American Economic Journal: Applied Economics 2016, 8(1): 129–149
http://dx.doi.org/10.1257/app.20130369

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

By Jessica Goldberg*

I use a field experiment to estimate the wage elasticity of employment in the day labor market in rural Malawi. Once a week for 12 consecutive weeks, I make job offers for a workfare-type program to 529 adults. The daily wage varies from the tenth to the ninetieth percentile of the wage distribution, and individuals are entitled to work a maximum of one day per week. In this context (the low agricultural season), 74 percent of individuals worked at the lowest wage, and consequently the estimated labor supply elasticity is low (0.15), regardless of observable characteristics. (*JEL* C93, J22, J31, O15, O18, R23)

Labor is a critical resource for the poor in developing countries, and labor markets in these countries are very different from those in industrialized countries. While 1.65 billion people worldwide are employed for regular wages, another 1.5 billion people, including most working adults in developing countries, are self-employed or participate in the casual day-labor market (World Bank 2013). Despite the importance of casual labor markets in developing countries, we lack evidence about how labor supply is determined in these settings.

Understanding labor supply can inform public policy in areas such as wage-setting for public works programs. The most widely known public cash-for-work program is India's National Rural Employment Guarantee Scheme, which employed almost 45 million day laborers in 2008–2009 alone, but similar programs exist in 29 sub-Saharan African countries (McCord and Slater 2009), including Malawi, the setting for this paper.

Proponents often describe these programs as self-targeting to poor beneficiaries through low wages (Besley and Coate 1992) and high time costs to participate. The opportunity cost of time is greater for wealthier households, and therefore poor households are more likely to select into these programs. Alatas et al. (2013) show

* Department of Economics, University of Maryland, 3115 G Tydings Hall, 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. Institutional Review Board (IRB) approval was obtained from the University of Michigan. Field work would not have been possible without Geoffrey Mdumuka, Lonnie Mwamlima, Kingsley Naravato, and the staffs of the Lobi Horticultural Association and Lobi Extension Planning Area office. I thank Brian Jacob, David Lam, Jeff Smith, Dean Yang, Susan Godlonton, Erik Johnson, Sara LaLumia, Molly Lipscomb, Elias Walsh, Kathleen Beegle, Judy Hellerstein, Sebastian Galliani, John Ham, and various seminar and conference participants for their extremely helpful comments. Brian Quistorff, Matthew Staiger, and Tara Kaul provided excellent research assistance. All errors and omissions are my own.

† Go to http://dx.doi.org/10.1257/app.20130369 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

theoretically that under some circumstances, so-called "ordeal mechanisms," such as travel to a registration center, are not sufficient to improve targeting. In their experiment in Indonesia, reductions in the distance to a registration center does not change the ratio of wealthy to poor households who apply for a conditional cash transfer program.

In public works programs, self-selection is determined by willingness to do the work required for the wage offered by the program. Therefore, estimating the elasticity of labor supply is important in understanding who is likely to participate in such programs, and how changes in daily wages will affect targeting. There are few convincing estimates of labor supply elasticities in developing countries, and the existing evidence comes from observational data. Studies of informal rather than salaried work in developing countries dates back to Lewis (1954), which assumes that the supply of labor is perfectly elastic. More recently, empirical estimates of labor supply elasticities in rural markets in developing countries have generally supported an upward-sloping labor supply curve. Bardhan (1979) estimates upward sloping labor supply curves with what he characterizes as "very small" elasticities for rural households in West Bengal; Abdulai and Delgado (1999) estimates somewhat greater elasticities for husbands and wives in Ghana. Rosenzweig (1978) 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.[1] These papers are all identified from nontransitory changes in wages and estimate changes in the labor supply curve. As Oettinger (1999) demonstrates, there is substantial downward bias in OLS estimates of labor supply elasticities from observational data, as changes in wages reflect shifts in both labor demand and labor supply.

In contrast to the previous studies in developing countries, I conduct an experiment that randomizes wages for work on community agricultural development projects. The two previous experiments about labor supply (Dal Bó, Finan, and Rossi 2013; Fehr and Goette 2007) are conducted in different contexts—either job markets for professionals, or developed countries—than the rural labor market I study. Dal Bó, Finan, and Rossi (2013) randomize salaries for professional public sector jobs in Mexico, in order to learn whether higher salaries attract applicants with more desirable characteristics. The experiment I conduct in Malawi is most similar conceptually to the study of Swiss bicycle couriers by Fehr and Goette (2007). Their study, like mine, introduces exogenous and temporary shocks to wages in a market where workers can flexibly adjust their labor supply. However, there are many reasons to expect different results in my study than in theirs. First, the experiments differ in their designs in ways that lead Fehr and Goette to estimate different labor supply adjustments. Second, bicycle couriers in Zurich, Switzerland likely have different opportunity costs of time and labor supply elasticities than peasants in rural Malawi.

More generally, differences in institutions alone will generate different labor supply responses to wage changes in developed and developing countries. In developed countries, there is often little flexibility to adjust labor supply at the intensive

[1] I estimate an extensive margin elasticity, so the change in labor supply is entirely a substitution effect, and therefore a backward bending labor supply curve is not possible in my experiment.

margin—a constraint that leads to larger extensive margin than intensive margin wage elasticities. Even in particular developed-country labor markets where labor supply is unusually flexible, such as stadium vendors (Oettinger 1999) or taxi drivers (Camerer et al. 1997, Chou 2000, and Farber 2003), other institutional factors specific to those contexts will influence the elasticity of labor supply. In developing countries, though, spot markets facilitate adjustment along the days-worked margin. Financial markets in developed countries facilitate smoothing over time, but the absence of such markets in developing countries presents challenges to the intertemporal substitution that is central to a lifecycle model of labor supply. Additionally, if workers in developing countries are more risk averse, then a lifecycle model would predict their labor supply to be less elastic with respect to wages than that of less risk averse workers in developed countries (Chetty 2006). And in general, workers in developing countries may have higher marginal utility of consumption than their counterparts in developed countries, which in a lifecycle model would increase the behavioral response to a change in wages.

For a myriad reasons, then, existing studies of labor supply in developed countries, even well identified experimental or observational studies, are unlikely to be particularly informative about labor supply in developing countries. My experiment begins to fill this gap by providing well-identified estimates of one key margin of labor supply adjustment in a developing country: the probability of accepting employment on a given day, given an unanticipated transitory shock to wages. Along this margin, my results show that labor supply is inelastic. A 10 percent increase in wages leads to only a 1.5-to-1.6 percent increase in the probability of working, with no differences along dimensions of demographic heterogeneity, including gender. Survey data that I collected after the fourth, eighth, and twelfth week of the project suggest that the labor supply decision is driven by the immediate necessity to purchase basic commodities, and by constraints on work imposed by illness or funerals.

The experiment takes place during the agricultural off-season in Malawi, so my results must be interpreted in the context of a very low opportunity cost of time. Perhaps unsurprisingly, then, my results indicate high participation even at low wages, which in turn bounds the elasticity of labor supply from above. The experiment is a starting point for other experiments about labor supply in developing countries and the results can inform the design and targeting of workfare programs even if the point estimates are not generalizable to labor supply elasticities in other contexts.

The paper proceeds as follows. I describe the experiment in Section I and describe the data in Section II. I present the framework for estimates of my main parameter in Section III. I discuss the main results in Section IV. I account for the use of a schedule of wages by including results from permutation tests and specifications including lagged and leading wages in Section V. Section VI concludes.

I. Experimental Design

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 nonfood 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—called "ganyu" in Malawi—can play an important role in bringing in cash and coping with shocks. In the 2004 Integrated Household Survey (IHS), a nationally representative household survey collected by Malawi's National Statistics office as part of the World Bank's Living Standards Measurement Study program, 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 and are extremely low in rural areas; the tenth percentile of the wage distribution in rural areas is 40 Malawian kwacha (MK) (US \$0.29) per day, and the ninetieth percentile is MK 135 (US \$0.96) per day.[2] 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. Studying 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 randomize the wages offered to 529 adults in 10 villages in rural Malawi 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 1 day per week for 12 consecutive weeks. I partnered with a local community-based organization called the Lobi Horticultural Association (LHA) to identify a sample and appropriate work activities. In cooperation with local leaders and government extension workers in Dedza, Malawi, I identified villages that were within 20 kilometers of LHA's headquarters at the Lobi Extension Planning Area (EPA) office, to facilitate supervision by LHA officers and extension workers. To minimize the chance that participants in one village would learn about wages in other villages, only 1 village per group village headman[3] was included in the project.

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.

[2] The exchange rate in 2009 was \$1 USD = 139.9 MK.

[3] Villages 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 4 to 12 or more villages and may coordinate development policies and other activities across villages under his domain.

LHA leaders were instructed to recruit up to 30 households in each village for participation in the project.[4] 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.[5]

A. *Project Timing and Work Schedule*

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-farm or off-farm, though by some measures the market for day labor is not drastically different than at other times of the year.[6] The dry season is the time of year when individuals have the most food and most cash, and the lowest opportunity cost of working off-farm. That opportunity cost, though, is constant throughout the experimental period.

Participants were given the opportunity to work for pay for one day per week for 12 consecutive weeks. Each week, participants could either accept the offered wage and work for the full day, or reject the wage and not work at all. The workday was the same each week for each village. Participants were told at the outset that they were eligible for work through an employment project funded by an outside entity partnering with the local horticultural organization; the outside entity was responsible for the terms of employment. 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 in cash, immediately after they worked. Work activities were carefully monitored by government extension agents to ensure that within each village, the intensity and duration of work were the same from week to week.

The once-per-week design of the project is suitable for studying labor in a static framework. Whereas spillovers in the disutility of working from one period to the next are important in interpreting the results in Fehr and Goette (2007), they are unlikely to play a role in labor supply decisions for participants in my experiment. The spacing combined with the project's timing (when little other paid work was

[4] When more than 30 households were identified, all were invited to participate. The number of participating households per village thus ranges from 25 to 40.

[5] While having multiple participants per household complicates analysis that aggregates individuals' responses to changes in their 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-farm 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.

[6] During the dry season, only 12 percent of adults in rural areas report having done ganyu in the previous week according to the IHS. However, the corresponding figure for the wet season is only 13 percent. Moreover, while dry season wages are lower than wet season wages—MK 71/per day during the dry season compared to MK 84/day during the wet season—the mean wages for both seasons fall well within the range of wages studied in this experiment.

available) also limits the potential that the experiment affects market wages in the project villages, which would have introduced another parameter into the analysis of the labor supply decision. The six-day gap between each work period does provide individuals substantial opportunity to rearrange their other obligations in order to be able to work on this project while continuing to devote time to other productive activities. This ability to reduce the opportunity cost of accepting employment through my project is likely to overstate the *level* of employment at each wage, but does not have clear effects on the predicted 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, in midline and endline surveys, respondents report working for other employers during only 12 percent of the person-weeks covered by the experiment. Second, the effect of wages in my project on the probability of outside employment is very small, though it is statistically significant in some specifications. The pattern of outside employment is nonmonotonic in wages.[7] Third, 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 result in a significantly different point estimate of elasticity of employment. 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 smaller for the more comprehensive definition of employment.

My analysis includes village and week fixed effects, which account for time-invariant village determinants of labor supply and common time trends in labor supply, respectively. The village fixed effects absorb any differences in labor supply due to differences in the type of work activity or day of week (since type and day were constant within the village over the 12 weeks of the project) or village-specific characteristics such as the chief's level of support for the project. The week fixed effects account for common seasonal variation such as depletion of food stores. The fixed effects do not account for time-varying village-specific factors. For example, village and week fixed effects would not be sufficient if heavy rainfall affected some villages in some but not all weeks. Timing the experiment to take place during the dry season was a deliberate effort to minimize the impact of such aggregate shocks; indeed, there was no rainfall during the experiment.

B. *Wage Schedule*

Randomly assigned wages for this project range from MK 30/day (US \$0.21) to MK 140/day (US \$1.00), in increments of MK 10.[8] The wage range spans the

[7] About one-quarter of individuals obtain outside employment in weeks with wages of MK 30, 70, 110, or 140, and between 3 and 5 percent obtain outside employment in weeks with other wage levels.

[8] The 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.

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	1,020
Chimowa	100	60	120	30	110	70	140	80	130	90	50	40	1,020
Manase	60	120	30	110	70	140	80	130	90	50	40	100	1,020
Lasani	120	30	110	70	140	80	130	90	50	40	100	60	1,020
Njonja	30	110	70	140	80	130	90	50	40	100	60	120	1,020
Hashamu	110	70	140	80	130	90	50	40	100	60	120	30	1,020
Kachule	70	140	80	130	90	50	40	100	60	120	30	110	1,020
Msangu/Kalute	140	80	130	90	50	40	100	60	120	30	110	70	1,020
Kamwendo	80	130	90	50	40	100	60	120	30	110	70	140	1,020
Kunfunda	130	90	50	40	100	60	120	30	110	70	140	80	1,020
Average	88	93	88	86	84	87	88	84	81	80	81	80	

Notes: Wages expressed in Malawi kwacha. At the time of the project, \$1 USD = 139.9 MK.

tenth to ninetieth percentile of wages for day labor reported for adults in rural areas in Malawi's 2004 IHS. 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 and ensured the same total earnings potential in all villages.

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 reducing the effective sample size and introducing serial correlation in the wages. To address the issues related to the small number of clusters, I describe a randomization inference procedure that permutes the *schedule* of wages among the ten schedules included in the project in Section VA.

The negative serial correlation appears to have been undetected by participants and does not affect their labor supply. In Section VB, I provide evidence that neither lagged wages nor leading wages have any predictive power for current employment. In addition, 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; 3 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, 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 ten.

II. Data

In total, the project includes 529 individuals[9] in 298 households. I follow these individuals for 12 weeks, recording their participation in each week's work activity.

[9] One individual died after week 6 of the project, so the sample size in weeks 7–12 is 528.

This gives me 6,333 binary observations of individual labor supply. Because wages are assigned at the village-week level, I aggregate individual data to 120 village-weeks in the main analysis. The outcome of interest is the fraction of eligible participants in each village who work for the project in each week.

To supplement 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 twelfth weeks of the project (with each village surveyed six days following its fourth, eighth, and twelfth 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 (with 83 percent reporting both the wage and their own participation correctly).

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 preselected 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.[10] 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 worked an average of one day in the week before the survey or 2.7 days in the month before the survey.

III. 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 refer to as the elasticity of employment. This is a reduced-form estimate of an uncompensated, intertemporal parameter, but is not the structural Frisch elasticity. 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. It is calculated by aggregating the extensive-margin decisions of

[10] Including 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.

Table 2—Baseline Characteristics

	Mean	SD	*N*	10th	Median	90th
Male	0.40	0.49	529			
One male and one female in household	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 household	2.25	0.97	495	1	2	3
Number of children in household	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

Notes: Figures in the top half of the table are from administrative records. Figures in the bottom half of the table come from the baseline survey conducted with individuals at baseline.

individuals within a village, so it reflects the substitution effect but cannot speak to the income effect of changes in wages.

I focus on the daily participation decision of adults who are in the labor force,[11] in response to temporary changes in wages. The corresponding elasticity measures the change in the probability of working on a given day for a change in that day's wage. Oettinger (1999) 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 employment on a given day for registered stadium vendors is between 0.55 and 0.65. Barmby and Dolton (2009) estimate the same wage elasticity for workers on an archeological dig in Syria in the 1930s, and find an elasticity of 0.035. Both Oettinger (1999) and Barmby and Dolton (2009) interpret their estimates as intertemporal elasticities of substitution, where workers experience anticipated, transitory shocks to wages and substitute between labor and leisure accordingly. Standard economic theory predicts larger responses to temporary changes in wages than to permanent changes, so the high-frequency experimental variation would be expected to overstate the magnitude of the labor supply response. Thus, the high frequency changes and short persistence of wages suggest that if anything, the small labor supply elasticity I estimate is an upper bound on what would have been detected with greater persistence in wages.

[11] In 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 first time an individual with no previous work experience works during this project.

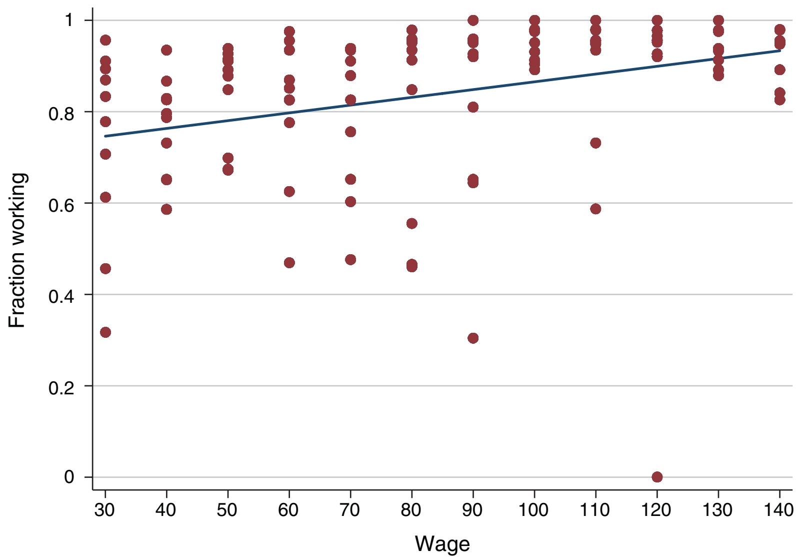

Figure 1. Fraction Working at Each Wage (*Wages in MK*)

Notes: Figure 1 plots 120 observations of employment at the village-week level. Each point represents the fraction of respondents who work in a given week. Wages are expressed in Malawi kwacha. The regression line represented the fitted values from a regression of fraction worked on wage: $labor_{tv} = \alpha + \beta \ln(wage_{tv}) + \nu_{tv}$.

IV. Results

A. *Level of Labor Supply*

This experiment was conducted in the agricultural off-season, and the high levels of employment observed in the experiment are consistent with a low opportunity cost of time. 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, nearly 74 percent of respondents worked. While this high base has a strong seasonal component, employment at low wages is characteristic of the market for ganyu in Malawi around the year. 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.

It is not possible to examine labor supply at the very bottom of the wage distribution using my data, because the IRB committee at the University of Michigan did not allow wages below MK 10/day to be included in the experiment. The experiment was not designed to study the individual determinants of working at a given wage, but rather the change in the probability of working conditional on the wage offered. Still, I examine the correlation between observable characteristics and individual labor supply at very low wages before turning to estimates of the elasticity of employment. I consider nine characteristics measured in the baseline survey: gender (indicator for male), household size, years of education, married (indicator), age (in years), number of rooms in the home, acres of land owned, tin roof on the home (indicator), and lack of any previous paid work experience (indicator). Results are shown in columns 1 and 2 of Table 3.

Table 3—Correlation between Employment and Observable Characteristics

	Individual indicator for working for a wage of							
	MK 30		MK 30 or MK 40		MK 130 or MK 140		Total days worked	
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Male	−0.067***	−0.075***	−0.062*	−0.065*	−0.027	−0.029	−0.464**	−0.524**
	(0.014)	(0.012)	(0.034)	(0.034)	(0.017)	(0.018)	(0.126)	(0.131)
Household size	−0.090	−0.122	0.003	0.012	0.016	0.024	0.122	0.233
	(0.084)	(0.080)	(0.056)	(0.063)	(0.020)	(0.022)	(0.351)	(0.325)
Years of education	−0.010	−0.006	−0.003	0.000	−0.001	−0.001	−0.058**	−0.030
	(0.008)	(0.004)	(0.002)	(0.002)	(0.001)	(0.001)	(0.021)	(0.018)
Married	0.052	0.032	0.018	0.009	−0.001	−0.002	0.143	0.054
	(0.059)	(0.054)	(0.043)	(0.045)	(0.011)	(0.010)	(0.172)	(0.213)
Age	0.001	0.000	0.001	0.001	0.000	0.000	0.006	0.007
	(0.002)	(0.002)	(0.001)	(0.001)	(0.000)	(0.000)	(0.005)	(0.006)
Number of rooms	−0.013	0.002	−0.010	−0.006	0.001	0.002	−0.129	−0.079
	(0.038)	(0.028)	(0.020)	(0.019)	(0.005)	(0.005)	(0.095)	(0.103)
Acres of land	0.067**	0.041**	0.017*	0.011	0.003	0.002	0.179	0.154
	(0.023)	(0.017)	(0.009)	(0.008)	(0.003)	(0.003)	(0.112)	(0.116)
Tin roof	−0.017	−0.069	−0.041	−0.043	−0.018	−0.018	−0.401*	−0.484**
	(0.064)	(0.065)	(0.027)	(0.028)	(0.021)	(0.022)	(0.197)	(0.162)
No previous work experience	−0.070	−0.044	−0.010	−0.013	0.016	0.013	0.296	0.136
	(0.040)	(0.054)	(0.047)	(0.052)	(0.009)	(0.011)	(0.186)	(0.193)
Village effects		x		x		x		x
Observations	488	488	488	488	488	488	488	488
Mean of dependent variable	0.74	0.74	0.93	0.93	0.99	0.99	10.19	10.19
Adjusted R^2	0.01	0.19	0.00	0.02	0.00	0.01	0.04	0.13
p-value: covariates jointly 0	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Notes: This table reports results from OLS estimates. Standard errors are clustered at the village level. Unit of observation is individual, sample is all individuals for whom baseline data are available. The last row presents the p-value from an F-test that the coefficients for the covariates male, household size, years of education, married, age, number of rooms, acres of land, tin roof, and no previous work experience are jointly equal to zero.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Women are 6 to 7 percentage points more likely to work at a wage of MK 30 than men, and each additional acre of land owned by a household increases the probability of working by 4 to 7 percentage points. Other characteristics are not significantly correlated with the probability of working at the lowest wage. The patterns are unchanged when including village fixed effects, and persist when considering employment at either of the lowest two wages as the dependent variable (shown in columns 3 and 4 of Table 3). For comparison, in columns 5 and 6 of Table 3, I examine the correlation between working for either of the highest two wages included in the experiment and the same baseline characteristics. None of the observed characteristics are significantly associated with the probability of working at a high wage. Results predicting total employment, measured as the total number of days worked, are shown in columns 7 and 8.

For all four outcomes, I reject that the correlations between the baseline characteristics and the outcome of interest are jointly zero. However, the explanatory power of the regressions is low. I avoid interpreting the R^2 for the binary dependent variable models in columns 1 to 6, but note that baseline characteristics explain only 4 percent of the variation in the number of days worked in column 7, and even including village fixed effects raises the adjusted R^2 to only 0.13 in column 8. I return to baseline characteristics when exploring heterogeneous labor supply responses in Section IVC, but the results from Table 3 already suggest that heterogeneity is not likely to be important in this sample.

B. *Point Estimate of the Elasticity of Employment*

To account for the village-level randomization, I estimate the elasticity of employment from data aggregated to 120 village-week observations. I run ordinary least squares regressions of the form

$$labor_{tv} = \alpha + \beta \ln(wage_{tv}) + \nu_{tv}. \tag{1}$$

The coefficient β is the marginal effect of a one log-point, or approximately one percent, change in wages on the fraction of individuals in village v working in a given week.

The marginal effect is not an elasticity, but it is easily transformed into one using the standard formula,

$$\epsilon_e = \frac{\partial Q}{\partial P} \times \frac{P}{Q}. \tag{2}$$

Because I am using log-wages as the independent variable, I compute $\epsilon_e = \frac{\beta}{mean(labor)}$.

In Table 4, I begin by regressing the average employment in village v in week t on the log wage without any additional controls. I find that a one percent increase in wages is associated with a 12.4 percentage-point increase in fraction of participants working. This effect is significantly different from zero at the 99 percent confidence level, using p-values from the wild-t bootstrap procedure suggested by Cameron, Gelbach, and Miller (2008) with 1,000 replications. The elasticity corresponding to the marginal effect reported 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 has small effects on the magnitude of the coefficient or associated elasticity. With village and week fixed effects, the coefficient $\beta = 0.135$, and the associated elasticity is 0.16.

The high level of labor supply at the lowest wage included in the experiment is an empirical result in itself, but it does constrain the maximum possible elasticity that could have been detected in the experiment. In order to calculate the maximum possible elasticity conditional on the observed level of employment at the wage of MK 30, I use the observed labor supply pattern at the lowest wage, and construct a counterfactual where all participants work at MK 140. The marginal effect of log wages on labor supply under that counterfactual is $\overline{\beta} = 0.171$, which corresponds to an elasticity of 0.20. The p-value for the hypothesis test that the coefficient from the regression in column 4 of Table 4 is equal to the counterfactual coefficient $\overline{\beta} = 0.171$ is 0.35.

Table 4—Elasticity of Employment with Respect to Wages

	Fraction working in each village-week			
Dependent variable	(1)	(2)	(3)	(4)
ln(*wage*)	0.124**	0.124**	0.135**	0.135**
	(0.038)	(0.039)	(0.034)	(0.035)
Village effects		x		x
Week effects			x	x
p-value from clustered SEs	0.0093	0.0114	0.0030	0.0040
p-value from wild bootstrap	0.0150	0.0160	0.0020	0.0010
p-value from RI	0.0120	0.0120	0.0050	0.0050
Observations	120	120	120	120
Adjusted R^2	0.11	0.13	0.34	0.38
Mean of dependent variable	0.85	0.85	0.85	0.85
Elasticity	0.15	0.15	0.16	0.16

Notes: This table reports results from OLS estimates. Standard errors clustered at the village level reported in parentheses. Unit of observation is village × week, sample is all individuals. p-value from 999 wild bootstrap iterations calculated against a null hypothesis of $\beta = 0$. p-value from RI calculated from $10! - 1$ permutations of the village wage schedule. The RI p-value is the fraction of permutations in which the true coefficient falls within the α tail of the distribution.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

C. *Heterogeneity*

Because individuals with different characteristics may differ in their opportunity cost of working, marginal utility of consumption, or institutional constraints to adjusting their labor supply, there may be heterogeneity in the elasticity of labor supply. The predictions developed in the context of a lifecycle model—specifically, that individuals with higher marginal utility of consumption will be more elastic in their supply of labor, and those who are more risk averse will be less elastic—do not carry over to the static model that is relevant when income is consumed in the same period it is earned. In static models with standard utility functions, increases in nonlabor income lead to larger labor supply elasticities. I examine heterogeneity by characteristics of individuals that may be correlated with higher nonlabor (or, in this case, outside-the-experiment) income: gender, land ownership, household size, asset ownership, and education.

For this analysis, I compute the median level of each characteristic within village, and then aggregate labor supply to the village-week level separately for those above and below the median (or separately for men and women). In Table 5 I report elasticities for each subgroup, and test that the effect of wages on labor supply for the above-median and below-median groups is the same. I cannot reject that wages have equal effects on labor supply for individuals above or below the median of each characteristic, or for men and women. This may be explained by the overall homogeneity of my sample, which by construction includes only poor households in rural areas who are already participating in causal labor markets.

Table 5—Subgroup Elasticity of Employment with Respect to Wages

Dependent variable	Fraction working in each village-week (1)	(2)	(3)	(4)	(5)
		Above median:			
	Female	Land owned	Household size	Assets owned	Education
Panel A					
ln(*wage*)	0.126***	0.131**	0.125**	0.139**	0.126**
	(0.035)	(0.038)	(0.027)	(0.044)	(0.041)
p-value from wild bootstrap	0.002	0.001	0.002	0.005	0.008
Mean of dependent variable	0.86	0.86	0.85	0.85	0.85
		Below median:			
	Male	Land owned	Household size	Assets owned	Education
Panel B					
ln(*wage*)	0.142***	0.131**	0.144**	0.138**	0.135**
	(0.037)	(0.034)	(0.041)	(0.038)	(0.032)
p-value from wild bootstrap	0.001	0.002	0.002	0.000	0.001
Mean of dependent variable	0.81	0.85	0.85	0.85	0.86
Observations	120	120	120	120	120
p-value for equality of coefficients	0.25	0.99	0.54	0.92	0.68

Notes: This table reports results from OLS estimates. Standard errors clustered at the village level are reported in parentheses. Unit of observation is village × week. Sample in panel A is individuals with above-median baseline levels of the indicated characteristic, within their village. The sample in panel B is individuals with below-median baseline levels. All specifications include village and week fixed effects. *p*-value from 999 wild bootstrap iterations are calculated against a null hypothesis of $\beta = 0$.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

D. *Self-Reported Explanations for Labor Supply*

In midline and endline surveys, respondents were asked to list up to three reasons for working in weeks that they worked, or three reasons for not working in weeks they did not work. Reasons for working were grouped into four categories: because of the wage,[12] 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, appears 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.

[12] Used only when the respondent's literal answer was "because of the wage" or "because the wage was good."

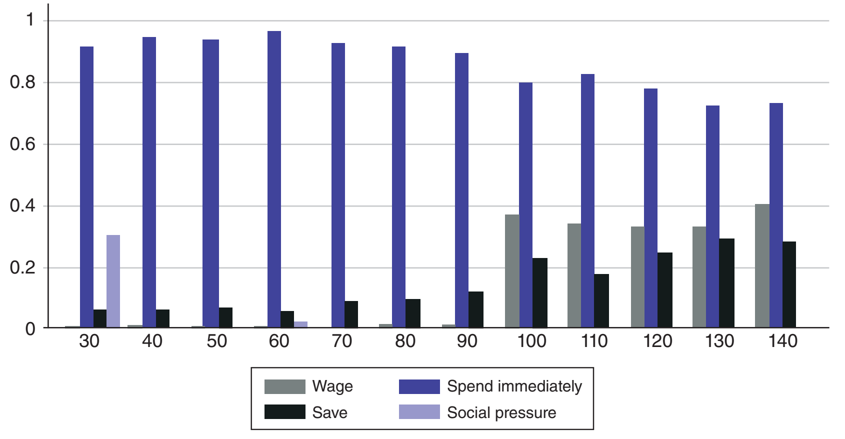

FIGURE 2. SELF-REPORTED REASONS FOR WORKING

Notes: Data are from surveys collected after weeks 4, 8, and 12. Unit of observation is the individual-week, and sample includes responses corresponding to the 4,173 individual-weeks in which respondents correctly recalled the wage and that they had worked. Responses 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. Respondents could list multiple reasons, so answers may not sum to 1.

Reasons for not working were grouped into six categories: because of the wage, because the respondent was occupied with other work, because money was not needed, because of a funeral, because of illness (of 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 as measured 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.

V. Accounting for Randomization of the Wage Schedule

A. *Randomization Inference*

While wages vary at the village-week level, they were randomized at the village level. That is, each village was assigned to one of ten possible schedules of wages, $S_v \in \{S_1, S_{10}\}$. Once a schedule was assigned, week-to-week variation within village was deterministic. I use the randomization inference method introduced by Fisher (1935) and discussed by Rosenbaum (2002) to test the null hypothesis that the true effect of wages on labor supply is zero.[13] Under this maintained null hypothesis, labor supply in village v in week t would have been the same if the wage in the

[13] Recent papers in development economics that make use of this method include Cohen and Dupas (2010) and Iyer (2010).

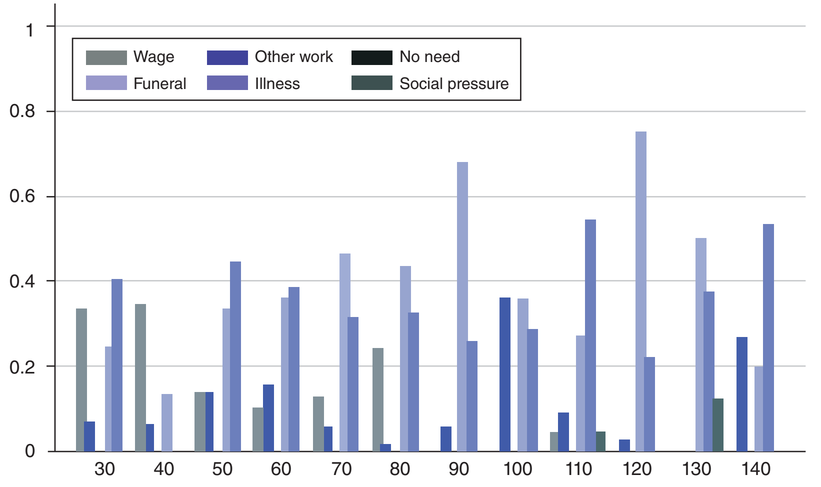

Figure 3. Self-Reported Reasons for Not Working

Notes: Data are from surveys collected after weeks 4, 8, and 12. Unit of observation is the individual-week, and sample includes responses corresponding to the 927 individual-weeks in which respondents correctly recalled the wage and that they had not worked. Responses 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. Respondents could list multiple reasons, so answers may not sum to 1.

village had been some w_{-tv} rather than the actual wage w_{tv}, so the counterfactual outcome is known (and equal to the observed outcome).

I permute the *schedule* of wages by considering the ten factorial possible assignments of ten villages to ten wage schedules. For each permutation, I compute the effect of the counterfactual wage schedule on labor supply. I collect $10! - 1$ coefficients from these permutations, and compare the observed coefficient (and elasticity) to the distribution of coefficients that would have been obtained under every possible counterfactual assignment. The randomization inference p-values represent the fraction of permutations in which the true coefficient falls within the α tail of the distribution of coefficients, under the null hypothesis that the true effect of wages on labor supply is zero.

I report p-values from this exercise for the main results in Table 4. Despite the small number of villages, I robustly reject that the true effect of wages on labor supply is zero: the randomization inference p-values are 0.0120 in both specifications without week fixed effects (columns 1 and 2) and 0.0050 in the specifications with week fixed effects (columns 3 and 4).

B. *Robustness to Lagged and Leading Wages*

If participants detected and reacted to the negative serial correlation in wages, then that feature of the wage schedule would affect both the interpretation of the elasticity 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

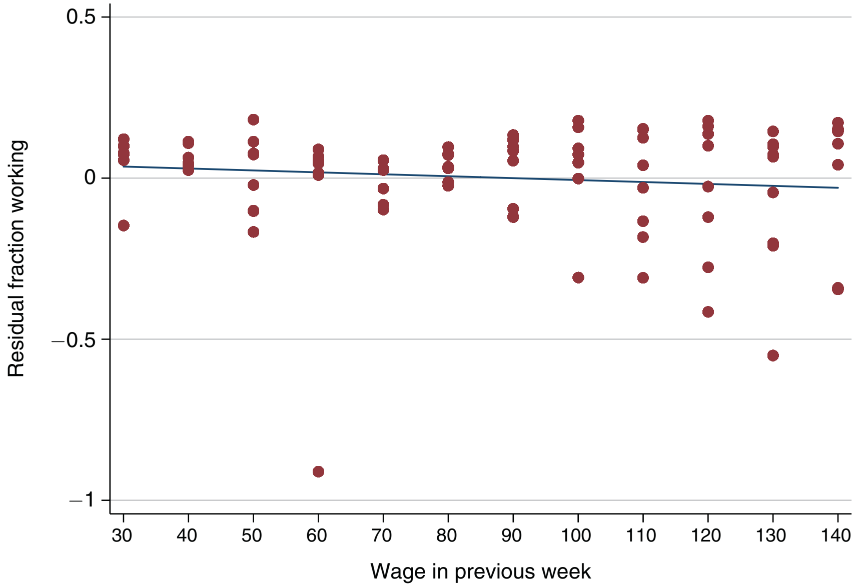

Figure 4. Residualized Fraction Working (*Previous week's wages in MK*)

Notes: Figure 4 plots 120 observations of residual employment at the village-week level. Wages are expressed in Malawi kwacha. I first obtain the residual from the regression of fraction worked on wage: $labor_{tv} = \alpha + \beta \ln(wage_{tv}) + \nu_{tv}$. I then plot this residual against the wage in the previous week, and include a regression line from the fitted values of the regression $\nu_{tv} = \delta + \gamma \ln(wage_{t-1,v}) + \omega_{tv}$.

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 announced change in current period wages.

Graphically, I create plots analogous to Figure 1 by plotting residualized labor supply against lagged and leading wages, respectively. That is, I plot the residual from equation (1) against wages in the previous week in Figure 4 and against wages from the subsequent week in Figure 5. In contrast to the fitted line in Figure 1, the regression lines in both of the new graphs are essentially flat (with slopes of -0.0065 and -0.0004, respectively).

More formally, I check whether participants react to future wages by estimating regression specifications including future and past wages, respectively. For this exercise, I have to limit the number of weeks included in the analysis. The left-hand panel of Table 6 includes weeks 1 to 11. Column 1, included for reference, is the same specification as Table 4 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 village and week fixed effects. In the right-hand panel of Table 6, 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 the coefficients are all close to zero. Additionally, I fail to reject the hypothesis that the coefficients on four weeks

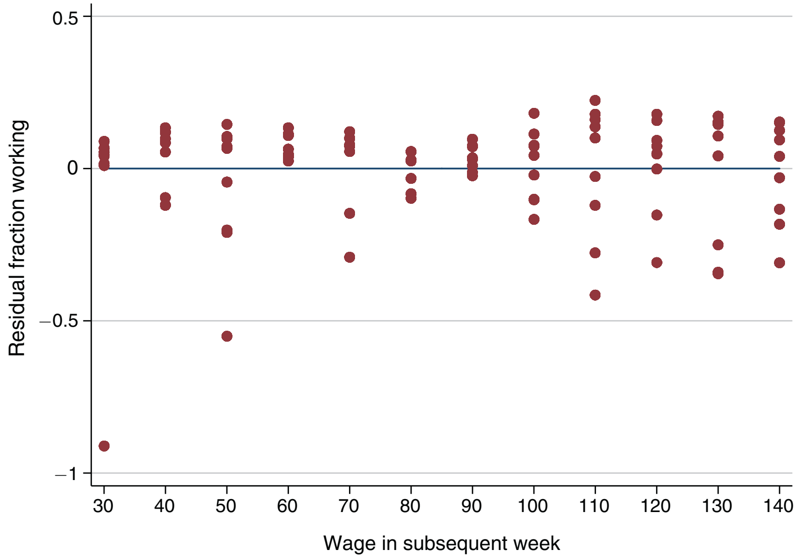

Figure 5. Residualized Fraction Working (*Subsequent week's wages in MK*)

Notes: Figure 5 plots 120 observations of residual employment at the village-week level. Wages are expressed in Malawi kwacha. I first obtain the residual from the regression of fraction worked on wage: $labor_{tv} = \alpha + \beta \ln(wage_{tv}) + \nu_{tv}$. I then plot this residual against the wage in the subsequent week, and include a regression line from the fitted values of the regression $\nu_{tv} = \delta + \gamma \ln(wage_{t+1,v}) + \omega_{tv}$.

of future wages are jointly equal to zero. I interpret the results in 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.

I run a similar robustness check that includes lagged wages in order to exclude the possibility that changing expectations about wages over the course of the project affected labor supply. Table 7 includes wages in past weeks, using specifications analogous to those for future weeks in Table 6. As before, the left-hand panel of Table 7 uses 11 weeks of labor supply choices (covering weeks 2–12 of the project) and incorporates 1 week of lagged wages, and the right-hand panel uses 8 weeks of labor supply data (weeks 5–12) and 4 lags. The magnitude of the effect of current wages on the probability of working is very similar with and without lagged wages, and those lagged wages themselves are not predictive of employment. None of the coefficients on lagged wages are statistically different from zero, and as in the specifications with future wages, I fail to reject the joint hypothesis that the coefficients on four weeks of lagged wages are jointly equal to zero.

VI. Conclusion

My paper builds upon a scant previous literature using experiments to study labor supply, and provides the first experimental evidence about labor supply in the rural labor markets typical of many developing countries. By experimentally varying wages offered for casual day labor to participants in ten villages in Malawi, I am

TABLE 6—ELASTICITY OF EMPLOYMENT WITH RESPECT TO FUTURE WAGES

	Individual × day indicator for working					
	Weeks 1 to 11			Weeks 1 to 8		
Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
ln(*wage*)	0.127**	0.123**	0.136**	0.146**	0.115**	0.122*
	(0.042)	(0.039)	(0.035)	(0.054)	(0.051)	(0.055)
	[0.017]	[0.011]	[0.000]	[0.032]	[0.053]	[0.032]
$\ln(wage_{t+1})$		−0.015	−0.008		0.031	0.033
		(0.046)	(0.040)		(0.084)	(0.076)
		[0.802]	[0.929]		[0.778]	[0.682]
$\ln(wage_{t+2})$					0.018	0.016
					(0.046)	(0.025)
					[0.818]	[0.705]
$\ln(wage_{t+3})$					−0.058	−0.034
					(0.037)	(0.045)
					[0.855]	[0.914]
$\ln(wage_{t+4})$					0.020	0.041
					(0.042)	(0.038)
					[0.667]	[0.483]
Village effects			x			x
Week effects			x			x
Observations	110	110	110	80	80	80
Adjusted R^2	0.11	0.10	0.36	0.11	0.08	0.28
Mean of dependent variable	0.84	0.84	0.84	0.82	0.82	0.82
Elasticity	0.15	0.15	0.16	0.18	0.14	0.15
p-value: leads jointly 0					0.96	0.93

Notes: This table reports results from OLS estimates. Clustered standard errors (clustered at the village level) are in parentheses. *p*-value from 999 wild bootstrap iterations are calculated against a null hypothesis of $\beta = 0$ in brackets. The last row reports the wild bootstrap *p*-value from a joint test that the coefficients on all four leads of wages are 0. Unit of observation is village × week, sample is all individuals.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

able to obtain a causal estimate of the effect of a change in wages on the probability of working. This elasticity of employment is between 0.15 and 0.16 in my sample.

I find highly inelastic labor in a context where overall participation is very high. Because my experiment takes place during the agricultural off season, take-up of the work opportunity provided through my project exceeds 70 percent at even the lowest wage offered. While women are somewhat more likely to participate than men, other observable characteristics do not predict the probability of working.

Not only are differences in the level of labor supply mostly unexplained by baseline characteristics, but also, there is little heterogeneity in the elasticity of employment. I estimate the effect of wages on the probability of working separately for men and women, and for those above and below the median for characteristics that may proxy for the marginal utility of consumption or opportunity cost of time: land ownership, household size, asset ownership, and education. For each characteristic, there are neither economically nor statistically significant differences in the effect of wages on the probability of working for each subgroup. While most of these dimensions of heterogeneity have not been investigated in previous work about

Table 7—Elasticity of Employment with Respect to Past Wages

	Individual × day indicator for working					
	Weeks 2 to 12			Weeks 5 to 12		
Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
ln(*wage*)	0.137**	0.120**	0.137**	0.160***	0.169***	0.169***
	(0.040)	(0.039)	(0.035)	(0.030)	(0.034)	(0.033)
	[0.010]	[0.017]	[0.000]	[0.003]	[0.005]	[0.000]
ln($wage_{t-1}$)		−0.052*	−0.027		0.012	0.016
		(0.027)	(0.031)		(0.010)	(0.016)
		[0.885]	[0.943]		[0.310]	[0.427]
ln($wage_{t-2}$)					−0.001	−0.000
					(0.019)	(0.018)
					[0.962]	[0.990]
ln($wage_{t-3}$)					0.010	0.009
					(0.019)	(0.009)
					[0.944]	[0.917]
ln($wage_{t-4}$)					0.001	0.000
					(0.017)	(0.016)
					[0.940]	[0.992]
Village effects			x			x
Week effects			x			x
Observations	110	110	110	80	80	80
Adjusted R^2	0.13	0.14	0.43	0.48	0.46	0.63
Mean of dependent variable	0.85	0.85	0.85	0.89	0.89	0.89
Elasticity	0.16	0.14	0.16	0.18	0.19	0.19
p-value: lags jointly 0					0.92	0.69

Notes: This table reports results from OLS estimates. Clustered standard errors (clustered at the village level) are in parentheses. *p*-value from 999 wild bootstrap iterations are calculated against a null hypothesis of $\beta = 0$ in brackets. The last row reports the wild bootstrap *p*-value from a joint test that the coefficients on all four lags of wages are 0. Unit of observation is village × week, sample is all individuals.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

labor supply in developing countries, the similarity between men's and women's elasticity of employment 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 and socioeconomic 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 4, 8, and 12, I collect survey data about recollection of wages and work history, as well as reasons for working or not 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 programs in Malawi and other developing countries. In theory, the low wages offered by these programs lead to self-targeting by beneficiaries and reduce the need for complicated screening procedures (Besley and Coate 1992). Inelastic labor force participation in my experiment casts doubt on whether this sort of self-selection occurs in Malawi, at least when programs are offered to apparently poor households. It may be that the vast majority of households that participate in the market for ganyu are too poor to select out of employment opportunities even at low wages, or that credit constraints generate inefficient labor supply patterns.

REFERENCES

Abdulai, Awudu, and Christopher Delgado. 1999. "Determinants of nonfarm earnings of farm-based husbands and wives in northern Ghana." *American Journal of Agricultural Economics* 81 (1): 117–30.

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi. 2013. "Ordeal Mechanisms in Targeting: Theory and Evidence from a Field Experiment in Indonesia." National Bureau of Economic Research (NBER) Working Paper 19127.

Bardhan, Pranab K. 1979. "Labor Supply Functions in a Poor Agrarian Economy." *American Economic Review* 69 (1): 73–83.

Barmby, Tim, and Peter Dolton. 2009. "What lies beneath? Effort and incentives on archaelogical digs in the 1930's." Unpublished.

Besley, Timothy, and Stephen Coate. 1992. "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs." *American Economic Review* 82 (1): 249–61.

Camerer, Colin, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *Quarterly Journal of Economics* 112 (2): 407–41.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–27.

Chetty, Raj. 2006. "A New Method of Estimating Risk Aversion." *American Economic Review* 96 (5): 1821–34.

Chou, Y. K. 2000. "Testing Alternative Models of Labor Supply: Evidence from Taxi-Drivers in Singapore." University of Melbourne Department of Economics Working Paper 768.

Cohen, Jessica, and Pascaline Dupas. 2010. "Free Distribution or Cost Sharing? Evidence from a Randomized Malaria Prevention Experiment." *Quarterly Journal of Economics* 125 (1): 1–45.

Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi. 2013. "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service." *Quarterly Journal of Economics* 128 (3): 1169–1218.

Farber, Henry S. 2003. "Is Tomorrow Another Day? The Labor Supply of New York Cab Drivers." National Bureau of Economic Research (NBER) Working Paper 9706.

Fehr, Ernest, and Lorenz Goette. 2007. "Do workers work more if wages are high? Evidence from a randomized field experiment." *American Economic Review* 97 (1): 298–317.

Fisher, Ronald. 1935. *The Design of Experiments.* London: Oliver and Boyd.

Goldberg, Jessica. 2016. "Kwacha Gonna Do? Experimental Evidence about Labor Supply in Rural Malawi: Dataset." *American Economic Journal: Applied Economics*. http://dx.doi.org/10.1257/app.20130369.

Iyer, Lakshmi. 2010. "Direct versus Indirect Colonial Rule in India: Long-Term Consequences." *Review of Economics and Statistics* 92 (4): 693–713.

Lewis, W. Arthur. 1954. "Economic development with unlimited supplies of labour." *Manchester School of Economic and Social Studies* 22 (2): 139–91.

McCord, Anna, and Rachel Slater. 2009. *Overview of Public Works Programmes in Sub-Saharan Africa*. Overseas Development Institute. London, September.

Oettinger, Gerald S. 1999. "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors." *Journal of Political Economy* 107 (2): 360–92.

Rosenbaum, Paul R. 2002. *Observational Studies*. 2nd ed. New York: Springer-Verlag.

Rosenzweig, Mark R. 1978. "Rural Wages, Labor Supply, and Land Reform: A Theoretical and Empirical Analysis." *American Economic Review* 68 (5): 847–61.

World Bank. 2013. *World Development Report 2013: Jobs*. Washington, DC: World Bank.