Experimental Evidence about Earning, Saving, and Borrowing Money in Rural Malawi

by

Jessica A. Goldberg

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Public Policy and Economics)
in The University of Michigan
2011

Doctoral Committee:

Professor Jeffrey A. Smith, Co-Chair
Associate Professor Dean C. Yang, Co-Chair
Professor Brian A. Jacob
Professor David A. Lam

ACKNOWLEDGEMENTS

I am deeply grateful to the members of my dissertation committee – Jeff Smith, Dean Yang, Brian Jacob, and David Lam – for their teaching and mentorship throughout my time in graduate school. Xavier Gine has been a superb co-author and mentor. Erik Johnson and Elias Walsh are always the first people I turn to with questions or ideas, and they have shared every step of the journey through graduate school with me. Susie Godlonton's friendship and perspective on research have been invaluable.

Additionally, Rebecca Thornton, John Bound, Rob Garlick, Emily Beam, Caroline Theoharides, Todd Pugatch, Brian Kovak, and seminar participants at the University of Michigan have provided valuable suggestions about my research. Sara LaLumia and Molly Lipscomb have offered wonderful advice and encouragement, and Liz Foster has helped me at every stage. My parents, Peter and Betsy Goldberg, have been tremendously supportive and I am glad they were able to visit me in Malawi and see the country I've come to love; my sister Shelly went above and beyond for me this year especially.

My work in Malawi would not have been possible without the assistance of a dedicated team. Thanks especially to Niall Keleher, Santhosh Srinivasan, Lonnie and Lutamayo Mwamlima, Kingsley Nalivata, T/A Kachere, Timothy Chimlomo, Grieven Nchocholo, and the staff and members of the Lobi Horticultural Association.

My research has been supported with grants from the Rackham Graduate School, the Center for International and Comparative Studies, the Center for Afro-American and African Studies, the World Bank, and USAID.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figure

LIST OF APPENDICES

Appendix

ABSTRACT

Experimental Evidence about Earning, Saving, and Borrowing Money in Rural Malawi

by

Jessica A. Goldberg

Co-Chairs: Jeffrey Smith and Dean Yang

In the first chapter of this dissertation, I estimate the wage elasticity of working in the day labor market in rural Malawi using panel data from a unique field experiment. Though employment in daily wage markets is important to individuals and governments in developing countries, there is little evidence about labor supply in those markets. My estimates are from a field experiment in which 529 adults from ten different villages are offered a day's work once per week for 12 consecutive weeks, with wages ranging from MK 30 ($US 0.21) to MK 140 ($US 1.00) per day. I find that the elasticity of employment is between 0.15 and 0.17, with no significant differences between men and women.

The second chapter asks whether informal insurance networks or social norms about sharing income can explain the high marginal propensities to consume that are common in developing countries. I employ a field experiment to distinguish between the use of windfall money when receipt of the money is known to others in the community versus when it is private information. I find that immediate spending is 35 percent higher for individuals who receive money in a public setting compared to

those who received money privately. This spending pattern is consistent with a seven percent tax on surplus income in a simple model where a fraction of money that is not spent immediately must be shared with others in the social network.

The third chapter, written jointly with Xavier Gine and Dean Yang, describes the results of a randomized field experiment in Malawi examining borrower responses to being fingerprinted when applying for loans. This intervention improved the lenders ability to implement dynamic repayment incentives, allowing it to withhold future loans from past defaulters while rewarding good borrowers with better loan terms. As predicted by a simple model, fingerprinting led to substantially higher repayment rates for borrowers with the highest ex ante default risk, but had no effect for the rest of borrowers. We provide unique evidence that this improvement in repayment rates is accompanied by behaviors consistent with less adverse selection and lower moral hazard.

CHAPTER I

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

1.1 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 pubic sector works program (*McCord and Slater* (2009)).

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* (1993) or *Rosenzweig* (1978)).

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* (1954) 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* (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) estimate somewhat greater elasticities for husbands and wives in Ghana; and *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. *Kochar* (1999) and *Rose* (2001) 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* (2007), 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* (1999) 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* (2009) 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

[1] Negative 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* (1978).

and other 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.* (1997), *Chou* (2000), and *Farber* (2003) 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.* (2010) 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 1.2 and describe the data in Section 1.3. I present the framework and results for estimates of my main parameter, the elasticity of employment, in Section 1.4. I highlight the methodological differences between my research and the previous literature about the elasticity of labor supply in developing countries in Section 1.5, 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 1.6. In Section 1.7, I test the hypothesis that obstacles to saving money lead to inflexible labor supply. I conclude in Section 1.8.

1.2 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 headman[2] 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

[2]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 four to 12 or more villages and may coordinate development policies and other activities across villages under his domain.

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 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 1.1 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]

[3]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.

[4]The 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

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 1.4.4, 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]

1.3 Data

In total, the project includes 529 individuals[6] 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. Additionally, I

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.

[5]Additional 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, perhaps the government rate for day labor (MK 110) or the local market rate.

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

have records of major community events that may affect participation, especially funerals held in the village.

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.

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 1.2 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.

1.4 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 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

[7] 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.

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.

1.4.1 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* (1993) 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 as intertemporal elasticities of working for individuals in a daily labor market.

Heckman *Heckman* (1993) 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:

$$E(H|W,Y,\nu) \tag{1.1}$$

$$E(H|W,Y,H>0) \tag{1.2}$$

$$E(H|W,Y) = E(H|W,Y,H>0) \times Pr(H>0|W,Y) \tag{1.3}$$

$$Pr(H>0|W,Y) \tag{1.4}$$

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 (1.1) 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 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 (1.4) 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* (1993)) 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* (2002)). 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 (1.2) and (1.4) above:

$$Pr(H > 0 | W, Y, \text{in daily labor market}) \tag{1.5}$$

This labor supply function combines elements of the intensive margin elasticity of hours worked for participants in Heckman's equation (1.2) by conditioning on participation, and of the extensive margin probability of participating in Heckman's equation (1.4) 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 $\frac{\partial Pr(H>0|W,Y,\text{in daily labor market})}{\partial W} \times \frac{W}{H}$.[8] Oettinger *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 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 (1.5) 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

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

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 (1.5). 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.

1.4.2 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 1.1. 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.[9] Despite this discontinuity, I focus on the elasticity of employment across the 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.[10] If the correct model is one that allows for a discontinuity at MK 100, then my estimates overstate the elasticity of employment

[9] Recall that the government's set rate for day labor is MK 110, so this discontinuity does not suggest a reference point corresponding to the government's wage rate.

[10] Using 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.

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 $labor_{itv} = \alpha + \beta ln(wage_{tv}) + \nu_i tv$. 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 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 $\epsilon_e = \frac{\beta}{mean(labor)}$. This elasticity corresponds to the extensive margin elasticity from labor supply equation 1.5 above.

In Table 1.3, 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.

1.4.3 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 1.4 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-Huber-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.[11]

[11] *Angrist and Pischke* (2009) 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* (2009) 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 $\epsilon_e = \frac{\beta}{mean(labor)}$ 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. (2008) 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^0_h.

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 $labor_{itv} = \alpha + \beta ln(wage_{tv}) + \nu$ 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 $\nu_{\mathbf{rh}}$ from imposing each of 101 values of the null hypothesis H^0_h from 0 to 1.

For each vector of residuals $\hat{\nu_{\mathbf{rh}}}$, 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 $lab\hat{o}r_{itv}$ by adding the residual to the observed outcome for each observation. I then estimate $lab\hat{o}r_{itv} = \tilde{\alpha} + \tilde{\beta} ln(wage_{tv}) + u$ 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 1.2 and 1.3 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.

1.4.4 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 A.1 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 1.3 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 A.1, 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 A.1. 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 A.2 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 A.3, 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 A.2 and A.3 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 A.4 correspond to the same-numbered columns in the main results table, Table 1.3. 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 1.4 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 1.5 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.

1.5 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*Abdulai and Delgado* (1999) and West Bengal*Bardhan* (1979), 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* (1979) 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

[12] *Rosenzweig* (1978) found a negative long-run 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.

differences between my preferred estimates and the Bardhan-style 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. I can then assess whether lack of external validity explains why the elasticities I present in section 1.4 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 1.5. 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 (1.1) 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]

[13] I also estimate the elasticity of labor supply by drawing six *consecutive* weeks of data for each

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* (1979) (0.20 to 0.29) and *Abdulai and Delgado* (1999) (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.

1.6 Gender

A long literature suggests that women supply labor more elastically than men in developed countries (e.g. *Killingsworth* (1983), *Heckman* (1993)). Previous work in developing countries is also consistent with women supplying labor more elastically than men in India *Rosenzweig* (1978) and Ghana*Abdulai and Delgado* (1999). In Ta-

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.

bles 1.6 and 1.7, 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 1.6 as in Table 1.3. 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 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 sep-

[14]Malawi 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.

[15]This 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.

arately 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 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 1.8 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.

1.7 Savings

The estimates shown in Table 1.3 indicate highly inelastic decisions about working for participants in a daily labor market. Inelastic labor supply could result from obstacles to saving money, making labor supply more responsive to marginal utility of consumption than to wages. If so, the ability to save should result in more elastic supply of labor. To test this hypothesis, I implement a cross-cutting randomization of savings accounts. I do not find any evidence that access to savings accounts affects the probability of working or the elasticity with respect to wages.

Since individuals within a household are likely to share resources, I randomized savings accounts at the household level. Within each village, respondents in half of the households were offered the chance to open a savings account with the LHA Savings and Credit Cooperative (SACCO), an affiliate of the regional government's SACCO. Randomization was conducted in the field, with one representative per household drawing a bottle cap from an envelope. No participants had accounts before this project, apparently because of lack of information about account availability and account opening procedures. All "winners" chose to open accounts. To be eligible for an account, individuals must be members of LHA. I paid the MK 150 ($US 1.10) membership dues for all participants in this project, including those who were not assigned to receive a savings account.[16] I collect information about deposits into SACCO accounts in the seven days following the respondent's assigned work day for all 12 weeks. The data about deposits come from the LHA SACCO files. I collect data about deposits only during the 12 weeks of the project, and therefore I am limited to examining very short-run effects of savings accounts.

I run regressions of the form $labor_{itv} = \alpha + \beta_1 ln(wage_{tv}) + \beta_2 account_i + \beta_3 (account_i \times ln(wage_{tv})) + \nu$, where "account" is an indicator variable that equals one for individuals in households that were randomly assigned to receive an account with the LHA SACCO, and equals zero otherwise. If access to savings accounts leads to more elastic labor force participation, then β_3 will be positive.

In columns (1) and (2) of Table 1.9, I use individual labor supply as the outcome of interest. The binary "account" equals one for all members of households who were assigned to receive savings accounts. All of the specifications in Table 1.9 use village and week fixed effects, as in columns (4) of the previous tables. Using individual fixed effects precludes separately identifying the effect of savings accounts on the outcomes

[16] I provided account-opening assistance, including application forms, to individuals assigned to receive accounts only. The local SACCO office did not have extra forms and did not open any accounts for members of the control group.

of interest. Neither the main effect of having an account, in column (1), nor the interaction between savings accounts and wages, in column (2), is statistically different from zero, and the point estimates themselves are very close to zero. Note that for those assigned to receive savings accounts, the elasticity of employment includes the main and interaction effects, so $\epsilon_e = \frac{\beta_1+\beta_3}{mean(labor)}$. The estimated elasticity when accounting for savings accounts is 0.17, similar to results in Table 1.3. In columns (3) and (4), the dependent variable is an indicator for whether both members of two-person households worked. Again, neither the main nor the interaction effects of savings accounts are statistically significant.

If being assigned to receive a savings account does not actually increase savings but only shifts savings from home to the SACCO, then the results in Table 1.9 would not be surprising. I do not have sufficient data to compare the total savings levels of those who received accounts and those who did not, but I can examine non-experimental data about accumulated savings to test whether the level of savings with the SACCO affects employment among the sample assigned to receive an account. In Table 1.10, I examine the relationship between savings accumulated in LHA accounts up to but not including week t on the effect of labor supply in week t. The sample is limited to those assigned to receive savings accounts. Accumulated savings in LHA accounts are zero by definition in all weeks for all respondents who did not receive accounts, so including the full sample reduces variation in the independent variable. Focusing on the subsample of those who received accounts should yield a more precise estimate of accumulated savings among the relevant population. Ninety-four of the 147 households assigned to receive savings accounts had made at least one deposit by week 12 (and 87 of those households had made at least one deposit by week six). Households that made at least one deposit had deposited an average of MK 752 as of week 12.

I estimate $labor_{itv} = \alpha + \beta_1 ln(wage_{tv}) + \beta_2 savings_{it-1} + \beta_3 (savings_{it-1} \times ln(wage_{tv})) + \nu$. Note that "$savings_{it-1}$" is predetermined as of week t, but not randomly assigned. As in Table 1.9, columns (1) and (2) refer to individuals' supply of labor, and columns (3) and (4) to an indicator for both members of two-person households working. The elasticity for savers in Column (2) includes the main and interaction effects of wages and $\epsilon_e = \frac{\beta_1+\beta_3}{mean(labor)}$, and the coefficients in the household regressions are not transformed into standard elasticities. I use individual or household fixed effects to control for omitted characteristics of respondents that might simultaneously determine past savings and present labor supply.

Neither the main effect of accumulated savings or the interaction between accumulated savings and wages are statistically different from zero in any of the specifications. The estimated elasticity of individual labor supply controlling for accumulated savings is 0.15, consistent with the results for individuals in Table 1.3. When including the interaction between accumulated savings and wages, the estimated elasticity of individual labor supply is 0.24. While this is substantively higher than my earlier estimates for individuals, the difference between this and earlier estimates is not statistically significant. Together, the results in Tables 1.9 and 1.10 do not support the hypothesis that lack of access to savings technology constrains the elasticity of the supply of casual wage labor, though they do not rule out that in a different context, such as a longer time horizon or a setting where the risks of saving cash at home were higher, access to savings accounts might affect labor supply.

1.8 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 530 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.

One potential explanation for highly inelastic supply of labor is that people face obstacles to saving their wages. If income cannot be transferred from one period to another, it is rational to supply labor in response to marginal utility of consumption rather than wages. I test this hypothesis by randomly assigning half of participating households to receive savings accounts with a local savings and credit cooperative. The effect of access to savings on the supply of labor and on the elasticity of the supply of labor is a precisely estimated zero. This suggests that inability to save wages does not cause the highly inelastic employment patterns observed in this sample.

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 1.1: 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 1.2: 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 1.3: Elasticity of employment w.r.t. wages

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Individual*day indicator for working					
Ln(wage)	0.127***	0.127***	0.140***	0.140***	0.127***	0.140***
	(0.033)	(0.033)	(0.032)	(0.032)	(0.033)	(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	0.15	0.17	0.17	0.15	0.17
	(0.040)	(0.040)	(0.040)	(0.040)	(0.040)	(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 1.4: Different methods for computing standard errors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable:		Individual*day indicator for working		Village*day average employment			Individual*day indicator for working	
Ln(wage)	0.127***	0.127***	0.127***	0.128***	0.126***	0.127**	0.127***	0.127
	(0.010)	(0.010)	(0.029)	(0.031)	(0.030)	(0.035)	(0.033)	
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
							(0.040)	
							3.75	[0.826, 3.101]
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	collapsed and robust	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

Table 1.5: Elasticity of labor supply from time-aggregated cross sectional data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:	Number of days worked out of 12 weeks					Number of days worked out of 6 weeks	
Ln(average accepted wage)	8.641***	8.527***	8.795***				
	(0.436)	(0.478)	(0.574)				
Ln(village average accepted wage, all weeks)				10.586***	9.059***		
				(0.928)	(1.167)		
Ln(village average accepted wage, half of weeks)						1.633*	1.527**
						(0.841)	(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 1.6: Elasticity of men's employment w.r.t. wages

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Individual*day indicator for working					
Ln(wage)	0.139***	0.139***	0.157***	0.157***	0.139***	0.157***
	(0.032)	(0.032)	(0.035)	(0.035)	(0.032)	(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	0.17	0.19	0.19	0.17	0.19
	(0.043)	(0.043)	(0.047)	(0.047)	(0.043)	(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 1.7: Elasticity of women's employment w.r.t. wages

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Individual*day indicator for working					
Ln(wage)	0.119**	0.119**	0.129***	0.129***	0.119**	0.129***
	(0.035)	(0.035)	(0.032)	(0.032)	(0.035)	(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	0.14	0.15	0.15	0.14	0.15
	(0.041)	(0.041)	(0.038)	(0.039)	(0.041)	(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 1.8: Dry and wet season elasticities from IHS data

Panel A. Dry Season

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	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

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	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 1.9: 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***	0.140***	0.287***	0.269***
	(0.033)	(0.027)	(0.044)	(0.051)
Account	-0.017	-0.026	-0.041	-0.195
	(0.013)	(0.097)	(0.030)	(0.191)
Account*Ln(wage)		0.002		0.035
		(0.020)		(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		0.17	
	(0.041)		(0.042)	
Elasticity (no account)		0.17		0.16
		(0.034)		(0.035)
Elasticity (account)		0.17		0.18
		(0.049)		(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 1.10: 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***	0.203	0.248***	0.420
	(0.030)	(0.126)	(0.056)	(0.263)
Ln(savings)	0.011	0.075	0.022	0.157
	(0.011)	(0.098)	(0.026)	(0.213)
Ln(savings)*Ln(wage)		-0.015		-0.031
		(0.021)		(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		0.14	
	(0.034)		(0.033)	
Elasticity (no savings)		0.23		0.24
		(0.141)		(0.149)
Elasticity (with savings)		0.21		0.22
		(0.118)		(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 1.1: Fraction working at each wage (wages in MK)

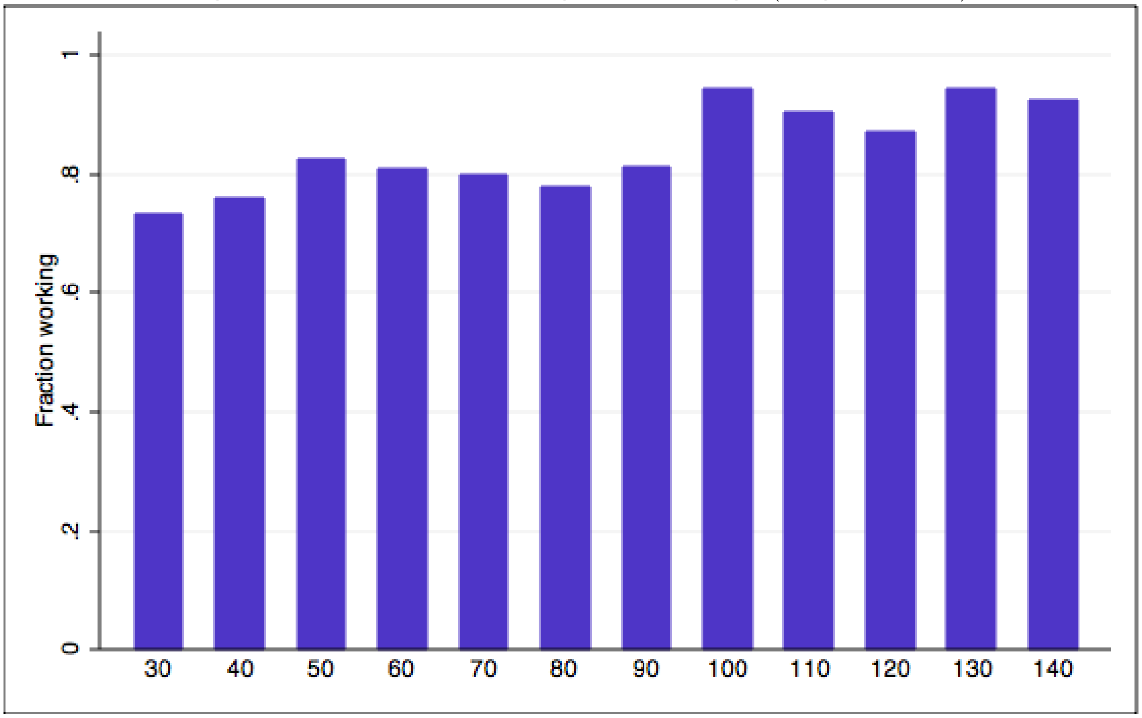

Figure 1.2: Rejection rate for null hypotheses about β from bootstrap-t procedure

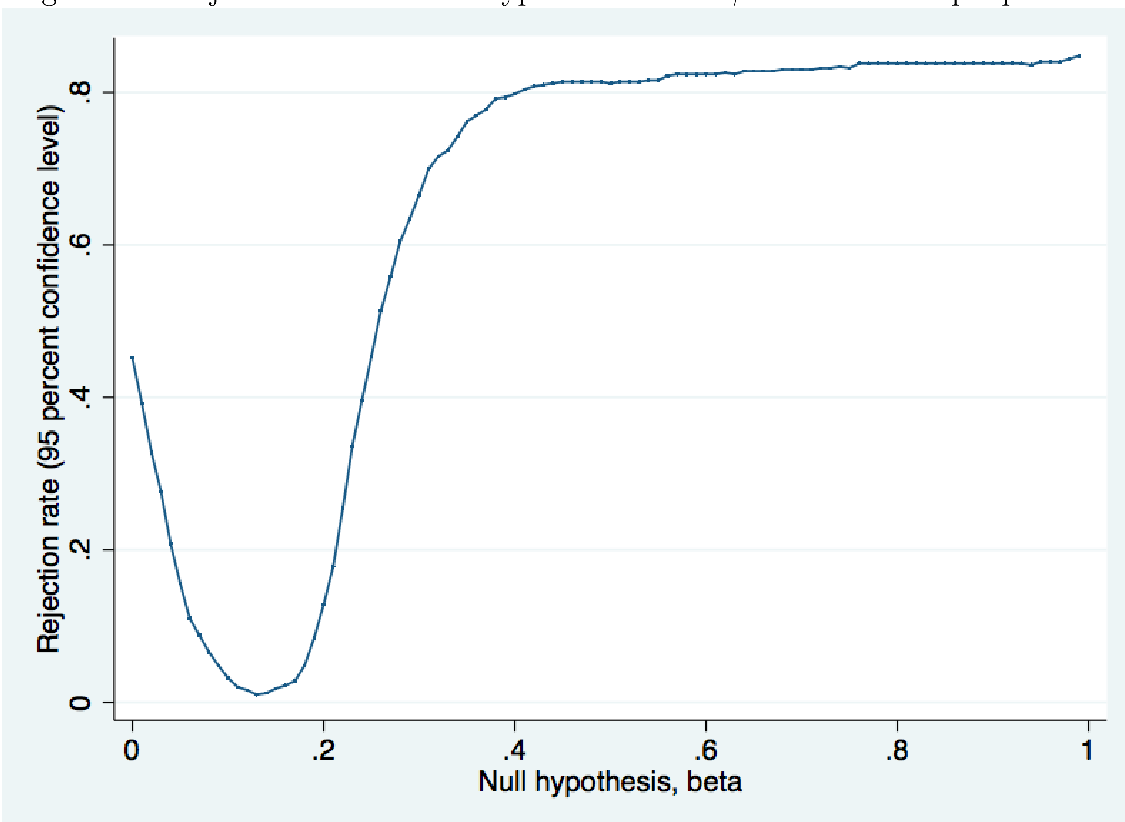

Figure 1.3: Rejection rate for null hypotheses about ϵ_e from bootstrap-t procedure

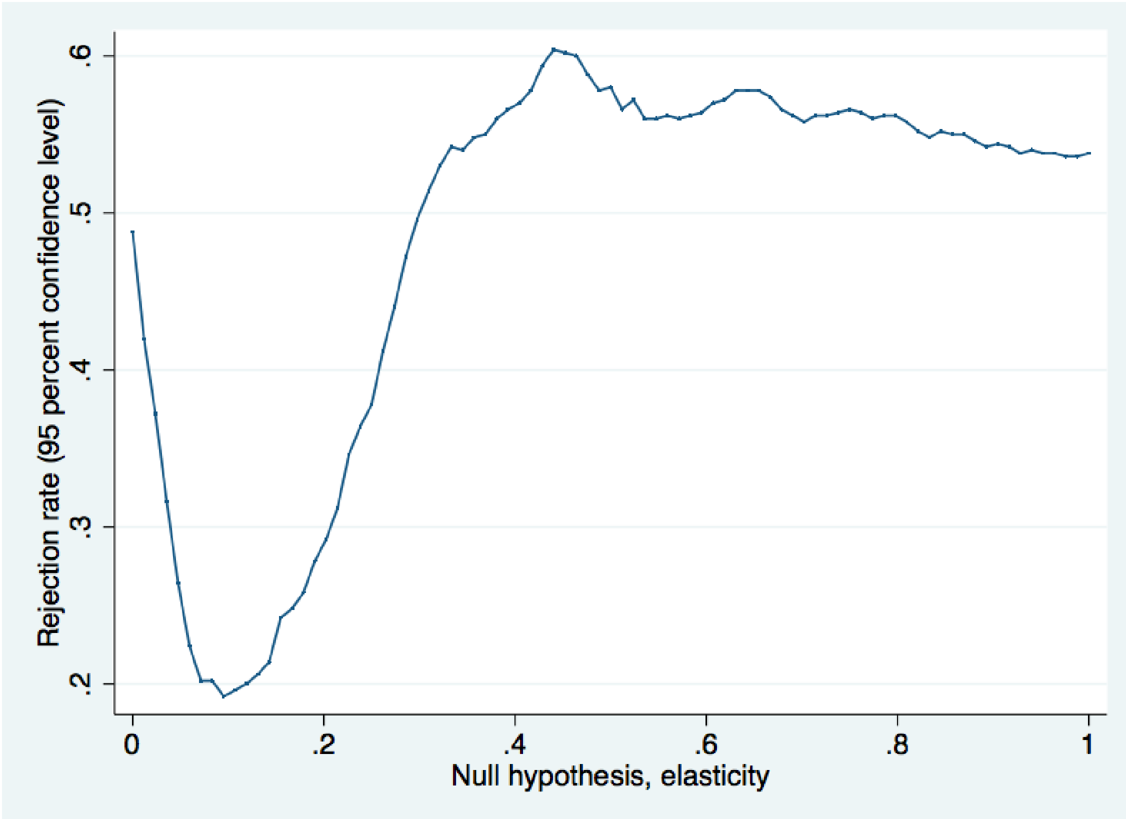

Figure 1.4: Self-reported reasons for working

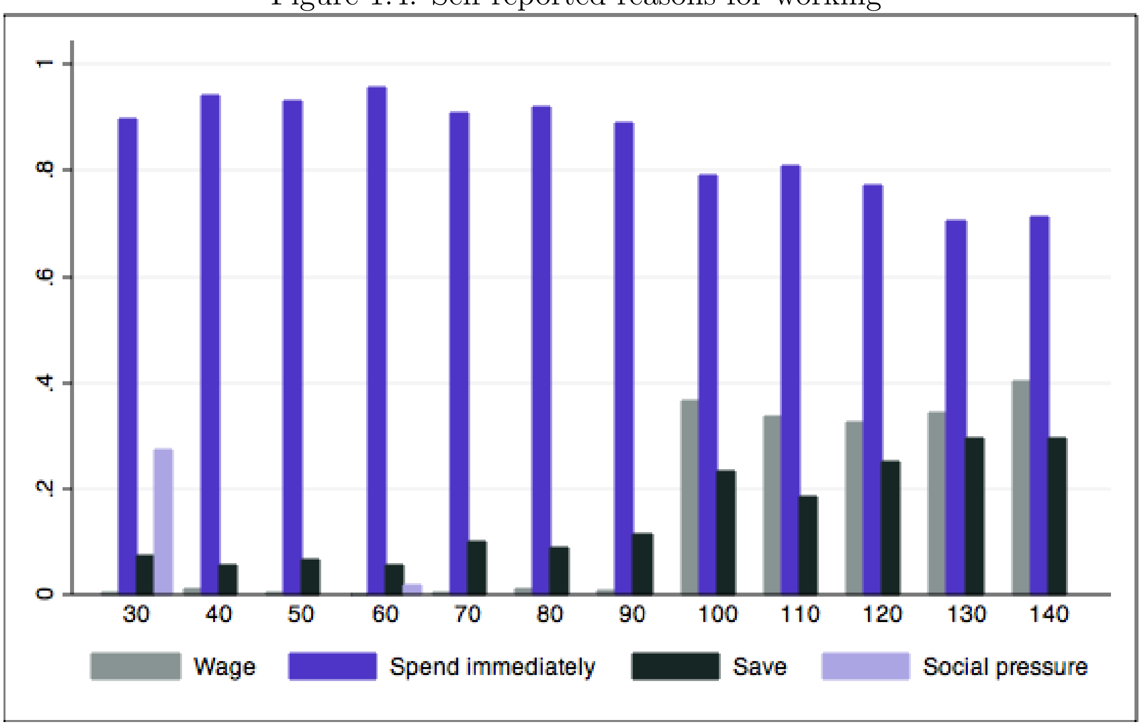

Figure 1.5: Self-reported reasons for not working

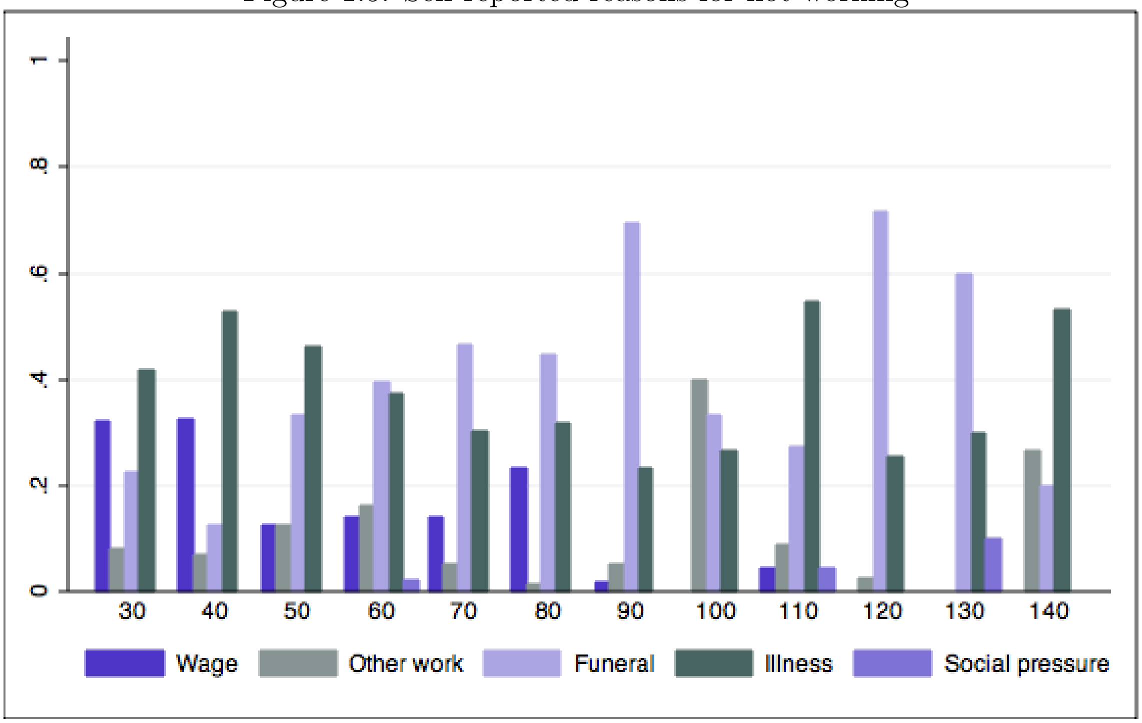

CHAPTER II

The Lesser of Two Evils: The Roles of Social Pressure and Impatience in Consumption Decisions

2.1 Introduction

In rural economies, income is concentrated at harvest time rather than distributed throughout the year. Expenditures on some household necessities obviously take place throughout the year, but major purchases of agricultural inputs are typically concentrated before planting season. Despite the lag between receipt of income and demand for major purchases, poor farmers often appear to have high marginal propensities to consume, spending most of their income soon after receiving it and relying on loans to purchase inputs for the subsequent season. In Malawi, only 18 percent of farmers who received loans to grow the cash crop paprika reported saving money to use for the next season's inputs, though larger fractions save for precautionary reasons or to smooth consumption (BFIRM baseline survey, October 2007). This behavior is often attributed to high discount rates or hyperbolic discounting, but may in fact be explained by constraints on consumption that are not included in standard models. I show that a high marginal propensity to consume is also consistent with a model in which a tax is imposed on income that is observed by others in the community but not immediately converted into illiquid or non-taxable assets.

High or hyperbolic discount rates predict a high marginal propensity to consume, but do not explain differences between timing of consumption when income is public information and when it is not. My model, though, predicts more money is spent immediately when income is received in public than when it is received in private. I test the model using a field experiment that assigns some individuals to receive money in public settings, and other individuals to receive money in private settings. I run two lotteries in 158 agriculture clubs in central Malawi. One lottery, and its winner, is publicly announced to the whole group. The other lottery is private, and only its winner knows that a second lottery was held. This experimental variation in information about cash assets allows me to distinguish between the private, first-best use of unexpected income, and the second-best use of such income when constrained by obligations to other members of the community. If public knowledge about income does not constrain its use, then spending patterns will be the same for winners of the public and private lotteries. On the other hand, if public winners face a different budget constraint than private winners, they will spend their prize more quickly in order to evade the obligation to share their income. Indeed, I find that those who receive money under the public condition spend 35 percent more windfall income within the first week of receiving it than those who receive money under the private condition.

The notion that social norms constrain consumption and savings decisions is not new. *Townsend* (1994) develops a model in which informal transfers within a village smooth consumption against idiosyncratic shocks. *Ligon et al.* (2002) show that with unenforceable commitments to village networks, informal insurance schemes do not fully insure against these shocks. Enforcement, then, constrains the private first-best allocation of income under some circumstances. Social anthropologists expand upon this notion, explaining that the pressure to share cash with family or neighbors is pervasive, and that people respond to these pressures by spending income quickly: "Consequently, on those infrequent occasions when they were able to earn money, they often made wasteful or ill-considered expenditures just to keep friends from borrowing it" *Maranz* (2001) (p. 18). *Baland et al.* (2007) document similar behavior among micro-finance customers in Cameroon who borrow money despite having savings of equal or greater value than their loans, thus incurring unnecessary interest costs for projects that could be self-financed. They report, "Excess borrowing is purposefully used by some members to signal financial difficulties to their relatives in search of financial help. Reimbursement obligations are then used as an argument to discourage such demands" (p. 9). *Comola and Fafchamps* (2010) demonstrate that gift-exchange between rural Tanzanian households is sometimes involuntary rather than an optional reciprocal strategy for coping with idiosyncratic shocks.

I extend this literature by using experimental variation in whether windfall income is received in a public or private setting to formally test the effect of social pressure on consumption decisions. My paper proceeds as follows. I describe the model in Section 2.2. I explain the experiment in Section 2.3, and the data in Section 2.4. I present results in Section 2.5. Section 2.6 concludes.

2.2 A Simple Model

Individuals with high discount rates consume income rapidly and have low marginal propensities to save. Hyperbolic discounters apply an additional penalty to consumption in future periods, which also leads to high marginal propensities to consume. In a two-period model with windfall income, consuming the windfall income in the initial period rather than spreading it between the two periods would be consistent with a high discount rate or hyperbolic discounting. Here, I develop a model in which the same behavior is explained by pressure to share, which is an economic constraint, rather than the discount rate, which is a behavioral parameter. In this model, exponential discounters and even people who are perfectly patient will consume income rapidly in order to avoid sharing their income when others observe their income. The combination of rapid consumption and obligatory transfers to people outside of the household results in high expenditures in the period in which income is received.

In each period, an individual's wealth consists of savings from the previous period and income earned in the current period. Both savings and income can be either public or private – that is, known either to the whole community or known only to the individual. Let w denote wealth, s savings, c consumption and y income. The superscript u will denote that it is public, and v that it is private. The subscript t indexes time periods. We have thus

$$
\begin{aligned}
w_t^u &= y_t^u + (1+r)s_{t-1}^u \\
w_t^v &= y_t^v + (1+r)s_{t-1}^v
\end{aligned}
$$

In addition, individuals are expected to share surplus income with others in their social network. Surplus income is public income that exceeds immediate immediate consumption needs. This sharing norm exists to redistribute money in a socially efficient way, from those with low marginal utility of consumption to those with high marginal utility of consumption. In practice, full redistribution is not achieved because individuals have incentives to reduce the amount they share, which they can justify by signaling an artificially high marginal utility of consumption. Individuals signal a high marginal utility of consumption by spending income immediately and visably, on goods for themselves or their own households. A fraction τ of public income that is not seen to be consumed immediately – that is, $max(0, w_t^u - c_t)$ – must be given to others in the network.

Spending income immediately demonstrates that a household has legitimate consumption needs. Conceptually, it is a strategy similar to non-cash-constrained families taking out loans, as a way to signal inability to give money to relatives*Baland et al.* (2007). In contrast, holding cash indicates a relatively low marginal utility of consumption, such that it might be socially optimal to reallocate money to someone else in the network. Individuals can choose to signal that they have high needs and should therefore be exempt from contributions to the network by spending their income immediately. Even though this rapid spending reduces contributions to the group, it escapes sanction because negative shocks are not perfectly observed, and thus others in the network are not sure whether the rapid spending represented genuine or strategic need to consume. Social sanctions can only be enforced for income that was known to others, so holding cash from w^u is subject to sharing but holding cash from w^v is not.

Consider a two period model where individuals will consume all their income in the second period. For simplicity, assume $r = 0$, and let $\theta = 1 - \tau$. Subsume y_0 into w_0 which is given and let $y_1 = y$, also given. There are two possibilities:

First if $c_0 < w_0^u$ then individuals contribute others in their networks, based on the amount of first period public income that is not consumed in the first period:

$$\begin{aligned} w_1^u &= \theta(w_0^u - c_0) + y^u \\ w_1^v &= w_0^v + y^v \\ c_1 &= \theta(w_0^u - c_0) + y^u + w_0^u + y^v \\ &= y^u + y^v + \theta w_0^u + w_0^v - \theta c_0 \end{aligned} \tag{2.1}$$

Second if $c_0 \geq w_0^u$ then individuals signal that they need all of the income they have received by consuming it immediately and do not make any contributions:

$$\begin{aligned}
w_1^u &= y^u \\
w_1^v &= (w_0^v + w_0^v - c_0) + y^v \\
c_1 &= y^u + (w_0^u + w_0^v - c_0) + y^v \\
&= y^u + y^v + w_0^u + w_0^v - c_0 \qquad (2.2)
\end{aligned}$$

In addition, I assume that there is no borrowing so $c_0 \leq w_0^u + w_0^v$. The set of possible combinations of c_0 and c_1 is shown in figure 2.1. The blue line corresponds to equation 2.1 and the green to equation 2.2. The solid portion of the line represents combinations of c_0 and c_1 that are actually possible.

An individual wants to maximize utility, $u(c_0, c_1) = \ln c_0 + \beta \ln c_1$. Maximizing utility such that equation 2.1 holds, I find the point (call it c_0^*) where an indifference curve of the utility function is tangent to the blue line. If I maximize such that equation 2.2 holds, I find the point (c_0^{**}) where an indifference curve is tangent to the green line.

The solution to this model has four cases, which I describe in detail in Appendix B. The sharing norm binds in two cases. In one, the sharing norm causes individuals to consume more of public income in period 0 than they would have otherwise, driving the marginal propensity to consume from public income above the marginal propensity to consume out of private income. At the extreme, individuals are driven to a second case, the corner solution. At the corner solution, they consume all of public income and none of private income in period 0.

Either of these cases are plausible if a large share of income is publicly observed and the fraction τ that must be contributed to others is large. It leads to high marginal propensities to consume, behaviors that are observationally equivalent to high discount rates or hyperbolic discounting. However, the high marginal propensity to consume is explained by the obligation to share income that others know about, rather than a high discount rate. In other words, this model provides an alternate explanation for behavior that may appear to result from non-standard discount rates under conditions about the observability of income and the importance of sharing norms that are thought to be important features of developing countries.

Importantly, the model predicts different levels of period 0 consumption when income is observed and when it is not. To test this, I conduct lotteries in clubs of farmers in Malawi to generate exogenous shocks to public and private income. If the model holds, I expect to see that a shock to public income is consumed immediately while a shock to private income is not.

2.3 Experimental Design

Enforcing sharing norms on savings or assets requires public information about the level of savings or assets, even when the "tax" is a contribution to an informal insurance network. When information about an individual's income is public, that individual faces constraints that lead to a higher marginal propensity to consume immediately after receiving income and more money given to others in the social network. To test those implications of the simple optimization model above and determine the causal effect of public information about individuals' choices on their financial decision-making, I study the allocation of windfall income under different information conditions. In principle, my experiment is similar to *Ashraf* (2009). While Ashraf focuses on information conditions that affect within-household bargaining, I am interested in community-level dynamics.

Maranz's description of social pressure to share turns on the assumption that "if someone receives money, those people who are socially close will know it." *Maranz* (2001) It is exactly that assumption that I manipulate in a simple field experiment involving 1,553 farmers in central Malawi. I run 316 lotteries in 158 agricultural clubs in central Malawi in May 2008. These clubs of approximately 10 members each were formed in late 2007 for the purpose of receiving extension services and borrowing through group liability schemes, and the lotteries are conducted when the clubs are assembled for meetings and surveys related to their loans for the coming agricultural season. The lotteries are facilitated by trained, Chichewa-speaking enumerators.

In each club, one lottery is "public." The lottery and the amount of the prize are announced to the group. Farmers each draw a ticket from a bag, and the farmer whose ticket has a star is declared the winner. The enumerator records the winning farmer's name and awards him his cash prize in front of the entire group. Everyone present knows there was a lottery, who won the lottery, and the value of the prize. Winning this prize is an increase in public income, w_0^u.

The second lottery is "private." Before meeting with each club, a winner and several alternate winners for this lottery are randomly selected. In the case that the designated winner is not present or won the public lottery, the prize is awarded to the highest-ranked alternate. The group is not told about the second lottery. Instead, the winner is informed privately, while responding to the baseline survey. The winner is assured that no one else in the community has won money in secret, and that no one else has been told that he (the private winner) received a prize. Because the supplemental survey for lottery winners is brief and completing the baseline survey takes longer for some group members than others, it is unlikely that the time to complete the lottery questionnaire signals anything out of the ordinary to other group members. Also, all questionnaires are administered in a private setting, out of sight and hearing of other members of the group. In other words, I take every reasonable precaution to ensure that the private lottery is indeed private, and that winners feel secure that no one in the group knows of their prize. Winning a private lottery is an increase in private income, w_0^v.

The prizes for the public and private lotteries are identical, MK 2500 ($17.86 US, at an exchange rate of MK 140 = $1 US) paid immediately in cash. That sum is roughly equivalent to one-tenth of average annual per capita cash income in Malawi, and will buy 25 kg of fertilizer or five chickens. Since the public and private lottery winners are randomly chosen, any differences between how they choose to use their prize can be attributed solely to the impact of their communities' knowledge of their income.

2.4 Data

My final sample is of 315[1] lottery winners, half of whom won in "public" settings and the other half of whom won under "private" conditions. I have data from four surveys. Surveys were administered on two occasions, in May and August 2008, at

[1]One individual who won a private lottery declined the prize.

a central meeting location were all group members were asked to gather. Chichewa-speaking enumerators conducted one-on-one interviews with respondents.

In each club, the May surveys were conducted on the same day as the lotteries. The baseline survey was administered to all 1,553 members of groups where lotteries were conducted. The supplemental lottery survey asked the 315 winners to list the ways in which they would use their prize money, and then indicate when each transaction would take place and who would be the beneficiary. In August, a subset of 81 participating clubs were revisited[2]. Some 627 members of those clubs were present and were administered a follow-up questionnaire about assets, savings, and the recent harvest. At least one of the lottery winners was present in 77 of those clubs; in total, 114 lottery winners were administered a supplemental survey asking how they had actually used their prize money.

Table 2.1 presents summary statistics for baseline characteristics of the public and private lottery winners. Public and private winners do not differ significantly in their gender, age, or years of education. The apparent difference in land ownership is due to outliers and becomes insignificant when trimming the top one percent of land holdings. Including or excluding these baseline characteristics does not affect the sign or significance of subsequent results. However, public and private winners

[2]Budget constraints precluded revisiting all clubs.

do differ substantially in their likelihood of being resurveyed in August. I examine this apparently selective attrition in Table 2.2.

Not all clubs were resurveyed in August, but since each club has one private and one public winner, equal response rates for the two types of lottery winners were desirable. Public lottery winners (41.4%) were about as likely as all respondents (40.4%) to appear in the August sample. Only 32.3 percent of private lottery winners were resurveyed, however. Among lottery winners, those who did and did not respond to the August survey were about equal in gender, age, and land owned. However, those who did not respond have significantly fewer years of education than those who did. It is possible that the more educated face a higher opportunity cost of time, though unlikely because the sample consists entirely of farmers who do not do regular wage labor and the August survey took place in the lull period between the harvest and the next season's planting.

Public and private winners are balanced on their baseline characteristics. Ultimately, I cannot explain the differences in response rates for public and private winners. The most plausible selective attrition story that is related to the lottery type is that private winners were concerned that their prize could be exposed to other group members during another encounter with the survey team. People who were concerned about their prize being disclosed to others were probably those who used the prize differently than they would have if others knew about it, so this story biases me towards finding no results.

My outcomes of interest span two concepts: when prize money is spent, and how it is spent. For each concept, I have two sets of measures. The measures from the April supplemental surveys tell me about anticipated use of the prize money, while the outcomes from the August surveys tell me about the realized use of the prize money. To explore timing of prize use, I aggregate spending by date. I have measures of spending the same day as the lottery, within one week of the lottery, within the same month as the lottery (May), and in each of three subsequent months (June, July, and August). I focus primarily on money spent immediately, meaning the week as the lottery, though the results are not sensitive to using the narrower same-day restriction. The reason for choosing the same-week rather than same-day measure is that the lotteries, surveys, and related research activities took most of the day and left little time for winners to spend their prize money. Also, market days happen once per week in most villages, so the primary opportunity to spend money occurs at a weekly interval.

My analysis of how prize money is used divides spending into five categories:

consumption by the winner, consumption by others in the winner's household, consumption by persons not in the winner's household, investment or purchase of durable goods for the household, and savings. These categories are mutually exclusive and exhaustive. I include purchase of agricultural inputs such as fertilizer and pesticides, purchase of livestock, and purchase of building materials in the "investment" category. Results for analysis of these categories are not sensitive to alternative definitions of investment, such as removing livestock.

2.5 Results

2.5.1 Timing of expenditures

I compare the timing of expenditures by those who won public lotteries to those who won private lotteries, and show that winners of public lotteries spend money more rapidly. Since public and private lottery winners were randomly selected and balanced at baseline, differences in the timing of their expenditures are caused by differences in treatment: winners of public lotteries were exposed to social pressure to share income, while winners of private lotteries were not.

Table 2.3 shows the difference in expenditures within one week of the lottery for public and private winners. Columns (1) through (4) use data about anticipated expenditures collected at the time of the lotteries for all winners. Note first that all lottery winners anticipated using a large amount of their prize money very quickly. In the full sample, winners of private lotteries anticipate spending MK 1,815, or 72 percent of their prize money, within one week of the lottery. Winners of public lotteries anticipate spending MK 2,051 (MK 2,041 controlling for baseline characteristics including gender, age, household size, education, land owned, physical attributes of the house, and the household's durable assets) in the same time period. The difference of MK 236 (MK 226 controlling for baseline characteristics) is statistically significant and represents a 13 percent increase relative to the immediate expenditures of winners of private lotteries. Limiting the analysis to the subsample of respondents who were resurveyed three months after the lotteries, the magnitude of the difference between public and private lottery winners' expenditures persists, though the standard errors are larger.

I use data from the follow-up survey about actual expenditures in columns (5) and (6). The differences in actual spending between public and private lottery winners are large and statistically significant. Within one week of the lotteries, private winners had spent MK 985, or 39 percent of their prize. Public winners, though, spent MK 1,334 (MK 1,327 controlling for baseline characteristics). The difference of MK 349 (MK 342) is statistically significant and large relative to the spending of private lottery winners within the same time frame. Receiving prize money in a public setting induced individuals to spend one third more of their money in the first week than they would have had they received the money in private. This is strong evidence that the context in which money is received affects the time over which it is spent. Spending money rapidly as a reaction to social sharing norms reduces the ability to smooth consumption across time or to adjust the amount shared to unanticipated negative shocks realized after income is received.

2.5.2 Strategic spending and sharing

Table 2.3 shows that gross expenditures immediately after the lotteries were higher for those who won in public settings than in private settings.[3] This is strong evidence that social pressure affects consumption decisions. In Table 2.4, I look more directly at strategic spending and sharing. Strategic spending is the difference between immediate consumption of public lottery winners compared to private lottery winners. It has three components: expenditures for the winner him/herself, expenditures for the winner's household, and expenditures on durable (relatively illiquid) goods for the winner's household. The model predicts that public winners spend strategically to reduce the amount of money they are obligated to give to others.

[3]Estimates in this and subsequent tables include baseline covariates. Including the covariates does not affect the results.

As shown in columns (1) to (3), private winners spend somewhat more than public winners on themselves, their households, and durable goods. However, these differences are not statistically significant. Column (4) is the sum of the three previous columns; the coefficient on the indicator for winning a public lottery in column (4) is the total amount of strategic spending. Receiving money in public causes lottery winners to consume MK 246 more of their prize in the first week than they would have if they had received money privately. This difference is not statistically significant, but the magnitude is large and in the predicted direction. MK 246 is ten percent of the total prize money, and a 25 percent increase in immediate spending relative to winners of public lotteries. It is an amount equivalent to about two-and-a-half day's wages, and enough to purchase a live chicken or pay for a course of medicine to treat malaria.

Column (5) measures expenditures on or gifts to others outside the household. This is the transfer payment $\tau \times (w_0^u - c_0)$. Winners of public lotteries spend MK 97 more on people outside of their households than winners of private lotteries. This confirms the model's basic assumption that observable income is subject to different sharing norms than unobservable income. Note that we can calibrate the model above to estimate τ:

$$\tau = \frac{transfers}{w_0^u - c_0} \tag{2.3}$$

In this case, $\tau = 0.07$. Winners were obligated to share with others seven percent of income they did not spend on their own immediate consumption.[4]

Thus far, my discussion has focused on differences in immediate consumption that are caused by social pressure to share when income is observed by others in the network. The model described in Section 2 also predicts differences in *overall* consumption and sharing between public and private winners. Winners of public lotteries are expected to share more money total, and therefore consume less total, than winners of private lotteries. Winners of public lotteries may also weight spending on their households towards durable goods as a way of smoothing consumption, and are likely to hold less cash or liquid assets at the end of the follow up period.

Most of these predictions are not borne out in the data from the follow-up survey. The categories in Table 2.5 correspond to those in Table 2.4, but Table 2.5 covers spending in three months following the lotteries. Public winners spend approximately the same amount as private winners on themselves. They spend MK 186 less than

[4]Since this experiment studies use of windfall income, results may not generalize to earned income. *Jakiela* (2009) finds a higher tax rate imposed on windfall than "earned" income in a laboratory experiment in Kenya.

private winners on consumption goods for their households, but MK 196 more on durables. Neither difference is statistically significant, though the reallocation is consistent with an effort to smooth consumption by winners of public lotteries.

More surprisingly, over the three month horizon, there is no meaningful difference in the amount of money given to people outside the household. Winners of private lotteries give away MK 190 over three months, and winners of public lotteries give away MK 186. Winners of public lotteries give away more money immediately than winners of private lotteries, but less in subsequent months. Receiving money in public affects the timing of consumption and sharing, but not the total level. This result is not explained by the model described in Section 2. One possible explanation is that social obligations are absolute instead of relative to income, and that efficiency concerns dictate collecting money from each person as soon as practical. Public winners therefore paid their "tax" immediately, while private winners paid small installments over time. Alternatively, it may imply that information about winning the prize in the private condition became public over time, and private winners were subject to taxation as information was revealed. Either of these alternative explanations still have welfare implications for receiving money in public compared to in private, since both imply constraints on individuals' flexibility in smoothing consumption.

2.5.3 Ability to predict spending

Recall that in Table 2.1, I show that public and private winners differed in their *predicted* as well as *realized* spending within one week of the lottery. While winners of public lotteries reported more immediate spending in both surveys, both public and private lottery winners actually spent less within one week than they anticipated.

I analyze errors in prediction of immediate expenditures in Table 2.6. Column 1 estimates the equation $Y_{i,August} = \alpha + \beta Y_{i,May} + \epsilon_i$. If individuals perfectly predict their spending, then α equals zero and β equals one. Alternatively, if the prediction contains no information about actual spending, then β equals zero. The coefficient is not statistically different from zero, consistent with predictions about spending within one week containing no information about actual spending in that time period. Column 2 adds an indicator for whether the individual won a public or private lottery, and an interaction between winning a public lottery and the predicted level of spending in the first week. The statistically significant coefficient on the interaction term suggests that public and private winners differ in their ability to predict spending in the week following the lottery. Indeed, private winners actually spend 0.32 kwacha for every one kwacha they predicted spending within one week of the lottery. For winners of public lotteries, though, the correlation between predicted and actual spending within one week is zero. Public winners faced apparently unexpected constraints in allocating their prize money relative to private winners. This might indicate uncertainty surrounding the social norms or tax rate governing sharing of windfall income from an unusual source.

In Tables 2.7 and 2.8, I examine the ability to predict spending in each of the five categories. The estimates in Table 2.7 are category-by-category OLS regressions of realized spending on anticipated spending and a constant. As in column (1) of table 5, there are two interesting tests of each coefficient in table 8. A coefficient of zero indicates that the predicted level of spending in a given category contained no information about actual spending in that category. A coefficient of one means that the prediction was perfect. For each category of spending, I reject the hypothesis that lottery winners perfectly anticipated their spending with their May survey responses. Anticipations do have some predictive value for all categories except spending on one's self, however. The strongest correlation between predicted and actual spending is for money shared with people outside of the household. For every kwacha lottery winners anticipated sharing with others, they actually shared MK 0.40.

Table 2.8 asks whether winners of public lotteries differed from winners of private lotteries in their ability to predict spending across the five categories. The regressions in this table include an indicator for whether the individual won a public lottery and an interaction term. As before, interaction terms that are significantly different from zero indicate that public winners differed from private winners in their ability to predict spending. The coefficient on the predicted level of spending is the marginal spending by August for each kwacha predicted in May for private lottery winners. The sum of the coefficients on the predicted level and the interaction term captures the same concept for public lottery winners.

Winners of public and private lotteries are not significantly different in their ability to predict spending in any of the categories, but there are interesting differences in the categories for which each group of winners is able to make meaningful predictions. Public lottery winners could predict about three times more of the variance in their spending on durable goods than private lottery winners. For each kwacha public lottery winners anticipated spending on durables, they actually spent MK 0.37, compared to MK 0.13 for private lottery winners. Though the difference is not statistically different from zero, public lottery winners' predictions explained a statistically significant fraction of the variance in their spending on durables and private lottery winners' predictions did not. In contrast, private lottery winners were three-and-a-half times more accurate at predicting the amount of money they would give to others. For each kwacha they anticipated sharing, private lottery winners actually gave MK 0.67 to others. Public lottery winners, though, actually gave away only 0.19 kwacha for each kwacha they anticipated sharing.

The differences in ability to predict investment are consistent with public winners reacting to a perceived threat of taxation, and with the results about the timing of use of the prize money in the previous section. Recall that public winners invested more money than private winners in the week immediately after the lottery. The heightened ability to predict investment might be because public winners planned to use investment - which converts income from taxable cash into non-taxable goods - to protect their money from others' claims. The inability to predict spending on others may suggest that the anticipation of a social tax has a bigger effect on spending than actual enforcement of such a tax. Public winners anticipated and acted to protect their income from a threat that, in fact, was not realized. This is a puzzling outcome in equilibrium, but is consistent with gradual revelation of private winners' prizes between May and August, and consequent catch-up in sharing by those private winners.

2.6 Conclusion

I use a simple experiment of allocating money to members of agricultural clubs in public and private lotteries to measure the impact of public information on farmers' anticipated use of their prizes. While all winners spend a large share of their prize money in the one-week period immediately following the lottery, those who won money in a public setting have 35 percent higher expenditures in that short window. The tendency to spend quickly could suggest very high discount rates, but the difference between the rapid spending by individuals who win the money publicly and that of those who win privately requires additional explanation. Strong sharing norms may constitute a tax on income that others know about. Then, spending such income quickly, before others lay claim to it, is the rational optimizing behavior of an individual with a standard discount rate facing a budget constraint with parameters such that a tax on public income is binding. The short-run spending patterns in my data suggest a tax of seven percent on surplus income. Over a longer horizon, however, the use of income received in public is statistically indistinguishable from income received in private. The welfare implications of such a finding are ambiguous. Sharing norms may be a constraint that force individuals to accept a second-best solution, where spending quickly limits the ability to shop for better prices or leads to hasty decisions that are regretted in the future. Even though identical fractions of public and private income are given to people outside the household over a longer horizon, the accelerated giving and spending when income is received in public reduces consumption smoothing and the ability to adjust sharing to unanticipated negative shocks. However, sharing norms may also provide an important means of insuring against idiosyncratic shocks. Public information about income may increase the enforceability of these informal insurance networks, therefore increasing the consumption-smoothing benefits they provide. Nonetheless, it is clear from my data that it is important to model this additional constraint when studying consumption decisions of individuals when income is easily observed and such norms are likely to be present. These individuals may exhibit high marginal propensities to consume despite not having high or hyperbolic discount rates.

Tables

Table 2.1: Sample means, public and private winners

	(1) Male	(2) Age	(3) Years of Education	(4) Land Owned	(5) In August Sample
Public	0.947	43.49	5.933	4.520	0.414
	(0.018)	(1.133)	(0.297)	(0.205)	(0.040)
Private	0.923	44.65	6.130	5.067	0.323
	(0.022)	(1.090)	(0.290)	(0.238)	(0.038)
Observations	306	303	304	311	315
p-value: public=private	0.351	0.478	0.603	0.073	0.021

Means for winners of public and private lotteries.

Table 2.2: Sample means, attriters and non-attriters

	(1) Male	(2) Age	(3) Years of Education	(4) Land Owned
In August Sample	0.947	44.83	5.491	4.948
	(0.024)	(1.273)	(0.365)	(0.251)
Not in August Sample	0.927	43.62	6.349	4.701
	(0.020)	(0.931)	(0.282)	(0.209)
Observations	306	303	304	311
p-value: in=out	0.521	0.444	0.063	0.448

Means for baseline respondents observed in August and not observed in August.

Table 2.3: Spending within one week of lottery

	(1)	(2)	(3)	(4)	(5)	(6)
		Predicted			Realized	
Public	236.097**	225.642*	293.754	351.604*	348.711*	342.715*
	(113.878)	(117.319)	(189.479)	(197.181)	(186.395)	(191.255)
Covariates		x		x		x
Observations	294	294	114	114	114	114
Mean for private winners	1815.07	1815.07	1715.29	1715.29	985.10	985.10
R^2	0.01	0.07	0.02	0.11	0.03	0.11

OLS estimates. All standard errors are clustered at the club level.
Columns 1 and 2 include data from all lottery winners. Columns 3 to 6 include data from winners who were interviewed at follow-up. Covariates are age, gender, education category dummies, household size, number of children, amount of land owned, index of housing quality, index of livestock owned, numeracy score, indicator for self reported risk taking, transfers received during the previous season, and transfers given during the previous season.
* p<0.10, ** p<0.05, *** p<0.001

Table 2.4: Immediate spending and sharing

Dependent variable:	(1) Self	(2) HH	(3) Investment	(4) Total (1) to (3)	(5) Non-HH	(6) Savings
Public	125.009	16.275	104.750	246.035	96.680*	-49.841
	(106.797)	(134.955)	(114.407)	(192.429)	(52.176)	(75.031)
Covariates	x	x	x	x	x	x
Observations	114	114	114	114	114	114
Mean for private winners	213.73	476.67	247.65	938.04	47.06	88.24
R^2	0.09	0.10	0.13	0.10	0.11	0.12

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.
See footnote for Table 2.3 for a list of covariates. * p<0.10, ** p<0.05, *** p<0.001

Table 2.5: Eventual spending and sharing

Dependent variable:	(1) Self	(2) HH	(3) Investment	(4) Total (1) to (3)	(5) Non-HH	(6) Savings
Public	-7.576	-185.549	195.575	2.449	-3.799	6.709
	(137.427)	(175.013)	(175.213)	(165.384)	(92.580)	(116.582)
Covariates	x	x	x	x	x	x
Observations	114	114	114	114	114	114
Mean for private winners	515.69	979.61	559.02	2054.31	190.20	150.98
R^2	0.13	0.17	0.11	0.10	0.04	0.06

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.
See footnote for Table 2.3 for a list of covariates. * p<0.10, ** p<0.05, *** p<0.001

Table 2.6: Relationship between actual and predicted spending within one week of lottery

	(1) Realized	(2) Realized
Predicted	0.121	0.320**
	(0.101)	(0.140)
Public		995.320**
		(342.079)
Public × Predicted		-0.372**
		(0.160)
Covariates	x	x
Observations	114	114
R^2	0.10	0.15
p-value: public+interaction = 0		0.00

OLS estimates. All standard errors are clustered at the club level.
Sample includes winners interviewed at follow up.
* p<0.10, ** p<0.05, *** p<0.001

Table 2.7: Actual versus predicted spending and sharing

Dependent variable:	(1) Self	(2) HH	(3) Investment	(4) Total (1) to (3)	(5) Non-HH	(6) Savings
Predicted Self	-0.083					
	(0.146)					
Predicted HH		0.208*				
		(0.107)				
Predicted Investment			0.232*			
			(0.126)			
Predicted Total (1) to (3)				0.159**		
				(0.079)		
Predicted Non-HH					0.404*	
					(0.242)	
Predicted Save						0.343**
						(0.159)
Covariates	x	x	x	x	x	x
Observations	114	114	114	114	114	114
R^2	0.13	0.20	0.13	0.14	0.09	0.17

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.
See footnote for Table 2.3 for a list of covariates. * p<0.10, ** p<0.05, *** p<0.001

Table 2.8: Actual versus predicted spending and sharing, by lottery type

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Self	HH	Investment	Total (1) to (3)	Non-HH	Savings
Public	-69.909	-197.688	103.191	-99.552	41.095	64.141
	(155.185)	(225.152)	(206.052)	(335.153)	(97.372)	(103.882)
Predicted Self	-0.220					
	(0.139)					
Predicted Self × Public	0.302					
	(0.296)					
Predicted HH		0.247*				
		(0.146)				
Predicted HH × Public		-0.043				
		(0.196)				
Predicted Investmen			0.134			
			(0.129)			
Predicted Investment × Public			0.239			
			(0.181)			
Predicted Total (1) to (3)				0.119		
				(0.124)		
Predicted Total (1) to (3) × Public				0.064		
				(0.157)		
Predicted Non-HH					0.674*	
					(0.349)	
Predicted Non-HH × Public					-0.485	
					(0.403)	
Predicted Save						0.395*
						(0.217)
Predicted Save × Public						-0.107
						(0.310)
Covariates	x	x	x	x	x	x
Observations	114	114	114	114	114	114
R^2	0.14	0.22	0.16	0.14	0.10	0.17
P(predicted + interaction = 0)	0.75	0.17	0.04	0.08	0.44	0.21

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.
See footnote for Table 2.3 for a list of covariates. * p<0.10, ** p<0.05, *** p<0.001

Figures

Figure 2.1: Possible combinations of c_0 and c_1

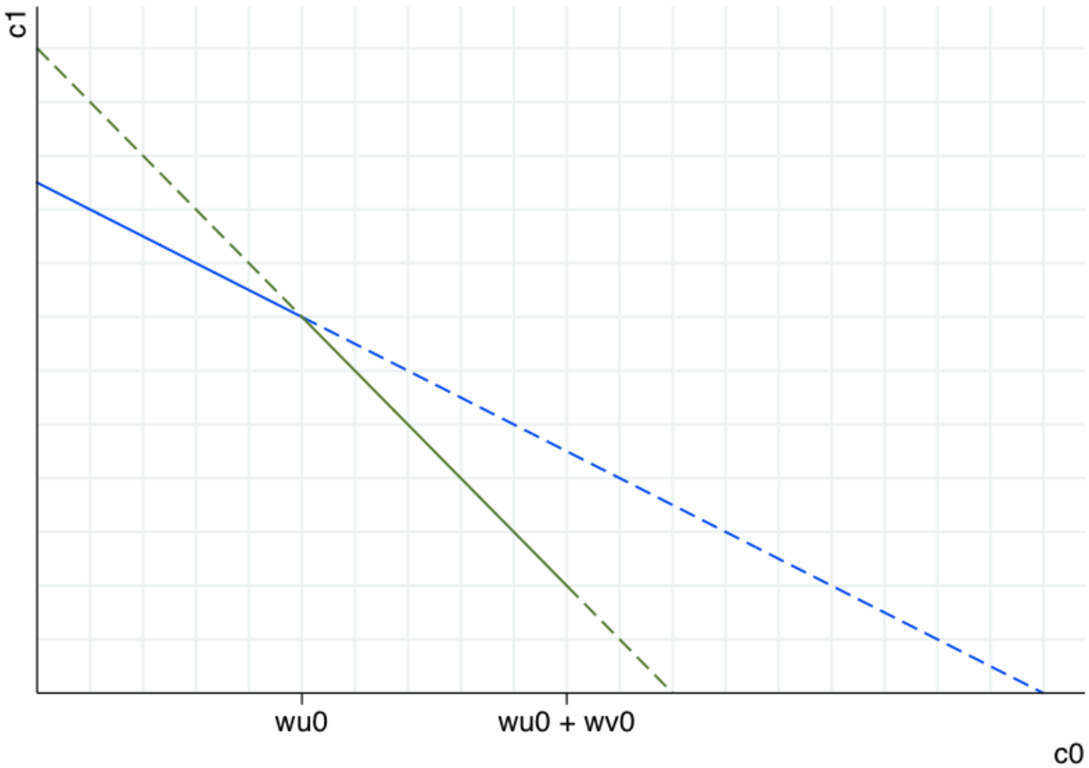

CHAPTER III

Identification Strategy: A Field Experiment on Dynamic Incentives in Rural Credit Markets

3.1 Introduction

Lending in low-income countries is notoriously difficult. Clients typically lack adequate collateral and lenders often have limited information about the profitability of their customers. Information asymmetries coupled with costly enforcement of repayment severely limits the profitability of lenders. The problem is particularly acute in agriculture because the nature of production precludes the use of many of the mechanisms used in microfinance. For example, lenders cannot schedule frequent repayments because cash flows are only received after harvest, several months after the loan is taken. In addition, all farmers need cash at the same time, so allowing some farmers to borrow only after others have repaid their loans is problematic because some farmers would end up receiving credit when they do not need it. Even if all clients were allowed to borrow at the same time, joint liability may be ineffective if most production shocks are covariate. Finally, and perhaps most importantly, lenders may lack the ability to deny access to future loans to defaulting clients in the absence of a national system that allows individuals to be uniquely identified.

When this happens, loan defaulters can often avoid sanction by simply applying for new loans under different identities. Lenders respond by limiting the supply of credit, due to the inability to sanction unreliable borrowers and, conversely, to reward reliable borrowers with expanded credit. As a result, many smallholder farmers are severely constrained by the inability to finance crucial inputs such as fertilizer and improved seeds, particularly for export crops.[1]

In this paper we implement a randomized field experiment to estimate the impact of biometric identification (fingerprinting) in a contextrural Malawi characterized by a lack of a unique identification system and limited access to credit.

According to the 2006 Doing Business Report, Malawi ranked 109 out of 129 countries in terms of private credit to GDP, a frequently-used measure of financial development. Malawi also gets the lowest marks in the depth of credit information

[1]The following quote from 1973 by Robert McNamara when he was the World Bank president exemplifies this view: "The miracle of the Green Revolution may have arrived, but for the most part, the poor farmer has not been able to participate in it. He simply cannot afford to pay for the irrigation, the pesticide, the fertilizer For the small holder operating with virtually no capital, access to capital is crucial."

index" which proxies for the amount and quality of information about borrowers available to lenders. Using more micro data, 74 percent of cash crop farmers in our baseline survey had not borrowed from a bank or microfinance institution in the last 10 years.

In the experiment, smallholder farmers organized in groups of 15-20 members applied for agricultural input loans to grow paprika and were randomly allocated to either a control group or a treatment group where each member had a fingerprint collected as part of the loan application. Unlike conventional ID cards or passports, a fingerprint is an effective personal identifier because it is unique to and embodied in each person, so it cannot be forgotten, lost or stolen. Thus, fingerprinting customers would allow lenders to construct credit histories and use them to withhold new loans from past defaulters. In essence, fingerprinting can make the threat of future credit denial more credible.

To guide the empirical strategy, we develop a simple two period model in the spirit of *Stiglitz and Weiss* (1983) that incorporates both adverse selection and moral hazard and show that "dynamic incentives," that is, the ability to deny credit in the second period based on the first period repayment performance, can reduce both types of asymmetric information problems and therefore raise repayment. Adverse selection problems can be mitigated because riskier individuals that would otherwise default may now take out smaller loans (or avoid borrowing altogether) to preserve access to credit in the future.[2] In addition, borrowers may have greater incentives to ensure that agricultural production is successful, either by exerting more effort or by diverting fewer resources away from production (lower moral hazard). Also intuitively, the model predicts that the impact of "dynamic incentives" will be largest for the riskiest individuals.

Consistent with the predictions of the model, fingerprinting led to substantially higher repayment rates for the subgroup of farmers with the highest ex-ante default risk. By contrast, fingerprinting had no impact on repayment for farmers with low ex ante default risk. While we cannot separate the effect of moral hazard and adverse selection on repayment, we collect unique additional evidence that points to the presence of both informational problems. Fingerprinting leads farmers to choose smaller loan sizes, consistent with a reduction in adverse selection. In addition, high-default-risk farmers who are fingerprinted also divert fewer inputs away from the contracted crop (paprika), which we interpret as a reduction in moral hazard. When we compare these benefits to fairly conservative costs of implementation, we find that adoption of fingerprinting is cost-effective (the benefit-cost ratio is 2.27).

The key contribution of the paper is that, to our knowledge, it is the first randomized field experiment examining the impact of a technology that improves the effectiveness of dynamic incentives in a credit market. Our analysis is further distinguished by the fact that, in addition to measuring impacts on borrowing decisions and repayment (using the lenders administrative data), we also estimate impacts on specific behaviors related to moral hazard using a detailed follow-up survey of borrowers.

Substantively, our intervention most closely resembles the promise of a future lower interest rate conditional on current loan repayment in *Karlan and Zinman* (2009),

[2] In this paper we use the term "adverse selection" to mean ex-ante selection effects deriving from borrowers hidden information. We acknowledge that such selection may occur on the basis of either unobserved risk type (emphasized in the model) or unobserved anticipated effort (as highlighted by *Karlan and Zinman* (2009)).

henceforth "KZ", who find evidence of moral hazard and weaker evidence of adverse selection in an experiment with a South African provider of consumer loans. Our experiment differs from KZs in several important respects. First, our experiment is concerned with lending for productive investment in rural areas, while KZs involves consumer loans for urban customers. Second, our experiment estimates the impact of manipulating the ability to impose dynamic incentives via a technological innovation. KZ, on the other hand, measures the impact of informing borrowers of the existence of a dynamic incentive. Third, we implement a follow-up survey of borrowers to provide additional insight into the specific behaviors that are changed by the intervention and that result in higher repayment. KZ, by contrast, relies exclusively on the lenders administrative data for analysis and so cannot shed light on what borrower behaviors may have changed.

The fourth and final key difference is in the timing of the intervention relative to the borrowing decision. In KZ, the dynamic incentive is announced after clients have agreed to borrow (and all loan terms have been finalized). As a result, differences in repayment can only be due to moral hazard. In our case, the lenders ability to use dynamic incentives (due to fingerprinting) is revealed before agents decide to borrow. Consequently, the composition of borrowers and the choice of loan terms may change as well. Because potential borrowers cannot repeatedly be surprised, an estimate of the impact of dynamic incentives that are revealed prior to the customers borrowing decision is the more relevant policy parameter.[3]

To be clear, because we informed the lender which clubs had been fingerprinted,

[3]In principle, one could fingerprint borrowers at different points in time along the loan cycle to identify various asymmetric information problems. For example a subset of borrowers (group 1) could be fingerprinted before loan decisions are made, then another group (group 2) immediately after loans are granted but before funds are invested into production and a yet another group (group 3) could be fingerprinted once production has taken place but before repayment. A final group of borrowers would not be fingerprinted (group 0). With full compliance, that is, when all subjects agree to be fingerprinted, one could then measure adverse selection by comparing group 1 and 2; ex-ante moral hazard by comparing 2 and 3 and strategic default by comparing 3 and 0. Given the number of farmers in our study, it was infeasible to implement this design because power calculations suggested we could have at best two groups. Our study therefore consists of groups 0 and 1.

loan officers could have changed their behavior towards treated and control clubs in response to this information. For example, they could have devoted more time to monitoring and enforcing repayment from control clubs, since fingerprinted clubs were already subject to dynamic incentives. We provide convincing evidence to the contrary: approval decisions and subsequent monitoring of clubs by loan officers did not differ across treated and control clubs. As a result, we interpret our findings as emerging solely from borrowers responses.

By documenting impacts on behaviors related to adverse selection and moral hazard, our findings contribute to a burgeoning empirical literature that tests claims made by contract theory and measures the prevalence of asymmetric information (see *Chiappori and Salanie* (2000) for a review). A number of recent papers provide empirical evidence of the existence and impacts of asymmetric information in credit markets, in both developed and developing countries. *Ausubel* (1999) uses a large-scale randomized trial of direct-mail pre-approved solicitations from a major US credit card company and finds evidence of higher risk individuals selecting less favorable credit cards, consistent with adverse selection. *Klonner and Rai* (2009) exploit the introduction of a cap in bidding ROSCAS of South India and find higher repayment rates in earlier rounds attributable to changes in the composition of bidders, consistent with lower adverse selection. *Visaria* (2009) documents the positive impact of expedited legal proceedings on loan repayment among large Indian firms, even among loans that originated before the reform, consistent with a reduction in moral hazard. **?** find that incomplete information about fishermens ability in coastal India limits their access to credit for technology adoption. *Edelberg* (2004) also develops a model of adverse selection and moral hazard that is taken to US data from the Survey of Consumer Finance and finds evidence consistent with both informational problems.[4]

[4] *Ligon et al.* (2002) write down competing models of risk-sharing that are taken to the data and find evidence of limited commitment. In a paper similar in spirit, *Paulson et al.* (2006) estimate

The paper is also related to a framed experiment conducted by *Gine et al.* (2010) in Peru that shows that dynamic incentives can be important. In addition, there is a theoretical and empirical literature on the impact of credit bureaus that are also related to this paper. The exchange of information about borrowers should theoretically reduce adverse selection (*Pagano and Jappelli* (1993)) and moral hazard (*Badilla and Pagano* (2000)). Empirically, *Janvry et al.* (forthcoming) study the introduction of a credit bureau in Guatemala and find that it did contribute to efficiency in the credit market. Finally, the paper is related to the literature motivated by the rise in personal bankruptcies in the US in the last decades (*Livshits et al.* (2010)).

The remainder of this paper is organized as follows. Section 3.2 describes the experimental design and survey data and Section 3.3 presents a simple model of loan repayment. Section 3.4 describes the regression specifications, and Section 3.5 presents the empirical results. Section 3.6 provides additional discussion and robustness checks. Section 3.7 presents the benefit-cost analysis of introducing biometric technology, and Section 3.8 concludes.

3.2 Experimental design and survey data

The experiment was carried out as part of the Biometric and Financial Innovations in Rural Malawi (BFIRM) project, a cooperative effort among Cheetah Paprika Limited (CP), the Malawi Rural Finance Corporation (MRFC), the University of Michigan, and the World Bank. CP is a privately owned agri-business company established in 1995 that offers extension services and high-quality inputs to smallholder farmers via an out-grower paprika scheme.[5] The farmer receives extension services and a package of seeds, pesticides and fungicides at wholesale rates in exchange for

structurally competing models of credit markets in Thailand and find moral hazard to be important.

[5]Extension services consist of preliminary meetings to market paprika seed to farmers and teach them about the growing process, additional group trainings about farming techniques, individual support for growers provided by the field assistants, and information about grading and marketing the crop.

the commitment to sell the paprika crop to CP at harvest time. Although CP is by far the largest paprika purchaser in the country, it does not provide credit to farmers because of the risks involved in contract enforcement.[6] CP has a staff of six extension officers and 15 field assistants in the locations chosen for the study. The staff maintain a database of all current and past paprika growers and handles the logistics of supplying farmers with the package of inputs as well as the purchase of the crop.

MRFC is a government-owned microfinance institution and is the largest provider of rural finance, with a nationwide outreach of 210,000 borrowers in 2007. MRFC provided financing for the in-kind loan package for 1/2 to 1 acre of paprika. The loan did not include any cash to purchase inputs. Instead, borrowers took an authorization form from MRFC to a pre-approved agricultural input supplier who provided the inputs to the farmer and billed MRFC at a later date. The loan amount was roughly 17,000 Malawi Kwacha (approximately $120), varying slightly by location. Sixty percent of the loan went towards fertilizer (one 50 kilogram bag of D-compound fertilizer and two 50 kilogram bags of CAN fertilizer); the rest went toward the CP input package: thirty-three percent covered the cost of nine bags of pesticides and fungicides (2 Funguran, 2 Dithane, 2 Benomyl, 1 Cypermethrin, 1 Acephate and 1 Malathion) and the remaining seven percent for the purchase of 0.4 kilograms of seeds.[7] While all farmers that took the loan were given the CP package, farmers had the option to borrow only one of the two available bags of CAN fertilizer. Expected yield for farmers using the package with two bags of CAN fertilizer on one acre of land was between 400 and 600 kg, compared to 200 kg with no inputs.[8]

In keeping with standard MRFC practices, farmers were expected to raise a 15

[6]In 2007, CP purchased approximately eighty-five percent of the one thousand tons of paprika produced annually in Malawi.

[7]The loan amount varied across locations because of modest differences in the transport cost for fertilizer. The cost of the CP package was the same in all locations.

[8]Yield is computed under the conservative assumption that farmers will divert one 50 Kg bag of CAN fertilizer towards maize cultivation. While larger quantities of inputs would result in higher output for experienced paprika-growers, the package described here was designed by extension experts to maximize expected profits for novice, small-holder growers.

percent deposit, and were charged interest of 33 percent per year (or 30 percent for repeat borrowers). Within a group, take-up of the loan was an individual decision, but the subset of farmers who took up the loan was told that they were jointly liable for each others loans. In practice, however, joint liability schemes in Malawi are seldom enforced.[9]

At the time of the study, the vast majority of farmers in the sample had no access to formal-sector credit. In our baseline survey, only 6.7% of farmers had any formal-sector loans in the previous year. Among this small number of farmers with formal-sector credit, MRFC was the largest single lender, providing 34% of loans (more than twice the share of the next largest single lender).[10] Farmers therefore had a strong interest in maintaining good credit history with MRFC so as to maintain access to what would likely be their primary source of formal credit going forward.

In the absence of fingerprinting, identification of farmers relies on the personal knowledge of loan officers (who may also rely on local informants such as village and locality leaders). While loan officers could build up reliable knowledge of borrowers over time, this identification "technology" is imperfect. Loan officers are sometimes promoted and routinely rotated to other localities. Among the 11 loan officers who handle our study areas, the median number of years at the branch is only two, while the median number of years working for the lender is 13.[11] In the absence of an independent mechanism for identifying borrowers, the institutional memory is lost when the loan officer is transferred to another location. Even when loan officers remain in a given location over time, the large number of borrowers can lead them to

[9]See *Gine and Yang* (2009) for another example of limited enforcement of joint liability loans.

[10]Across study areas, access to formal credit varies from 4% to 10%. In Dedza, the region with highest access to formal loans, MRFC provides almost half of these formal loans.

[11]Because soft information about borrowers is important, one may be surprised by the high turnover rate among credit officers. However, MRFC management, like that of many other lenders, rotates credit officers for a number of reasons. For example, rotation is believed to help keep morale high given that they work in difficult environments. It is also thought to minimize corruption and collusion of credit officers with borrowers against the bank. Promotion of successful individuals within the organization also leads to replacement of credit officers at the local level and some loss of soft information on borrowers.

make mistakes in identification. In this project, loan officers issued an average of 104 loans, and also handled other loan customers not associated with the project.

The timeline of the experiment is presented in Figure 3.1. In July 2007, CP asked farmers in the study areas to organize themselves into clubs of 15 to 20 members to accommodate MRFCs group lending rules.[12] Most of these clubs were already in existence, primarily to ease delivery of Cheetah extension services and collection of the crop. Our study sample consists of 249 clubs with approximately 3,500 farmers in Dedza, Mchinji, Dowa and Kasungu districts.

Farmer clubs in the study were randomly assigned to be fingerprinted (the treatment group) or not (the control group), with an equal probability of being in either group. During the baseline survey and fingerprinting period (August and September 2007), CP staff provided a list of paprika growing clubs in each locality to be visited in each week, and randomization of treatment status was carried out after stratifying by locality and week of club visit. The stratification thus ensured that each credit officer handled roughly the same number of treatment and control clubs.[13]

Club visits began with private administration of the baseline survey to individual farmers, and were followed by a training session. Both treatment and control groups were given a presentation on the importance of credit history in ensuring future access to credit. The training emphasized that defaulters would face exclusion from future borrowing, while borrowers in good standing could be rewarded with larger loans in the future. Then, in treatment clubs only, individual participants fingerprints were collected. Our project staff explained how their fingerprint uniquely identified them for credit reporting to all major Malawian rural lenders, and that future credit providers would be able to access the applicants credit history simply by checking his

[12]A typical CP group has between 15 and 30 farmers and is organized around a paprika collection point. MRFCs lending groups have at most 20 farmers, so most of the CP groups participating in the study had to be split to be able to access MRFCs loans.

[13]There are 16 localities or "extension planning areas" (EPAs) in the study. EPAs are administrative boundaries set up for the delivery of agricultural services by Malawis agriculture ministry.

or her fingerprint.[14] Appendix C provides the script used during the training. See Appendix D for further technical details on the biometric technology used.

After fingerprints were collected, a demonstration program was used to show participants that the laptop computer was now able to identify an individual with only his or her fingerprint. One farmer was chosen at random to have his right thumb scanned again, and the club was shown that the individuals name and demographic information (entered earlier alongside the original fingerprint scan) subsequently was retrieved by the computer program. During these demonstration sessions all farmers whose fingerprints were re-scanned were correctly identified. The control group was not fingerprinted, but as mentioned previously, also received the same training emphasizing the importance of ones credit history and how it influences ones future credit access.[15]

The baseline survey administered prior to the training and the collection of fingerprints included questions on individual demographics (education, household size, religion), income generating activities and assets including detailed information on crop production and crop choice, livestock and other assets, risk preferences, past and current borrowing activities, and past variability of income. Summary statistics from the baseline survey are presented in Table 3.1, and variable definitions are provided in Appendix E.[16]

After the completion of the survey, credit history training, and fingerprinting of the treatment group, the names and locations of the members that applied for loans along with their treatment status were handed over to MRFC loan officers so that they could screen and approve the clubs according to their protocols. Among other

[14] Our team of enumerators encountered essentially no opposition to fingerprint collection, perhaps due to the novelty of the technology.

[15] It should be clear that, because we provided education on the importance of credit history to our control group as well, we can estimate neither the impact of fingerprinting without such education, nor the impact of the credit history education alone.

[16] These survey data were collected prior to the farmers being informed about the role of biometrics in the project and their treatment status, to ensure that farmers survey answers were not influenced by knowledge of the nature of the experiment.

standard factors, MRFC conditions lending on the clubs successful completion of 16 hours of training. MRFC approved loans for 2,063 out of 3,206 customers (in 121 out of 239 clubs). Of the customers approved for loans, some failed to raise the required down payment and others opted not to borrow for other reasons. The final sample consists of 1,147 loan customers from 85 clubs.[17] These loan customers received loan packages with an average value of MK 16,913 (US$117).[18]

During the months of July and August, farmers harvested the paprika crop and sold it to CP at predefined collection points. CP then transferred the proceeds from the sale to MRFC who then deducted the loan repayment and credited the remaining post-repayment proceeds to an individual farmers savings account. This garnishing of the proceeds for loan repayment, therefore essentially allows MRFC to "seize" the paprika crop when farmers sell to CP (and for most farmers it is the only sales outlet).[19] Farmers could also make loan repayments directly to MRFC at their branch locations or during credit officer visits to their villages; this occurred, for example, among the small number of farmers who sold to paprika buyers other than CP.

We also implemented a follow-up survey of farmers in August 2008, once crops had been sold and income received. The sample size of this follow-up survey is 1,226 in total (borrowers plus non-borrowers), among whom 520 were borrowers.[20] The formal loan maturity (payment) date was September 30, 2008. Some additional

[17]While a natural question at this point is whether selection into borrowing was affected by treatment status, treatment and control groups did not differ in their rates of MRFC loan approval or the fraction of farmers who ended up with a loan (as will be detailed in the results section below).

[18]All conversions of Malawi kwacha to US dollars in this paper assume an exchange rate of MK145/US$.

[19]Proceeds from other types of crops of course cannot be seized in this way to secure loan repayment because MRFC does not have analogous garnishing arrangements with other crop buyers.

[20]The follow-up sample is smaller than the sample of baseline borrowers because for budget reasons we could not visit each borrowing household at their place of residence. Instead, we invited study participants to come to a central location at a certain date and time to be administered the follow-up interview. Not all farmers attended the meeting where the follow-up survey was administered, but as we discuss below in Section 3.6. (see Appendix Table G.3), there is no evidence of selective attrition related to treatment status. For the full sample as well as the borrower subsample, in no regression is fingerprinting or fingerprinting interacted with predicted repayment statistically significantly associated with attrition from the survey.

payments were made after the formal due date; MRFC reports that there is typically no additional loan repayment two months past the due date for agricultural loans. In the empirical analysis we obtain our dependent variables from the August 2008 survey data as well as administrative data from MRFC on loan take-up, amount borrowed, and repayment.

3.2.1 Balance of baseline characteristics across treatment vs. control groups

To confirm that the randomization across treatments achieved balance in terms of pre-treatment characteristics, Table 3.2 presents the means of several baseline variables for the control group as reported prior to treatment, alongside the difference vis-a-vis the treatment group (mean in treatment group minus mean in control group). We also report statistical significance levels of the difference in treatment-control means. These tests are presented for both the full baseline sample and the loan recipient sample. Overall, we find balance between the two groups in both the full baseline sample and the loan recipient sample. In the full baseline sample, the difference in means for the treatment and control groups is not significant for any of the 11 baseline variables. In the loan recipient sample, for 10 out of these 11 baseline variables, the difference in means between treatment and control groups is not statistically significantly different from zero at conventional levels, and so we cannot reject the hypothesis that the means are identical across treatment groups. For only one variable, the indicator for the study participant being male, is the difference statistically significant (at the 10% level): the fraction male in the treatment group is 6.6 percentage points lower than in the control group.[21]

[21]It will turn out, however, that the regression results to come are not substantially affected by the inclusion or exclusion in the regressions of a large set of control variables (including the "male" indicator).

3.3 A simple model of borrower behavior

To study how dynamic incentives affect borrower behavior, we develop a simple model of risk-neutral agents that incorporates both moral hazard and adverse selection. By virtue of the experiment, the credit contract is kept fixed, so our goal here is not to solve for the optimal contract in the presence of both information asymmetries (*Guesnerie et al.* (1988) or *Chassagnon and Chiappori* (1997) for risk averse agents), but rather to derive the agents optimal behavior with and without dynamic incentives.

Agents (or farmers) decide how much to borrow for cash crop inputs and how much to invest. We assume that they do not have collateral or liquid assets, so the maximum they can invest in cash crop production is the loan amount.

We introduce the possibility of adverse selection by allowing farmers to differ in the probability p (unobserved by the lender) that cash crop production is successful. Production is given by $f_S(b)$ when successful and by $f_F(b)$ when it fails, which happens with probability $1 - p$. The amount b denotes total cash crop inputs invested. We assume that $f_j(b), j \in \{S, F\}$, satisfies the usual properties $f_j(0) = 0$, $f_j'(b) > 0$, and $f_j''(b) < 0$.

We model moral hazard by allowing borrowers to divert inputs instead of investing them in cash crop production. If they decide to divert, they earn q per unit of input diverted, which can be interpreted as the secondary market price for inputs or the expected return if these inputs are invested in another crop. Given the arrangement to buy the cash crop (paprika) in the experiment, we assume that the lender can only seize cash crop production but not the proceeds from diverted inputs. To simplify matters, we assume that the choice of diversion is binary, that is, either all or nothing is diverted.[22]

[22]One can extend the model to the case where diversion is a continuous variable but the intuition is already captured in the simpler version presented.

We consider first the case where identification of clients is not possible, so borrowers can obtain a fresh loan even if they have defaulted in the past by simply using a different identity. Lenders cannot use dynamic incentives and are thus forced to offer the same one season contract every period, as they cannot tailor the terms of the contract to individual credit histories. Though in practice loan officers may recognize clients by sight, loan officers may resign or be transferred and so the new loan officer will not know the clients. Even if loan officers remain on the job, clients could borrow from a different branch or from a different lender altogether.

We then consider the case with biometric technology which provides the lender the ability to use dynamic incentives by denying credit to past defaulters. In this situation, borrowers face a tradeoff between diverting inputs away from cash crop production but jeopardizing chances of a loan in the future versus ensuring repayment of the current loan and therefore securing a loan in the future.

In both cases, the credit contract offered by the lender is given by a loan amount b and gross interest rate R. We assume that the loan size b can take on two values, b_L and b_H where $b_L < b_H$.[23] We also assume that even when cash crop production fails, the borrower has enough funds to cover loan repayment provided that the small amount b_L is borrowed and inputs are not diverted. More formally, $f_F(b_L) = b_L R$. This implies that if the borrower chooses to invest the large amount b_H in paprika production but the crop fails, then the borrower defaults because by concavity of $f_F(\cdot)$, $f_F(b_H) < b_H R$. Finally, we assume that if the crop succeeds, the large loan size yields higher farm profits than the smaller loan size. If we let $y_S(b_k) = f_S(b_k) - b_k R$, for $k \in \{L, H\}$ denote net profits from successful cash crop production, this assumption can be expressed as $y_S(b_H) > y_S(b_L)$.[24]

[23]This assumption is in accord with the actual details of the loan package, where the most important determinant of loan size is whether the farmer chooses to have the loan fund one vs. two bags of CAN fertilizer. We can think of b_H including two bags, and b_L only one.

[24]Using similar notation, the previous assumption implies that when the crop fails, farm profits are larger under the smaller loan size: $y_F(b_H) < y_F(b_L)$.

We assume that there are two periods and no discounting, although the model could easily be extended to an infinite horizon setting with discounting. The timing within a period follows the set-up of the field experiment: the borrower first learns whether the lender can use dynamic incentives; then the borrower decides how much to borrow and whether to divert inputs; then paprika production takes place; the loan is repaid if sufficient funds are available and finally the borrower consumes any remaining income.

In what follows, we take the credit contract as given and characterize optimal borrower behavior with and without dynamic incentives. Then we briefly discuss the optimality of the credit contract and compare the predictions of the model to those of other models in the literature.

3.3.1 Borrower behavior without dynamic incentives

Since the lender offers the same contract in each period, lifetime optimization coincides with period-by-period optimization. In a given period, the borrower chooses how much to borrow b and whether to divert inputs D by solving the following problem:

$$v(p) = \max_{b \in \{b_L, b_H\}} \{max_{D\in\{0,1\}} Dqb_H + (1-D)py_S(b_H), \max_{D\in\{0,1\}} Dqb_L + (1-D)py_S(b_L)\}$$

The dependency of net income from borrowing v on p is made explicit. If the borrower diverts, consumption is qb because the bank cannot seize income, but if the borrower invests in paprika production, consumption only takes place when production is successful as the bank seizes all output if paprika production fails.

Now let p_D be the success probability that leaves a borrower with the larger loan size b_H indifferent between diverting the inputs or investing them in paprika production. More formally, $qb_H = p_D y_S(b_H)$ as plotted in Figure 3.2.

If $p < p_D$, the solution to the problem when dynamic incentives are absent is to always borrow the large amount b_H and to divert all inputs ($D = 1$). If $p \geq p_D$, the borrower also borrows the large amount b_H but does not divert and therefore repays with probability p. Expected net income in a period $v(p)$ is

$$v(p) = qb_H \text{ if } p < p_D \text{ and } v(p) = py_S(b_H) \text{ if } p \geq p_D \tag{3.1}$$

3.3.2 Borrower behavior with dynamic incentives

In this case, the lender will only provide credit in period two to borrowers that have successfully repaid in period one. Because there are only two periods, in the last period the lender cannot provide additional incentives to elicit repayment, so the optimization problem that borrowers face is the same as the period-by-period optimization when dynamic incentives were absent. Borrowers maximize their lifetime utility by solving the following problem in period one:

$$V(p) = \max_{b \in \{b_L, b_H\}} \{ \max_{D \in \{0,1\}} Dqb_H + (1-D)p[y_S(b_H) + v(p)] \max_{D \in \{0,1\}} Dqb_L + (1-D)py_S(b_L) + v(p) \}$$

where again the dependency of V and v on p is made explicit. Net income v in period two is derived in equation (3.1). If the lower amount b_L is chosen, the borrower can always repay the loan and so net income from borrowing $v(p)$ in period two is assured. If, on the other hand, the higher amount b_H is chosen, then the borrower will obtain $v(p)$ in period two only if there is no diversion ($D = 0$) and paprika production is successful in period one. Income from not borrowing is normalized to zero.

It is easy to see that with dynamic incentives, diversion of inputs in the first period is never optimal. A borrower with a high probability of success $p \geq p_D$ would not divert in the absence of penalties, so he would certainly not do it when the lender can impose penalties. More formally, because $py_S(b_H) > qb_H$ if $p \geq p_D$, it follows that

$p[y_S(b_H) + v(p)] > qb_H$ since $v(p) > 0$.

When $p < p_D$, borrowers choose to divert in the absence of dynamic incentives. When dynamic incentives are in place, they can increase lifetime utility by choosing the lower amount in the first period. They then secure a loan in the second period which can then be diverted to achieve the same utility as if they had diverted in the first period. In addition, if cash crop production succeeds, then they also consume in the first period.[25]

We now study the choice of loan amount in the first period. Let p_{B0} be the probability of success that leaves a borrower with success probability $p \geq p_D$ indifferent between the two loan amounts. If success probability is such that $p_D < p < p_{B0}$, then the borrower chooses b_L to ensure loan repayment, but if the probability is high enough, so that $p_D < p_{B0} < p$ he then chooses b_H. The subscript 0 denotes the fact that in the absence of dynamic incentives the borrower would not divert because $p \geq p_D$. Probability p_{B0} can be written as

$$p_{B0} = \frac{y_S(b_L)}{y_S(b_H)} \tag{3.2}$$

Now let p_{B1} be analogous to p_{B0} for borrowers with success probability $p < p_D$. Here the subscript 1 indicates that the borrower would divert in the absence of dynamic incentives. If success probability satisfies $p < p_{B1} < p_D$, the borrower will choose the smaller loan amount b_L and if $p_{B1} < p < p_D$ the larger amount b_H. It is easy to show that p_{B1} satisfies

$$qb_H(1 - p_{B1}) = p_{B1}[y_S(b_H) - y_S(b_L)] \tag{3.3}$$

[25] While this result is immediate without discounting, it can be obtained with discounting provided the discount rate is low enough.

or, after some algebra and substitutions,

$$p_{B1} = \frac{p_D}{p_D + 1 - p_{B0}} \tag{3.4}$$

As it turns out, depending on the magnitude of $y_S(b_L)$, $y_S(b_H)$, and qb_H only $p_D > p_{B1}$ or $p_D < p_{B0}$ will hold, because $p_D > p_{B1}$ is true if and only if $p_D > p_{B0}$.[26] So either p_{B0} or p_{B1} is relevant. There are three cases, which we label (i), (ii), and (iii), distinguished by the size of the gains from input diversion (qb_H) relative to those from successful cash crop production, $y_S(b_H)$ and $y_S(b_L)$.

The first case is where (i) $qb_H > y_S(b_H)$, in which the gains from diversion are higher than the gains from cash crop production even when the high loan amount is taken and production is successful. In this case, $p_D > 1 > p_{B1} > p_{B0}$ and p_{B0} becomes irrelevant because $p_D < p_{B0}$ is violated. Intuitively, $p_D > 1$ means that there are no borrowers who would repay without dynamic incentives, because the gains from diversion are higher than the gains from cash crop production even for borrowers with the highest success probabilities; p_{B0} is irrelevant because there are no farmers for whom $p > p_D$. In the first period with dynamic incentives, borrowers with $p \geq p_{B1}$ take the larger loan and those for whom $p < p_{B1}$ take the smaller loan size.

The second and probably most interesting case is where (ii) $y_S(b_H) > qb_H > y_S(b_L)$, in which the gains from diversion (relative to cash crop production) are intermediate. In this case, in the absence of dynamic incentives, some borrowers (those with highest success probabilities, for whom $p > p_D$) will choose to produce rather than divert, which others with lower success probabilities will divert rather than produce. In this case we have $1 > p_D > p_{B1} > p_{B0}$,[27] and so p_{B0} is irrelevant (those with $p > p_D$ always choose the larger loan in the first period). In the first period with

[26] This is easy to see using the expression for p_{B1} derived in equation (3.4).

[27] To see this, divide inequalities in (ii) by $y_S(b_H)$ and recall $qb_H = p_D y_S(b_H)$ and expression (3.4).

dynamic incentives, borrowers with $p \geq p_{B1}$ take the larger loan and those for whom $p < p_{B1}$ take the smaller loan size.

The third case is where (iii) $y_S(b_L) > qb_H$, in which the gains from diversion are small relative to the gains from successful cash crop production, even when the small loan size is taken. Here, $1 > p_{B0} > p_{B1} > p_D$ so that p_{B1} now becomes irrelevant (because all individuals with $p < p_D$ will take the smaller loan size in the first period with dynamic incentives). Now it is those borrowers for whom $p > p_D$ that show variation in loan size in the first period with dynamic incentives: those with $p \geq p_{B0}$ take the larger loan and those for whom $p < p_{B0}$ take the smaller loan size.

Figure 3.2 is drawn assuming Case (ii) holds. It plots p_{B0} and p_{B1}, and because $p_D > p_{B0}$, p_{B0} is irrelevant. Probability p_{B1} is shown as the intersection of the left hand side and right hand side of the equality in (3.3) above.

For each regime (with and without dynamic incentives), Table 3.3 reports the first period optimal choices of loan size and whether to divert as well as repayment rate as a function of the borrowers success probability.

Interestingly, dynamic incentives have different effects on the optimal choices of borrowers depending on their probability of success. For example, borrowers with relatively low probability of success are most affected by the introduction of dynamic incentives. They choose the higher loan amount and to divert it all without dynamic incentives but borrow the lower amount and invest it in cash crop production when dynamic incentives are introduced. As a result, their repayment rate changes from zero to one once incentives are introduced.

Borrowers with relatively high probability of success are the least affected, since they never divert inputs and always choose the higher loan amount, except for in Case (i) where they would divert without incentives and not divert with incentives.

Borrowers with an intermediate value of the probability of success will, upon introduction of dynamic incentives, change either the diversion or the loan size decisions depending on the parameter values and functional forms. In Case (ii) they always choose the higher loan amount but move from diversion to no diversion when incentives are introduced. In Case (iii), they never divert but incentives lead them to move from the higher to the lower loan amount.

3.3.3 Discussion

If the lender sets gross interest rate R to break even, and the individual probability of success $p \in [0, 1]$ is drawn from the density function $G(p)$, then R satisfies

$$ib_H = [1 - G(p_D)][E(p|p \geq p_D)Rb_H + (1 - E(p|p \geq p_D))f_F(b_H)] \tag{3.5}$$

where i is the deposit rate and $E(p|p \geq p_D) = \int_{p_D}^{1} pdG(p)$.

Notice that the bank breaks-even whenever $p_D < 1$, otherwise all borrowers would divert and the bank would be unable to collect repayment. As a result, there is no interest rate R such that case (i) considered before is an equilibrium.

Depending on the parameters, a separating equilibrium may exist where the lender maximizes borrower welfare subject to breaking even by offering a menu of loan sizes and gross interest rates. Borrowers with low probability of success p may either borrow the large amount and default or borrow the lower amount and produce (again depending on the parameters), borrowers with intermediate probability of success will borrow the lower amount and produce and borrowers with high probability of success will borrow the large amount and produce.

When dynamic incentives are introduced, the lender can follow a strategy similar to *Stiglitz and Weiss* (1983) or *Boot and Takor* (1994). In words, the lender could lower the interest rate associated with the lower loan size b_L in the second period below the per period break even interest rate (thereby making a loss) but raise it in the first period so as to satisfy the break even constraint intertemporally. This may be optimal because in the first period the borrower has the added incentive of the promise of a loan in the future, a loan that will be ever more attractive the lower is the interest rate charged.

If collateral was available, then a menu of interest rates and collateral could always be offered in both periods (*Bester* (1985)). But as *Boot and Takor* (1994) point out, dynamic incentives can be more efficient than static incentives like collateral. As in their model, the value of long-term contracting does not arise from the ability to learn the borrower type (in their model all agents are equal) nor from improved risk-sharing (in both models agents are risk neutral). Long term relations are valuable because the lender has the ability to punish defaulters and to reward good borrowers.

Because repayment is higher with dynamic incentives, lenders could lower the interest rate and as a result borrowers might borrow more. The lender should also be willing to extend more credit if dynamic incentives can be used. As a result, overall borrowing could increase, although borrowers with low probability of success may still borrow less to ensure future access to loans. This increase in borrowing is also predicted by the more macro literature that tries to explain the increase in personal bankruptcies over the last few decades as a result of improvements in information technology available to lenders for credit decisions (see for example *Livshits et al.* (2010); *Narajabad* (2010); and *Sanchez* (2009)).

The source of heterogeneity in the model is the probability of success p. If there was heterogeneity in the discount rate, then dynamic incentives would only be relevant for agents that are patient (i.e. with low enough discount rate). In this alternative model, if borrowers prefer to divert in the absence of dynamic incentives, repayment would be low without fingerprinting and would only increase for agents with low discount rate when fingerprinting is introduced.

In many multi-period models of limited commitment and asymmetric information, agents are not allowed to save because they could borrow and default and then live in autarky from reinvesting the savings (*Bulow and Rogoff* (1989)). In *Boot and Takor* (1994), the agent has no incentive to save because the long-term contract provides better-than-market interest rates. In this model without dynamic incentives, agents with high probability of success will not find it profitable to default and save for period 2 either, even if a savings technology were available at rate i. But if the probability is low enough, in particular if p is such that

$$p < \frac{(i-1)qb_H}{y_S(b_L)}$$

then agents would borrow the higher amount b_H in period one, divert and hence default and save it into period 2 to earn $i > 0$. When dynamic incentives are allowed, then the same argument of *Boot and Takor* (1994) applies and so agents would prefer to borrow again in the second period, even if savings technology were available.

3.4 Regression specification

Because the treatment is assigned randomly at the club level, its impact on the various outcomes of interest (say, repayment) can be estimated via the following regression equation:

$$Y_{ij} = \alpha + \beta B_j + \gamma X_{ij} + \epsilon_{ij} \tag{3.6}$$

where Y_{ij} = repayment outcome for individual i in club j (e.g., equal to 1 if repaying in full and on time, and 0 otherwise), B_j is biometric identification (1 if fingerprinted and 0 if not), and X_{ij} is a vector of club and individual farmer characteristics collected at baseline. ϵ_{ij} is a mean-zero error term. Treatment assignment at the club level creates spatial and other correlation among farmers within the same club, so standard errors must be clustered at the club level (*Moulton* (1986)). Inclusion of the vector X_{ij} of baseline characteristics can reduce standard errors by absorbing residual variation. In our case, we include the baseline characteristics reported in Table 3.1, as well as indicators for the two stratification variables (locality/EPA fixed effects and week of loan offer fixed effects) and all interactions between the dummy variables for locality and week of loan offer.

The coefficient β on the biometric treatment status indicator is the impact of being fingerprinted on the dependent variable of interest.

We also examine the interactions between the randomized treatment and a particular baseline characteristic: a measure of the ex-ante probability of repayment. Examining this dimension of heterogeneity is a test of the theoretical models prediction that the impact of dynamic incentives on repayment is negatively related with the ex-ante repayment rate (what the repayment rate would have been in the absence of dynamic incentives): borrowers who, without the dynamic incentive, would have had lower repayment will see their repayment rates rise more when the dynamic incentive is introduced.[28] To test this question, we estimate regression equations of the following form:

$$Y_{ij} = \alpha + \rho(B_j \times D_{ij}) + \beta B_j + \gamma X_{ij} + \epsilon_{ij} \tag{3.7}$$

D_{ij} is a variable representing the individuals predicted likelihood of repayment (its main effect is included in the vector X_{ij}). The coefficient ρ on the interaction term $B_j \times D_{ij}$ reveals the extent to which the impact of biometric identification on repayment varies according to the borrowers predicted repayment.

To implement equation (3.7) examining heterogeneity in the effect of fingerprinting, we construct an index of predicted repayment. This involves creating what is essentially a credit score" for each borrower in the sample on the basis of the relationship between baseline characteristics (some of which may not be observable to

[28]While in the model the single dimension of borrower heterogeneity is the probability of success, p, we have no way to estimate this directly for our full borrowing sample. Note that the repayment rate is monotonic in p, making it a good proxy for p. While in principle one could apply the procedure in Appendix F with crop output as the dependent variable, in practice this would limit us because crop output is only observed in the smaller subsample of borrowers (N=520). The repayment rate, on the other hand, comes from administrative data and so is available for the entire borrowing sample.

the lender) and repayment in the control (non-fingerprinted) group. (See Appendix F for details on the construction of the predicted repayment variable. Appendix Table G.1 presents the auxiliary regression results used in construction of the predicted repayment variable.) This index is either interacted linearly with the treatment indicator, or it is converted into indicators for quintiles of the distribution of predicted repayment in the absence of fingerprinting and then interacted with the treatment indicator.[29] In all regression results where the treatment indicator is interacted with predicted repayment, we report bootstrapped standard errors because the predicted repayment variable is a generated regressor.[30]

3.5 Empirical results: impacts of fingerprinting

This section presents our experimental evidence on the impacts of fingerprinting on a variety of inter-related outcomes. We examine impacts on loan approval and borrowing decisions, on repayment outcomes, and on intermediate farmer actions and outcomes that may ultimately affect repayment.

Tables 3.4 through 3.8 will present regression results from estimation of equations (3.6) and (3.7) in a similar format. In each table, each column will present regression results for a given dependent variable. Panel A will present the coefficient on treatment (fingerprint) status from estimation of equation (3.6).

Then, to examine heterogeneity in the effect of fingerprinting, Panels B and C will present results from estimation of versions of equation (3.7) where fingerprinting is interacted linearly with predicted repayment (Panel B) or with dummy variables for

[29] In other results that are analogous to the analysis of Table 3.2 (available from authors on request), we show that there is balance in key baseline characteristics across treatment and control observations within each quintile of predicted repayment.

[30] We calculate standard errors for regressions in the form of equation (3.7) from 200 bootstrap replications. In each replication, we re-sample borrowing clubs from our original data (which preserves the original club-level clustering), compute predicted repayment based on the new sample, and re-run the regression in question using the new value of predicted repayment for that replication. See **?** for details.

quintiles of predicted repayment (Panel C). In both Panels B and C the respective main effects of the predicted repayment variables are also included in the regression (but for brevity the coefficients on the predicted repayment main effects will not be presented). In Panel C, the main effect of fingerprinting is not included in the regression, to allow each of the five quintile indicators to be interacted with the indicator for fingerprinting in the regression. Therefore, in Panel C the coefficient on each fingerprint-quintile interaction should be interpreted as the impact of fingerprinting on borrowers in that quintile, compared to control group borrowers in that same quintile.

Finally, in Tables 3.4 through 3.8 the mean of the dependent variable in a given column, for the overall sample as well for each quintile of predicted repayment separately, are reported at the bottom of each table.

3.5.1 Loan approval, take-up, and amount borrowed

The first key question to ask is whether fingerprinted farmers were more likely to have their loans approved by the lender, or were more likely to take out loans, compared to the control group. This question is important because the degree of selectivity in the borrower pool induced by fingerprinting status affects interpretation of any effects on repayment and other outcomes. Although loan officers were told which clubs had been fingerprinted in September 2007 when loan applications were due, they do not appear to have used this information in their loan approval decisions. Since biometric technology can be seen as a substitute for loan officer effort, one would expect loan officers to have better knowledge about non-fingerprinted clubs. However, this is not what we find.

Appendix Table G.2 combines the reports from all loan officers collected in August 2008 as well as borrower responses in the August 2008 follow-up survey. Loan officers were first asked about the specific treatment status of five clubs randomly selected from the sample of clubs for which they were responsible. They were then asked whether they knew the secretary or president of the club and finally they were asked to estimate the number of loans given out in each club. The first row of the table shows that loan officers had very little knowledge about the actual treatment status of clubs. Only 54 percent of the fingerprinted clubs are reported correctly as being fingerprinted and an even lower 22 percent of non-fingerprinted clubs are reported correctly as such. Pure guesswork would yield an accuracy rate of 50 percent. This evidence alone suggests that loan officers did not take into account treatment status in their interactions with the clubs.

Loan officers know club officers roughly half of the time, and on average misreport the number of loans disbursed to a club by 1.5 loans. More importantly, there are no statistical differences in the reporting accuracy of fingerprinted clubs compared to non-fingerprinted ones. Borrower reports in the last three rows of the table paint a similar picture. Loan officers are no more likely to visit non-fingerprinted clubs to collect repayment compared to fingerprinted clubs, and as a result, members of non-fingerprinted clubs report talking the same number of times to loan officers as do members of fingerprinted clubs. Finally, they all report finding it relatively easy to contact the loan officer.

The evidence in the table indicates that loan officers did not respond to the treatment. Therefore, any impacts of the treatment should be interpreted as emerging solely from borrowers responses to being fingerprinted.

Because loan officers did not take treatment status into account, it is not surprising that fingerprinting had no effect on loan approval. We also find no effect on loan-take-up by borrowers, perhaps because clubs were formed with the expectation of credit availability and fingerprinting did not act as a strong enough deterrent to borrowing to affect farmers decisions at the extensive margin. Columns 1 and 2 of Table 3.4 present results from estimation of equations (3.6) and (3.7) for the full baseline sample where the dependent variables are, respectively, an indicator for the lenders approving the loan for the given farmer (mean 0.63), and an indicator for the farmer ultimately taking out the loan (mean 0.35).[31]

There is no evidence that the rate of loan approval or take-up differs substantially across the treatment and control groups on average: the coefficient on fingerprinting is not statistically different from zero in either columns 1 or 2, Panel A.

There is also no indication of selectivity in the resulting borrowing pool across subgroups of borrowers with different levels of predicted repayment. The coefficient on the interaction of fingerprinting with predicted repayment is not statistically significantly different from zero in either columns 1 or 2 of Panel B. When looking at interactions with quintiles of predicted repayment (Panel C), while the fingerprint-quintile 2 interaction is positive and significantly different from zero at the 10% level in the loan approval regression, none of the interaction terms with fingerprinting are significantly different from zero in the loan take-up regression.

While there is no indication that the pool of ultimate borrowers was itself substantially affected by fingerprinting, it does appear that conditional on borrowing fingerprinted borrowers took out smaller loans. In Column 3 of Table 3.4, the dependent variable is the total amount borrowed in Malawi kwacha. Panel A indicates that loans of fingerprinted borrowers were MK 697 smaller than loans in the control group on average, a difference that is significant at the 10% level.

Inspecting the coefficients on the interactions of fingerprinting with predicted repayment, it appears that this effect is confined exclusively to borrowers in the lowest quintile of expected repayment. Differences between fingerprinted and non fingerprinted borrowers are small and not significant in quintiles two through four, but in quintile one, where fingerprinted borrowers take out loans that are smaller by MK

[31]Not all farmers who were approved for the loan ended up taking out the loan. Anecdotal evidence indicates that a substantial fraction of non-take-up among approved borrowers resulted when borrowers failed to raise the required deposit (amounting to 15% of the loan amount).

2,722 (roughly US$19) than those in the corresponding quintile in the control group, the difference is marginally significant (the t-statistic is 1.63). This result is in accord with the theoretical models prediction that the "bad" borrowers (those whose repayment rates would be lowest in the absence of dynamic incentives) will respond to the imposition of a dynamic incentive by voluntarily reducing their loan sizes. We view this result voluntarily lower borrowing amounts on the part of fingerprinted borrowers in the lowest quintile as evidence that fingerprinting reduces adverse selection in the credit market, albeit on a different margin than is usually discussed in the credit context.

The existing literature tends to emphasize that improved enforcement should lead low-quality borrowers to be excluded from borrowing entirely in other words, the improvement of the borrower pool operates on the extensive margin of borrowing. Our result here that low-quality borrowers (those in the lowest quintile of predicted repayment) voluntarily take out smaller loans leads the overall loan pool in money terms to be less weighted towards the low-quality borrowers, but in this case the improvement in the borrowing pool operates on the intensive margin of borrowing, rather than the extensive margin.

Interpretation of subsequent differences in the repayment rates (discussed below) should keep this result in mind. Improvements in repayment among fingerprinted borrowers (particularly among those in the lowest quintile) may in part result from their decisions to take out smaller loans at the very outset of the lending process and improve their eventual likelihood of repayment.

3.5.2 Loan repayment

How did fingerprinting affect ultimate loan repayment? Columns 1-3 of Table 3.5 present estimated effects of fingerprinting for the loan recipient sample on three outcomes: outstanding balance (in Malawi kwacha), fraction of loan paid , and an indicator for whether the loan is fully paid, all by September 30, 2008 (the official due date of the loan, after which the loan is officially past due). The next three columns (columns 4-6) are similar, but the three variables refer to eventual" repayment as of the end of November 2008. The lender makes no attempt to collect past-due loans after November of each agricultural loan cycle, so the eventual repayment variables represent the final repayment status on these loans.

Results for all loan repayment outcomes are similar: fingerprinting improves loan repayment, in particular for borrowers expected ex ante to have poorer repayment performance. Coefficients in Panel A indicate that fingerprinted borrowers have lower outstanding balances, higher fractions paid, and are more likely to be fully paid on-time as well as eventually (and the coefficient in the regression for fraction paid on-time is statistically significant at the 10% level).

In Panel B, the fingerprinting-predicted repayment interaction term is statistically significantly different from zero (at least at the 5% level) in all regressions. The effect of fingerprinting on repayment is larger the lower is the borrowers ex ante likelihood of repayment. In Panel C, it is evident that the effect of fingerprinting is isolated in the lowest quintile of expected repayment, with coefficients on the fingerprint-quintile 1 interaction all being statistically significantly different from zero at the 5% or 1% level and indicating beneficial effects of fingerprinting on repayment (lower outstanding balances, higher fraction paid, and higher likelihood of full repayment). Coefficients on other fingerprint-quintile interactions are all smaller in magnitude and not statistically significantly different from zero (with the exception of the negative coefficient on the fingerprint-quintile 5 interaction for fraction paid, which is odd and may simply be due to sampling variation).

The magnitudes of the repayment effect found for the lowest predicted-repayment quintile are large. The MK7,202.65 effect on eventual outstanding balance amounts to 40% of the average loan size for borrowers in the lowest predicted-repayment quintile.

While outstanding balance should mechanically be lower due to the lower loan size in the lowest predicted-repayment quintile, the effect is almost three times the size of the reduction in loan size, so by itself lower loan size cannot explain the treatment effect on repayment. The 31.7 percentage point increase in eventual fraction paid and the 39.6 percentage point increase in the likelihood of being eventually fully paid are also large relative to bottom quintile percentages of 81% and 68% respectively.

3.5.3 Intermediate outcomes that may affect repayment

In this section we examine decisions that farmers make throughout the planting and harvest season that may contribute to higher repayment among fingerprinted farmers. The dependent variables in the remaining results tables (Tables 3.6-3.8) are available from a smaller subset of loan recipients (N=520) who were successfully interviewed in the August 2008 follow-up survey round. To help rule out the possibility that selection into the 520-observation August 2008 follow-up survey sample might bias the regression results for that sample, Column 2 of Appendix Table G.3 examines selection of loan recipients into the follow-up survey sample. The regressions are analogous in structure to those in the main results tables (Panels A, B, and C), and the dependent variable is a dummy variable for attrition from the baseline (September 2007) to the August 2008 survey. There is no evidence of selective attrition related to treatment status: in no case is fingerprinting or fingerprinting interacted with predicted repayment statistically significantly associated with attrition from the survey.

Appendix Table G.4 presents regression results for repayment outcomes that are analogous to those in Table 3.5, but where the sample is restricted to this 520-observation sample. The results confirm that the repayment results in the 520-observation sample are very similar to those in the overall loan recipient sample, in terms of both magnitudes of effects and statistical significance levels.

3.5.3.1 Land area allocated to various crops

One of the first decisions that farmers make in any planting season (which typically starts in November and December) is the proportion of land allocated to different crops. Table 3.6 examines the average and heterogeneous impact of fingerprinting on land allocation; the dependent variables across columns are fraction of land used in maize (column 1), 7 cash crops (columns 2-8), and all cash crops combined (column 9).

Why might land allocation to different crops respond to fingerprinting? As discussed in the context of the theoretical model, non-production of paprika is a form of moral hazard, since the lender can only feasibly seize paprika output (in collaboration with the paprika buyer, Cheetah Paprika) and not other types of crop output. By not producing paprika (or producing less), the borrower is better able to avoid repayment on the loan. Therefore, by improving the lenders dynamic incentives, fingerprinting may discourage such diversion of inputs and land to other crops, as farmers face increased incentives to generate cash profits that are sufficient for loan repayment.

While none of the effects of fingerprinting in Table 3.6 (either overall in Panel A or in interaction with predicted repayment in Panels B and C) are statistically significant at conventional levels, there is suggestive evidence that there is an impact of fingerprinting on land allocation for borrowers in the first predicted-repayment quintile. In this group, the effect of fingerprinting on land allocated to paprika (column 5, first row of Panel C) is marginally significant (with a t-statistic of 1.63) and positive, indicating that fingerprinting leads farmers to allocate 8.3 percentage points more land to paprika. This effect is roughly half the size of the paprika land allocation in the lowest quintile of predicted repayment.

It is worth considering that the effect on land allocated to paprika may be smaller than it might be otherwise because farmers began preparing and allocating land earlier in the agricultural season than our treatment. If land is less easily reallocated than other inputs from one crop to another, then we would anticipate smaller short run effects on land allocation than on the use of inputs such as fertilizer and chemicals (to which we now turn). In the long run, when farmers incorporate the additional cost of default due to fingerprinting into their agricultural planning earlier in the season, we might find larger impacts on land allocation.

3.5.3.2 Inputs used on paprika

After allocating land to different crops, the other major farming decision made by farmers is input application. Non-application of inputs on the paprika crop facilitates default on the loan and is therefore another form of moral hazard, again since only paprika output can feasibly be seized by the lender.

It is worth keeping in mind that input application takes place later in the agricultural cycle than land allocation, and agricultural inputs are more fungible than land. Also, inputs are added multiple times throughout the season, so farmers can incorporate new information about the cost of default into their use of inputs but cannot change land allocation after planting. Thus, we may expect use of inputs to respond more quickly to the introduction of fingerprinting than would allocation of land.

Table 3.7 examines the effect of fingerprinting on the use of inputs on the paprika crop. The dependent variables in the first five columns (all denominated in Malawi kwacha) are applications of seeds, fertilizer, chemicals, man-days (hired labor), and all inputs together. Columns 6 and 7 look at, respectively, manure application (denominated in kilograms because this input is typically produced at home and not purchased) and the number of times farmers weeded the paprika plot. We view the manure and weeding dependent variables as more purely capturing labor effort exerted on the paprika crop, while the other dependent variables capture both labor effort and financial resources expended.

The results for paid inputs (columns 1-5) indicate that particularly for farmers with lower likelihood of repayment fingerprinting leads to higher application of inputs on the paprika crop. In Panel B, the coefficients on the fingerprint-predicted repayment interaction are all negative in sign, and the effects on the use of fertilizer and paid inputs in aggregate are statistically significantly different from zero. In Panel C, the coefficient on the fingerprint-quintile 1 interaction is positive and significantly different from zero at the 5% confidence level for spending on seeds and is marginally significant for spending on fertilizer (t-statistic 1.44) and for all paid inputs (t-statistic 1.55). The negative and significant impact on use of paid labor in the fourth quintile is puzzling and may be attributable to sampling variation.

Results for inputs not purchased in the market are either nonexistent or ambiguous. No coefficient is statistically significantly different from zero in the regressions for manure (column 6) or times weeding (column 7).

It is worth asking whether the impact of fingerprinting seen in Table 3.7 means that farmers are less likely to divert input to use on other crops, or, alternatively, less likely to sell or barter the inputs for their market value. To address this, we examined the impact of fingerprinting on use of inputs on all crops combined. Results were very similar to Table 3.7s results for input use on the paprika crop only (results are available from the authors on request). This suggests that in the absence of fingerprinting, inputs were not used on other non-paprika crops. (If fingerprinting simply led inputs to be substituted away from non-paprika crops to paprika, the estimated impact of fingerprinting on input use on all crops would be zero.) It therefore seems most likely that fingerprinting made farmers less likely to dispose of the inputs via sale or barter.

In sum: for borrowers with a lower likelihood of repayment, fingerprinting leads to increased use of marketable inputs in growing paprika. While this effect is at best only marginally significant for borrowers in the lowest predicted repayment quintile, the magnitudes in that quintile are substantial. For the lowest predicted-repayment subgroup, fingerprinted farmers used MK6,540 more paid inputs in total, which is substantial compared to the mean in the lowest predicted-repayment subgroup of MK7,440.

3.5.3.3 Farm profits

Given these effects of fingerprinting on intermediate farming decisions such as land allocation and input use, what is the effect on agricultural revenue and profits? Columns 1-3 of Table 3.8 present regression results where the dependent variables are market crop sales, the value of unsold crops, and profits (market sales plus value of unsold crops minus value of inputs used), all denominated in Malawi kwacha. The magnitudes of the overall impacts of fingerprinting on value of sales, unsold harvest, and total profits (Panel A), and in the bottom two quintiles (Panel C) are large and positive, but the effects are imprecisely estimated and none are statistically significantly different from zero.

To help deal with the problem of outliers in the profit figures, column 4 presents regression results where the dependent variable is the natural log of agricultural profits.[32] The effect of fingerprinting in the bottom quintile of predicted repayment is positive but not statistically significant (t-statistic 1.11).

In sum, then, it remains possible that increased use of paid inputs led ultimately to higher revenue and profits among fingerprinted farmers in our sample, but the imprecision of the estimates prevents us from making strong statements about the impact of fingerprinting on farm profits.

[32]For seven (7) observations profits are zero or negative, and in these cases ln(profits) is replaced by 0. These observations are not driving the results, as results are essentially identical when simply excluding these 7 observations from the regression.

3.6 Discussion and additional analyses

In sum, the results indicate that for the lowest predicted-repayment quintile, fingerprinting leads to substantially higher loan repayment. In seeking explanations for this result, we have provided evidence that for this subgroup fingerprinting leads farmers to take out smaller loans, devote more land to paprika, and apply more inputs on paprika.

We view these results so far as indicating that for the farmers with the lowest ex ante likelihood of repaying their loans fingerprinting leads to reductions in adverse selection and ex-ante moral hazard. The reduction in adverse selection (a reduction in the riskiness of the loan pool) comes about not via the extensive margin of loan approval and take-up, but through farmers decisions to take out smaller loans if they are fingerprinted (the intensive margin of loan take-up).

Ex-ante moral hazard is the problem that borrower behavior that is unobserved to the lender may be detrimental for repayment. We interpret changes in intermediate outcomes and behaviors such as increased land use and input application for paprika as reductions in ex-ante moral hazard. We believe that the most likely scenario is that in the absence of fingerprinting, borrowers in the lowest predicted-repayment subgroup were not using the paprika inputs received as part of the loan for paprika production. Rather, they are most likely to have sold them in the market or bartered them away. Then when such borrowers were fingerprinted, they became more likely to use the inputs as intended, expanding land allocated to paprika and using the inputs on that crop as the loan required. Below we provide a test of whether increased repayment as a result of fingerprinting in part reflects reductions in ex-post moral hazard. We also report results of a test of the positive correlation property that reveals the presence of asymmetric information (*Chiappori and Salanie* (2003) and *Chiappori et al.* (2006)). Finally, at the end of this section, we summarize the results of additional robustness checks that are presented in greater detail in Appendix F.

3.6.1 Evidence for a reduction in ex-post moral hazard

Reductions in ex-ante moral hazard may help encourage higher loan repayment by improving farm output so that farmers have higher incomes with which to make loan repayments. Reductions in adverse selection reduced loan sizes for the "bad" borrowers also help increase repayment performance. But a question that remains is whether any of the increase in repayment is due to reductions in ex-post moral hazard. In other words, are there reductions in strategic or opportunistic default by borrowers, holding constant loan size and farm profits?

We investigate this by running regressions where repayment outcomes are the dependent variables, but where we include as independent variables in the regression controls for agricultural profits and the total originally borrowed. Results are reported in Table 3.9.[33] The profits and total borrowed variables are flexibly specified as indicators for the borrower being in the 1st through 10th decile of the distribution of the variable (one indicator is excluded in each resulting group of 10 indicators, so there so there are 18 additional variables in each regression.)

We cannot reject the hypothesis that fingerprinting has no effect on eventual repayment (columns 4-6) once we control for agricultural profits and original loan size. Coefficient estimates that were previously statistically significant (in Appendix Table G.4) are now uniformly smaller in magnitude and not statistically significantly different from zero. Indeed, the previously significant coefficients on the fingerprint $\times$ quintile 1 interaction across the columns are roughly cut in half. Results are similar for repayment by the due date (columns 1-3), with the exception of the regression for "Balance, Sept. 30" where the linear interaction term and the interaction term with quartile 1 of predicted repayment remain statistically significant at the 5% and 10% levels, respectively. Even in this latter cases, however, the coefficient magnitudes are

[33] We limit ourselves to the 520-observation sample because of the need to control for profits, which was only observed among those in the August 2008 survey. These results should therefore be compared with Appendix Table G.4, which is also for the 520-observation sample.

reduced substantially vis-a-vis the corresponding estimates in Appendix Table G.4.

All told, we view these results as providing no strong support for the idea that a reduction in ex-post moral hazard increases in repayment even conditional on amount borrowed and agricultural profits is also an important contributor to the increased repayment we observe among fingerprinted farmers in the lowest predicted-repayment quintile.

3.6.2 Test of the positive correlation property

Following several recent articles that use data from insurance markets to test for the presence of asymmetric information (*Chiappori and Salanie* (2003) and *Chiappori et al.* (2006)), the predictions of the theoretical model of Section 3.3 can be used to perform a similar test. In the insurance market context, many models of adverse selection and possibly moral hazard that assume competitive insurance markets predict a positive correlation between coverage and the probability of the event insured, conditioning on the information available to the insurer. In our context, the test involves a positive correlation between loan size and default.

In order to test this prediction, multiple loan contracts must coexist in equilibrium, but according to the model (see Table 3.3), all agents should borrow the high amount when dynamic incentives cannot be used, and so there should be no correlation. With dynamic incentives however, both high and low loan sizes (b_L and b_H) will be taken and so the correlation can be tested. Using data on the loan size and default at maturity date, we find, as expected, no correlation for borrowers in the control group (t-stat = 1.13), but find a strong positive correlation in the treatment group (t-stat=3.30). In the treatment group, a MK1,000 increase in the loan amount is associated with a decrease in the probability of default (not being fully paid at the loan due date) of roughly 3 percentage points.

3.6.3 Additional robustness checks

Here we summarize additional robustness checks that are presented in greater detail in Appendix G. Appendix G also includes regression tables for all results discussed below.

3.6.3.1 Impact of fingerprinting in the full sample

Most results presented so far are for the subsample of farmers who took out a loan. We have argued that when restricting ourselves to this subsample, estimated treatment effects are not confounded by selection concerns because treatment has no statistically significant effect on selection into borrowing, either on average or in interaction with predicted repayment (Table 3.4, column 2). That said, one may raise a concern about statistical power: 95% confidence intervals around the point estimates in Table 3.4, column 2 admit non-negligible effects of treatment on selection into borrowing. The concern would be that there was in fact selection into borrowing in response to fingerprinting, which would cloud the interpretation of our results. For example, one might worry that that fingerprinting led borrowers in quintile 1 of predicted repayment to be on average different from control group borrowers in quintile 1 (along various observed and unobserved dimensions) in ways that make them more likely to repay, to devote land to paprika, and to use fertilizer on paprika.

Analyses of the full sample of farmers, without restricting the sample only to borrowers, can help address such concerns about selection bias. Estimated effects of treatment (and interactions with predicted repayment) would then represent effects of being fingerprinted on average across treated individuals, whether or not the individual took out a loan. While such an analysis makes little sense for outcomes specific to loans such as repayment (as in the outcomes of Table 5), we carry out this analysis for the other examined variables from the August 2008 follow-up survey, namely land use, input use, and profits (the outcomes in Tables 3.6, 3.7, and 3.8 respectively).

As it turns out, full-sample regression results are very similar to those from the borrower-only regressions. The general pattern is for coefficients that were significant before to remain statistically significant, but to be only around half the magnitude of the coefficients in the borrowing sample regressions. This reduction in coefficient magnitude is consistent with effect sizes in the full sample representing a weighted average of no effects for nonborrowers and nonzero effects for borrowers (slightly less than half of individuals in the full sample are borrowers). All in all, we conclude that selection into borrowing is not driving the treatment effect estimates of Tables 3.6, 3.7, and 3.8.

3.6.3.2 Results with "simple" predicted repayment regression

Results discussed so far examining heterogeneity in treatment effects construct the predicted repayment variable using the regression in column 3 of Appendix Table G.1. The right-hand-side of this regression contains farmer-level characteristics as well as all interactions between locality and week of initial loan offer fixed effects.

Because the baseline farmer-level characteristics listed in Appendix Table G.1 are the most readily interpretable, we check the robustness of the results to constructing predicted repayment using only baseline farmer-level characteristics. The alternative predicted repayment regression is that of column 3 of Appendix Table G.1, except that (locality)×(week of initial loan offer) fixed effects are dropped. This regression is then used to predict repayment for the full sample, and the predicted repayment variable is interacted with treatment to examine heterogeneity in the treatment effect.

Regression results are very similar when using this simpler index of predicted repayment. Overall, the general conclusion stands: fingerprinting has more substantial effects on repayment and activities on the farm for individuals with lower predicted repayment, even when repayment is predicted using only a restricted set of baseline farmer-level variables.

3.6.3.3 Results with predicted repayment coefficients obtained from partition of control group

In heterogeneous treatment effect results presented so far, there may be a concern that for idiosyncratic reasons control farmers in some geographic areas could have unusually low repayment rates compared to treatment farmers in the same areas. If this were the case, then the main analyses we have conducted so far might mechanically find a positive effect of treatment in cohorts where control group farmers had idiosyncratically low repayment rates.

We address this type of concern in two ways. First, we point to the robustness check just described above, where we find that results are very similar when the predicted repayment index is estimated without locality$\times$(week of initial loan offer) fixed effects. These results reveal that the patterns of treatment effect heterogeneity we emphasize are not simply an artifact of inclusion of the locality fixed effects (and interactions with week of initial loan offer) in the predicted repayment regression.

Second, we gauge the extent to which our main results diverge from those of an alternative approach that involves partitioning the control group into two parts: one part used to generate coefficients in the predicted repayment regression, and the other part used as a counterfactual for the treatment group in the main regressions. Because observations used to generate coefficients in the auxiliary predicted repayment regression are not then used as counterfactuals for the treatment observations, this approach avoids the possibility that our results arise mechanically from overfitting the repayment model.

Due to sampling variation, different randomly-determined partitions of the control group will yield different results, so we conduct this exercise 1,000 times and then examine the distribution of the regression coefficients generated. We focus our attention on coefficients on the interaction between the treatment indicator and the indicator for quintile 1 of predicted repayment (in Panel C) for the dependent variables of

Tables 3.4 to 3.8.

We find that in all cases the quintile 1 interaction term coefficient in Tables 3.4 to 3.8 falls within the 95 percent confidence interval of the coefficients generated in the partitioning exercise. Furthermore, whenever the interaction term coefficient is statistically significantly different from zero in Tables 3.4 to 3.8, the 95 percent confidence interval of the coefficients generated in the partitioning exercise does not include a coefficient of zero or of the opposite sign.

We therefore conclude that our main results are not mechanically driven by idiosyncratically low repayment among some control farmers in certain localities.

3.7 Benefit-cost analysis

The analysis so far has estimated the gains to the financial institution (MRFC) from using fingerprinting to identify new borrowers as part of the process of loan screening. These gains need to be weighed against the costs of fingerprinting. In this section, we present a benefit-cost analysis of biometric fingerprinting of borrowers. The analysis is most valid for institutions similar in characteristics to those of our partner institution, MRFC, but we have made the elements of the calculation very transparent so that they can be easily modified for other institutions with different characteristics.

The benefit-cost calculation is presented in Table 3.10. The uppermost section of the table is the calculation of benefits per individual fingerprinted. At the suggestion of MRFC, we assume that all new loan applicants are fingerprinted, and that 50% of applicants are approved for loans. Based on our experimental results we assume that the increase in repayment due to fingerprinting is confined to the first quintile (20% of borrowers), and that for this subgroup fingerprinting causes an increase in repayment amounting to 31.7% of the loan balance (from column 5 of Table 3.5). We assume that the total amount to be repaid is MK15,000 on average. Total benefit per individual fingerprinted is therefore MK475.50 (US$3.28).

The next section of the table calculates cost per individual fingerprinted. There are three general types of costs. First, equipment costs need to be amortized across farmers fingerprinted. We assume each equipment unit (a laptop computer and external fingerprint scanner) costs MK101,500,[34] and is amortized over three years, for annual cost of each equipment package of MK33,833. Twelve of these equipment packages (two for each of six branches) will be required to fingerprint MRFCs borrowers throughout the country. With an estimated 5,000 new loan applicants per year, each of these equipment units will be used to fingerprint 417 farmers on average. The equipment cost per farmer fingerprinted is therefore MK81.20.

The second type of cost is loan officer time. We estimate that it takes 5 minutes to fingerprint a customer and enter his or her personal information into the database. At a salary of MK40,000 per month and 173.2 work hours per month, this comes out to a cost of MK19.25 per customer fingerprinted.

The third type of cost is the transaction cost per fingerprint checked, MK108.75 (US$0.75). We assume here that MRFC hires a private firm to provide the fingerprint identification services, in which case the fingerprint database is stored on the firms server overseas and batches of fingerprints to be checked are sent electronically by MRFC to the firm during loan processing season. Lists of identified defaulters are sent back to MRFC with fast turnaround. In consultation with a U.S. private firm that provides such services, we were given a range of $0.03-$0.75 per fingerprint identification transaction. Per-fingerprint transaction costs are higher when the client has a relatively low number of transactions per year, and MRFCs 5,000 transactions per year is considered low, so we conservatively assume the transaction cost per fingerprint at the higher end of this range, $0.75 (MK108.75).

[34]This is the actual cost of each equipment unit we purchased for the project, which included a laptop computer ($480), an extra laptop battery ($120), a laptop carrying case ($20), and an external fingerprint scanner ($80).

Summing up these three types of costs, total cost per individual fingerprinted is MK209.20. The net benefit per individual fingerprinted is therefore MK266.30 (US$1.84), and the benefit-cost ratio is an attractive 2.27.[35]

For several reasons, this benefit-cost calculation is likely to be quite conservative. First of all, under reasonable circumstances some of the individual costs could be brought down considerably. The cost for equipment units could fall substantially if a fingerprinting function were integrated into equipment packages that had multiple functionalities, such as the hand-held computers that MRFC is considering providing for all of its loan officers. Transaction costs for fingerprint checking could fall due to volume discounts if the lending institution banded together with other lenders to channel all their fingerprint identification through a single service provider (in the context of a credit bureau, for example).

In addition, there are other benefits to the lending institution that this benefit-cost calculation is not capturing. The impact of fingerprinting on loan repayment may become larger in magnitude over time as the lenders threat of enforcement becomes more credible. We have also assumed that all the benefits come from fingerprinting new loan customers (the subject of this experiment), but there may also be increases in repayment among existing customers who are fingerprinted (on which this experiment does not shed light). Finally, there may be broader benefits that are not captured by the lending institution, such as increased income due to more intensive input

[35] An alternative is for a lending institution to purchase its own fingerprint matching software and do fingerprint identification in-house instead of subcontracting this function to an outside firm. This would eliminate the $0.75 (MK108.75) transaction cost per fingerprint checked. According to a U.S. fingerprint identification services firm we consulted, the initial fixed cost of installing an off-the-shelf fingerprint matching software system is in the range of $15,000 to $50,000 (depending on specifications), with an annual maintenance cost of 10-20% of the initial fixed cost. In addition, there would be personnel costs for staff to operate the system. Assuming an initial fixed cost of $15,000, maintenance cost of 10% of the original fixed cost, and an additional full-time staff member to run the system costing the same as a current MRFC loan officer, NPV is lower when fingerprint identification is done in-house than when this function is contracted out (which is why Table 3.10s calculation assumes contracting out). But with a high enough annual volume of transactions (perhaps in the context of a credit bureau in which many or all of Malawis lenders participate), in-house fingerprint identification could make economic sense.

application by fingerprinted farmers.[36]

3.8 Conclusion

We conducted a field experiment where we randomly selected a subset of potential loan applicants to be fingerprinted, which improved the effectiveness of dynamic repayment incentives for these individuals. For all the recent empirical work on microcredit markets in developing countries, to our knowledge this is the first randomized field experiment of its kind, and the first to shed light (thanks to a detailed follow-up survey of borrowers) on the specific behaviors germane to the presence of asymmetric information problems. Consistent with a simple model of asymmetric information in credit markets, we find heterogeneous effects of being fingerprinted, with the strongest effects among borrowers expected (ex ante) to have the worst repayment performance. Fingerprinting leads these "worst" borrowers to raise their repayment rates dramatically, partly as a result of voluntarily choosing lower loan sizes as well as devoting more agricultural inputs to the cash crop that the loan was intended to finance. The treatment-induced reduction in loan size represents a reduction in adverse selection, while the increased use of agricultural inputs on the cash crop represents a reduction in ex-ante moral hazard.

The short-term improvements in repayment estimated in this paper may indeed be smaller than the effects that would be found over a longer horizon. First of all, borrowers assessments of the effectiveness of the technology and the credibility of the threat to withhold credit would likely rise over time as they gained further exposure to the system, observed that their past credit performance was being correctly retrieved by the lender, and saw that credit history information was indeed being shared with other lenders. In addition, the lender should be able to selectively allocate credit to

[36] Unfortunately, our estimates of the impact of fingerprinting on profits are too imprecise to say whether profits definitely increased due to this intervention.

the pool of good-performing borrowers over time, further improving overall repayment performance of the borrowing pool. Finally, because there is less risk involved for the lender, the credit contract terms could be made more attractive to borrowers, which may further improve repayment.[37]

By revealing the presence of specific asymmetric information problems and the behaviors that result from them, this papers findings can help guide future theoretical work on rural credit markets. To be specific, models of credit markets in contexts similar to rural Malawi should allow for adverse selection on the intensive margin of loan take-up (i.e., the choice of loan size), as well as ex-ante moral hazard (actions during the production season that may affect farm profits). On the other hand, our results suggest that it may be less important for models to incorporate ex-post moral hazard (strategic or opportunistic default), since we find no evidence of it in this context.

Our results also have implications for microlending practitioners, by quantifying the benefits from exploiting a commercially-available technology to raise repayment rates. Beyond improving the profitability and financial sustainability of microlenders, increased adoption of fingerprinting (or other identification technologies) can bring additional benefits if lenders are thereby encouraged to expand the supply of credit, and if this expansion of credit supply has positive effects on household well-being. Credit expansions enabled by improved identification technology may be particularly large in previously underserved areas, such as the rural sub-Saharan context of our experiment, where problems with personal identification are particularly severe.

Another potential implication of this research is that in the absence of an alternative national identification system, fingerprints could serve as the unique identifier that allows individual credit histories to be stored and accessed in a cross-lender credit

[37]After learning about the benefits of biometric technology, MRFC applied for a grant from a donor agency to finance the purchase of handheld devices and software to mainstream the collection of biometric information from all its clients. OIBM, a competitor that operates in mostly urban areas, collects an electronic fingerprint from every borrower.

bureau. It has been noted that a key obstacle to establishment of credit bureaus is the lack of a unique identification system (*Conning and Udry* (2005), *Fafchamps* (2004)). Our results indicate that borrowers (particularly the worst borrowers) do perceive fingerprinting as an improvement in the lenders dynamic enforcement technology, and so support the use of fingerprints as an identifier in a national credit bureau.

As is the case with all field experiments, it is important to replicate this study in other contexts to gauge the external validity of the results. In addition to conducting similar studies in other rural sub-Saharan African contexts, it is also crucial to gauge the extent to which impacts of fingerprinting-enabled dynamic incentives are different in urban areas or areas with greater access to microcredit, for example. As mentioned above, the effects of fingerprinting on repayment could very well rise over time, and so future studies should monitor effects beyond a single loan cycle. Future work should also make sure to examine responses by the lender, such as changes in the credit contract, approval rates or in loan officer monitoring. While in our case loan officers did not behave differently towards treated borrowers, in other contexts, perhaps under different loan officer incentives, this may not be the case. We view these and related questions as promising areas for future research.

Tables

Table 3.1: Summary statistics

	Mean	Standard Deviation	10th Percentile	Median	90th Percentile	Observations
Baseline Characteristics						
Male	0.80	0.40	0	1	1	1147
Married	0.94	0.24	1	1	1	1147
Age	39.96	13.25	24	38	59	1147
Years of Education	5.35	3.50	0	5	10	1147
Risk Taker	0.56	0.50	0	1	1	1147
Days of Hunger Last Year	6.05	11.05	0	0	30	1147
Late Paying Previous Loan	0.13	0.33	0	0	1	1147
Income SD	27568.34	46296.41	3111.27	15556.35	57841.34	1147
Years of Experience Growing Paprika	2.22	2.36	0	2	5	1147
Previous Default	0.02	0.14	0	0	0	1147
No Previous Loans	0.74	0.44	0	1	1	1147
Predicted repayment	0.79	0.26	0.33	0.90	1.02	1147
Take-up						
Approved	0.99	0.08	1	1	1	1147
Any Loan	1.00	0.00	1	1	1	1147
Total Borrowed (MK)	16912.60	3908.03	13782	16100	20136.07	1147
Land Use						
Fraction of Land used for Maize	0.43	0.16	0.28	0.40	0.63	520
Fraction of land used for Soya/Beans	0.15	0.16	0.00	0.11	0.38	520
Fraction of land used for Groundnuts	0.13	0.12	0.00	0.11	0.29	520
Fraction of land used for Tobacco	0.08	0.12	0.00	0.00	0.27	520
Fraction of land used for Paprika	0.19	0.13	0.00	0.18	0.36	520
Fraction of land used for Tomatoes	0.01	0.03	0.00	0.00	0.00	520
Fraction of land used for Leafy Vegetables	0.00	0.02	0.00	0.00	0.00	520
Fraction of land used for Cabbage	0.00	0.01	0.00	0.00	0.00	520
Fraction of Land used for all cash crops	0.57	0.16	0.38	0.60	0.72	520
Inputs						
Seeds (MK, Paprika)	247.06	348.47	0	0	560	520
Fertilizer (MK, Paprika)	7499.85	7730.05	0	5683	18200	520
Chemicals (MK, Paprika)	671.31	1613.13	0	0	2500	520
Man-days (MK, Paprika)	665.98	1732.99	0	0	2400	520
All Paid Inputs (MK, Paprika)	9084.19	8940.13	0	8000	19990	520
KG Manure, Paprika	90.84	313.71	0	0	250	520
Times Weeding, Paprika	1.94	1.18	0	2	3	520
Outputs						
KG Maize	1251.30	1024.36	360	1080	2160	520
KG Soya/Beans	83.14	136.86	0	40	200	520
KG Groundnuts	313.89	659.34	0	143	750	520
KG Tobacco	165.47	615.33	0	0	400	520
KG Paprika	188.14	396.82	0	100	364	520
KG Tomatoes	30.56	126.29	0	0	0	520
KG Leafy Vegetables	29.94	133.24	0	0	0	520
KG Cabbage	12.02	103.79	0	0	0	520
Revenue and Profits						
Market sales (MK)	65004.30	76718.29	9800	44000	137100	520
Profits (market sales + value of unsold crop - cost of inputs, MK)	117779.20	303100.80	33359	95135	261145	520
Value of Unsold Harvest (Regional Prices, MK)	80296.97	288102.70	24645	70300	180060	520
Repayment						
Balance, Sept. 30	2912.91	6405.77	0	0	13981	1147
Fraction Paid by Sept. 30	0.84	0.33	0	1	1	1147
Fully Paid by Sept. 30	0.74	0.44	0	1	1	1147
Balance, eventual	2080.86	5663.98	0	0	9282	1147
Fraction Paid, eventual	0.89	0.29	0	1	1	1147
Fully paid, eventual	0.79	0.41	0	1	1	1147

Table 3.2: Tests of balance in baseline characteristics between treatment and control group

	Full baseline sample		Loan recipient sample	
	Mean in control group	Difference in treatment (fingerprinted) group	Mean in control group	Difference in treatment (fingerprinted) group
Variable:				
Male	0.81	-0.036 (0.022)	0.80	-0.066* (0.037)
Married	0.92	-0.004 (0.011)	0.94	0.003 (0.016)
Age	39.50	0.019 (0.674)	39.96	-0.088 (1.171)
Years of education	5.27	-0.046 (0.175)	5.35	-0.124 (0.272)
Risk taker	0.57	-0.033 (0.032)	0.56	0.013 (0.051)
Days of hunger in previous season	6.41	-0.647 (0.832)	6.05	-0.292 (1.329)
Late paying previous loan	0.14	0.005 (0.023)	0.13	0.030 (0.032)
Standard deviation of past income	25110.62	1289.190 (1756.184)	27568.34	-1158.511 (2730.939)
Years of experience growing paprika	2.10	0.096 (0.142)	2.22	0.299 (0.223)
Previous default	0.03	-0.002 (0.010)	0.02	0.008 (0.010)
No previous loan	0.74	-0.006 (0.027)	0.74	-0.020 (0.041)
P-value for test of joint significance	0.91		0.66	
Observations	3206		1147	

Each row presents mean of a variable in the baseline (September 2008) survey in the control group, and the difference between the treatment group mean and the control group mean of that variable (standard error in parentheses). Differences and standard errors calculated via a regression of the baseline variable on the treatment group indicator; standard errors are clustered at the club level.

Table 3.3: Borrower behavior under various theoretical cases, with and without dynamic incentives

	Without Dynamic Incentives			With Dynamic Incentives		
Case(i): $qb_H > y_S(b_H)$						
	$p < p_{B1}$	$p \geq p_{B1}$		$p < p_{B1}$	$p \geq p_{B1}$	
Loan size b	b_H	b_H		b_L	b_H	
Diversion D	1	1		0	0	
Repayment Rate	0	0		1	p	
Case(ii): $y_S(b_H) > qb_H > y_S(b_L)$						
	$p < p_{B1}$	$p_{B1} \leq p < p_D$	$p \geq p_D$	$p < p_{B1}$	$p_{B1} \leq p < p_D$	$p \geq p_D$
Loan size b	b_H	b_H	b_H	b_L	b_H	b_H
Diversion D	1	1	0	0	0	0
Repayment Rate	0	0	p	1	p	p
Case(iii): $y_S(b_L) > qb_H$						
	$p < p_D$	$p_D \leq p < p_{B0}$	$p \geq p_{B0}$	$p < p_D$	$p_D \leq p < p_{B0}$	$p \geq p_{B0}$
Loan size b	b_H	b_H	b_H	b_L	b_L	b_H
Diversion D	1	0	0	0	0	0
Repayment Rate	0	p	p	1	1	p

Table 3.4: Impact of fingerprinting on loan approval, loan take-up, and amount borrowed

	(1)	(2)	(3)
Sample:	All Respondents		Loan Recipients
Dependent variable:	Approved	Any Loan	Total Borrowed (MK)
Panel A			
Fingerprint	0.038	0.051	-696.799*
	(0.053)	(0.044)	(381.963)
Panel B			
Fingerprint	0.207	0.108	-2812.766
	(.161)	(.145)	(2371.685)
Predicted repayment * fingerprint	-0.219	-0.074	2630.653
	(.197)	(.168)	(2555.167)
Panel C			
Fingerprint * Quintile 1	0.093	0.075	-2721.780
	(.115)	(.111)	(1666.068)
Fingerprint * Quintile 2	0.180*	0.102	-258.179
	(.096)	(.086)	(828.500)
Fingerprint * Quintile 3	-0.030	0.061	-458.924
	(.082)	(.073)	(596.109)
Fingerprint * Quintile 4	-0.001	-0.037	-101.028
	(.086)	(.082)	(575.968)
Fingerprint * Quintile 5	-0.017	0.039	-400.620
	(.100)	(.089)	(784.509)
Observations	3206	3206	1147
Mean of dependent variable	0.63	0.35	16912.60
Quintile 1	0.58	0.29	17992.53
Quintile 2	0.64	0.36	17870.61
Quintile 3	0.71	0.44	16035.10
Quintile 4	0.70	0.47	15805.54
Quintile 5	0.59	0.30	16886.56

Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels.
Each column presents estimates from three separate regressions: main effect of fingerprinting in Panel A, linear interaction with predicted repayment in Panel B, and interactions with quintiles of predicted repayment in Panel C. All regressions include locality×week of initial loan offer fixed effects, baseline characteristics (male, five-year age categories, one-year education categories, and marriage), and baseline risk indicators (dummy for self-reported risk-taking, days of hunger in the previous season, late payments on previous loans, standard deviation of income, years of experience growing paprika, dummy for default on previous loan, and dummy for no previous loans). Panel B regressions include the main effect of the level of predicted repayment, and Panel C regressions include dummies for quintile of predicted repayment main effects. Standard errors on Panel A coefficients are clustered at the club level, while those in Panels B and C are bootstrapped with 200 replications and club-level resampling.

Table 3.5: Impact of fingerprinting on loan repayment

	(1)	(2)	(3)	(4)	(5)	(6)
Sample:	Loan recipients					
Dependent variable:	Balance, Sept. 30	Fraction Paid by Sept. 30	Fully Paid by Sept. 30	Balance, Eventual	Fraction Paid Eventual	Fully Paid Eventual
Panel A						
Fingerprint	-1556.383*	0.073*	0.096	-996.430	0.045	0.085
	(824.174)	(0.040)	(0.062)	(754.301)	(0.036)	(0.058)
Panel B						
Fingerprint	-15174.149***	0.716***	0.842***	-9727.739**	0.438**	0.602***
	(2743.271)	(.110)	(.178)	(4199.085)	(.184)	(.224)
Predicted repayment * fingerprint	16930.139***	-0.799***	-0.928***	10855.103**	-0.489**	-0.643***
	(3047.515)	(.121)	(.196)	(4499.549)	(.196)	(.243)
Panel C						
Fingerprint * Quintile 1	-10844.169***	0.499***	0.543***	-7202.647**	0.317**	0.396**
	(2681.861)	(.127)	(.147)	(2969.045)	(.136)	(.156)
Fingerprint * Quintile 2	-1104.582	0.066	0.163	-1028.696	0.060	0.170
	(2025.425)	(.105)	(.160)	(1871.298)	(.097)	(.148)
Fingerprint * Quintile 3	-307.761	0.005	-0.004	-297.918	0.002	0.007
	(966.586)	(.048)	(.091)	(901.013)	(.045)	(.087)
Fingerprint * Quintile 4	818.275	-0.037	-0.045	775.231	-0.035	-0.028
	(942.466)	(.046)	(.078)	(883.076)	(.044)	(.075)
Fingerprint * Quintile 5	1674.419	-0.078*	-0.084	1404.812	-0.061	-0.050
	(1022.895)	(.046)	(.074)	(951.535)	(.043)	(.071)
Observations	1147	1147	1147	1147	1147	1147
Mean of dependent variable	2912.91	0.84	0.74	2080.86	0.89	0.79
Quintile 1	6955.67	0.62	0.52	4087.04	0.81	0.68
Quintile 2	4024.05	0.77	0.63	3331.17	0.81	0.67
Quintile 3	1571.44	0.92	0.83	1301.79	0.93	0.84
Quintile 4	877.80	0.95	0.85	781.59	0.95	0.87
Quintile 5	1214.19	0.94	0.85	950.29	0.95	0.88

See notes for Table 3.4.

Table 3.6: Impact of fingerprinting on land use

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable: Fraction of land used for	Maize	Soya/Beans	Groundnuts	Tobacco	Paprika	Tomatoes	Leafy Vegetables	Cabbage	All cash crops
Panel A									
Fingerprint	0.001	0.015	-0.012	-0.004	0.005	-0.001	-0.003	-0.001	-0.001
	(0.019)	(0.019)	(0.016)	(0.016)	(0.014)	(0.003)	(0.003)	(0.001)	(0.019)
Panel B									
Fingerprint	-0.009	-0.025	-0.025	-0.033	0.079	0.009	0.006	-0.003	0.009
	(.092)	(.094)	(.060)	(.062)	(.064)	(.010)	(.015)	(.004)	(.092)
Predicted repayment * fingerprint	0.013	0.049	0.016	0.036	-0.092	-0.013	-0.011	0.003	-0.013
	(.101)	(.105)	(.068)	(.066)	(.073)	(.013)	(.016)	(.005)	(.101)
Panel C									
Fingerprint * Quintile 1	-0.061	-0.013	-0.008	-0.012	0.083	0.005	0.007	-0.002	0.061
	(.066)	(.063)	(.052)	(.050)	(.051)	(.008)	(.012)	(.003)	(.066)
Fingerprint * Quintile 2	0.065	0.019	-0.014	-0.019	-0.035	-0.005	-0.010	-0.002	-0.065
	(.052)	(.042)	(.041)	(.030)	(.037)	(.008)	(.008)	(.002)	(.052)
Fingerprint * Quintile 3	-0.012	0.002	-0.009	0.004	0.009	0.008	-0.002	-0.001	0.012
	(.044)	(.045)	(.033)	(.022)	(.038)	(.008)	(.007)	(.002)	(.044)
Fingerprint * Quintile 4	0.008	0.015	-0.026	0.009	-0.003	-0.002	-0.003	0.002	-0.008
	(.041)	(.040)	(.034)	(.021)	(.037)	(.009)	(.007)	(.003)	(.041)
Fingerprint * Quintile 5	-0.005	0.043	-0.001	-0.001	-0.018	-0.012	-0.005	-0.002	0.005
	(.044)	(.040)	(.036)	(.023)	(.034)	(.009)	(.006)	(.003)	(.044)
Observations	520	520	520	520	520	520	520	520	520
Mean of dependent variable	0.43	0.15	0.13	0.08	0.19	0.01	0.00	0.00	0.57
Quintile 1	0.44	0.07	0.13	0.18	0.17	0.01	0.01	0.00	0.56
Quintile 2	0.49	0.10	0.13	0.13	0.15	0.00	0.00	0.00	0.51
Quintile 3	0.42	0.21	0.12	0.03	0.20	0.01	0.00	0.00	0.58
Quintile 4	0.42	0.19	0.12	0.04	0.21	0.01	0.01	0.00	0.58
Quintile 5	0.40	0.17	0.14	0.04	0.23	0.01	0.01	0.00	0.60

See notes for Table 3.4.

Table 3.7: Impact of fingerprinting on agricultural inputs used on paprika crop

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:	Seeds (MK)	Fertilizer (MK)	Chemicals (MK)	Man-days (MK)	All Paid Inputs (MK)	KG Manure	Times Weeding
Panel A							
Fingerprint	74.107	733.419	345.328*	-395.501**	757.354	29.649	0.019
	(47.892)	(1211.905)	(190.262)	(181.958)	(1389.230)	(32.593)	(0.147)
Panel B							
Fingerprint	262.116*	11115.814**	466.677	411.043	12255.650**	52.882	0.182
	(146.417)	(5660.459)	(594.037)	(579.097)	(5987.210)	(144.033)	(.466)
Predicted repayment * fingerprint	-234.438	-12946.332**	-151.316	-1005.720	-14337.806**	-28.970	-0.203
	(183.931)	(6245.378)	(701.923)	(732.887)	(6700.416)	(161.334)	(.591)
Panel C							
Fingerprint * Quintile 1	188.703**	5871.126	374.260	106.406	6540.496	78.234	0.445
	(95.018)	(4062.716)	(406.741)	(347.367)	(4210.469)	(111.980)	(.367)
Fingerprint * Quintile 2	78.717	3597.540	244.449	-236.338	3684.368	27.058	-0.443
	(95.343)	(3026.725)	(414.863)	(454.498)	(3362.245)	(81.930)	(.338)
Fingerprint * Quintile 3	124.548	-585.618	500.669	-348.598	-309.000	58.670	-0.191
	(97.766)	(2250.453)	(427.366)	(458.033)	(2602.025)	(94.443)	(.333)
Fingerprint * Quintile 4	-10.190	-1790.213	283.962	-1065.690**	-2582.132	-25.080	-0.254
	(110.489)	(2503.022)	(430.040)	(537.142)	(2952.953)	(73.404)	(.348)
Fingerprint * Quintile 5	18.589	-2444.617	264.620	-315.018	-2476.427	21.879	0.564
	(110.367)	(2201.579)	(445.234)	(572.589)	(2635.638)	(93.481)	(.379)
Observations	520	520	520	520	520	520	520
Mean of dependent variable	247.06	7499.85	671.31	665.98	9084.19	90.84	1.94
Quintile 1	174.13	6721.24	401.30	143.48	7440.15	97.39	1.47
Quintile 2	140.00	6080.46	620.67	238.94	7080.08	39.25	1.55
Quintile 3	269.90	8927.65	674.48	836.98	10709.00	105.73	2.05
Quintile 4	292.07	7649.51	715.08	936.29	9592.95	93.23	2.24
Quintile 5	340.18	8078.58	892.05	1065.18	10375.99	118.13	2.28

See notes for Table 3.4.

Table 3.8: Impact of fingerprinting on revenue and profits

Dependent variable:	(1) Market sales (Self Report, MK)	(2) Value of Unsold Harvest (Regional Prices, MK)	(3) Profits (market sales + value of unsold harvest - cost of inputs, MK)	(4) Ln(profits)
Panel A				
Fingerprint	7246.174	5270.320	14509.457	0.060
	(8792.055)	(14879.349)	(16679.311)	(0.095)
Panel B				
Fingerprint	69102.211	-29468.424	24207.068	0.651
	(49177.370)	(85252.270)	(90535.890)	(.423)
Predicted repayment * fingerprint	-77131.415	43317.493	-12092.441	-0.737
	(51232.390)	(103316)	(108112.600)	(.501)
Panel C				
Fingerprint * Quintile 1	30766.147	7940.835	31915.287	0.401
	(36850.940)	(50587.570)	(63206.880)	(.363)
Fingerprint * Quintile 2	41981.091	6364.782	45650.027	0.283
	(33084.250)	(75026.680)	(81848.520)	(.264)
Fingerprint * Quintile 3	-20925.441	-14911.454	-26932.651	-0.202
	(17938.730)	(59934.020)	(63400.760)	(.227)
Fingerprint * Quintile 4	-12785.841	7481.854	3609.228	-0.038
	(14733.930)	(57096.050)	(60385.110)	(.231)
Fingerprint * Quintile 5	1053.151	33336.147	34125.843	-0.054
	(15282.460)	(71891.840)	(74254.990)	(.240)
Observations	520	520	520	520
Mean of dependent variable	65004.30	80296.97	117779.16	11.44
Quintile 1	60662.57	82739.24	121222.50	11.36
Quintile 2	89028.25	29995.27	91652.71	11.55
Quintile 3	57683.74	96247.91	123242.30	11.44
Quintile 4	61088.27	104927.50	136467.50	11.45
Quintile 5	56593.43	85817.08	115172.50	11.39
Mean of dependent variable (US $)	464.32	573.55	841.28	n.a.

See notes for Table 3.4.

Table 3.9: Ex post moral hazard

Dependent variable:	(1) Balance, Sept. 30	(2) Fraction Paid by Sept. 30	(3) Fully Paid by Sept. 30	(4) Balance, Eventual	(5) Fraction Paid Eventual	(6) Fully Paid Eventual
Panel A						
Fingerprint	-102.571	-0.005	0.000	424.455	-0.031	-0.003
	(775.942)	(0.039)	(0.072)	(565.064)	(0.028)	(0.058)
Panel B						
Fingerprint	-8282.762*	0.320	0.399	-3537.222	0.082	0.173
	(4698.647)	(.237)	(.298)	(5140.761)	(.237)	(.280)
Predicted repayment * fingerprint	9794.172**	-0.389	-0.478	4743.330	-0.134	-0.211
	(4903.271)	(.246)	(.326)	(5378.012)	(.247)	(.308)
Panel C						
Fingerprint * Quintile 1	-7589.870*	0.304	0.314	-4443.517	0.149	0.221
	(4479.864)	(.211)	(.248)	(4252.255)	(.187)	(.216)
Fingerprint * Quintile 2	2964.264	-0.164	-0.142	2679.579	-0.151	-0.132
	(2231.033)	(.118)	(.166)	(1950.827)	(.104)	(.149)
Fingerprint * Quintile 3	-597.239	0.035	0.049	-358.930	0.024	0.050
	(1105.313)	(.062)	(.108)	(978.83)	(.055)	(.100)
Fingerprint * Quintile 4	419.174	-0.026	-0.058	763.678	-0.042	-0.052
	(1000.512)	(.057)	(.113)	(909.161)	(.052)	(.102)
Fingerprint * Quintile 5	460.280	-0.019	-0.023	732.027	-0.029	-0.005
	(1134.012)	(.062)	(.119)	(955.498)	(.052)	(.099)
Observations	520	520	520	520	520	520
Mean of dependent variable	2071.21	0.89	0.79	1439.16	0.92	0.83
Quintile 1	6955.67	0.62	0.52	3472.29	0.83	0.71
Quintile 2	4024.05	0.77	0.63	2610.41	0.85	0.75
Quintile 3	1571.44	0.92	0.83	476.63	0.97	0.91
Quintile 4	877.80	0.95	0.85	661.79	0.96	0.86
Quintile 5	1214.19	0.94	0.85	311.66	0.98	0.93

See notes for Table 3.4.

Table 3.10: Benefit-cost analysis

Benefit			
	(a) Increase in repayment due to fingerprinting in Quintile 1	4,755.00	Malawi kwacha
	(b) Quintile 1 as share of all borrowers	20.0%	
	(c) Borrowers as share of all fingerprinted	50%	
	(d) Total benefit per individual fingerprinted [= (a)*(b)*(c)]	475.50	Malawi kwacha
Cost			
	(e) Cost per equipment unit	101,500	Malawi kwacha
	(f) Equipment amortization period	3	years
	(g) Annual equipment amortization [= (e) / (f)]	33,833	
	(h) Fingerprinted individuals per equipment unit	417	individuals
	(i) Equipment cost per farmer [= (g) / (h)]	81.20	Malawi kwacha
	(j) Loan officer time cost per farmer	19.25	Malawi kwacha
	(k) Transaction cost per fingerprint checked	108.75	Malawi kwacha
	(l) Total cost per individual fingerprinted [= (i) + (j) + (k)]	209.20	Malawi kwacha
(m) Net benefit per fingerprinted farmer [= (d) - (l)]		266.30	Malawi kwacha
(n) Benefit-cost ratio [= (d) / (l)]		2.27	

Assumptions:

Exchange rate:	145	MK/US$
Loan size	15,000	Malawi kwacha
Increase in share of loan repaid due to fingerprinting in Quintile 1	31.7%	
Cost per equipment unit (laptop computer + fingerprint scanner)	700	USD
Number of equipment units	12	
Fingerprinting time per individual	5	minutes
Monthly salary of MRFC loan officer	40,000	Malawi kwacha
Hours worked per month by MFRC loan officer	173.2	hours

Figures

Figure 3.1: Experimental timeline

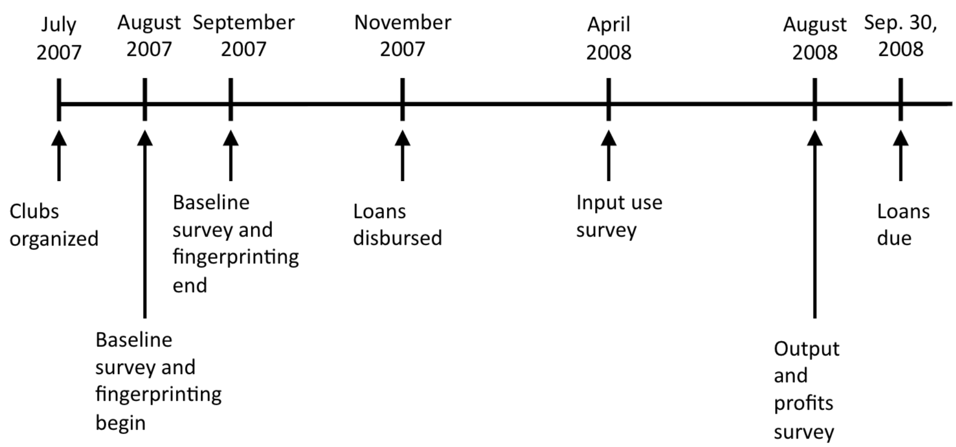

Figure 3.2: Optimal behavior as a function of p

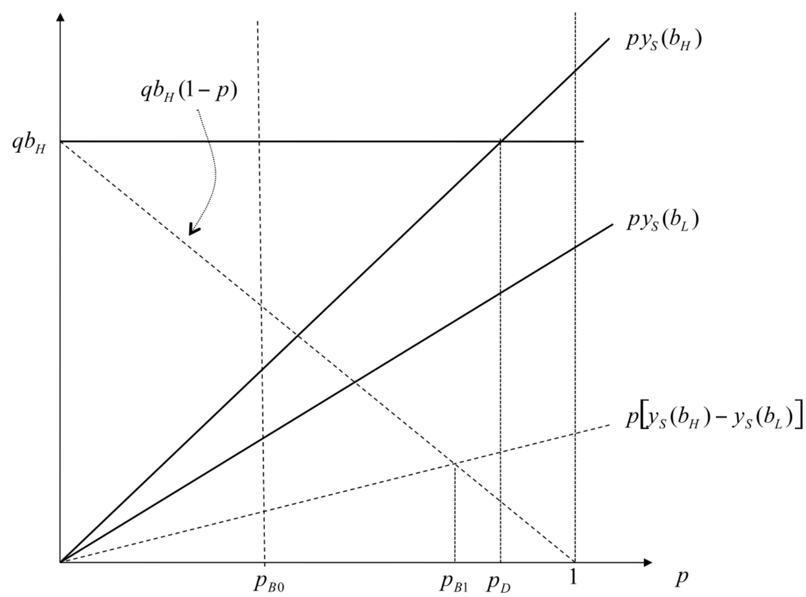

APPENDICES

APPENDIX A

Robustness checks

Table A.1: 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***	0.125***	0.142***	0.149**	0.120**	0.133**
	(0.037)	(0.035)	(0.034)	(0.047)	(0.045)	(0.055)
Ln($wage_{t+1}$)		-0.018	-0.010		0.027	0.029
		(0.044)	(0.037)		(0.080)	(0.066)
Ln($wage_{t+2}$)					0.016	0.013
					(0.048)	(0.027)
Ln($wage_{t+3}$)					-0.047	-0.028
					(0.037)	(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	0.15	0.17	0.18	0.15	0.16
	(0.046)	(0.043)	(0.044)	(0.060)	(0.057)	(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 A.2: 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***	0.123***	0.143***	0.166***	0.176***	0.175***
	(0.036)	(0.034)	(0.033)	(0.027)	(0.031)	(0.028)
Ln($wage_{t-1}$)		-0.057**	-0.029		0.011	0.014
		(0.025)	(0.026)		(0.011)	(0.016)
Ln($wage_{t-2}$)					-0.005	-0.004
					(0.019)	(0.017)
Ln($wage_{t-3}$)					0.009	0.009
					(0.018)	(0.011)
Ln($wage_{t-4}$)					-0.002	-0.002
					(0.018)	(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	0.15	0.17	0.19	0.20	0.20
	(0.044)	(0.042)	(0.041)	(0.034)	(0.038)	(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 A.3: Elasticity of employment w.r.t. the average of past wages

	(1)	(2)	(3)
Sample:		Weeks 2 to 12	
Dependent variable:		Individual*day indicator for working	
Ln(wage)	0.141***	0.152***	0.140***
	(0.036)	(0.035)	(0.034)
Ln($\overline{wage_{t-1}}$)		0.075	-0.143
		(0.065)	(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	0.18	0.17
	(0.044)	(0.043)	(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 A.4: Probit estimates of the elasticity of employment w.r.t. wages

	(1)	(2)	(3)	(4)
Dependent variable:	Individual*day indicator for working			
Ln(wage)	0.120***	0.119***	0.136***	0.135***
	(0.030)	(0.030)	(0.030)	(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	0.14	0.16	0.16
	(0.037)	(0.037)	(0.038)	(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.
* $p<0.10$, ** $p<0.05$, *** $p<0.001$

APPENDIX B

Solution to the social pressure model

Recall that

$$\begin{aligned} w_t^u &= y_t^u + (1+r)s_{t-1}^u \\ w_t^v &= y_t^v + (1+r)s_{t-1}^v \end{aligned}$$

where u denotes public and v denotes private. A fraction τ of public, liquid assets ("surplus income") held at the end of period 1 must be contributed to the social network.

There are four different cases as illustrated in figure B.1.

Case 1 $c_0^* < w_0^u$ and $c_0^{**} < w_0^u$. Thus $c_0 = c_0^*$ is the optimal point allowed.

Case 2 $c_0^* > w_0^u$ and $c_0^{**} > w_0^u$. Thus $c_0 = c_0^{**}$ is the optimal point allowed.

Case 3 $c_0^* > w_0^u$ and $c_0^{**} < w_0^u$. Thus $c_0 = w_0^u$ is the optimal point allowed.

Case 4 $c_0^* > w_0^u + w_0^v$ and $c_0^{**} > w_0^u + w_0^v$. Thus $c_0 = w_0^u + w_0^v$ is the optimal point allowed.

Figure B.1: Different cases of the model

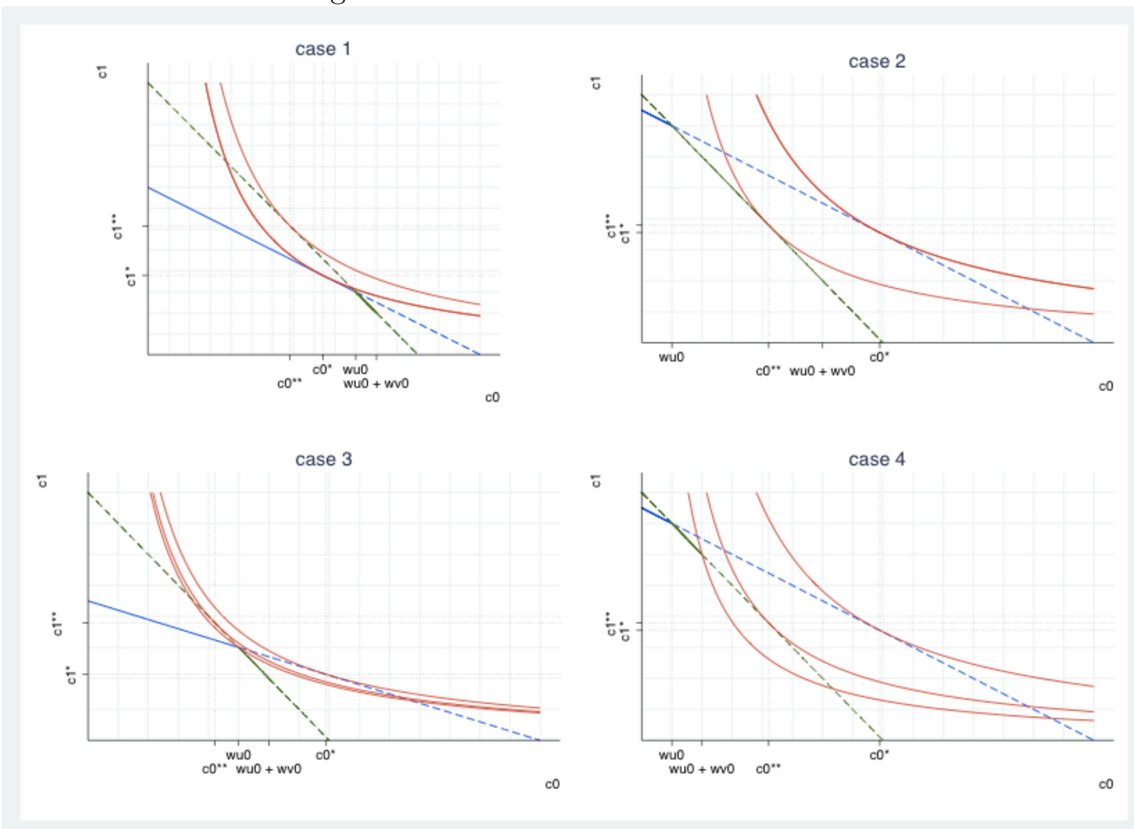

In case 1, the income effect dominates the substitution effect. A shock to first period private income increases first period consumption by more than a shock to first period public income. In case 2, the sharing obligation for unconsumed public income is not binding because individuals want to consume more than their public income in the first period. The effects of public and private shocks on first period consumption are symmetric. In case 3, the sharing norm binds: individuals would like to save some of their public income, but the substitution effect dominates the income effect and forces them to the corner solution. Here, first period consumption increases one-for-one with first period public income, while first period private income has no effect on first period consumption. The paper focuses on the behavior predicted by cases two and three of the model. In case 4, the no-borrowing constraint binds and individuals can consume only as much as they receive in the first period.

First, I maximize utility with respect to the sharing budget constraint, such that equation 2.1 holds, to find the utility-maximizing level of consumption c_0^*.

$$\max_{c_0} \ln c_0 + \beta \ln(y^u + y^v + \theta w_0^u + w_0^v - \theta c_0)$$

$$\frac{\partial F}{\partial c_0} = \frac{1}{c_0} + \beta \frac{-\theta}{y^u + y^v + \theta w_0^u + w_0^v - \theta c_0} = 0$$

$$\beta\theta c_0 = y^u + y^v + \theta w_0^u + w_0^v - \theta c_0$$

$$c_0^* = \frac{1}{\theta(1+\beta)}(y^u + y^v + \theta w_0^u + w_0^v)$$

I sustitute into the budget constraint (equation 2.1) to find c_1^*:

$$c_1^* = \frac{\beta}{(1+\beta)}(y^u + y^v + \theta w_0^u + w_0^v)$$

Next, I maximize utility with respect to the no-sharing budget constraint, such that equation 2.2 holds, and call the utility-maximizing level of consumption under this constraint c_0^{**}.

$$\max_{c_0} \ln c_0 + \beta \ln(y^u + y^v + w_0^u + w_0^v - c_0)$$

$$\frac{\partial F}{\partial c_0} = \frac{1}{c_0} + \beta \frac{-1}{y^u + y^v + w_0^u + w_0^v - c_0} = 0$$

$$\beta c_0 = y^u + y^v + w_0^u + w_0^v - c_0$$

$$c_0^{**} = \frac{1}{1+\beta}(y^u + y^v + w_0^u + w_0^v)$$

Again, I substitute this solution into the budget constraint (equation 2.2) to find c_1^{**}:

$$c_1^{**} = \frac{\beta}{(1+\beta)}(y^u + y^v + w_0^u + w_0^v)$$

The derivatives of interest are $\partial c_0/\partial w_0^u$ and $\partial c_0/\partial w_0^v$.

	Case 1	Case 2	Case 3	Case 4
c_0	$\frac{1}{\theta(1+\beta)}\left(y^u + y^v + \theta w_0^u + w_0^v\right)$	$\frac{1}{1+\beta}\left(y^u + y^v + w_0^u + w_0^v\right)$	w_0^u	$w_0^u + w_0^v$
$\frac{\partial c_0}{\partial w_0^u}$	$\frac{1}{1+\beta}$	$\frac{1}{1+\beta}$	1	1
$\frac{\partial c_0}{\partial w_0^v}$	$\frac{1}{\theta(1+\beta)}$	$\frac{1}{1+\beta}$	0	1

c_1	$\frac{\beta}{(1+\beta)}\left(y^u + y^v + \theta w_0^u + w_0^v\right)$	$\frac{\beta}{(1+\beta)}\left(y^u + y^v + w_0^u + w_0^v\right)$	$y^u + y^v + w_0^v$	$y^u + y^v$
$\frac{\partial c_1}{\partial w_0^u}$	$\frac{\beta\theta}{1+\beta}$	$\frac{\beta}{1+\beta}$	0	0
$\frac{\partial c_1}{\partial w_0^v}$	$\frac{\beta}{1+\beta}$	$\frac{\beta}{1+\beta}$	1	0

As argued in the paper, case 3 is plausible when a large fraction of income is observed and when liquid assets must be shared with others in the social network, which is common in many rural and developing economies.

APPENDIX C

Biometrics training script

C.1 Benefits of good credit

Having a record of paying back your loans can help you get bigger loans or better interest rates.

Credit history works like trust. When you know someone for a long time, and that person is honest and fair when you deal with him, then you trust him. You are more likely to help him, and he is more likely to help you. You might let him use your hoe (or something else that is important to you), because you feel sure that he will give it back to you. Banks feel the same way about customers who have been honest and careful about paying back their loans. They trust those customers, and are more willing to let them borrow money.

MRFC already gives customers who have been good borrowers a reward. It charges them a lower interest rate, 30 percent instead of 33 percent. That means that for the loan we have described today, someone who has a good credit history would only have to pay back 8855, instead of 8971.[1]

[1] Loan amounts mentioned in the script are lower than actual loan amounts observed in the data because fertilizer prices rose somewhat in the time between the initial intervention (in Aug-Sep 2007) and loan disbursement (Nov 2007).

Another way that banks might reward customers they trust is by letting them borrow bigger amounts of money. Instead of 7700 MK to grow one acre of paprika, MRFC might lend a trusted customer 15400, to grow two acres.

To earn trust with the bank, and get those rewards, you have to be able to prove to the bank that you have taken loans before and paid them back on time. You can do that by making sure that you give the bank accurate information when you fill out loan applications. But if you call yourself John Jacob Phiri one year, and Jacob John Phiri the next year, then the bank might not figure out that you are the same person, so they wont give you the rewards you have earned.

C.2 Costs of bad credit

But trust can be broken. If your neighbor borrows your radio and does not give it back or it gets ruined, then you probably wouldnt lend him anything else until the radio had been replaced.

Banks work the same way. If you take a loan and break the trust between yourself and the bank by not paying back the loan, then the bank wont lend to you again. This is especially true if you have a good harvest but still choose not to pay back the loan.

When you apply for a loan, one of the things that a bank does to decide whether or not to accept your application is to look in its records to see if you have borrowed money before. If you have borrowed but not paid back, then you will be turned down for the new loan. This is like you asking your neighbors if someone new shows up in the village and asks you to work for him. You might first ask around to see if the person is fair to his employees and pays them on time. If you learn that the person does not pay his workers, then you wont work for him. Banks do the same thing by checking their records.

MRFC does not ever give new loans to people who still owe them money. And

MRFC shares information about who owes money with other banks, so if you fail to pay back a loan from MRFC, it can stop you from getting a new loan from OIBM or another lender, also.

Remainder of scripts is administered to fingerprinted clubs only

C.3 Biometric technology

Fingerprints are unique, which means that no two people can ever have the same fingerprints. Even if they look similar on a piece of paper, people with special training, or special computer equipment, can always tell them apart.

Your fingerprint can never change. It will be the same next year as it is this year. Just like the spots on a goat are the same as long as the goat lives, but different goats have different spots.

Fingerprints can be collected with ink and paper, or they can be collected with special machines. This machine stores fingerprints in a computer. Once your fingerprint is stored in the computer, then the machine can recognize you, and know your name and which village you come from, just by your fingerprint! The machine will recognize you even if the person who is using it is someone you have never met before. The information from the machines is saved in many different ways, so if one machine breaks, the information is still there. Just like when Celtels building burned, peoples phone numbers did not change.

Administer the following after all fingerprints have been collected

C.4 Demo

Now, I can figure out your name even if you dont tell me. Will someone volunteer to test me? *(Have a volunteer swipe his finger, and then tell everyone who it was).*

The bank will store information about your loans with your fingerprint. That means that bank officers will know not just your name, but also what loans you have taken and whether or not you have paid them back. They will be able to tell all of this just by having you put your finger on the machine.

Before, banks used your name and other information to find out about your credit history. But now they will use fingerprints to find out. This means that even if you tell the bank a different name, they will still be able to find all of your loan records. Names can change, but fingerprints cannot.

Having your fingerprint on file can make it easier to earn the rewards for good credit history that we talked about earlier. It will be easy for the bank to look up your records and see that you have paid back your loans before. It will also be easier to apply for loans, because there will be no new forms to fill out in the future!

But, having your fingerprint on file also makes the punishment for not paying back your loan much more certain. Even if you tell the bank a different name than you used before, or meet a different loan officer, or go to a different branch, the bank will just have to check your fingerprint to find out whether or not you paid your loans before. Having records of fingerprints also makes it easy for banks to share information. Banks will share information about your fingerprints and loans. If you dont pay back a loan to MRFC, OIBM will know about it!

APPENDIX D

Details on biometric fingerprinting technology

In consultation with MRFCs management, fingerprint recognition was chosen over face, iris or retina recognition because it is the cheapest, best known and most widely used biometric identification technology. Fingerprinting technology extracts features from impressions made by the distinct ridges on the fingertips and has been commercially available since the early 1970s.

Loan applicants from fingerprinted clubs had the image of their right thumb fingerprint captured by an optical fingerprint scanner attached to a laptop. To maximize accuracy, farmers washed their thumbprints prior to scanning, and the scanner was also cleaned after each impression. During collection, about 2 per cent of farmers had the left thumbprint recorded (instead of the right) because the right thumbprint was worn out. (Many farmers grow tobacco, which involves thumb usage during seedling transplantation that can wear out a thumbprint over many years.)

Upon scanning, the fingerprint image was enhanced and added to the borrower database. We purchased the VeriFinger 5.0 Software Development Kit from Fulcrum Biometrics and had a programmer develop a data capture program that would allow the user to (i) enter basic demographic information such as the name, address, village, loan size and the unique BFIRM identifier, (ii) capture the fingerprint with the scanner and (iii) review the fingerprint alongside the demographic information.

APPENDIX E

Variable definitions

Data used in this paper come from two surveys: a baseline conducted in August-September 2007 and a follow-up survey about farm outputs and other outcomes conducted in August 2008. We also used administrative data about loan take-up and repayment, obtained from MRFCs internal records.

E.1 Baseline characteristics (from baseline survey)

Male equals 1 for men and 0 for women.

Married equals 1 for married respondents and 0 for respondents who are single, widowed, or divorced.

Age is respondents age in years. In regressions, we use dummies for 5-year age categories rather than a continuous measure of age.

Years of education is years of completed schooling, and is top-coded at 13. In regressions, we use dummies for years of completed schooling, rather than a continuous measure of education.

Risk taker equals 1 for respondents who report that they frequently take risks, and 0 for respondents who do not.

Days of hunger last year is the number of days in the 2006-2007 season that individuals reduced the number of meals they ate per day.

Late paying previous loan equals 1 for respondents who report paying back a previous loan late, and 0 for respondents who do not.

Income SD is the standard deviation of income between the self-reported best and worst incomes of the 5 most recent years.

Years of experience growing paprika is the self reported number of seasons in which the respondent has grown paprika before the season studied in this project.

Previous default equals 1 for respondents who report that they have defaulted on a previous loan and 0 otherwise.

No previous loans equals 1 for respondents who report that they have not had any other loans from formal financial institutions (including micro lenders, savings and credit cooperatives, and NGO schemes) and 0 otherwise.

E.2 Take-up and repayment (from administrative data)

Approved equals 1 if the respondent was approved by MRFC for a loan and 0 otherwise.

Any loan equals 1 if the respondent borrowed money from MRFC and 0 otherwise (this could differ from Approved if the respondent chose not to take out the loan after it was approved by MRFC).

Total borrowed is the amount owed to MRFC, in Malawi kwacha (MK 145 = $US 1). This includes the loan principal and 33 percent interest charged by MRFC.

Balance is the unpaid loan amount remaining to be paid to MRFC. The balance includes principal and accumulated interest, and is reported in MK.

Fraction paid is the amount paid on the loan, divided by the total borrowed defined above.

Fully paid equals 1 if the respondent has completely repaid the loan and 0 if there is an outstanding balance.

We examine different versions of the variables *Balance*, *Fraction paid*, and *Fully paid* that vary by the date at which loan repayment status is measured. One set of variables refers to loan repayment status as of September 30, 2008, which is the formal due date of the loan. Another set of variables refers to “eventual repayment as of the end of November 2008. MRFC considers loan repayment status at the end of November 2008 as the final repayment status of the loan, and makes no subsequent attempts to collect loan repayments after that point.

E.3 Land use and inputs (from follow-up survey)

Fraction of land used for various crops is the land used for the given crop, divided by total land cultivated.

Seeds is the value of paprika seeds used by the respondent, in MK.

Fertilizer is the value of all chemical fertilizer used by the respondent on the paprika crop, in MK.

Chemicals is the value of all pesticides and herbicides used by the respondent on the paprika crop, in MK.

Man-days is the amount of money spent on hired, non-family labor for the paprika crop, in MK.

All paid inputs is the total amount of money spent on inputs for the paprika crop, in MK. Mathematically, it is the sum of Seeds, Fertilizer, Chemicals, and Man-days defined above.

KG manure is the kilograms of manure applied to the paprika crop.

Times weeding is the number of times the paprika crop was weeded, by the respondent or hired labor.

E.4 Output, revenue and profits (from follow-up survey)

KG of various crops is the self-reported kilograms harvested of each crop.

Market sales is the amount of MK received from any sales of maize, soya, groundnuts, tobacco, paprika, tomatoes, leafy vegetables, and cabbage between April and August, which encompasses the entire main harvest and selling season for these crops.

Profits is the value of Market sales, plus the value of unsold crop estimated based on the farmers reported quantity, valued at district average price reported by the EPA office (*Value of unsold harvest*, defined below), minus *All paid inputs* as defined above.

Value of unsold harvest is the value, in MK, of the difference between the kg harvested and the kg sold of each crop. We use district average prices, as reported by the EPA office.

APPENDIX F

Construction of predicted repayment variable

To construct the predicted repayment variable, we first limit the sample to borrowers in the control group (N=563), and run a regression of a repayment outcome (fraction of loan repaid by September 30, 2008) on various farmer- and club-level baseline characteristics. Conceptually, the resulting index will be purged of any bias introduced by effects of fingerprinting on repayment because it is constructed using coefficients from a regression predicting repayment for only the control (non-fingerprinted) farmers.

Appendix Table G.1 presents results from this exercise. Statistically significant results in column 1, which only includes farmer-level (individual) variables on the right-hand-side, indicates that older farmers and those who do not self-identify as risk-takers have better repayment performance on the loan. Inclusion of a complete set of fixed effects for (locality)×(week of initial loan offer) interactions raises the R-squared substantially (from 0.05 in column 1 to 0.46 in column 2). The explanatory power of the regression is marginally improved further in column 3 (to an R-squared of 0.48) when age and education are specified as categorical variables (instead of being entered linearly).

We then take the coefficient estimates from column 3 of Appendix Table G.2 and predict the fraction of loan repaid for the entire sample (both control and treatment observations). This variable, which we call "predicted repayment", is useful for analytical purposes because it is a single index that incorporates a wide array of baseline information (at the individual and locality level) correlated with repayment outcomes. In the loan-recipient subsample, predicted repayment has a mean of 0.79, with standard deviation 0.26. As expected, predicted repayment is highly skewed, with median predicted repayment of 0.90. Predicted repayment reaches 100 percent at the 84th percentile.

APPENDIX G

Additional robustness checks

G.1 Impact of fingerprinting in full sample

Analyses of the full sample of farmers, without restricting the sample only to borrowers, can help address concerns about selection bias. Appendix Tables G.5, G.6, and G.7 present results from regressions analogous to Tables 3.6, 3.7, and 3.8, respectively, with the difference that the regressions include all 1,226 individuals interviewed in the follow-up survey (borrowers plus nonborrowers).

Full-sample regression results in Appendix Tables G.5, G.6, and G.7 are very similar to those from the borrower-only regressions. As discussed in the main text, the general pattern is for coefficients that were significant before to remain statistically significant, but to be only around half the magnitude of the coefficients in the borrowing sample regressions. This reduction in coefficient magnitude is consistent with effect sizes in the full sample representing a weighted average of no effects for nonborrowers and nonzero effects for borrowers.

To be specific, in the land-use full-sample regressions (Appendix Table G.5), fingerprinting leads farmers in quintile 1 of predicted repayment to devote 5 percentage points more of their land to paprika (significant at the 5% level). In the inputs regressions (Appendix Table G.6), the interaction of fingerprinting with predicted repayment in Panel B is negative and significant at the 10% level in the regressions for fertilizer and all paid inputs, as in Table 3.7. The fingerprinting×(quintile 1) interaction term is also positive and statistically significant at the 5% or 10% level for all input types in the table except for man-days. Results in the profits regressions of Appendix Table G.7 are similar to corresponding ones in Table 3.8, but as before they are not statistically significantly different from zero.

G.2 Results with "simple" predicted repayment regression

We discuss here robustness of treatment effect heterogeneity results to constructing the predicted repayment variable when excluding the locality×(week of initial loan offer) fixed effects. Compared with the predicted repayment regression used in the main results (column 3, Appendix Table G.1), when (locality)×(week of initial loan offer) fixed effects are dropped the R-squared of the regression falls from 0.48 to 0.08.

This simpler regression is then used to predict repayment for the full sample, and the predicted repayment variable is interacted with treatment to examine heterogeneity in the treatment effect. Results from this exercise are presented in Appendix Tables G.8 through G.12, which should be compared (respectively) to the main Tables 3.4 through 3.8.

Results are very similar when using this simpler index of predicted repayment. For example, the coefficients on the interaction between linear predicted repayment and fingerprinting in Panel B remain large in magnitude and retain statistical significance in the repayment and inputs regressions (Appendix Tables G.9 and G.11, respectively). In Panel C, where fingerprinting is interacted with quintiles of predicted repayment, a slight difference vis--vis previous results is that typically the significant interaction term is (fingerprinting)×(quintile 2) rather than the interaction with quintile 1. The main pattern that fingerprinting has more substantial effects on repayment and activities on the farm for individuals with lower predicted repayment is robust to using this simpler predicted repayment regression.

G.3 Results with predicted repayment coefficients obtained from partition of control group

This section describes our approach to estimating predicted repayment using a partition of the control group separate from a partition used as a counterfactual for the treatment group in the main regressions. We conduct this exercise 1,000 times, where in each replication we first randomly select 50% of the control group for inclusion in the auxiliary regression to predict repayment. We then predict repayment for the other half of the control group and the full treatment group. Finally, we estimate the heterogeneous effects of treatment on repayment, land use, input use, and farm profits using equation (3.7) on a sample that includes the full treatment group and the half of the control group not randomly chosen for the auxiliary regression.

We report the 95 percent confidence interval for coefficients obtained from this procedure in Appendix Table G.14. We focus on results for the interaction between the treatment indicator and the indicator for quintile 1 of predicted repayment. Panel A of Appendix Table G.14 corresponds to Table 3.5; Panel B corresponds to Table 3.6; and so on. The coefficient and standard error reported are the original estimates and bootstrap replications using the full sample, as described previously.

In every case, the coefficient from the estimate using the full sample falls within the 95 percent confidence interval from the procedure using the partitioned sample. Furthermore, in every case where the original coefficient is significant, all coefficients in the 95 percent confidence interval of the partitioning exercise have the same sign as the coefficient in the main regressions of the paper, and the confidence interval never includes zero.

Tables

Table G.1: Auxilary regression for predicting loan repayment

	(1)	(2)	(3)
Dependent variable:	Fraction Paid by Sept. 30		
Male	0.080	0.061	0.058
	(0.073)	(0.048)	(0.048)
Married	-0.071	-0.091	-0.101
	(0.060)	(0.044)**	(0.046)**
Age	0.004	0.001	
	(0.001)***	(0.001)	
Years of education	-0.005	-0.003	
	(0.005)	(0.004)	
Risk taker	-0.078	0.008	0.013
	(0.041)*	(0.031)	(0.031)
Days of Hunger in previous season	0.001	-0.000	-0.001
	(0.002)	(0.001)	(0.001)
Late paying previous loan	-0.058	-0.084	-0.084
	(0.071)	(0.046)*	(0.047)*
Standard deviation of past income	-0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
Years of experience growing paprika	0.005	0.007	0.007
	(0.013)	(0.011)	(0.011)
Previous default	0.088	0.128	0.097
	(0.163)	(0.079)	(0.078)
No previous loan	-0.012	0.015	0.013
	(0.062)	(0.032)	(0.034)
Constant	0.729	0.949	0.982
	(0.114)***	(0.072)***	(0.090)***
Locality * week of initial loan offer fixed effects	–	Y	Y
Dummy variables for 5-year age groups	–	–	Y
Dummy variables for each year of education	–	–	Y
Observations	563	563	563
R-squared	0.05	0.46	0.48

Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels.
Sample is non-fingerprinted loan recipients from the September 2008 baseline survey.
All standard errors are clustered at the club level.

Table G.2: Impact of fingerprinting on loan officer knowledge and behavior

	Means			P-value of T-test of (2)=(3)	Num. of obs.
		All			
	(1)	(2)	(3)		
Loan officer reports					
Knows treatment status of club (1=yes)	0.37	0.54	0.22	0.16	51
Knows identity of club officers (1=Yes)	0.47	0.46	0.48	0.88	51
Abs. diff. between actual and officer report of number of loans	1.6	1.3	1.9	0.47	50
Borrower reports					
Number of times loan officer visited club to request loan repayment	0.35	0.41	0.27	0.41	396
Number of times borrower spoke to loan officer since April 2008	2.62	2.57	2.68	0.74	450
Difficulty in locating loan officer (1=easy 2=moderate 3=difficult)	1.2	1.17	1.24	0.32	453

The first three rows present loan officer reports about knowledge of clubs and treatment status collected in August 2008. The last three rows present borrower reports about interactions with the loan officer collected in the follow-up survey of August 2008.

Table G.3: Impact of fingerprinting on attrition from sample

Dependent variable:	Indicator for attrition from September 2008 baseline survey to August 2009 survey	
	(1)	(2)
Sample:	All respondents	Loan recipients
Panel A		
Fingerprint	-0.057	-0.086
	(0.036)	(0.070)
Panel B		
Fingerprint	-0.042	-0.134
	(.107)	(.197)
Predicted repayment * fingerprint	-0.020	0.059
	(.128)	(.225)
Panel C		
Fingerprint * Quintile 1	-0.023	-0.148
	(.075)	(.136)
Fingerprint * Quintile 2	-0.074	0.035
	(.071)	(.109)
Fingerprint * Quintile 3	-0.069	-0.106
	(.068)	(.105)
Fingerprint * Quintile 4	-0.086	-0.109
	(.076)	(.124)
Fingerprint * Quintile 5	-0.080	-0.115
	(.071)	(.128)
Observations	3206	1147
Mean of dependent variable	0.63	0.55
Quintile 1	0.58	0.59
Quintile 2	0.57	0.54
Quintile 3	0.63	0.58
Quintile 4	0.60	0.50
Quintile 5	0.70	0.52

Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels.
Each column presents estimates from three separate regressions: main effect of fingerprinting in Panel A, linear interaction with predicted repayment in Panel B, and interactions with quintiles of predicted repayment in Panel C. All regressions include locality× week of initial loan offer fixed effects, baseline characteristics (male, five-year age categories, one-year education categories, and marriage), and baseline risk indicators (dummy for self-reported risk-taking, days of hunger in the previous season, late payments on previous loans, standard deviation of income, years of experience growing paprika, dummy for default on previous loan, and dummy for no previous loans). Panel B regressions include the main effect of the level of predicted repayment, and Panel C regressions include dummies for quintile of predicted repayment main effects. Standard errors on Panel A coefficients are clustered at the club level, while those in Panels B and C are bootstrapped with 200 replications and club-level resampling.

Table G.4: Impact of fingerprinting on loan repayment

	(1)	(2)	(3)	(4)	(5)	(6)
Sample:	Loan recipients included in August 2009 survey					
Dependent variable:	Balance, Sept. 30	Fraction Paid by Sept. 30	Fully Paid by Sept. 30	Balance, Eventual	Fraction Paid Eventual	Fully Paid Eventual
Panel A						
Fingerprint	-1743.102*	0.073*	0.085	-875.314	0.031	0.060
	(885.950)	(0.044)	(0.069)	(670.297)	(0.032)	(0.057)
Panel B						
Fingerprint	-15386.752***	0.684***	0.759***	-8931.946*	0.362	0.390
	(3782.488)	(.196)	(.213)	(5162.708)	(.237)	(.257)
Predicted repayment * fingerprint	17012.956***	-0.761***	-0.841***	10046.221*	-0.413*	-0.411
	(4018.014)	(.206)	(.240)	(5446.717)	(.250)	(.284)
Panel C						
Fingerprint * Quintile 1	-12684.695***	0.566***	0.599***	-8016.543*	0.334*	0.373*
	(4085.065)	(.195)	(.198)	(4347.488)	(.195)	(.201)
Fingerprint * Quintile 2	1699.375	-0.098	-0.071	1799.143	-0.104	-0.090
	(2125.301)	(.111)	(.168)	(1914.282)	(.101)	(.152)
Fingerprint * Quintile 3	-690.017	0.038	0.052	-586.977	0.032	0.062
	(973.012)	(.055)	(.105)	(871.625)	(.050)	(.097)
Fingerprint * Quintile 4	443.620	-0.029	-0.065	549.532	-0.033	-0.034
	(924.169)	(.053)	(.113)	(821.086)	(.047)	(.103)
Fingerprint * Quintile 5	212.990	-0.006	0.007	289.061	-0.008	0.044
	(978.124)	(.054)	(.110)	(804.733)	(.045)	(.092)
Observations	520	520	520	520	520	520
Mean of dependent variable	2071.21	0.89	0.79	1439.16	0.92	0.83
Quintile 1	6955.67	0.62	0.52	3472.29	0.83	0.71
Quintile 2	4024.05	0.77	0.63	2610.41	0.85	0.75
Quintile 3	1571.44	0.92	0.83	476.63	0.97	0.91
Quintile 4	877.80	0.95	0.85	661.79	0.96	0.86
Quintile 5	1214.19	0.94	0.85	311.66	0.98	0.93

See notes for Table G.3

Table G.5: Impact of fingerprinting on land use

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sample:	Loan recipients included in August 2009 survey								
Dependent variable: Fraction of land used for	Maize	Soya/Beans	Groundnuts	Tobacco	Paprika	Tomatoes	Leafy Vegetables	Cabbage	All cash crops
Panel A									
Fingerprint	-0.010	-0.004	-0.001	-0.003	0.018	0.000	-0.001	0.000	0.010
	(0.012)	(0.012)	(0.009)	(0.009)	(0.011)	(0.002)	(0.002)	(0.001)	(0.012)
Panel B									
Fingerprint	-0.007	-0.045	0.005	-0.005	0.044	0.004	0.006	-0.000	0.008
	(.035)	(.037)	(.029)	(.029)	(.029)	(.004)	(.008)	(.001)	(.035)
Predicted repayment * fingerprint	-0.004	0.058	-0.007	0.003	-0.037	-0.005	-0.009	0.001	0.003
	(.044)	(.045)	(.035)	(.033)	(.038)	(.006)	(.009)	(.002)	(.044)
Panel C									
Fingerprint * Quintile 1	0.001	-0.046*	-0.004	-0.004	0.050**	0.002	0.001	0.000	-0.001
	(.027)	(.027)	(.021)	(.022)	(.022)	(.003)	(.006)	(.001)	(.027)
Fingerprint * Quintile 2	-0.029	0.025	0.030	-0.003	-0.022	0.000	0.001	-0.001	0.030
	(.033)	(.031)	(.023)	(.023)	(.025)	(.005)	(.005)	(.001)	(.033)
Fingerprint * Quintile 3	-0.015	-0.010	-0.011	-0.008	0.044*	0.004	-0.004	0.001	0.016
	(.030)	(.028)	(.021)	(.018)	(.026)	(.005)	(.005)	(.002)	(.030)
Fingerprint * Quintile 4	-0.006	0.017	-0.012	0.002	0.002	-0.001	-0.003	0.002	0.007
	(.029)	(.030)	(.024)	(.017)	(.028)	(.006)	(.005)	(.002)	(.029)
Fingerprint * Quintile 5	0.003	0.021	-0.016	-0.009	0.003	-0.005	-0.002	-0.001	-0.008
	(.033)	(.028)	(.023)	(.020)	(.027)	(.005)	(.004)	(.002)	(.033)
Observations	1226	1226	1226	1226	1226	1226	1226	1226	1226
Mean of dependent variable	0.46	0.16	0.12	0.09	0.15	0.01	0.01	0.00	0.54
Quintile 1	0.46	0.11	0.13	0.16	0.12	0.00	0.01	0.00	0.54
Quintile 2	0.49	0.12	0.13	0.13	0.12	0.00	0.00	0.00	0.51
Quintile 3	0.45	0.22	0.12	0.03	0.17	0.01	0.01	0.00	0.55
Quintile 4	0.44	0.21	0.12	0.04	0.19	0.01	0.01	0.00	0.56
Quintile 5	0.47	0.17	0.11	0.05	0.17	0.01	0.01	0.00	0.52

See notes for Table G.3

Table G.6: Impact of fingerprinting on agricultural inputs used on paprika crop

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Sample:	Loan recipients included in August 2009 survey						
Dependent variable:	Seeds (MK)	Fertilizer (MK)	Chemicals (MK)	Man-days (MK)	All Paid Inputs (MK)	KG Manure	Times Weeding
Panel A							
Fingerprint	50.470*	737.425	271.094**	-167.300	891.689	3.965	0.054
	(29.015)	(640.797)	(105.270)	(116.390)	(757.696)	(27.703)	(0.120)
Panel B							
Fingerprint	38.748	3603.741**	346.463	183.782	4172.733**	112.813	0.365
	(77.354)	(1781.483)	(223.419)	(177.176)	(1939.315)	(95.394)	(.307)
Predicted repayment * fingerprint	16.381	-4005.460*	-105.322	-490.610*	-4585.011*	-152.107	-0.435
	(100.259)	(2254.326)	(300.036)	(289.354)	(2537.854)	(118.904)	(.427)
Panel C							
Fingerprint * Quintile 1	95.528**	2271.800*	279.824*	66.218	2713.370**	112.236*	0.436*
	(45.098)	(1261.271)	(159.893)	(128.375)	(1368.638)	(65.572)	(.233)
Fingerprint * Quintile 2	55.135	1865.945	325.669	-127.295	2119.453	-63.420	-0.301
	(59.103)	(1439.897)	(208.452)	(270.537)	(1629.287)	(68.651)	(.297)
Fingerprint * Quintile 3	89.214	-576.071	314.824	-191.217	-363.251	-80.829	-0.181
	(66.460)	(1419.720)	(249.577)	(293.996)	(1650.901)	(75.873)	(.311)
Fingerprint * Quintile 4	3.099	-1269.499	259.021	-629.191*	-1636.570	-30.861	-0.104
	(74.887)	(1533.243)	(266.207)	(353.683)	(1905.085)	(71.051)	(.329)
Fingerprint * Quintile 5	62.079	1.107	253.485	-156.877	159.795	26.081	0.305
	(77.898)	(1402.017)	(233.795)	(322.488)	(1698.568)	(83.457)	(.322)
Observations	1226	1226	1226	1226	1226	1226	1226
Mean of dependent variable	185.56	3948.93	362.92	396.56	4893.98	83.63	1.54
Quintile 1	129.13	2335.99	182.64	152.25	2800.01	85.47	1.20
Quintile 2	132.03	2924.01	277.92	178.55	3512.50	59.29	1.28
Quintile 3	198.08	5481.54	426.67	593.28	6699.57	125.01	1.78
Quintile 4	237.16	5837.92	543.52	726.95	7345.56	83.53	1.91
Quintile 5	239.52	4786.95	516.97	579.91	6123.35	80.35	1.83

See notes for Table G.3

Table G.7: Impact of fingerprinting on revenue and profits

	(1)	(2)	(3)	(4)
Sample:		Loan recipients included in August 2009 survey		
Dependent variable:	Market sales (Self Report, MK)	Value of Unsold Harvest (Regional Prices, MK)	Profits (market sales + value of unsold harvest - cost of inputs, MK)	Ln(profits)
Panel A				
Fingerprint	1617.616	2720.001	3548.584	0.016
	(4178.351)	(27456.635)	(28440.566)	(0.074)
Panel B				
Fingerprint	20134.232	18828.141	31797.814	0.196
	(15678.99)	(86516.02)	(91020.54)	(.223)
Predicted repayment * fingerprint	-25875.572	-22509.908	-39476.163	-0.252
	(18176.9)	(94749.57)	(100277.2)	(.257)
Panel C				
Fingerprint * Quintile 1	8754.046	47355.641	52526.195	0.068
	(36850.940)	(50587.570)	(63206.880)	(.363)
Fingerprint * Quintile 2	18385.539	-51226.844	-38430.247	0.167
	(33084.250)	(75026.680)	(81848.520)	(.264)
Fingerprint * Quintile 3	-18623.896	2028.603	-13189.753	-0.154
	(17938.730)	(59934.020)	(63400.760)	(.227)
Fingerprint * Quintile 4	-6705.782	1150.111	-785.509	0.027
	(14733.930)	(57096.050)	(60385.110)	(.231)
Fingerprint * Quintile 5	-2174.627	-5799.395	-9926.923	-0.092
	(15282.460)	(71891.840)	(74254.990)	(.240)
Observations	1226	1226	1226	1226
Mean of dependent variable	53965.29	86793.08	119870.13	11.28
Quintile 1	48912.14	103543.10	138101.00	11.23
Quintile 2	70582.23	60989.97	109699.20	11.33
Quintile 3	44931.14	86190.55	108497.00	11.27
Quintile 4	54127.28	98467.02	125928.20	11.34
Quintile 5	47991.75	84126.01	109740.80	11.26
Mean of dependent variable (US $)	385.47	619.95	856.22	n.a.

See notes for Table G.3

Table G.8: Impact of fingerprinting on loan approval, loan take-up, and amount borrowed

	(1)	(2)	(3)
Sample:	All Respondents	All Respondents	Loan Recipients
Dependent variable:	Approved	Any Loan	Total Borrowed (MK)
Panel A			
Fingerprint	0.038	0.051	-696.799*
	(0.053)	(0.044)	(381.963)
Panel B			
Fingerprint	-0.090	-0.016	-692.208
	(.151)	(.144)	(2398.497)
Predicted repayment * fingerprint	0.158	0.083	-5.664
	(.171)	(.175)	(2700.889)
Panel C			
Fingerprint * Quintile 1	0.013	0.064	76.320
	(.073)	(.062)	(871.746)
Fingerprint * Quintile 2	0.030	0.032	-1237.509
	(.069)	(.062)	(704.973)
Fingerprint * Quintile 3	0.068	0.026	-1675.665
	(.069)	(.067)	(680.063)
Fingerprint * Quintile 4	0.019	0.047	-389.588
	(.069)	(.064)	(655.217)
Fingerprint * Quintile 5	0.053	0.087	-224.950
	(.072)	(.067)	(598.228)
Observations	3277	3277	1147
Mean of dependent variable	0.63	0.35	16912.60
Quintile 1	0.58	0.29	17992.53
Quintile 2	0.64	0.36	17870.61
Quintile 3	0.71	0.44	16035.10
Quintile 4	0.70	0.47	15805.54
Quintile 5	0.59	0.30	16886.56

See notes for Table G.3. These regressions use the “simple” measure of predicted repayment.

Table G.9: Impact of fingerprinting on loan repayment

	(1)	(2)	(3)	(4)	(5)	(6)
Sample:	Loan recipients					
Dependent variable:	Balance, Sept. 30	Fraction Paid by Sept. 30	Fully Paid by Sept. 30	Balance, Eventual	Fraction Paid Eventual	Fully Paid Eventual
Panel A						
Fingerprint	-1556.383*	0.073*	0.096	-996.430	0.045	0.085
	(824.174)	(0.040)	(0.062)	(754.301)	(0.036)	(0.058)
Panel B						
Fingerprint	-11352.022***	0.556***	0.656***	-6991.700**	0.321*	0.417*
	(3680.308)	(.174)	(.249)	(3553.984)	(.171)	(.232)
Predicted repayment * fingerprint	12085.689***	-0.595***	-0.692**	7396.859*	-0.340*	-0.409
	(4171.199)	(.193)	(.271)	(3947.945)	(.187)	(.251)
Panel C						
Fingerprint * Quintile 1	-2544.854	0.134*	0.159	-1201.900	0.059	0.093
	(1565.447)	(.077)	(.100)	(1499.982)	(.073)	(.097)
Fingerprint * Quintile 2	-3435.559**	0.157**	0.208**	-2590.627*	0.123*	0.180*
	(1561.394)	(.076)	(.101)	(1470.999)	(.072)	(.099)
Fingerprint * Quintile 3	-1825.518	0.065	0.115	-1139.191	0.024	0.115
	(1271.12)	(.066)	(.098)	(1215.847)	(.062)	(.094)
Fingerprint * Quintile 4	-512.427	0.033	0.011	-541.084	0.035	0.046
	(1130.745)	(.056)	(.086)	(1011.273)	(.050)	(.077)
Fingerprint * Quintile 5	313.466	-0.012	-0.001	322.623	-0.008	0.003
	(998.939)	(.046)	(.073)	(904.357)	(.041)	(.068)
Observations	1147	1147	1147	1147	1147	1147
Mean of dependent variable	2912.91	0.84	0.74	2080.86	0.89	0.79
Quintile 1	6955.67	0.62	0.52	4087.04	0.81	0.68
Quintile 2	4024.05	0.77	0.63	3331.17	0.81	0.67
Quintile 3	1571.44	0.92	0.83	1301.79	0.93	0.84
Quintile 4	877.80	0.95	0.85	781.59	0.95	0.87
Quintile 5	1214.19	0.94	0.85	950.29	0.95	0.88

See notes for Table G.3. These regressions use the "simple" measure of predicted repayment.

Table G.10: Impact of fingerprinting on land use

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable: Fraction of land used for	Maize	Soya/Beans	Groundnuts	Tobacco	Paprika	Tomatoes	Leafy Vegetables	Cabbage	All cash crops
Panel A									
Fingerprint	0.001	0.015	-0.012	-0.004	0.005	-0.001	-0.003	-0.001	-0.001
	(0.019)	(0.019)	(0.016)	(0.016)	(0.014)	(0.003)	(0.003)	(0.001)	(0.019)
Panel B									
Fingerprint	-0.175	-0.022	0.048	0.015	0.101	0.015	0.024	-0.005	0.175
	(.126)	(.100)	(.092)	(.069)	(.089)	(.018)	(.019)	(.006)	(.126)
Predicted repayment * fingerprint	0.215	0.045	-0.073	-0.023	-0.117	-0.019	-0.032	0.005	-0.215
	(.139)	(.113)	(.108)	(.079)	(.103)	(.021)	(.021)	(.007)	(.139)
Panel C									
Fingerprint * Quintile 1	-0.021	0.021	-0.008	-0.002	0.009	0.000	0.003	-0.003	0.021
	(.057)	(.047)	(.035)	(.033)	(.037)	(.006)	(.007)	(.003)	(.057)
Fingerprint * Quintile 2	-0.033	-0.011	-0.055	0.013	0.080**	0.008	-0.001	-0.002	0.033
	(.050)	(.043)	(.038)	(.031)	(.038)	(.008)	(.007)	(.004)	(.050)
Fingerprint * Quintile 3	0.007	-0.001	0.033	-0.011	-0.018	-0.007	-0.002	-0.001	-0.007
	(.046)	(.042)	(.036)	(.028)	(.040)	(.009)	(.007)	(.003)	(.046)
Fingerprint * Quintile 4	0.021	0.020	0.018	-0.022	-0.031	-0.002	-0.006	0.001	-0.021
	(.039)	(.036)	(.033)	(.024)	(.036)	(.009)	(.007)	(.002)	(.039)
Fingerprint * Quintile 5	0.021	0.034	-0.056*	0.007	0.004	-0.004	-0.006	-0.001	-0.021
	(.033)	(.034)	(.033)	(.022)	(.029)	(.008)	(.006)	(.001)	(.033)
Observations	520	520	520	520	520	520	520	520	520
Mean of dependent variable	0.43	0.15	0.13	0.08	0.19	0.01	0.00	0.00	0.57
Quintile 1	0.44	0.07	0.13	0.18	0.17	0.01	0.01	0.00	0.56
Quintile 2	0.49	0.10	0.13	0.13	0.15	0.00	0.00	0.00	0.51
Quintile 3	0.42	0.21	0.12	0.03	0.20	0.01	0.00	0.00	0.58
Quintile 4	0.42	0.19	0.12	0.04	0.21	0.01	0.01	0.00	0.58
Quintile 5	0.40	0.17	0.14	0.04	0.23	0.01	0.01	0.00	0.60

See notes for Table G.3. These regressions use the "simple" measure of predicted repayment.

Table G.11: Impact of fingerprinting on agricultural inputs used on paprika crop

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:	Seeds (MK)	Fertilizer (MK)	Chemicals (MK)	Man-days (MK)	All Paid Inputs (MK)	KG Manure	Times Weeding
Panel A							
Fingerprint	74.107	733.419	345.328*	-395.501**	757.354	29.649	0.019
	(47.892)	(1211.905)	(190.262)	(181.958)	(1389.230)	(32.593)	(0.147)
Panel B							
Fingerprint	244.347	12905.576**	1051.749	138.645	14340.317**	163.231	0.561*
	(258.741)	(5258.075)	(996.875)	(1418.959)	(6018.261)	(197.135)	(.937)
Predicted repayment * fingerprint	-207.279	-14820.444**	-860.116	-650.360	-16538.199**	-162.645	-0.660*
	(317.171)	(5932.685)	(1171.106)	(1699.914)	(6898.503)	(230.615)	(1.110)
Panel C							
Fingerprint * Quintile 1	168.071*	2182.842	535.269	-474.586	2411.596	-19.357	0.017
	(100.037)	(2408.763)	(339.744)	(478.934)	(2729.886)	(74.719)	(.360)
Fingerprint * Quintile 2	197.376*	5738.649**	468.314	-133.943	6270.395**	126.483	0.129
	(102.322)	(2554.428)	(348.92)	(534.438)	(2881.465)	(86.882)	(.326)
Fingerprint * Quintile 3	-66.591	140.304	472.752	-591.554	-45.088	33.042	0.089
	(106.405)	(2603.701)	(421.368)	(566.615)	(3002.609)	(77.734)	(.350)
Fingerprint * Quintile 4	7.974	-1593.374	279.405	-193.117	-1499.112	16.910	0.047
	(94.280)	(2187.822)	(418.937)	(510.335)	(2578.145)	(76.278)	(.340)
Fingerprint * Quintile 5	100.395	-837.790	162.736	-567.242	-1141.901	-13.431	-0.090
	(112.185)	(2135.513)	(384.179)	(528.735)	(2548.739)	(77.566)	(.293)
Observations	520	520	520	520	520	520	520
Mean of dependent variable	247.06	7499.85	671.31	665.98	9084.19	90.84	1.94
Quintile 1	174.13	6721.24	401.30	143.48	7440.15	97.39	1.47
Quintile 2	140.00	6080.46	620.67	238.94	7080.08	39.25	1.55
Quintile 3	269.90	8927.65	674.48	836.98	10709.00	105.73	2.05
Quintile 4	292.07	7649.51	715.08	936.29	9592.95	93.23	2.24
Quintile 5	340.18	8078.58	892.05	1065.18	10375.99	118.13	2.28

See notes for Table G.3. These regressions use the “simple” measure of predicted repayment.

Table G.12: Impact of fingerprinting on revenue and profits

Dependent variable:	(1) Market sales (Self Report, MK)	(2) Value of Unsold Harvest (Regional Prices, MK)	(3) Profits (market sales + value of unsold harvest - cost of inputs, MK)	(4) Ln(profits)
Panel A				
Fingerprint	7246.174	5270.320	14509.457	0.060
	(8792.055)	(14879.349)	(16679.311)	(0.095)
Panel B				
Fingerprint	105858.940	-77048.968	6406.374	1.143*
	(50814.47)	(162717.6)	(171947.4)	(.627)
Predicted repayment * fingerprint	-120067.879	100229.440	9866.066	-1.318*
	(55777.81)	(197387.4)	(208686.7)	(.744)
Panel C				
Fingerprint * Quintile 1	16476.436	-5726.186	7300.249	-0.055
	(21944.42)	(44155.49)	(51075.73)	(.253)
Fingerprint * Quintile 2	35444.339	-9950.472	16468.680	0.535**
	(22735.24)	(57621.38)	(62562.85)	(.269)
Fingerprint * Quintile 3	-1029.777	40054.423	49032.571	0.229
	(21583.30)	(58961.12)	(64835.04)	(.239)
Fingerprint * Quintile 4	-1653.502	-65999.076	-64498.215	-0.106
	(17323.32)	(66483.52)	(70231)	(.245)
Fingerprint * Quintile 5	-5201.902	51729.358	51882.714	-0.197
	(15736.72)	(63939.72)	(65743.91)	(.211)
Observations	520	520	520	520
Mean of dependent variable	65004.30	80296.97	117779.16	11.44
Quintile 1	60662.57	82739.24	121222.50	11.36
Quintile 2	89028.25	29995.27	91652.71	11.55
Quintile 3	57683.74	96247.91	123242.30	11.44
Quintile 4	61088.27	104927.50	136467.50	11.45
Quintile 5	56593.43	85817.08	115172.50	11.39
Mean of dependent variable (US $)	464.32	573.55	841.28	n.a.

See notes for Table G.3. These regressions use the "simple" measure of predicted repayment.

Table G.13: Auxilary regression for predicting loan repayment, no fixed effects

	(1)	(2)	(3)
Dependent variable:	Fraction Paid by Sept. 30		
Male	0.080	0.074	
	(0.073)	(0.071)	
Married	-0.071	-0.080	
	(0.060)	(0.065)	
Age	0.004		
	(0.001)***		
Years of education	-0.005		
	(0.005)		
Risk taker	-0.078	-0.072	
	(0.041)*	(0.043)*	
Days of Hunger in previous season	0.001	0.000	
	(0.002)	(0.001)	
Late paying previous loan	-0.058	-0.045	
	(0.071)	(0.067)	
Standard deviation of past income	-0.000	-0.000	
	(0.000)	(0.000)	
Years of experience growing paprika	0.005	0.004	
	(0.013)	(0.012)	
Previous default	0.088	0.062	
	(0.163)	(0.169)	
No previous loan	-0.012	-0.009	
	(0.062)	(0.061)	
Constant	0.729	1.006	
	(0.114)***	(0.108)***	
Locality * week of initial loan offer fixed effects	–	–	
Dummy variables for 5-year age groups	–	Y	
Dummy variables for each year of education	–	Y	
Observations	563	563	
R-squared	0.05	0.08	

Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels.
Sample is non-fingerprinted loan recipients from the September 2008 baseline survey.
All standard errors are clustered at the club level.

Table G.14: 95% confidence interval of Q1×Treatment interaction term from partitioning exercise

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Corresponds to Table 3.5									
Dependent variable:	Balance, Sept. 30	Fraction Paid by Sept. 30	Fully Paid by Sept. 30	Balance, Eventual	Fraction Paid Eventual	Fully Paid Eventual			
Coefficient: Fingerprint×Quintile 1	-10844.169***	0.499***	0.543***	-7202.647**	0.317**	0.396**			
Bootstrapped standard error	(2681.861)	(.127)	(.147)	(2969.045)	(.136)	(.156)			
95 percent confidence interval using half of control group in 1st stage	[-11279.51, -5421.963]	[.267, .52]	[.304, .629]	[-8038.583, -2882.11]	[.132, .361]	[.176, .518]			
Panel B: Corresponds to Table 3.6									
Dependent variable: Fraction of land used for	Maize	Soya/Beans	Groundnuts	Tobacco	Paprika	Tomatoes	Leafy Vegetables	Cabbage	All cash crops
Coefficient: Fingerprint×Quintile 1	-0.061	-0.013	-0.008	-0.012	0.083	0.005	0.007	-0.002	0.061
Bootstrapped standard error	(.066)	(.063)	(.052)	(.050)	(.051)	(.008)	(.012)	(.003)	(.066)
95 percent confidence interval using half of control group in 1st stage	[-.151, .054]	[-.064, .073]	[-.073, .062]	[-.093, .057]	[-.006, .131]	[-.005, .011]	[-.015, .011]	[-.008, .002]	[-.054, .151]
Panel C: Corresponds to Table 3.7									
Dependent variable:	Seeds (MK)	Fertilizer (MK)	Chemicals (MK)	Man-days (MK)	All Paid Inputs (MK)	KG Manure	Times Weeding		
Coefficient: Fingerprint×Quintile 1	188.703**	5871.126	374.260	106.406	6540.496	78.234	0.445		
Bootstrapped standard error	(95.018)	(4062.716)	(406.741)	(347.367)	(4210.469)	(111.980)	(.367)		
95 percent confidence interval using half of control group in 1st stage	[55.431, 285.169]	[1209.801, 9846.745]	[-.048, 946.45]	[-347.262, 358.159]	[1755.798, 10490.54]	[-99.285, 183.423]	[-.247, .839]		
Panel D: Corresponds to Table 3.8									
Dependent variable:	Market sales (Self Report, MK)	Value of Unsold Harvest (Regional Prices, MK)	Profits (market sales + value of unsold harvest - cost of inputs, MK)	Ln(profits)					
Coefficient: Fingerprint×Quintile 1	30766.147	7940.835	31915.287	0.401					
Bootstrapped standard error	(36850.940)	(50587.570)	(63206.880)	(.363)					
95 percent confidence interval using half of control group in 1st stage	[6749.634, 94461.64]	[-121480.1, 71494.77]	[-82988.08, 127890.7]	[-.073, .942]					

The coefficients reported in this table are for the interaction between the indicator for being in the bottom quintile of predicted repayment and being assigned to have a fingerprint collected when applying for a loan. They correspond to the coefficients for bottom quintile in Panel C of Tables 3.5, 3.6, 3.7, and 3.8. The standard errors are the bootstrapped standard errors reported in those tables. The confidence intervals are from 1000 replications of each regression where one half of the control group was randomly chosen for inclusion in the first stage regression, and the remaining half of the control group plus the full treatment group was preserved for inclusion in the second stage regression.

Table G.15: Robustness check of confidence intervals

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Corresponds to Table 3.5									
Dependent variable:	Balance, Sept. 30	Fraction Paid by Sept. 30	Fully Paid by Sept. 30	Balance, Eventual	Fraction Paid Eventual	Fully Paid Eventual			
Coefficient: Fingerprint×Quintile 1	-10844.169***	0.499***	0.543***	-7202.647**	0.317**	0.396**			
Bootstrapped standard error	(2681.861)	(.127)	(.147)	(2969.045)	(.136)	(.156)			
95 percent confidence interval using half of control group in 1st stage	[-14371.41, -1832.52]	[.079, .671]	[.087, .796]	[-11903.3, 597.541]	[-.035, .539]	[-.017, .685]			
Panel B: Corresponds to Table 3.6									
Dependent variable: Fraction of land used for	Maize	Soya/Beans	Groundnuts	Tobacco	Paprika	Tomatoes	Leafy Vegetables	Cabbage	All cash crops
Coefficient: Fingerprint×Quintile 1	-0.061	-0.013	-0.008	-0.012	0.083	0.005	0.007	-0.002	0.061
Bootstrapped standard error	(.066)	(.063)	(.052)	(.050)	(.051)	(.008)	(.012)	(.003)	(.066)
95 percent confidence interval using half of control group in 1st stage	[-.188, .14]	[-.178, .163]	[-.144, .096]	[-.151, .105]	[-.082, .167]	[-.013, .021]	[-.024, .016]	[-.014, .004]	[-.14, .188]
Panel C: Corresponds to Table 3.7									
Dependent variable:	Seeds (MK)	Fertilizer (MK)	Chemicals (MK)	Man-days (MK)	All Paid Inputs (MK)	KG Manure	Times Weeding		
Coefficient: Fingerprint×Quintile 1	188.703**	5871.126	374.260	106.406	6540.496	78.234	0.445		
Bootstrapped standard error	(95.018)	(4062.716)	(406.741)	(347.367)	(4210.469)	(111.980)	(.367)		
95 percent confidence interval using half of control group in 1st stage	[-98.689, 396.033]	[-4404.095, 12100.82]	[-486.832, 1619.236]	[-1071.866, 716.063]	[-3893.477, 13651.09]	[-305.1, 326.601]	[-.793, 1.1]		
Panel D: Corresponds to Table 3.8									
Dependent variable:	Market sales (Self Report, MK)	Value of Unsold Harvest (Regional Prices, MK)	Profits (market sales + value of unsold harvest - cost of inputs, MK)	Ln(profits)					
Coefficient: Fingerprint×Quintile 1	30766.147	7940.835	31915.287	0.401					
Bootstrapped standard error	(36850.940)	(50587.570)	(63206.880)	(.363)					
95 percent confidence interval using half of control group in 1st stage	[-21692.7, 155189.1]	[-176686.8, 144512.1]	[-166265.2, 229394]	[-.371, 1.357]					

See notes for Table G.14.

Table G.16: Robustness check of confidence intervals

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Corresponds to Table 3.5									
Dependent variable:	Balance, Sept. 30	Fraction Paid by Sept. 30	Fully Paid by Sept. 30	Balance, Eventual	Fraction Paid Eventual	Fully Paid Eventual			
Coefficient: Fingerprint×Quintile 1	-10844.169***	0.499***	0.543***	-7202.647**	0.317**	0.396**			
Bootstrapped standard error	(2681.861)	(.127)	(.147)	(2969.045)	(.136)	(.156)			
95 percent confidence interval using half of control group in 1st stage	[-9569.22, 887.283]	[-.036, .444]	[-.097, .643]	[-7312.281, 1637.288]	[-.084, .29]	[-.119, .565]			
Panel B: Corresponds to Table 3.6									
Dependent variable: Fraction of land used for	Maize	Soya/Beans	Groundnuts	Tobacco	Paprika	Tomatoes	Leafy Vegetables	Cabbage	All cash crops
Coefficient: Fingerprint×Quintile 1	-0.061	-0.013	-0.008	-0.012	0.083	0.005	0.007	-0.002	0.061
Bootstrapped standard error	(.066)	(.063)	(.052)	(.050)	(.051)	(.008)	(.012)	(.003)	(.066)
95 percent confidence interval using half of control group in 1st stage	[-.166, .181]	[-.1, .16]	[-.145, .081]	[-.11, .059]	[-.042, .113]	[-.011, .014]	[-.025, .013]	[-.009, .003]	[-.181, .166]
Panel C: Corresponds to Table 3.7									
Dependent variable:	Seeds (MK)	Fertilizer (MK)	Chemicals (MK)	Man-days (MK)	All Paid Inputs (MK)	KG Manure	Times Weeding		
Coefficient: Fingerprint×Quintile 1	188.703**	5871.126	374.260	106.406	6540.496	78.234	0.445		
Bootstrapped standard error	(95.018)	(4062.716)	(406.741)	(347.367)	(4210.469)	(111.980)	(.367)		
95 percent confidence interval using half of control group in 1st stage	[-110.111, 292.123]	[-4288.664, 10087.83]	[-218.602, 1044.806]	[-1329.363, 366.171]	[-4015.115, 10794.47]	[-121.158, 278.817]	[-.901, 1.016]		
Panel D: Corresponds to Table 3.8									
Dependent variable:	Market sales (Self Report, MK)	Value of Unsold Harvest (Regional Prices, MK)	Profits (market sales + value of unsold harvest - cost of inputs, MK)	Ln(profits)					
Coefficient: Fingerprint×Quintile 1	30766.147	7940.835	31915.287	0.401					
Bootstrapped standard error	(36850.940)	(50587.570)	(63206.880)	(.363)					
95 percent confidence interval using half of control group in 1st stage	[-12025.86, 179096.6]	[-142196.9, 92000.43]	[-119491.3, 197887.8]	[-.416, 1.069]					

See notes for Table G.14.

BIBLIOGRAPHY

BIBLIOGRAPHY

Abdulai, A., and C. Delgado (1999), Determinants of nonfarm earnings of farm-based husbands and wives in northern ghana, *American Journal of Agricultural Economics*, pp. 117–130.

Angrist, J. D., and J.-S. Pischke (2009), *Mostly Harmless Econometrics*, Princeton University Press.

Ashenfelter, O., K. Doran, and B. Schaller (2010), A shred of credible evidence on the long run elasticity of labor supply, *NBER Working Paper*.

Ashraf, N. (2009), Spousal control and intra-household decision making: An experimental study in the philippines, *American Economic Review*, *99*(4), 1245–1277.

Ausubel, L. M. (1999), Adverse selection in the credit card market, *Working paper*, University of Maryland.

Badilla, J., and M. Pagano (2000), Sharing default information as a borrower discipline device, *European Economic Review*, *44*(10), 1951–1980.

Baland, J., C. Guirkinger, and C. Mali (2007), Pretending to be poor: Borrowing to escape forced solidarity in credit cooperatives in cameroon, *Working Paper, Center for Research in Economic Development, University of Namur*.

Bardhan, P. (1979), Labor supply functions in a poor agrarian economy, *American Economic Review*, pp. 73–83.

Barmby, T., and P. Dolton (2009), What lies beneath? effort and incentives on archaelogical digs in the 1930's, *Working Paper*, pp. 1–38.

Bester, H. (1985), Screening vs. rationing in credit markets with imperfect information, *American Economic Review*, *75*(4), 850–855.

Boot, A. W. A., and A. V. Takor (1994), Moral hazard and secured lending in an infinitely repeated credit market game, *International Economic Review*, *35*, 899–920.

Bulow, J., and K. Rogoff (1989), A constant recontracting model of sovereign debt, *Journal of Political Economy*, *97*(1), 155–178.

Burtless, G., and J. A. Hausman (1978), The effect of taxation on labor supply: evaluating the gary negative income tax experiment, *Journal of Political Economy*, *86*(6), 1103–1130.

Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler (1997), Labor supply of new york city cabdrivers: one day at a time, *Quarterly Journal of Economics*, pp. 407–441.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008), Bootstrap-based improvements for inference with clustered errors, *Review of Economics and Statistics*, *90*(3), 414–427.

Chassagnon, A., and P. A. Chiappori (1997), Insurance under moral hazard and adverse selection: the case of pure competition, *Manuscript*, Departement et Laboratoire d'Economie Theorique et Appliquee.

Chiappori, P. A., and B. Salanie (2000), Testing assymetric information in insurance markets, *Journal of Political Economy*, *108*, 56–78.

Chiappori, P. A., and B. Salanie (2003), Testing contract theory: a survey of some recent work, in *Advances in Economics and Econometrics: Theory and Applications, Eight World Congress*, vol. III, edited by M. Dewatripont, L. P. Hansen, and S. Turnovsky, pp. 115–149, Cambridge University Press.

Chiappori, P. A., B. Jullien, B. Salanie, and F. Salanie (2006), Assymetric information in insurance: general testable implications, *Rand Journal of Economics*, *37*, 783–798.

Chou, Y. K. (2000), Testing alternative models of labor supply: evidence from taxi-drivers in singapore, *University of Melbourne Department of Economics Working Paper Series*, *768*, 1–39.

Comola, M., and M. Fafchamps (2010), Are gifts and loans between households voluntary?, *Centre for the Study of African Economies Working Paper Series*, (20).

Conning, J., and C. Udry (2005), Rural financial markets in developing countries, in *The Handbook of Agricultural Economics*, vol. 3, edited by R. E. Everson, P. Pingali, and T. P. Schultz, Elsevier Science.

Edelberg, W. (2004), Testing for adverse selection and moral hazard in consumer loan markets, *Finance and economics discussion paper series*, Board of Governors of the Federal Reserve System.

Fafchamps, M. (2004), *Market Institutions in Sub-Saharan Africa: Theory and Evidence*, MIT Press.

Farber, H. S. (2003), Is tomorrow another day? the labor supply of new york cab drivers, *NBER*, p. 43.

Fehr, E., and L. Goette (2007), Do workers work more if wages are high? evidence from a randomized field experiment, *American Economic Review*, *97*(1), 298–317.

Gine, X., and D. Yang (2009), Insurance, credit, and technology adoption: field experimental evidence from malawi, *Journal of Development Economics*, *89*, 1–11.

Gine, X., P. Jakiela, D. Karlan, and J. Morduch (2010), Microfinance games, *American Economic Journal: Applied Economics*, *2*(3), 60–95.

Guesnerie, R., P. Picard, and P. Rey (1988), Adverse selection and moral hazard with risk neutral agents, *European Economic Review*, *33*, 807–823.

Heckman, J. J. (1993), What has been learned about labor supply in the past twenty years?, *American Economic Review*, *83*(2), 116–121.

Jakiela, P. (2009), How fair shares compare: Experimental evidence from two cultures.

Janvry, D., A. C. McIntosh, and E. Sadoulet (forthcoming), The supply and demand side impacts of credit market information, *Journal of Development Economics*.

Karlan, D., and J. Zinman (2009), Observing unobservables: identifying information asymmetries with a consumer credit field experiment, *Econometrica*, *77*(6), 1993–2008.

Killingsworth, M. R. (1983), *Labor Supply*, Cambridge Surveys of Economic Literature, Press Syndicate of the University of Cambridge.

Klonner, S., and A. S. Rai (2009), Adverse selection in credit markets: evidence from bidding roscas, *Working paper*, Williams College.

Kochar, A. (1999), 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.

Lewis, W. A. (1954), Economic development with unlimited supplies of labour, *The Manchester School of Economic and Social Studies*, *XXII*(2), 139–191.

Ligon, E., J. P. Thomas, and T. Worrall (2002), Informal insurance arrangements with limited commitment: Theory and evidence from village economies, *Review of Economic Studies*, *69*(1), 209–244.

Livshits, I., J. MacGee, and M. Tertilt (2010), Accounting for the rise in consumer bankruptcies, *American Economic Journal: Macroeconomics*, *2*(2), 165–193.

Maranz, D. (2001), *African Friends and Money Matters: Observations from Africa*, *Publications in Ethnography*, vol. 37, SIL International.

McCord, A., and R. Slater (2009), Overview of public works programmes in subsaharan africa, *Tech. rep.*, Overseas Development Institute.

Meyer, B. (2002), Labor supply at the extensive and intensive margins: the eitc, welfare, and hours worked, *American Economic Review.*

Moulton, B. (1986), Random group effects and the precision of regression estimates, *Journal of Econometrics*, *32*(3), 385–397.

Narajabad, B. N. (2010), Information technology and the rise of household bankruptcy, job Market Paper, University of Texas at Austin.

Oettinger, G. S. (1999), An empirical analysis of the daily labor supply of stadium vendors, *Journal of Political Economy*, *107*(2), 360–392.

Pagano, M., and T. Jappelli (1993), Information sharing in credit markets, *Journal of Finance*, *48*, 1693–1718.

Paulson, A. L., R. M. Townsend, and A. Karaivanov (2006), Distinguishing limited commitment from moral hazard in models of growth with inequality, *Journal of Political Economy.*

Rose, E. (2001), Ex ante and ex post labor supply response to risk in a low-income area, *Journal of Development Economics*, *64*, 371–388.

Rosenzweig, M. (1978), Rural wages, labor supply, and land reform: A theoretical and empirical analysis, *The American Economic Review*, pp. 847–861.

Sanchez, J. M. (2009), The role of information in the rise of consumer bankruptcies, *Working Paper 09-4*, Federal Reserve Bank of Richmond.

Stiglitz, J. E., and A. Weiss (1983), Incentive effects of terminations: applications to the credit and labor markets, *American Economic Review*, *73*(5), 912–927.

Townsend, R. M. (1994), Risk and insurance in village india, *Econometrica*, *62*(3), 539–591.

Visaria, S. (2009), Legal reform and loan repayment: the microeconomic impact of debt recovery tribunals in india, *American Economic Journal: Applied Economics*, pp. 59–81.