percent increase in wages leads to a 1.5 to 1.7 percent increase in the probability of working, with no differences between men and women. My results stand in contrast to the common finding in developing and

developed countries that women’s labor supply is more elastic than men’s (see, for example, Heckman[14] or Rosenzweig[22]).

My experimental approach improves identification relative to the techniques used in previous estimates of the elasticity of labor supply in developing countries. The early literature about economic growth in developing countries followed Lewis[17] in assuming that the supply of labor is perfectly elastic. More recently, empirical estimates of labor supply elasticities in developing countries have generally supported an upward-sloping labor supply curve. Bardhan[4] estimates upward sloping labor supply curves with what he characterizes as “very small” elasticities for rural households in West Bengal; Abdulai and Delgado[1] estimate somewhat greater elasticities for husbands and wives in Ghana; and Rosenzweig[22] estimates that the long-run labor supply curve for women in India slopes up, while the long-run labor supply curve for men is backward bending. Kochar[16] and Rose[21] study the response of labor supply to weather shocks in India, supporting the hypothesis that poor households increase the level of their wage labor to cope with negative shocks to non-labor income. These papers, like most of the research about labor supply in developed countries, rely on econometric identification strategies or structural models to obtain causal estimates.

To my knowledge, the only previous study to randomize pre-tax wages is Fehr and Goette[11], which randomly assigns bicycle messengers to receive a 25% increase in commissions for deliveries for four weeks.¹ My experiment, which includes a larger sample and a much wider range of wages, not only provides a unique source of exogenous variation in wages for the most common labor arrangement in Malawi and other developing countries with large rural populations, but also connects the development literature to the more recent literature about day labor markets in developed countries. Oettinger[20] studies the attendance decisions of registered stadium vendors and finds that the elasticity of working on a given day with respect to that day’s expected wage is between 0.55 and 0.65. Barmby and Dolton[5] estimate that the elasticity of working on a given day of a 1938 archeological dig in Syria was 0.035.

An important characteristic of casual wage labor markets in developing countries is that labor supply is extremely flexible on a short-term basis. In the United States and other

¹Negative income tax experiments in the 1970s in the United States typically find small effects of post-tax wages on labor supply. See, for example, Burtless and Hausman[6].

developed countries, many employees are constrained to work either full-time or not at all, with little opportunity to adjust their number of weeks per year. Camerer et al.[7], Chou[9], and Farber[10] take advantage of a notable exception to rigid labor markets by studying the relationship between hours worked and the implied hourly wage for taxi cab drivers. Camerer et al. and Chou find a puzzling result: taxi drivers to work fewer hours on more profitable days, implying a downward-sloping supply curve. They explain this result through so-called “target earning” behavior: taxi drivers set a goal for daily earnings and stop work when they reach their goal. Using a richer data set and a different approach to imputing hourly wages, Farber finds that taxi drivers work longer hours when hourly wages are higher, the standard upward-sloping labor supply curve. Ashenfelter et al.[3] return to the taxi cab driver puzzle and study changes in hours worked in response to exogenous changes in fares. They estimate the elasticity of labor supply in response to a long run change in wages to be -0.20. Though these papers are based on data from the United States and Singapore, they are some of the only papers in the labor supply literature to study a situation comparable to that in Malawi’s market for ganyu, where labor supply can be freely adjusted in the short run.

My paper proceeds as follows. I describe the experiment in Section 2 and describe the data in Section 3. I present the framework and results for estimates of my main parameter, the elasticity of employment, in Section 4. I highlight the methodological differences between my research and the previous literature about the elasticity of labor supply in developing countries in Section 5, and use supplemental data from Malawi’s 2004 IHS to provide a context for the similarity of men’s and women’s elasticities in Section 6. I conclude in Section 7.

2 Experimental Design

I randomize the wages that 529 adults in 10 villages in rural Malawi are offered for doing manual labor on agricultural development projects. Project participants are recruited from households who have done similar paid work in the past year. They are offered a job one day per week for 12 consecutive weeks. The job is the same each week, but wages change. Each week, participants can either accept the offered wage and work for the full day, or reject the wage and not work at all. Wages are announced one week in advance, and the MK 30 to

MK 140 range spans the 10th to 90th percentile of wages for day labor reported for adults in rural areas in Malawi's 2004 Integrated Household Survey (IHS).

Casual wage labor arrangements are common in Malawi, a small, extremely poor country in southeastern Africa. Fifty-two percent of Malawians consume less than a minimum subsistence level of food and non-food items, according to the 2006 World Bank Poverty and Vulnerability Assessment, and 28 percent fall below the PPP-adjusted \$1/day threshold. While on-farm production is the dominant source of income and use of time for the rural poor, day labor can play an important role in bringing in cash and coping with shocks. In the 2004 IHS, 28 percent of those living in rural areas report doing some ganyu within the last year and 21 percent reported doing some ganyu in the previous seven days. Wages vary seasonally and geographically; the 10th percentile of the wage distribution in rural areas is MK 40 (\$US 0.21) per day, and the 90th percentile is MK 135 (\$US 0.96) per day. My study takes place in Lobi, a rural area in the Central Region, along Malawi's western border with Mozambique. Lobi was chosen as the study area because it has a typical market for labor with both private and public employers, including the national Public Works Programme. Working in an area where some people already perform ganyu helps in defining a sample of individuals already participating in the relevant market and makes it more likely that people will treat the work offered through the project as a routine business decision rather than a special opportunity subject to non-economic considerations.

I partnered with a local community-based organization called the Lobi Horticultural Association (LHA) to identify a sample and appropriate work activities and, in a cross-cutting randomization, provide access to savings accounts with LHA's savings and credit cooperative (SACCO) for half of the participating households. In cooperation with local leaders and government extension workers in Dedza, Malawi, I identified 10 villages that were within 20 kilometers of LHA's headquarters, situated at the Lobi Extension Planning Area offices. The villages were chosen to be near enough to LHA's office to make it easy for people who received savings accounts to access those accounts. To minimize the chance that participants in one village would learn about wages in other villages, only one village per group village headman² was included in the project.

²Villages are led by a traditional leader known as the headman. A higher-ranking traditional leader known as the "group village headman" presides over clusters of four to 12 or more villages and may coordinate

Within each village, LHA leaders and extension workers chose a work activity. These activities were by design labor intensive, unskilled, and had public rather than private benefits. To be consistent with local standards, “one ganyu,” or a day’s work, lasted four hours. Activities included clearing and preparing communal land for planting, digging shallow wells to be used for irrigation, and building compost heaps to be used to fertilize communal land. Within each village, the activity was the same for all 12 weeks. The amount of effort was held constant by objective standards from week to week: participants had to dig the same number of cubic feet or hoe the same number of linear feet each week. Since all analyses incorporate village fixed effects, differences between activities across villages do not affect the results.

Up to 30 households in each village were invited to participate in the project. Qualifying households had to have at least one adult member who had performed ganyu within the last year. Up to two adults per household – usually but not always the head of household and his spouse – were invited to participate. While having multiple participants per household complicates the analysis of an individual’s response to a change in his own wages because household income is not held constant, it allows me to identify the elasticity with respect to the change in wages that is relevant in this context. Much of the literature in labor economics considers changes in wages for a single member of a household, holding constant income for other household members. That is the relevant parameter in developed countries or urban areas, where household members often participate in different job markets. However, it is not relevant in rural areas in developing countries, where adults have homogenous work opportunities. In Malawi, men and women perform similar on- and off-farm labor. Men and women may participate in the government’s Public Works Programme, which pays individuals in poor households to work on community infrastructure projects such as road construction. Allowing multiple adults per household to participate in this project is akin to studying the effect of a transitory change in the prevailing village wage for unskilled labor.

Participating households were given the opportunity to work for pay on their village’s activity one day per week for 12 consecutive weeks. The workday was the same each week for each village, so that village fixed effects also control for day-of-week effects. Participants were told at the outset that the project would last 12 weeks, that the work would be the development policies and other activities across villages under his domain.

same each week, that the wage would be different each week, and that they could work as many or as few days as they chose without penalty. Work was supervised by government agricultural extension agents. Wages were announced one week in advance, and in each village, a foreman was responsible for communicating the wage to all participants in the village. Participants were paid immediately, in cash, after they worked. Payments were made by a three-person team that included one Chichewa-speaking research assistant who handled money and recorded attendance, one government extension worker who supervised the community project, and one local foreman who helped identify participants to ensure that only pre-selected participants were included. Work activities were carefully monitored to ensure that within each village, the intensity and duration of work was the same from week to week.

The once-a-week design of the project was intended to minimize general equilibrium effects and to ensure that regular village activities were not unduly disrupted. Also, spreading the project over 12 weeks, rather than 12 consecutive days, allowed additional time for participants to experience positive and negative shocks, and thus for me to observe the supply of labor in response to these shocks. A disadvantage of the design is that the six-day gap between each employment offer gives individuals substantial opportunity to rearrange their other obligations in order to be able to work on this project without reducing their time in other productive activities. This ability to minimize the opportunity cost of accepting employment through my project is likely to overstate the *level* of employment, but does not have clear effects on the elasticity.

Intertemporal elasticities of substitution typically are interpreted as substitution between labor and leisure. Because my experiment offers employment for one out of seven days, individuals could instead substitute work on my project for other wage employment. I argue, however, that respondents' behavior is more consistent with substitution between labor and leisure than labor for different employers. First, the effect of wages in my project on the probability of outside employment is very small, though it is statistically significant in some specifications. Second, using an alternate definition of labor supply that counts individuals as working if they work either for my project or for another employer during the week does not change the estimated elasticity of working. If individuals were substituting away from other wage work into employment on my project, we would expect that the effect of project

wages would be lower on the more comprehensive definition of employment. The lack of an effect on outside employment is consistent with the notion that demand for labor is scarce during the dry, unproductive time of year when my project took place. My interpretation is also consistent with the limited literature on employment in daily wage markets: despite similar gaps between employment opportunities for stadium vendors, Oettinger interprets his estimates as intertemporal elasticities of substitution of labor for leisure.

The project took place in June, July, and August, months that fall between the harvest and planting seasons in Malawi and come during the country's dry season. This is a time of year with low marginal productivity either on- or off-farm. It is the time of year when individuals have the most food and most cash. Importantly, I can be confident that the opportunity cost of time was constant throughout the experimental period. Labor supply elasticities may vary seasonally, and the estimates from this experiment are not necessarily valid for a different time of year, when the opportunity cost of time is higher.

Wages for this project range from MK 30/day (\$US 0.21) to MK 140/day (\$US 1.00), in increments of MK 10.³ Table 1 shows the schedule of wages, which alternated high and low wages over the 12-week duration of the project, then shifted the schedule forward in order to have 10 separate schedules that followed the same pattern of increases and decreases. Using 10 different wage schedules creates *village* \times *week* variation that allows me to control for village and time fixed effects separately. The shifted schedule (as opposed to i.i.d. randomized wages) means that each village has the same total earnings potential and that averages across villages, within week, are approximately constant. Since it is possible that participants will consider relative wages, the schedule is designed such that each village faces the same number of wage increases and decreases. After randomly allocating each village to a wage-schedule, I allowed LHA leaders and government extension workers to determine the day of the week on which villages would be visited.⁴

³The wages are based on outcomes from a pilot study I conducted in March 2009, where 77 percent of participants worked for the lowest offered wage of MK 70, and 96 percent worked for the highest offered wage of MK 120.

⁴The list of villages given to LHA leaders and extension workers reflected the randomization, i.e. the village randomly selected as "village one" was listed first, the village randomly selected as "village two" was second, etc. The LHA leaders and extension workers retained that ordering in many cases when deciding which villages to visit on which days of the week. Since I use village fixed effects, and since the wage schedule is exogenous in each village, the relationship between day-of-week and wage schedule does not compromise the results.

Randomizing the villages' starting points in the wage schedule rather than separately assigning wages for each village-week was ultimately a trade off that insured against poorly distributed wages in a small sample at the cost of introducing serial correlation in the wages. This correlation appears to have been undetected by participants, however. In section 4.3, I provide evidence that neither lagged wages nor leading wages have any predictive power for current employment. Participants did not adjust their employment to anticipated future wages or exhibit learning about the wage process based on past wages.⁵

3 Data

In total, the project includes 529 individuals⁶ in 298 households. I follow these individuals for 12 weeks, recording their participation in each week's work activity. This gives me 6333 binary observations of individual labor supply.

To complement the administrative data, I use data from four surveys: a baseline survey and three follow-up surveys. The baseline survey was conducted at the outset, before participants were told about the nature of the project or the activities involved. It contains demographic and socioeconomic characteristics of respondents and information about their previous work history. The three follow-ups were conducted after the fourth, eighth, and 12th weeks of the project (with each village surveyed 6 days following its 4th, 8th, and 12th assigned work day). These follow-up surveys first ask respondents to recall their own participation and the wages over the previous four weeks, then ask about reasons for working or not working each week. The recall questions verify that participants are reasonably accurate in describing their participation in the project, and enhance my confidence in their self-reported reasons for working or not working in specific weeks.

⁵Additional survey evidence supports the notion that participants did not detect the negative serial correlation in wages. The survey conducted after work for week eight had been completed and wages for week nine had been announced asked participants, "what do you think the wage will be next week?" and "what do you think the wage will be in two weeks?" Eighty percent of participants knew the correct wage for their village in week nine; three percent answered but gave an incorrect wage; 17 percent said that they did not know the wage for week nine. This is clear evidence that wage changes were properly communicated to participants one week in advance. In contrast, fewer than one percent of those surveyed in week eight knew the correct wage for week 10. When asked, "what will the wage be in two weeks?" eight percent answered but gave an incorrect wage; 92 percent said that they did not know the wage for week 10. It seems reasonable to assume that participants' expectations of wages after the anticipated change in week $t + 1$ would revert to some constant level.

⁶One individual died after week six of the project, so the sample size in weeks 7-12 is 528.

Of the 529 individuals included in the project, 370 respondents are spouses living in 185 households. Another 74 are women in households where both project participants are women, and 18 are men in households where both project participants are men. The remaining 67 are individuals who are the only participants in their households. The survey team was able to interview 495 participants the week before the project began. Respondents in pre-selected households who were not available during the survey period were nonetheless allowed to participate in the study, to avoid creating a sample biased towards those with low opportunity cost of time. Table 2 presents baseline characteristics for participants in this project. The majority of the sample are married women.⁷ Participants have attended an average of four years of school and live in households with approximately two adults and three children. Respondents own an average of 1.8 acres of land; their houses have an average of two rooms; and only 16 percent of respondents have tin roofs on their houses. They work an average of one day in the week before the survey or 2.7 days in the month before the survey.

4 Elasticity of employment

I estimate a change in the probability of working on a given day with respect to a change in that day's wages, a parameter I will call the elasticity of employment. This is a reduced-form estimate of an uncompensated, intertemporal parameter, but it differs from the familiar Frisch elasticity or the elasticity of labor force participation in ways I explain in the next section. The change in the probability of working captures the relevant margin of choice in the market for day labor in poor rural economies, where individuals work either a full day or not at all but may choose their number of days with considerably more flexibility than is common in developed countries. I estimate that the elasticity of employment is between 0.15 and 0.17. These estimates are robust to alternative specifications using different combinations of village, week, and individual fixed effects; the marginal effects from OLS and probit specifications are virtually identical. Including wages for previous or future weeks does not change the point estimates of the elasticity with respect to the current week's wage, and my inferences are robust to several alternative methods of computing standard errors.

⁷Including widowed men and women or those whose spouses are disabled or permanently unavailable for work was a preference of my partner organization. All of my results are robust to limiting the sample to the 370 respondents who are married and whose spouses are also participating in the project.

4.1 Theoretical framework

Three key dimensions of labor supply elasticities discussed in the literature are the margin of choice of labor supply, the anticipation of the wage change, and the persistence of the wage change. Heckman[14] provides a useful taxonomy of the different labor supply margins in his 1993 review of the literature; the most important consideration is whether variation in labor supply is at the intensive or extensive margin. Each of the labor supply functions that Heckman describes can be estimated for different types of variation in wages: anticipated or unanticipated changes, and permanent or temporary changes. The standard intertemporal elasticity of substitution applies to trade-offs between labor and leisure in response to an anticipated, temporary change in wages. I will argue that the wage changes induced by my experiment are anticipated, temporary changes, and that my estimates should be interpreted intertemporal elasticities of working for individuals in a daily labor market. The estimates in this paper are reduced-form rather than structural parameters. The underlying model is a static optimization problem, appropriate given that the labor supply of individuals in this sample is separable across days.⁸

Heckman [14] describes four different labor supply functions, where H represents labor supply (in days or hours), W represents wages, Y represents non-labor income, and ν represents other variables that affect labor supply. These labor supply functions are:

$$E(H|W, Y, \nu) \tag{1}$$

$$E(H|W, Y, H > 0) \tag{2}$$

$$E(H|W, Y) = E(H|W, Y, H > 0) \times Pr(H > 0|W, Y) \tag{3}$$

$$Pr(H > 0|W, Y) \tag{4}$$

When H is properly defined to represent a margin at which individuals can choose to adjust their labor supply, the elasticity of labor supply at the intensive margin comes from the derivative of expression (1) with respect to W : $\epsilon_{intensive} = \frac{\partial E(H|W, Y, \nu)}{\partial W} \frac{W}{H}$. In situations where individuals cannot adjust their supply of labor at the intensive margin and instead have to choose between working a fixed number of hours (or days, or weeks) and not working, or

⁸See Ham[12] and Ham and Reilly[13] for a discussion of static versus lifecycle models of labor supply.

when only the binary participation decision is observed, we may estimate the *extensive* margin elasticity or the elasticity of participation from the derivative of expression (4) with respect to W : $\epsilon_{extensive} = \frac{\partial Pr(H>0|W,Y)}{\partial W} \frac{W}{H}$. Theoretically, the marginal effect of wages on labor supply at the intensive margin may be larger or smaller than the marginal effect of wages on labor supply at the extensive margin. Empirically, “Participation (or employment) decisions generally manifest greater responsiveness to wage and income variation than do hours-of-work equations for workers,” (Heckman [14]) based on empirical estimates for developed countries.

While the elasticity of labor supply at the intensive margin has received more attention in the empirical literature in developed countries, there are many instances where the extensive margin elasticity is the policy relevant parameter. For example, the change in aggregate supply of labor by single women due to the expansion of the Earned Income Tax Credit (EITC) in the 1990s was dominated by an increase in labor force participation (Meyer [19]). Understanding the impact of the EITC expansion, then, requires an estimate of the increase in labor force participation due to the policy change. In developing countries with large-scale public works programs, including Malawi’s \$40 million Community Livelihoods Support Fund and India’s National Rural Employment Guarantee Act, which makes over a billion people eligible for up to 100 days of work per year, understanding the change in the fraction of the population who would work under the program at different wages is of crucial importance.

The market for day labor, where individuals can work or not work for the prevailing wage each day, blurs the distinction between the intensive and extensive margin at the same time it makes clear the separation of participation versus employment. In a daily labor market the decision of $H = 0$ or $H > 0$ is made each day, and reflects movement between employment and unemployment but not between labor force participation and non-participation. Some people choose not to work on a given day because the prevailing wage is less than their opportunity cost, but would have worked had the day’s wage been higher. Thus, they are *in the market* for day labor even though they are not *employed* on a given day. Empirical estimates of the probability of working in a day labor market should condition on a different participation indicator than $H > 0$, and estimate a labor supply function that combines elements of equations (2) and (4) above:

$$Pr(H > 0|W, Y, \text{in daily labor market}) \tag{5}$$

This labor supply function combines elements of the intensive margin elasticity of hours worked for participants in Heckman’s equation (2) by conditioning on participation, and of the extensive margin probability of participating in Heckman’s equation (4) since the outcome of interest is the probability of positive hours of work. The corresponding elasticity, which I will call “the elasticity of working” is $\frac{\partial Pr(H>0|W,Y,\text{in daily labor market})}{\partial W} \times \frac{W}{H}$.⁹ Oettinger [20] calls this parameter the elasticity of participation in a daily labor market in his study of the labor supply of stadium vendors. He finds that the elasticity of working on a given day for registered stadium vendors is between 0.55 and 0.65. Barmby and Dolton (2010) also estimate the wage elasticity implied by equation (5) for workers on an archeological dig in Syria in the 1930s, and find an elasticity of 0.035.

Both Oettinger and Barmby and Dolton interpret their estimates as intertemporal elasticities of substitution, where workers experience anticipated, transitory shocks to wages and substitute between labor and leisure accordingly. Oettinger assumes that stadium vendors form expectations about future wages based on the popularity of the visiting team. Barmby and Dolton assume that serial correlation in the probability of unearthing valuable objects for which bonus payments are made allows archeological workers to form expectations based on past work.

Like Oettinger and Barmby and Dolton, I estimate changes in the probability of working on a given day among a sample of individuals who are known to be participants in the relevant labor market. My sample is restricted to households that have performed ganyu in the recent past, which satisfies the conditioning on labor market participation in equation (5). The margin of choice is at the level of a day, and because each participant is offered one day’s employment at each wage, the only possible values of H (measured in units of days) are 0 or 1. I represent that choice of employment with a binary variable in my empirical estimates.

⁹In my sample, 46 individuals had not done any paid work in the previous year. For these individuals, the estimated elasticity blurs the intensive and extensive margins because the *first* decision to work is also a decision to enter the labor market. All individuals work at least once over the 12 weeks of the project, so all do enter the labor market. My results are robust to dropping individuals who have not worked in the year before the project or to dropping observations corresponding to the the first time an individual with no previous work experience works during this project.

4.2 Point estimate of the elasticity of employment

I find that overall employment is high and the elasticity of employment is low, precisely estimated, and robust to many alternate specifications. I plot the fraction of the sample who work at each wage offer in Figure 1. At MK 30/day, the lowest wage in the sample, more than seventy percent of respondents worked. This high base has a strong seasonal component: marginal productivity at home or on one's own farm is low during the dry season, and there is very little demand for off-the-farm labor. However, employment at low wages is characteristic of the market for ganyu in Malawi. The lowest reported wages in the IHS are MK 10/day, and a quarter of those who do ganyu report receiving MK 40/day or less on average.

There is a marginally significant ($p = 0.10$) discontinuity in the probability of employment at a wage of MK 100/day.¹⁰ Despite this discontinuity, I focus on the elasticity of employment across the full range of wages rather than the change in the probability of working at MK 100. Much of the literature about labor supply in developing countries focuses on the elasticity of labor supply, so this choice facilitates comparisons between my results and previous research. Furthermore, the design of my experiment is not well-suited to identifying a non-linear change in the probability of working at MK 100. Because of the wage schedule I use, every wage of MK 100 or higher is an increase from the previous week's wage (except, of course, in the first week), and every wage of MK 90 or lower is a decrease from the previous week's wage. Therefore, it is not possible to determine whether a jump up in the probability of working at MK 100 is because of a reservation wage of MK 100, or because of a preference for wage increases.¹¹ If the correct model is one that allows for a discontinuity at MK 100, then my estimates overstate the elasticity of employment and my conclusion that the probability of working is inelastic with respect to wages would be strengthened.

In order to estimate the elasticity of working, I run ordinary least squares regressions of the form $labor_{itv} = \alpha + \beta \ln(wage_{tv}) + \nu_{itv}$. The coefficient β is the marginal effect of a one log-point, or approximately one-percent, change in wages on the probability that an individual works. The marginal effect is not an elasticity, but it is easily transformed into one

¹⁰The government's rate for day labor is currently set at MK 200, but was previously MK 110. The discontinuity does not suggest a reference point corresponding to the government's wage rate.

¹¹Using data from the first week only and relying on cross-village identification for variation, the probability of working for wages of MK 90 and lower is not statistically different from the probability of working for wages of MK 100 and higher.

using the standard formula, $\epsilon_e = \frac{\partial Q}{\partial P} \times \frac{P}{Q}$. Because I am using log-wages as the independent variable, I compute $\epsilon_e = \frac{\beta}{\text{mean}(\text{labor})}$. This elasticity corresponds to the extensive margin elasticity from labor supply equation 5 above.

In Table 3, I begin by pooling observations across weeks and villages without any additional controls. I find that a one-percent increase in wages is associated with a 12.8 percentage-point increase in the probability of working. This effect is significantly different from zero at the 99 percent confidence level. The elasticity corresponding to the estimate from the pooled data in Column (1) is 0.15. In columns (2), (3), and (4) respectively, I add fixed effects for village, week, and village and week together. Controlling for village and week separately or together does not change the magnitude of the coefficient or associated elasticity much. The elasticity in the specifications with week effects increases slightly to 0.17. In Column (5), I replace village and week fixed effects with individual fixed effects, controlling for unobserved time-invariant characteristics that are commonly thought to affect labor supply. Finally, I include individual and week fixed effects in Column (6). As before, this specification does not substantially alter the results: a one-percent increase in wages is associated with a 12.8 percentage-point increase in the probability of working, for an implied elasticity of employment equal to 0.17. The standard errors reported in this and subsequent tables come from 500 block-bootstrap replications, to allow for the possibility of persistent village-level shocks and address the small number of villages in the sample. See Appendix A for a detailed discussion of the standard errors.

4.3 Robustness checks

Given the schedule used to assign wages, the most plausible threat to the internal validity of my estimates would be that participants detected and reacted to the negative serial correlation in wages. If this were the case, it would affect both the interpretation of the elasticity as an intertemporal parameter, and the magnitude of the estimate. Respondents who understood that a low offer in week t implied a high offer in week $t + 1$ would exhibit larger elasticities than those who did not anticipate the wage in week $t + 1$. However, there is substantial evidence that participants did not detect the pattern in the wage schedule, and that they react only to the current, announced change in wages.

I check whether participants react to future wages by adding future wages to my basic

specification. To include future wages, I have to limit the sample accordingly. The left hand panel of Table A1 includes weeks one to 11. I first present a baseline specification for the subsample, then show specifications with future wages and with fixed effects. Column (1), included for reference, is the same specification as Table 3 column (1). The estimated elasticity when using the first 11 weeks of data barely differs from that for the full sample. Adding a measure of wages one week in the future does not change the estimated elasticity, and the coefficient on future wages is very small and not statistically different from zero in both column (2), which does not include fixed effects, and column (3), which includes individual and week fixed effects. In the right hand panel of Table A1, I further limit the sample in order to include more weeks of future wages. None of the coefficients on the measures of future wages are significant, and I also reject joint significance of the coefficients on future wages. I interpret this table as evidence that participants did not detect the negative serial correlation in the wages, and that their labor supply decision was based on current wages rather than anticipation of future wages.

Another challenge to the interpretation of my estimates as intertemporal parameters is that the underlying expectations about wages could have changed over the course of the experiment. Though I design the experiment to replicate typical market employment as much as possible by having regular employers supervise the work and distribute wages, and by using a task for which a wage market does exist, participants were aware that they were working for a “project” with the very non-standard feature of high-variance wages. At the beginning of the project, it is reasonable to assume that they expected a wage of MK 110 – the usual wage rate on government projects. The assumption is that the temporary, announced changes in wages for the project did not alter participants’ underlying expectations. If, however, expectations evolved in response to realized wage shocks, then the estimated elasticity would not be intertemporal in the standard sense of a change in labor supply in response to an anticipated temporary change from the long run expectation of wages.

The robustness of my estimates to week fixed effects provides some indication that changes in expectations – which would be correlated with time in the project – are not a major factor. For a more direct test, I include wages in past weeks, using specifications analogous to those for future weeks in Table A1. That past wages do not affect the probability of working and that the coefficient on current wages does not change when past wages are added to the

regression is consistent with two important aspects of participants' decisions about whether or not to work. First, those results suggest that expectations about future wages are not changing in response to past wages. Second, they suggest that there is no income effect of past earnings on the current employment decision. Instead, each week's choice about whether to work or not work can be interpreted as a response to the temporary change in the wage that week.

Indeed, the results in Table A2 support both hypotheses. As before, the left hand panel of the table uses 11 weeks of data and incorporates one additional week of wages, and the right hand panel uses eight weeks of data and four weeks of additional wages. The coefficient on wages in week $t - 1$ is significant when using one week of past wages with no fixed effects (column (2)), but none of the coefficients on wages one, two, three, or four weeks prior are individually or jointly significant in any of the other specifications.

In Table A3, I use the running average of wages in previous weeks as an alternative specification to study the effect of past wages on employment. Note that the construction of this additional variable is different for each week. In week two, the "average" of past wages in village v is simply the wage in village v in week 1. In week three, the average of past wages in village v is the average of wages in weeks one and two, and so on. The effect of past wages is not statistically significant either with or without week and individual fixed effects, and including this measure does not change the coefficient on current wages. Tables A2 and A3 provide strong evidence that changes in expectations are not affecting the magnitude of the elasticity or the interpretation of that elasticity as an intertemporal parameter.

My main specifications are ordinary least squares regressions even though the dependent variable is binary. I use OLS rather than maximum likelihood estimators in order to recover marginal effects estimates from specifications that include individual fixed effects, which is not possible with a conditional logit model. However, I present estimates from probit specifications without individual fixed effects in order to demonstrate that the OLS coefficients and probit marginal effects are nearly identical. Columns (1) through (4) in Table A4 correspond to the same-numbered columns in the main results table, Table 3. In subsequent analyses, I will use OLS specifications for ease of interpretation and to allow inclusion of individual fixed effects where appropriate.

An additional cause for concern is whether respondents reacted not to wages, but to

some other aspect of the experimental setting. A specific pitfall would be if labor supply was inelastic because respondents felt pressured to work despite the wage, or thought they would be eligible for some other benefit if they were perceived as “cooperative” or “hard-working.” I have evidence that this is not the case. Respondents listed up to three reasons for working in weeks that they worked, or three reasons for not working in weeks they did not work. Wages do not appear to be a major factor in the decision either to work or not to work. Reasons for working were grouped into four categories: because of the wage (used only when the respondent’s literal answer was “because of the wage” or “because the wage was good”), to get money to spend immediately, to get money to save, or because of social pressure or perceived benefits besides the wage. Figure 2 shows the fraction of individuals who mentioned each reason, aggregated across weeks for individuals who worked at each wage. Earning money to spend immediately is the dominant factor at all wage levels and is mentioned by over 70 percent of respondents, no matter what the wage. Social pressure to work, which includes being told to work by a local leader or government extension worker or anticipating some reward for cooperation, seems relevant only at the lowest wage, MK 30. The wage itself is mentioned by fewer than two percent of respondents for all wages less than MK 100, but by 30 percent or more of respondents at wages of MK 100 or higher.

Reasons for not working were grouped into six categories: because of the wage (again, used only when respondents specifically referenced bad wages), because the respondent was occupied with other work, because money was not needed, because of a funeral, because of illness (to the respondent or someone he/she was caring for), and because of social pressure not to work. Figure 3 shows the reasons for not working at each wage. Illnesses and funerals were the dominant causes of not working, which is consistent with the strong negative effect of funerals on labor supply in the administrative data. Wages were mentioned by fewer than 20 percent of respondents at all wage levels except for the lowest two, MK 30 and MK 40, and an unexplained spike at MK 80.

These self-reported data are consistent with the highly inelastic labor supply estimated in the previous section. Other factors dominate wages in the decision to work or not to work, even at very high or very low wage levels.

5 Comparison to previous reduced-form estimates

My data differ from data used in previous estimates of labor supply in three important ways. First, wages are randomly assigned. Second, I observe the full distribution of wage offers (for employment covered by my outcome variable), rather than only the average wage *accepted* by each individual. Third, I have panel rather than cross sectional data. In addition to these differences in data, I estimate the elasticity of labor supply at the extensive, rather than the intensive, margin. Even ignoring questions of external validity and using identically-structured data, my results would not match exactly those in the previous literature because I estimate a different parameter. I show that these differences in data and methodology account for my estimates being lower than those found for men and women in Ghana[1] and West Bengal[4], and that there is therefore no reason to suspect that the small elasticities I estimate indicate that Malawi is inherently different than in other developing countries.¹²

The data sets most commonly used in empirical analysis of labor supply are individual-level cross sections with measures of hours or days worked over some interval, average wages received over that interval, and a variety of individual background characteristics. There are two potential sources of bias from estimating in the wages reported in these cross sectional data. First, wages are endogenous and potentially correlated with unobservable characteristics that also affect the amount of labor supplied. Second, relying on the measure of wages received by respondents introduces selection bias because data are censored on the dependent variable.

To address the potential biases in using wages from cross sectional data, previous reduced-form work in both developed and developing countries has used measures of average market wages instead of individuals' own wages. Bardhan[4] is one example of this strategy in the development literature. I collapse my panel into a cross section that mimics the limitations of the commonly available data and use that data set to calculate an intensive margin elasticity that is directly comparable to those in the existing literature. This exercise is helpful in identifying the source of differences between my estimates and those in the previous literature. I focus on four major differences between my preferred estimates and the Bardhan-style

¹²Rosenzweig[22] found a negative elasticity of labor supply with respect to wages for men in India. Neither my experimental results nor those I will present from the time-aggregated cross section support backward bending labor supply curves in Malawi.

estimates to which analysts are limited when using cross sectional data without exogenous variation in wages. The first difference is context: there may be inherent differences between the labor markets in rural Malawi, West Bengal, and Ghana. The second difference is the parameter being estimated. I estimate an extensive margin elasticity, the change in the probability of working on a given day for people who have already selected into the market for day labor. Most of the literature focuses on an intensive margin, the change in hours (or days) worked. The point estimates of the elasticities at these two margins would be different even if estimated from the same data set. The third difference is in the distribution of wages: I observe the full distribution of wage offers, while most estimates have data on censored wages. The fourth difference is in the source of variation in wages. Wages are exogenous by design in my project, but endogenous in non-experimental cross sectional data.

Using the time-aggregated cross section allows me to hold constant the methodological issues, which are the second, third, and fourth differences discussed above. I can then assess whether lack of external validity explains why the elasticities I present in section 4 are lower than those in the previous literature about developing countries.

To construct the dependent variable, I add up the total number of days worked (which ranges from 0 to 12). This is the concept that Bardhan uses by taking the total number of days worked in the seven-day period covered by the survey of households in West Bengal that he analyzes. Note that this measure in my cross section is already more precise than normal in survey data, because it comes from administrative records rather than self-reports. Every individual in the sample worked at least two days, and, on average, individuals worked 10 days. Since every individual worked at least once, it is not possible to estimate the elasticity of labor force participation using the cross sectional data for this sample.

I construct three different measures of wages. First, I use the common “average wage” measure by taking the within-person across-week average *accepted* wage. This measure does not correct for endogenous wages or selection into employment at all. Also, because all wages that were offered in this experiment were accepted by at least some participants (and in practice, even the lowest wage was accepted 73 percent of the time it was offered) and all participants had the same distribution of wage offers, the individual average wage measures in the simulated cross section are endogenous but not censored on the dependent variable. Second, following Bardhan, I compute the “village average wage” as the within-village across-

week average accepted wage. However, the average offered wage is the same in each village in my sample by construction. Therefore, the “village average wage” measure varies across village because of supply side determinants. At least part of the variation in cross sectional data used in previous studies is due to differences in demand in different villages, however. To capture a village average wage measure that incorporates demand-side variation in offered wages, I construct a third measure of wages by sampling half of the weeks in each village. I randomly select six of the 12 weeks of data from each village and compute the within-village across-week average accepted wage for those six weeks. The corresponding outcome variable is the number of those six days that each individual worked.

I present the results from this exercise in Table 4. The dependent variable in this table is the scalar number of days worked during the project. The elasticity is interpreted as the percentage increase in days worked for a one-percent increase in wages, and comes from equation (1) in Section 4.2. Column (1) is a baseline specification with no additional controls. In this specification, a one percent increase in wages is associated with an 8.64 increase in days worked, for an elasticity of 0.86 (because average days worked is close to 10). Despite lack of individual covariates, the r-squared for this specification is very high, 0.81. In column (2), I add village fixed effects. In column (3), I add individual controls for gender, marriage status, age, and three measures of wealth: acres of land owned by the household, number of rooms in the house, and whether the house has a tin roof. The elasticities estimated in these two specifications are 0.85 and 0.86, respectively, and are not statistically different from the baseline specification. In column (4), I use average village wages as the key regressor and do not include any additional covariates. In column (5), I add the same individual covariates as in column (3). Village fixed effects are not separately identified with this measure of wages, so they are not included. The regressor of interest is average village wages. The elasticity is 1.05 without including individual covariates and 0.89 when including those covariates. Neither point estimate is statistically different from estimates using person-specific average wages. I cannot reject perfectly elastic labor supply at the intensive margin ($\epsilon = 1$) in any of the estimates in columns (1) through (5).

In columns (6) and (7), I use data from six randomly chosen weeks per village in order to preserve demand-side variation in offered wages. These results are from 1000 replications of choosing half of the weeks for each village, without replacement. On average, respondents

worked five of six possible days. The elasticity of labor supply with respect to this better-measured concept of average village wages is between 0.33 (without covariates) and 0.30 (with individual covariates).¹³

When I use a comparable data set to identify the intensive margin elasticity, my estimates are similar to or larger than elasticities estimated by Bardhan[4] (0.20 to 0.29) and Abdulai and Delgado[1] (0.32 for men and 0.66 for women). This suggests that the highly inelastic estimates in my preferred specifications that take advantage of the experimental design and estimate the change in the probability of working on a given day are not explained by inherent differences between the labor markets in rural Malawi and these other countries. Instead, a combination of the three types of methodological differences I discussed at the beginning of this section leads to much lower estimates than found in previous research. I would find higher elasticities using data from the same labor market if my data were subject to the biases in standard analysis of a non-experimental cross section.

6 Gender

A long literature suggests that women supply labor more elastically than men in developed countries (e.g. Killingsworth[15], Heckman[14]). Previous work in developing countries is also consistent with women supplying labor more elastically than men in India [22] and Ghana[1]. In Tables 5 and 6, I look at my experimental samples of men and women separately. On average, 81 percent of men work when offered employment. The estimated elasticity for men ranges between 0.16 and 0.19, with fixed effects added across columns in Table 5 as in Table 3. Results for women are strikingly similar. Some 86 percent of women work across the entire sample. Their elasticity with respect to wages falls between 0.14 and 0.15, estimates that are not statistically different from the estimated elasticities for men. In this section, I demonstrate that similar elasticities for men and women is a characteristic of the market for ganyu during Malawi's dry season rather than an artifact of my experimental design.

Just as there are many reasons that my point estimates of the elasticity of employment

¹³I also estimate the elasticity of labor supply by drawing six *consecutive* weeks of data for each village, because consecutive weeks is more closely analogous to the concept measured in cross sectional data. The elasticities from estimates using cross sectional data are 0.39 (without covariates) and 0.37 (with covariates). My preferred specification is the one using non-consecutive weeks, because the wage schedule mechanically reduces the across-village variation when using consecutive weeks.

differ from other estimates in the literature, there are many possible explanations for why the gender patterns in my results do not coincide with those in other studies. The most damaging explanation would be that the similar elasticities for men and women in my results are an artifact of my experimental design and do not reflect true labor supply patterns for Malawi. I test this using data from Malawi’s 2004 IHS survey and show that equal elasticities for men and women are typical of Malawi during the dry season, which is the time of year when my project took place. During the wet season, elasticities for women are higher than for men and therefore conform to the pattern found in the existing literature.

The IHS was administered across all 12 months and includes questions about supply and demand of ganyu. I exploit the variation in timing of survey administration to estimate labor supply elasticities for the wet and dry seasons separately using the survey data. The information about labor supply is somewhat limited: individuals are asked how many hours of ganyu they did in the past week, if they did any ganyu in the past 12 months, and how many days of ganyu they did in the past month. They are also asked how much they received for one day’s ganyu on average for all of the work they did in the past 12 months. There is no data about wages received for ganyu in the past week. Information about demand for ganyu is collected somewhat more precisely: individuals are asked about the amount of ganyu hired and the daily wage paid separately for the rainy (main agricultural season) and dry (off season) separately. I construct a measure of the average wage paid within a Traditional Authority (TA)¹⁴ in the wet and dry seasons respectively.¹⁵ Then, I regress labor supply in the previous week on the TA-level average wage for the corresponding season separately for the wet and dry seasons. I rely on the assumption that individual labor supply does not affect the market wage, and identification comes from across-location variation in wages. Employers do not report the characteristics of those hired to perform ganyu, and they do not report separate wages for men and women. I do not expect the point estimates from the IHS

¹⁴Malawi is divided into 350 administrative regions, which are called “Traditional Authorities” in rural areas and “wards” in urban areas. TAs are roughly the equivalent of counties in the United States; Malawi’s 28 “districts” are more organizationally similar to American states.

¹⁵This measure of wages captures employment by private individuals only. Wages paid by firms, nongovernmental organizations, or the government Public Sector Works Programme are not covered by the IHS survey. Wages paid by these employers are less likely to be seasonal because they are likely to hire at fixed rates, or for non-agricultural projects, or for agricultural work on irrigated land. Therefore, using a measure of wages that is limited to wages paid by private individuals should capture the key source of seasonal variation in ganyu wages.

data to match the point estimates from my experiment: wages in the IHS are endogenous and estimates using the IHS are likely biased. I am interested in comparing the pattern of elasticities by gender, not the point estimates.

Table 7 shows results from this exercise. Panel A contains results for the dry season, which includes June-November. Panel B contains results for the wet season, December-May. The sample is limited to the head of household and his or her spouse, if present, to match the selection criteria for my experiment. All regressions control for gender, age, household items score, housing quality score, land area farmed during the dry season, land area farmed during the wet season, amount of fertilizer used during the rainy season, education, and district of residence. Columns (1) to (3) capture the intensive margin elasticity from the regression of log hours on log wages. The elasticity during the dry season is 0.475, marginally different from zero. Estimates for men and women are imprecise but not significantly different from each other. During the wet season, however, the intensive margin elasticity falls by half and is not statistically different from zero. However, the separate estimates for men and women tell a different story. For men, the point estimate is -0.256, which, while not statistically different from zero, is consistent with previous findings that men's labor supply is either inelastic or in the backward-bending portion of the labor supply curve. Women have an elasticity of 0.639, significantly higher than men. During the rainy or high-productivity season, then, the pattern of men's and women's intensive margin elasticities in rural Malawi are consistent with evidence from other developing countries. During the dry season, though, gender differences are much harder to detect.

The extensive margin estimates in columns (4) to (6) are more comparable to estimates from my experiment. During the dry season, the elasticity of working in the past week for men and women combined is 0.27. Women have somewhat larger elasticities than men, but the difference between men and women is not statistically significant. In the wet season, though, the elasticity for women is 0.45, significantly different from zero, while the elasticity for men is -0.11 and not statistically significant. In other words, finding positive elasticities of working that are similar for men and women does not appear to be an artifact of my experimental design. The same pattern is present in nationally-representative survey data when looking at data from the same part of the agricultural season, though the estimates are less precise.

7 Conclusion

I use experimental variation in wages to study the effect of wages on the probability of working in the daily labor market in rural Malawi. This unique field experiment allows me to estimate a causal effect of wages on the probability of employment and to avoid the standard problems associated with simultaneous determination of supply and demand in cross sectional data about employment. I randomize wages at the village-week level, then offer employment to up to two adult members of pre-selected households in participating villages for one day per week for 12 weeks. The final sample consists of 529 individuals in 298 households, across ten villages. The panel of administrative outcomes allows me to use individual fixed effects in most specifications. I estimate that the elasticity of employment for individuals in this sample is between 0.15 and 0.17, and I robustly reject perfectly inelastic supply of labor in all specifications.

Two patterns in my results are distinct from those in the previous literature, and while my point estimates are unlikely to apply to other countries, these patterns may be more general. First, my point estimates of the elasticity of employment are very low relative to those from Ghana and West Bengal, but firmly reject the backward-bending labor supply curve that has been found for men in India (and in many developed countries). I show that my preferred estimates using the experimental panel are much lower than the estimates I would obtain using commonly available cross sectional data. This suggests that previous intensive margin estimates using cross sectional data may overstate the responsiveness to wages that actually characterizes the decision to work or not work on a given day. Second, I find that men and women have the same elasticity of employment. This finding is in stark contrast to the literature from both developing and developed countries that indicates a substantially higher elasticity of labor supply for women than men. The equality of men's and women's elasticities is not an artifact of the experimental design, but rather a characteristic of Malawi's labor market during the unproductive dry season. Further research to explore gender patterns in the seasonality of labor supply in countries with distinct wet and dry seasons is warranted, and has the potential to inform the design and targeting of public sector employment programs.

After weeks four, eight, and 12, I collect survey data about recollection of wages and work history, as well as reasons for working or not working. The data about recollection

of wages and work history confirm that respondents are accurate in their memory of the events, reporting both wages and past work accurately in 83 percent of the cases. I then use information from weeks in which respondents remembered the wage and whether they worked to examine self-reported reasons for working. At all wage levels, earning money to spend immediately is the most frequently reported reason for working, and funerals and illnesses are the dominant reasons for not working. Wages are cited by more than 20 percent of respondents as a reason for not working predominantly at very low wages (MK 30 and MK 40), and as a reason for working only at high wages of MK 100 or higher. These survey responses are consistent with the inelastic supply of labor observed in the administrative data.

Understanding the labor supply behavior of poor individuals is crucial for the design of public employment projects in Malawi and other developing countries. The Government of Malawi and the World Bank are spending \$40 million on a Community Livelihoods Support fund that uses public sector employment to meet dual goals: providing a safety net for poor individuals by offering employment, and improving infrastructure in the communities where those individuals live. Inelastic labor force participation makes it clear that there are stark tradeoffs between these goals when determining wage levels for the program. Malawi is not the only developing country with an interest in public employment programs: 29 countries in sub-Saharan Africa alone have such programs. The estimates I obtain from my experiment in Malawi not only contribute to the long and evolving literature about labor supply in developing countries, but also provide important parameters for understanding the impact of government and NGO programs that are already reaching millions of people.

Tables

Table 1: Weekly Wage Schedule (MK)

	1	2	3	4	5	6	7	8	9	10	11	12	Total
Kafotokoza	40	100	60	120	30	110	70	140	80	130	90	50	1020
Chimowa	100	60	120	30	110	70	140	80	130	90	50	40	1020
Manase	60	120	30	110	70	140	80	130	90	50	40	100	1020
Lasani	120	30	110	70	140	80	130	90	50	40	100	60	1020
Njonja	30	110	70	140	80	130	90	50	40	100	60	120	1020
Hashamu	110	70	140	80	130	90	50	40	100	60	120	30	1020
Kachule	70	140	80	130	90	50	40	100	60	120	30	110	1020
Msangu/Kalute	140	80	130	90	50	40	100	60	120	30	110	70	1020
Kamwendo	80	130	90	50	40	100	60	120	30	110	70	140	1020
Kunfunda	130	90	50	40	100	60	120	30	110	70	140	80	1020
Average	88	93	88	86	84	87	88	84	81	80	81	80	

Table 2: Baseline Characteristics

	Mean	SD	N	10th	Median	90th
Male	0.40	0.49	529			
One male and one female in HH	0.70	0.46	529			
Two female participants	0.14	0.35	529			
Two male participants	0.04	0.19	529			
One participant	0.13	0.33	529			
Married	0.80	0.40	495			
Years of education	4.33	3.15	493	0	4	8
Number of adults in HH	2.25	0.97	495	1	2	3
Number of children in HH	3.12	1.90	495	1	3	6
Tin roof	0.16	0.37	495			
Number of rooms	2.02	0.92	490	1	2	3
Acres of land	1.81	0.87	495	1	1.5	3
Days of paid work last week	1.02	1.59	495	0	0	3
Days of paid work last month	2.73	4.65	495	0	1	7

Table 3: Elasticity of employment w.r.t. wages

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
	Individual*day indicator for working					
Ln(wage)	0.127*** (0.033)	0.127*** (0.033)	0.140*** (0.032)	0.140*** (0.032)	0.127*** (0.033)	0.140*** (0.032)
Village effects		x		x		
Week effects			x	x		x
Individual effects					x	x
Observations	6333	6333	6333	6333	6333	6333
Mean of dependent variable	0.84	0.84	0.84	0.84	0.84	0.84
Elasticity	0.15 (0.040)	0.15 (0.040)	0.17 (0.040)	0.17 (0.040)	0.15 (0.040)	0.17 (0.040)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

Unit of observation is individual*week, sample is all individuals.

* p<0.10, ** p<0.05, *** p<0.001

Table 4: Elasticity of labor supply from time-aggregated cross sectional data

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Number of days worked out of 12 weeks	Number of days worked out of 12 weeks	Number of days worked out of 12 weeks	Number of days worked out of 12 weeks	Number of days worked out of 6 weeks	Number of days worked out of 6 weeks
Ln(average accepted wage)	8.641*** (0.436)	8.527*** (0.478)	8.795*** (0.574)				
Ln(village average accepted wage, all weeks)				10.586*** (0.928)	9.059*** (1.167)		
Ln(village average accepted wage, half of weeks)						1.633* (0.841)	1.527** (0.765)
Village effects		x	x				
Individual controls				x	x		x
Observations	529	529	488	529	488	529	488
Mean of dep. var.	10.09	10.09	10.18	10.09	10.18	5.02	5.07
Elasticity	0.86	0.85	0.86	1.05	0.89	0.33	0.30

Standard errors in columns (1) to (5) are clustered at the village level, and standard errors in columns (6) and (7) are from 1000 bootstrap replications.

Additional controls are gender, marital status, age, number of rooms, acres owned, and having a tin roof. Sample in columns (1), (2), (4), and (6) is all individuals.

Sample in columns (3), (5), and (7) is all individuals who answered baseline survey.

* p<0.10, ** p<0.05, *** p<0.001

Table 5: Elasticity of men's employment w.r.t. wages

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
	Individual*day indicator for working					
Ln(wage)	0.139*** (0.032)	0.139*** (0.032)	0.157*** (0.035)	0.157*** (0.035)	0.139*** (0.032)	0.157*** (0.035)
Village effects		x		x		
Week effects			x	x		x
Individual effects					x	x
Observations	2532	2532	2532	2532	2532	2532
Mean of dependent variable	0.81	0.81	0.81	0.81	0.81	0.81
Elasticity	0.17 (0.043)	0.17 (0.043)	0.19 (0.047)	0.19 (0.047)	0.17 (0.043)	0.19 (0.047)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

Unit of observation is individual*week, sample is all men.

* p<0.10, ** p<0.05, *** p<0.001

Table 6: Elasticity of women's employment w.r.t. wages

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
	Individual*day indicator for working					
Ln(wage)	0.119** (0.035)	0.119** (0.035)	0.129*** (0.032)	0.129*** (0.032)	0.119** (0.035)	0.129*** (0.032)
Village effects		x		x		
Week effects			x	x		x
Individual effects					x	x
Observations	3801	3801	3801	3801	3801	3801
Mean of dependent variable	0.86	0.86	0.86	0.86	0.86	0.86
Elasticity	0.14 (0.041)	0.14 (0.041)	0.15 (0.038)	0.15 (0.039)	0.14 (0.041)	0.15 (0.038)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

Unit of observation is individual*week, sample is all women.

* p<0.10, ** p<0.05, *** p<0.001

Table 7: Dry and wet season elasticities from IHS data

Panel A. Dry Season						
Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
	Ln(hours worked in past week)			Indicator for any work in past week		
	All	Men	Women	All	Men	Women
Ln(average paid wage in TA, dry season)	0.475*	0.576	0.380	0.033*	0.041	0.026
	(0.260)	(0.413)	(0.305)	(0.019)	(0.029)	(0.022)
Observations	1709	876	833	1709	876	833
Mean of dependent variable	2.04	3.04	0.99	0.12	0.17	0.07
Average wage	70.78	70.78	70.78	70.78	70.78	70.78
Elasticity				0.27	0.24	0.36

Panel B. Wet Season						
Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
	Ln(hours worked in past week)			Indicator for any work in past week		
	All	Men	Women	All	Men	Women
Ln(average paid wage in TA, rainy season)	0.232	-0.256	0.639*	0.017	-0.017	0.044*
	(0.293)	(0.455)	(0.369)	(0.021)	(0.033)	(0.027)
Observations	1805	887	918	1805	887	918
Mean of dependent variable	1.72	2.12	1.33	0.13	0.16	0.10
Average wage	83.78	83.78	83.78	83.78	83.78	83.78
Elasticity				0.13	-0.11	0.45

Data from the 2004 Malawi Integrated Household Survey.

OLS estimates. * p<0.10, ** p<0.05, *** p<0.001

Sample includes all adult heads of household or spouses who are at least 18 years old and who live in rural areas.

All estimates include controls for gender, age, household items score, housing quality score, land area farmed during the dry season, land area farmed during the wet season, amount of fertilizer used during the rainy season, and indicators for using any fertilizer during the rainy season, education category, and district of residence.

Table 8: Effect of savings accounts on elasticity of employment w.r.t. wages

	(1)	(2)	(3)	(4)
Dependent variable:	Individual	Individual	Household	Household
Ln(wage)	0.141*** (0.033)	0.140*** (0.027)	0.287*** (0.044)	0.269*** (0.051)
Account	-0.017 (0.013)	-0.026 (0.097)	-0.041 (0.030)	-0.195 (0.191)
Account*Ln(wage)		0.002 (0.020)		0.035 (0.041)
Village effects	x	x	x	x
Week effects	x	x	x	x
Observations	6285	6285	2748	2748
Mean of dependent variable	0.84	0.84	0.74	0.74
Elasticity	0.17 (0.041)		0.17 (0.042)	
Elasticity (no account)		0.17 (0.034)		0.16 (0.035)
Elasticity (account)		0.17 (0.049)		0.18 (0.051)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

In columns (1) and (2), unit of observation is individual*week, sample is all individuals.

In columns (3) and (4), unit of observation is HH*week, sample is HHs with two participants.

* p<0.10, ** p<0.05, *** p<0.001

Table 9: Effect of accumulated savings on elasticity of employment w.r.t. wages

	(1)	(2)	(3)	(4)
Dependent variable:	Individual	Individual	Household	Household
Ln(wage)	0.122*** (0.030)	0.203 (0.126)	0.248*** (0.056)	0.420 (0.263)
Ln(savings)	0.011 (0.011)	0.075 (0.098)	0.022 (0.026)	0.157 (0.213)
Ln(savings)*Ln(wage)		-0.015 (0.021)		-0.031 (0.046)
Individual effects	x	x		
Household effects			x	x
Observations	1666	1666	730	730
Mean of dep. variable	0.89	0.89	1.76	1.76
Elasticity	0.14 (0.034)		0.14 (0.033)	
Elasticity (no savings)		0.23 (0.141)		0.24 (0.149)
Elasticity (with savings)		0.21 (0.118)		0.22 (0.123)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

Sample restricted to households that received savings accounts.

In columns (1) and (2), unit of observation is individual*week.

In columns (3) and (4), unit of observation is household*week,

sample is restricted to households with two participants.

* p<0.10, ** p<0.05, *** p<0.001

Figures

Figure 1: Fraction working at each wage (wages in MK)




Figure 2: Self-Reported Reasons for Working

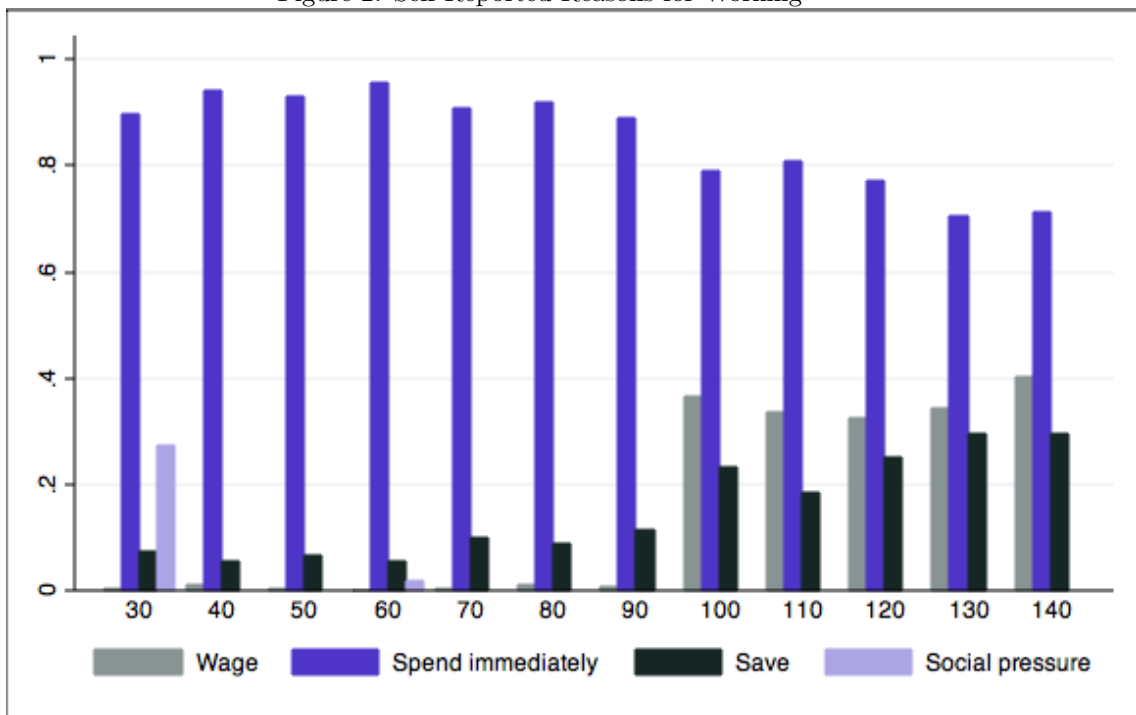
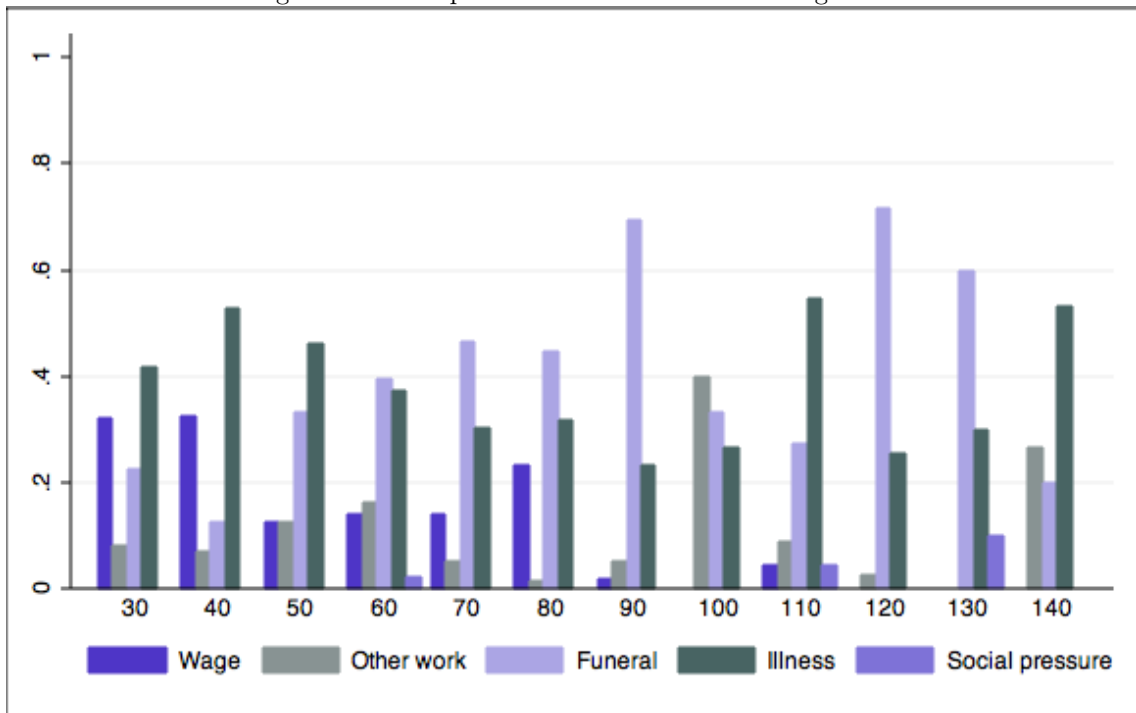


Figure 3: Self-Reported Reasons for Not Working



References

- [1] Awudu Abdulai and Christopher Delgado. Determinants of nonfarm earnings of farm-based husbands and wives in northern ghana. *American Journal of Agricultural Economics*, pages 117–130, Feb 1999.
- [2] Joshua D. Angrist and Jorn-Steffen Pischke. *Mostly Harmless Econometrics*. Princeton University Press, 2009.
- [3] O Ashenfelter, K Doran, and B Schaller. A shred of credible evidence on the long run elasticity of labor supply. *NBER Working Paper*, Jan 2010.
- [4] Pranab Bardhan. Labor supply functions in a poor agrarian economy. *American Economic Review*, pages 73–83, Mar 1979.
- [5] Tim Barmby and Peter Dolton. What lies beneath? effort and incentives on archaeological digs in the 1930’s. *Working Paper*, pages 1–38, Sep 2009.
- [6] Gary Burtless and Jerry A Hausman. The effect of taxation on labor supply: evaluating the gary negative income tax experiment. *Journal of Political Economy*, 86(6):1103–1130, Oct 1978.
- [7] Colin Camerer, Linda Babcock, George Loewenstein, and Richard Thaler. Labor supply of new york city cabdrivers: one day at a time. *Quarterly Journal of Economics*, pages 407–441, May 1997.
- [8] A Colin Cameron, Jonah B Gelbach, and Douglas L Miller. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427, Jul 2008.
- [9] Y K Chou. Testing alternative models of labor supply: evidence from taxi-drivers in singapore. *University of Melbourne Department of Economics Working Paper Series*, 768:1–39, Nov 2000.
- [10] Henry S Farber. Is tomorrow another day? the labor supply of new york cab drivers. *NBER*, page 43, May 2003.

- [11] Ernest Fehr and Lorenz Goette. Do workers work more if wages are high? evidence from a randomized field experiment. *American Economic Review*, 97(1):298–317, Apr 2007.
- [12] John Ham. Testing whether unemployment represents life-cycle labor supply. *Review of Economics Studies*, 54(4):559–578, October 1986.
- [13] John Ham and Kevin Reilly. Testing intertemporal substitution, implicit contract, and hours restriction models of the labor market using micro data. *American Economic Review*, pages 905–927, 2002.
- [14] James J Heckman. What has been learned about labor supply in the past twenty years? *American Economic Review*, 83(2):116–121, Sep 1993.
- [15] Mark R. Killingsworth. *Labor Supply*. Cambridge Surveys of Economic Literature. Press Syndicate of the University of Cambridge, 1983.
- [16] Anjini Kochar. Smoothing consumption by smoothing income: hours-of-work responses to idiosyncratic agricultural shocks in rural india. *The Review of Economics and Statistics*, 81(1):50–61, Sep 1999.
- [17] W Arthur Lewis. Economic development with unlimited supplies of labour. *The Manchester School of Economic and Social Studies*, XXII(2):139–191, 1954.
- [18] Anna McCord and Rachel Slater. Overview of public works programmes in sub-saharan africa. Technical report, Overseas Development Institute, 2009.
- [19] B Meyer. Labor supply at the extensive and intensive margins: the eitic, welfare, and hours worked. *American Economic Review*, Jan 2002.
- [20] Gerald S Oettinger. An empirical analysis of the daily labor supply of stadium vendors. *Journal of Political Economy*, 107(2):360–392, Dec 1999.
- [21] Elaina Rose. Ex ante and ex post labor supply response to risk in a low-income area. *Journal of Development Economics*, 64:371–388, Jan 2001.
- [22] Mark Rosenzweig. Rural wages, labor supply, and land reform: A theoretical and empirical analysis. *The American Economic Review*, pages 847–861, Dec 1978.

Appendices

A Standard errors

There are a range of potential challenges to calculating appropriate standard errors for the estimates in this paper. Anticipated problems with using unadjusted OLS standard errors are generic heteroskedasticity, correlation in outcomes at the village-week level, correlation in outcomes at the village level, and the relatively small number of clusters (villages) in the sample. Additionally, analytic standard errors for the elasticity cannot be computed since the joint distribution of $labor_{itv}$ and β is unknown. In this subsection, I discuss each of these issues and the method of calculating standard errors to address each issue in turn. I demonstrate that the block-bootstrap standard errors I use in the main results throughout the paper are conservative, and that results are robust to alternative ways of calculating standard errors.

Table A5 reports standard errors, p-values, and t-statistics for the coefficient on log wages from the regression $labor_{itv} = \alpha + \beta \ln(wage_{tv}) + \nu$ with no additional covariates. When possible, I include standard errors, p-values, and t-statistics for the elasticity of employment $\epsilon_e = \frac{\beta}{mean(labor)}$. I have deliberately omitted subscripts on the residual term ν ; I address various possibilities for the structure of the error term and techniques for dealing with them in the remainder of this section. The standard error of 0.010 in column (1) is unadjusted and included as a benchmark. The t-statistic for the test that $\beta = 0$ is 13.149, and the associated p-value is less than 0.001. The assumption underlying the standard errors in column (1) is that the residuals $\nu = \nu_{itv}$ are distributed i.i.d..

With the linear probability model, there is heteroskedasticity in the residuals such that the distribution of the residuals $\nu = \nu_{itv}$ is conditional on the regressors. In column (2), I allow for possible heteroskedasticity in the error terms by using heteroskedasticity-robust (Eicker-White) standard errors. The point estimate of the standard error on log wages is virtually unchanged and the t-statistic for the test that $\beta = 0$ declines slightly to 12.566.¹⁶

¹⁶Angrist and Pischke[2] point out that if the standard errors are in fact homoskedastic, the robust estimator is more biased than the conventional estimator. Their suggestion of using the maximum of the conventional and robust standard errors is unnecessary in my case, because the two estimators produce nearly identical standard errors.

A second concern is that there could be village-week correlation in outcomes. This could take the form of village-week specific shocks, such as an illness that affects one village in a single week. In this case, the residuals have the structure $\nu = \nu_{tv} + \nu_{itv}$ and village-week clustered standard errors are appropriate. I report these standard errors in column (3). The standard error for the coefficient on log wages increases to 0.029, for a t-statistic of 4.306. An alternative approach for addressing village-week correlation is to aggregate to 120 village-week observations. Angrist and Pischke[2] suggest showing that results are robust to analysis at the group level when the number of clusters is small. Since treatment is at the village level, this approach also makes clear the source of variation. In columns (4) and (5), the dependent variable is the fraction of participants in each village v who work in week t . I use Stata's `aweights` to weight by the square root of the number of participants per village. The standard error in column (4) is unadjusted, and the standard error in column (5) is robust to heteroskedasticity. As expected, the standard errors obtained from using village averages are not much different than the clustered standard errors, and conclusions about the magnitude of the elasticity of employment are robust to group-level analysis.

A third concern is that there could be village-level correlation in the outcomes. Village level correlation could come from persistent village-level shocks, such as an illness that strikes in one week and lingers or has effects in subsequent weeks, or could simply be that outcomes in villages are correlated because the people who live in the same village have many unobserved (but not time-invariant) characteristics that affect their employment probabilities in common. In either case, the residuals would have the structure $\nu = \nu_v + \nu_{itv}$. In this case, standard errors should be clustered at the village level. The village level is also the level of randomization, and since the regressor of interest varies only at the group level the impact of clustering is potentially large. The standard errors in column (6) are clustered at the village level. The standard error of β is 0.035; the t-statistic for the test that $\beta = 0$ is 3.600, and the p-value for that test is 0.006. As expected, clustering increases the magnitude of the standard errors. However, the point estimate of β remains significantly different from zero when using clustered standard errors.

The relatively small number of villages in my sample may be problematic if there are persistent village-level shocks. In column (7), I allow for persistent village-level shocks and address the small number of villages by calculating the standard errors from 500 block-

bootstrap replications. In this approach, first proposed by Hall, I re-sample villages with replacement and calculate the coefficient β in each replication. The standard error is the standard deviation of the coefficients from 500 replications. Bootstrapping is a common approach with a small number of units of randomization, as it simulates a larger sample of villages. The standard errors I obtain from the block-bootstrap procedure are very similar to those from clustering at the village level. The standard error of the coefficient on log wages is 0.033, for a t-statistic of 3.848 against the null hypothesis that $\beta = 0$.

With the block-bootstrap, I can also compute the point estimate of the elasticity $\epsilon_e = \frac{\beta}{\text{mean}(\text{labor})}$ in each replication and obtain a standard error for the elasticity. I do not report standard errors for the estimated elasticities in columns (1) to (6) because I cannot calculate a standard error for the elasticity analytically or with the delta method, since the joint distribution of the coefficients and the dependent variable is not known. The standard error for the elasticity of employment is 0.040. The 95 percent confidence interval for the elasticity is [0.072, 0.228], meaning that I reject perfectly inelastic labor supply but also reject elasticities higher than about 0.23.

Cameron et al.[8] demonstrate in a recent paper that the block bootstrap procedure produces downwardly-biased standard errors when the number of clusters is “small.” Their simulations are for data with six clusters; my 10 villages are few enough to merit consideration of their alternate procedure, a residual-swapping or “wild” bootstrap. A complication arises in implementing their procedure for my results: the method that they propose is a bootstrap-t procedure, not a procedure for estimating standard errors. When computing a t-statistic, though, it is necessary to propose a null hypothesis. For estimates of treatment effects, the null of zero is natural. For estimates of the effect of changes in wages on the probability of employment, though, the most interesting null hypothesis is not obvious. Therefore, I loop over 101 different possible values of H^0 from 0 to 1, in increments of 0.01, and calculate the bootstrap-t statistic associated with each of those possible null hypotheses H_h^0 .

I follow procedure 2a from Cameron et al.’s Appendix B. For each of 500 replications, I draw a sample of 10 villages with replacement. I estimate $\text{labor}_{itv} = \alpha + \beta \ln(\text{wage}_{tv}) + \nu$ using the bootstrap sample to obtain the point estimate of the coefficient. I also calculate the elasticity, $\hat{\epsilon}_e$, for each replication r . Then for each replication r , I calculate the restricted residuals $\nu_{\mathbf{r}h}$ from imposing each of 101 values of the null hypothesis H_h^0 from 0 to 1.

For each vector of residuals ν_{rh} , I follow Cameron et al.'s method of randomly swapping the sign of half of the residuals ν_{rhi} , then computing a new predicted outcome \widehat{labor}_{itv} by adding the residual to the observed outcome for each observation. I then estimate $\widehat{labor}_{itv} = \tilde{\alpha} + \tilde{\beta} \ln(wage_{tv}) + u$ and take the t-statistic for the test that $\tilde{\beta} = H_h^0$. I obtain 500 t-statistics for each of the 101 null hypotheses. I report the 95 percent confidence intervals for t-statistics of the tests that $\beta = 0$ and $\epsilon_e = 0$ in column (8). Reporting the statistic for the test of $\beta = 0$ is the standard convention in regression output and corresponds to the significance levels from block-bootstrapped standard errors that I report throughout this paper.

However, as discussed above, the tests that β and especially ϵ_e are zero are perhaps not the most relevant when estimating the elasticity of employment. Instead of taking a stand on the most appropriate null hypothesis, in Figures A1 and A2 I plot the rejection rate (t-statistics below -1.96 or above +1.96) against each possible value of ϵ_e^0 between 0 and 1. Rejection rates from the wild bootstrap procedure are lowest for null hypotheses of β and ϵ_e that approximate the confidence intervals from the clustered or block-bootstrapped standard errors.

My main results are robust to adjusting standard errors to allow for generic heteroskedasticity, village-week correlation in outcomes, and village level correlation in outcomes. The results also stand up to bootstrapping methods that take account of the small number of clusters in my data. The block-bootstrapped standard errors that I use throughout the paper are conservative in their magnitude and address both village level correlation in standard errors and the small number of villages in the sample.

B Appendix Tables

Table A1: Elasticity of employment w.r.t. future wages

	(1)	(2)	(3)	(4)	(5)	(6)
	Weeks 1 to 11			Weeks 1 to 8		
Dependent variable:	Individual*day indicator for working					
Ln(wage)	0.131*** (0.037)	0.125*** (0.035)	0.142*** (0.034)	0.149** (0.047)	0.120** (0.045)	0.133** (0.055)
Ln(wage _{t+1})		-0.018 (0.044)	-0.010 (0.037)		0.027 (0.080)	0.029 (0.066)
Ln(wage _{t+2})					0.016 (0.048)	0.013 (0.027)
Ln(wage _{t+3})					-0.047 (0.037)	-0.028 (0.044)
Ln(wage _{t+4})					0.029 (0.039)	0.039 (0.041)
Week effects			x			x
Individual effects			x			x
Observations	5805	5804	5804	4221	4217	4217
Mean of dependent variable	0.84	0.84	0.84	0.81	0.81	0.81
Elasticity	0.16 (0.046)	0.15 (0.043)	0.17 (0.044)	0.18 (0.060)	0.15 (0.057)	0.16 (0.071)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

Unit of observation is individual*week, sample is all individuals.

* p<0.10, ** p<0.05, *** p<0.001

Table A2: Elasticity of employment w.r.t. past wages

	(1)	(2)	(3)	(4)	(5)	(6)
	Weeks 2 to 12			Weeks 5 to 12		
Dependent variable:	Individual*day indicator for working					
Ln(wage)	0.141*** (0.036)	0.123*** (0.034)	0.143*** (0.033)	0.166*** (0.027)	0.176*** (0.031)	0.175*** (0.028)
Ln(wage _{t-1})		-0.057** (0.025)	-0.029 (0.026)		0.011 (0.011)	0.014 (0.016)
Ln(wage _{t-2})					-0.005 (0.019)	-0.004 (0.017)
Ln(wage _{t-3})					0.009 (0.018)	0.009 (0.011)
Ln(wage _{t-4})					-0.002 (0.018)	-0.002 (0.017)
Week effects			x			x
Individual effects			x			x
Observations	5813	5813	5813	4226	4226	4226
Mean of dependent variable	0.84	0.84	0.84	0.89	0.89	0.89
Elasticity	0.17 (0.044)	0.15 (0.042)	0.17 (0.041)	0.19 (0.034)	0.20 (0.038)	0.20 (0.034)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

Unit of observation is individual*week, sample is all individuals.

* p<0.10, ** p<0.05, *** p<0.001

Table A3: Elasticity of employment w.r.t. the average of past wages

	(1)	(2)	(3)
	Weeks 2 to 12		
Sample:	Individual*day indicator for working		
Dependent variable:	Individual*day indicator for working		
Ln(wage)	0.141*** (0.036)	0.152*** (0.035)	0.140*** (0.034)
Ln(\overline{wage}_{t-1})		0.075 (0.065)	-0.143 (0.087)
Week effects			x
Individual effects			x
Observations	5813	5813	5813
Mean of dependent variable	0.84	0.84	0.84
Elasticity	0.17 (0.044)	0.18 (0.043)	0.17 (0.042)

OLS estimates. Cluster bootstrapped standard errors (clustered at the village level).

Unit of observation is individual*week, sample is all individuals.

* p<0.10, ** p<0.05, *** p<0.001

Table A4: Probit estimates of the elasticity of employment w.r.t. wages

Dependent variable:	(1)	(2)	(3)	(4)
	Individual*day indicator for working			
Ln(wage)	0.120*** (0.030)	0.119*** (0.030)	0.136*** (0.030)	0.135*** (0.031)
Village effects		x		x
Week effects			x	x
Observations	6333	6333	6333	6333
Mean of dependent variable	0.84	0.84	0.84	0.84
Elasticity	0.14 (0.037)	0.14 (0.037)	0.16 (0.038)	0.16 (0.039)

Marginal effects (derivative at the mean) from probit estimates.

Cluster bootstrapped standard errors (clustered at the village level).

Unit of observation is individual*week, sample is all individuals.

* p<0.10, ** p<0.05, *** p<0.001

Table A5: Different methods for computing standard errors

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		Individual*day indicator for working	Individual*day indicator for working	average employment	Village*day average employment	Individual*day indicator for working	Individual*day indicator for working	Individual*day indicator for working
Ln(wage)	0.127*** (0.010)	0.127*** (0.010)	0.127*** (0.029)	0.128*** (0.031)	0.126*** (0.030)	0.127*** (0.035)	0.127*** (0.033)	0.127
t-statistic	13.149	12.566	4.306	4.062	4.152	3.600	3.848	[0.781, 3.128]
p-value	0.000	0.000	0.000	0.000	0.000	0.006		
Elasticity	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15
Observations	6333	6333	6333	120	120	6333	6333	6333
Mean of dependent variable	0.84	0.84	0.84	0.85	0.85	0.84	0.84	0.84
SE method	no adjustment	robust to heteroskedasticity	clustered at village-week level	collapsed to village-week level	collapsed and robust at village-week level	clustered at village level	block bootstrapped	wild bootstrapped

OLS estimates. Unit of observation is individual*week in columns (1) to (3) and (6) to (8), and village*week in columns (4) and (5). Sample is all individuals.
 * p<0.10, ** p<0.05, *** p<0.001

C Appendix Figures

Figure A1: Rejection rate for null hypotheses about β from bootstrap-t procedure




Figure A2: Rejection rate for null hypotheses about ϵ_e from bootstrap-t procedure

