Three Field Experiments in Development Economics

by

Lasse Florian Brune

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy (Economics) in the University of Michigan 2014

Doctoral Committee:

Associate Professor Dean Yang, Chair
Professor Charles C. Brown
Assistant Professor Stephen Leider
Associate Professor Melvin Stephens Jr.

Dedicated to my parents, to my brother, and to Sasha.

ACKNOWLEDGEMENTS

I am deeply grateful to my dissertation committee chair Dean Yang for his guidance and support during my time in graduate school, and to my dissertation committee members Mel Stephens, Steve Leider and Charlie Brown for their insightful comments throughout. I also thank Rebecca Thornton, Manuela Angelucci and Raj Arunachalam, Peter Hudomiet, Adithya Aladangady, Jason Kerwin, Laura Zimmermann, Sasha Brodsky and various University of Michigan seminar participants for helpful comments and advice (academic and real life).

In addition I would like to thank Sasha Brodsky, Peter Hudomiet, Adithya Aladangady, Ryan Monarch, Dan Marcin, Desmond Toohey, Eric Lewis, Cynthia Doniger, Christina DePasquale, and Breno Braga for sharing their life with me in Ann Arbor.

The research of this thesis was based on empirical work in Malawi and would not have been possible without the help of a number of people: I thank Ndema Longwe for outstanding fieldwork management, Niall Keleher and Lutamyo Mwamlima for mentoring and general advice, Lonnie Mwamlima for dedicated data entry management and Khorwani Zimba for logistical support. The field work for chapter 1 was made possible through the cooperation and support of the management and payroll processing staff of my anonymous partner tea firm in Mulanje. Moffat Kayembe and Carl Bruessow from Mulanje Mountain Conservation Trust helped make the field work for chapter 2 happen, and Esperanza Martinez Maldonado provided excellent research assistance.

I am grateful for research funding from the University of Michigan's Economics Department, the Michigan Institute for Teaching and Research in Economics (MITRE), Rackham Graduate School and the Center for International Business Education and Research; from IPA and Yale Savings and Payments Research Fund (funded by the Bill and Melinda Gates Foundation), and the University of Michigan Population Studies Center.

TABLE OF CONTENTS

LIST OF TABLES

Table

LIST OF FIGURES

Figure

LIST OF APPENDICES

Appendix

CHAPTER I

The Effect of Lottery-Incentives on Labor Supply: A Firm Experiment in Malawi

1.1 Introduction

The provision of incentives for workers is a central question for organizations. Employers use —and economists have studied— a wide range of compensation mechanisms to induce supply of the right amount and the right type of effort. In this paper I consider a particular incentive structure that is motivated by preferences for skewed payoffs that are observed in a variety of settings. I study the introduction of temporary "lottery" bonuses for manual workers of a large agricultural firm in Malawi to test how lottery incentives affect labor supply and how decisions evolve over time as workers gain experience with the outcomes of the lotteries.

Empirical evidence both from lab studies and real world behavior suggests a pronounced preference for positive skew for choices under uncertainty for many people. In the face of long-odds gambles lab study participants often are no longer risk averse but become risk-loving (Kahneman and Tversky, 1979*a*) even in the face of high stakes (Kachelmeier and Shehata, 1992). This observation was one of the motivations for the development of Prospect Theory (Kahneman and Tversky, 1992) which is now well accepted as a description of choices under uncertainty in the lab (Barberis, 2013).[1]

Outside the lab lotteries are popular and are a reliable source of government revenue; in betting markets long-shots are overvalued (Snowberg and Wolfers, 2010); and private households tend to under-diversify portfolios (Barberis and Huang, 2008). Lotteries are also commonly used directly as incentives in a variety of contexts. Firms often use sweepstakes for product promotions in order to incentivize "attention"; prize-linked savings accounts, in

[1] More recent theory developments focus on the role of salience (Bordalo, Gennaioli and Shleifer, 2012) to explain many of the same departures from Expected Utility Theory and more.

which savers give up safe interest returns in exchange for a chance to win big, are a common product in many countries and are used to incentivize savings (Guillen and Tschoegl, 2002); surveys often use lotteries in order to incentivize participation.

Preferences for long-odds gambles may also play a role for incentives in the context of labor markets. If such preferences are strong enough and persistent across time, then similarly to the examples above, labor contracts with lottery payoffs could be used to improve effects of incentives on effort in organizations. Workers may be more willing to provide effort under contracts with lottery elements compared to deterministic contracts valued at the same expected value as the lottery contract. By taking into account these particular features of workers' risk-preferences, firms could then increase profits. Explicit lottery contracts may be uncommon. But as Zabojnik (2002) points out, many real world contracts have skewed payoff structures which, according to standard theory for risk-averse expected utility maximizers, they should not have; e.g. stock options are more popular than one would expect and bonuses for sales staff are commonly tied to discrete targets with large payouts instead of being paid as a piece rate.

In contrast to many lab settings with one-shot decisions over gambles, if preferences for skewed payoffs are to be taken into account successfully by organizations in the context of labor contracts, then these preferences should not be one-off phenomena resulting from temporary decision mistakes; rather, as workers are exposed to the outcomes of the lotteries, effects need to persist over time. In addition, in contrast to decisions with immediate effects such as lab choices or the filling of short surveys, preferences for lottery payoffs should only affect labor supply if preferences translate into relatively more sustained continued effort, e.g. higher weekly attendance or increased productivity.

There are many reasons that have been suggested for why people may exhibit preferences for long-odd gambles. In this paper I will not attempt to distinguish between them in general, but I aim to separate two sets of fundamentally distinct categories of theories by observing how workers' choices in this experiments evolve over time. On the one hand, the observed choices could be the result of mere mistakes based on temporarily biased beliefs. It is possible that people place biased weights on the probability of a favorable — but very unlikely — outcome that are not equivalent to the true probabilities; however, with sufficient information and experience, in a repeated setting, the same people's beliefs may adjust and the preference for long-odd gambles would disappear. In fact, recent work in psychology finds that people no longer prefer lotteries over receiving the expected value for sure when small-stakes decisions are made repeatedly and feedback on the outcomes of draws is provided (Hertwig and Erev, 2009; Erev et al., 2010).

On the other hand, people may overweigh small probabilities of very favorable events not because of a temporary mistake but because of a stable preference for such gambles. These could have psychological foundations (see Barberis, 2013, p. 178 for a brief summary) or they might arise indirectly for economic reasons that are especially relevant in the context of this study, e.g. a combination of credit constraints and indivisibilities of expenditures such as for durable goods (Kwang, 1965). Lotteries may also permanently provide direct utility via general excitement or anticipation of winning something big (Conlisk, 1993) and regret aversion could lead people to work harder after having been assigned to lottery incentives in order to avoid regret over forgoing the chance to win big (Loomes and Sugden, 1982).

A number of studies test the effect of lottery incentives on survey responses as well as on health-related behavior. Volpp et al. (2008*a*) and Volpp et al. (2008*b*) find positive effects for medication adherence and weight loss, respectively, compared to a no-incentive condition – but not compared to a deterministic incentive The impact of lotteries on survey response rate is mixed (see e.g. Singer and Ye (2013) which includes a review of lottery incentives for web surveys). Most studies not only vary whether incentives are deterministic or lotteries but also vary other features such as whether the incentive is conditional on completion of the survey or not.This makes a clear interpretation difficult, or omits a fixed incentive group for comparison. An exception are two recent studies that explicitly compare lottery and fixed incentives by randomizing assignment to each condition and otherwise keeping requirements even. Gajic, Cameron and Hurley (2012) find "high" lottery incentives to be more cost-effective than fixed incentives for a sample of Ontario (Canada) residents. Halpern et al. (2011) find that lotteries did not improve response rates compared to no incentives at all for a sample of US physicians, while expected-value equivalent fixed incentives did.

In order to test how workers' labor supply responds to lottery incentives and to test if these effects endure over time, I partner with a large tea producer in Malawi to conduct a field experiment. I compare the incentive effects of two temporary bonus schemes — a fixed scheme and a lottery scheme with the same expected value — for a total of over 1,600 piece-rate workers who harvest tea leaves. Bonus schemes were re-assigned weekly for a period of up to 13 weeks and varied exogenously at the worker level. Eligibility requirements to receive the bonus were a combination of full weekly attendance and conditions on weekly total output. The requirements were identical for both bonus schemes. The (expected) values of the bonuses increased over time and ranged from between about 5% to about 15% of total weekly pay at the relevant margins for qualifying. Attendance and high productivity are important for the firm in order to maintain harvesting cycles that influence quality of the final product and the total potential output of the tea fields. High attendance and productivity also help with planning of production and reduce per-worker fixed costs.

I find that attendance at work was higher under both schemes and that the effect of the lottery bonus on the probability of full weekly attendance —a requirement for bonus eligibility and an explicit target variable of the incentive schemes— was about twice as large as the effect of the fixed bonus. The larger effect on attendance persists over time as workers gain experience with the lottery distribution. The bonus schemes affected workers' output only in the later stages of the experiment when bonus amounts and eligibility criteria were changed. Consistent with the attendance effects, the point estimate of the lottery bonus effect is larger than that of the fixed bonus effect during that part of the experiment, but I the coefficients are not statistically significantly different.

While the lottery bonus scheme did not have a strong differential effect on output, the results on attendance are significant for the firm. Attendance is an important in the production process. Tea fields need to be harvested in certain recurring intervals to ensure quality and maximize plant growth and falling behind the optimal schedule can result in significant reductions in quality and can inhibit growth. Estimates by the firm's management place the loss in sales prices in a typical main harvest season that are due to lower quality leaves from harvesting fields off their optimal cycle in the range of 10-15% of average final sales prices.

This paper is the first to empirically test the predictions implied by preferences for skewed payoffs in the context of compensation practices. In contrast to previous experiments with lottery one-shot incentives or choices in the lab, subjects in this experiment make repeated choices in the real world and face incentives of substantial value. The results suggests that at least in this context profit maximizing firms can improve worker compensation schemes by taking into account the behavioral "anomaly" of workers' preference for long shots at large gains.

1.2 Context and Field Experiment

1.2.1 Organization of production

This study was conducted in partnership with a large tea firm in Malawi.[2] The field experiment was designed and implemented in close collaboration with the firm's management in the second half of the 2012/2013 main season. The main season is characterized by more favorable weather conditions with higher temperatures and sufficient rainfall, and typically ranges from November through April. During off-season, plant growth is generally lower; in addition, the time is used for cutting down the plants ("pruning") and as a results about a third of fields does not produce leaves for harvesting until the beginning of the next main season.

[2]The partner firm wishes to stay anonymous due to business and political considerations.

The firm's tea growing and harvesting operations are overseen by estate managers. Estates are large (> 1000 workers) and are divided into divisions which are further divided into so-called "gangs". Each gang has at least one gang supervisor and several workers who help with organization such as the leaf clerk who records the weight of tea collected for each worker and a leaf inspector who performs quality control. Workers who hand pick –or "pluck"– tea leaves from the bushes are referred to as "pluckers". Gang sizes ranged in size from 30 to 100 pluckers.

This study was carried out with tea pluckers. Pluckers temporarily store the plucked tea leaves in baskets that are carried on their backs. Leaves are dropped off at the gang's weighing station where individual output is recorded several times per day. At the end of the work day the gang's harvested leaves are transported to the factory. There is no explicit teamwork involved in plucking tea. However, output of one worker is related to the output of another worker in the same gang, since on a given day in a given location the total amount of leaf that a gang can plucked is essentially fixed.

In contrast to crops that are harvested once or a few times per year, tea bushes grow continuously throughout the season and gangs of pluckers return to a fixed set of fields in regular intervals. A field is usually harvested for one or two days. Fields are scheduled to be visited in pre-set intervals (either 7 days or alternating 10/11 days depending on the pruning cycle of the field's tea bushes). Plucking tea on the correct schedule is essential for quality of output as leaves plucked too early or too late have undesirable characteristics that decrease the value of the final product. Plucking on the optimal cycle is also important for increasing total growth of the bushes of a field over time. Irregular attendance can disrupt the optimal plucking schedule, and managing plucking cycles is an important ongoing task for managers and gang supervisors who adjust working hours and the number of workers (temporarily) assigned to gangs to stay on schedule.

The schedule for each gang is set at the beginning of the main season between October or November and remains fixed until the end of the main season in April or May. The composition of each gang is also set at the beginning of the season and remains stable over time for the most part. A set of core workers is assigned to plots in each field for the rest of the season, and workers are at times temporarily added to cope with higher leaf growth. At the end of the season most workers are assigned to non-harvesting tasks for a transitional period as the the season winds down. After that a subset of workers is employed throughout the off-season, typically only for two or three days per week.

In the majority of gangs in this study tea pluckers are assigned plots at the beginning of the season to which they return when the gang returns on the same field in the next plucking interval. While pluckers are expected to finish their plots, assignments are in general not entirely rigid: pluckers who are finished with their plots can pluck in other plots; if a plucker falls behind, the supervisor will assign others to help finish the plot; and if a worker is absent on a given day, the supervisor will assign the plot to other workers for the day. In some divisions pluckers are not assigned fixed plots but pluck in rows, as supervisors coordinate the group. Both under the row organization and under the plot assignment workers have a daily minimum weight of 44kg that they are required to pluck. Repeated failure to do so will result in employment termination, but the handling is situation specific and depends on the supervisors. On the other hand, once a worker has completed the 44kg requirement, there is some pressure by supervisors to finish the assigned plots, but workers can in principle leave work. Workers have incentives to care about supervisors' assessment, beyond being good in relations with their superiors, because on average the more productive pluckers will be selected for off-season work and will more likely be employed for high-value tasks during the transitional period at the end of the season.

Workers usually arrive between 6.30am and 7.30am and typically work until between 2pm and 3.30pm in the main season. Workers arriving later than 7.30am are ineligible to work on the same day and are considered absent. The gang's end time of a given day is determined by the gang supervisor and depends on the crop that is available on a given day in the scheduled field. Usually the day ends when all tea bushes of the day's scheduled field have been plucked. Workers mostly arrive around the same time as the supervisor, but sometimes, e.g. when the leaf growth is high, workers start early to increase their day's output or to avoid daytime heat. Regular work days are Monday through Saturday. On Saturdays work ends earlier, usually by 12, and no later than 2.30pm, especially on paydays which are every second Saturday in the afternoon.

Tea pluckers are paid a constant piece rate for each kg of tea plucked.[3,4] Measurements are taken and announced every time a plucker drops off tea at the weighing station during the day. Most workers know the past *daily* total kg plucked for the current pay period; though many do not know the period's cumulative total. This aspect of worker's knowledge becomes important when discussing how workers' weekly output responded to the bonus incentives.

[3] The rate was MK 7.95 per kg since the 2011/2012 season and was increased to MK 9.77 per kg mid-season to keep up with substantial cost of living increases in the country. The average exchange rate during the time of the experiment was about MK 380 to US$1.

[4] As an additional incentive for attendance, during the months of February and March, workers who are not absent without permission for an entire week qualify for a 25kg of maize —the local staple food— to be purchased at reduced prices.

1.2.2 Data and sample characteristics

1.2.2.1 Data

This study uses administrative data provided by the tea firm that includes daily information about attendance, nature of task performed and daily total output. The data covers the 2012/2013 main season including the transition in and out of it, and ranges from calendar week 43 in Oct 2012 to calendar week 21 in May 2013. The experiment was phased in starting week 5 in February (a pilot took place with one gang in weeks 3 & 4) and ended uniformly in calendar week 17 in April of 2013. Details of the phase-in are discussed in the next section.

While in general workers are assigned to certain classes of jobs for the entire season, e.g. tea plucking or general duty, there can be temporary assignments to other classes of jobs; e.g. non-regular pluckers can help with plucking, or regular pluckers could be assigned, for example, to weeding. Since the bonus incentives introduced as part of this study only applied to pluckers, and since workers cannot choose which tasks they are assigned to, I work only with a worker's observations from weeks in which the worker was plucking on at least one day of the week or weeks in which a worker was absent the entire week but had worked as a plucker at least one day in the previous week. The analysis in this paper focuses on data associated with all of the 1,678 workers that have worked as pluckers at least once during the time of the experiment. Collapsed to the week level, the individual worker data covers a total of 38,295 worker-weeks of which 15,603 fall into weeks when the respective workers' gangs were part of the bonus experiment.

Table 1.1 presents summary statistics of key administrative variables. For selected variables the statistics are presented separately for the duration of the bonus experiment and for the time before and after the experiment that is covered by the data. While the entire data set is used during the analysis, identification of the bonus scheme effects comes from the periods in which gangs were part of the bonus experiment, i.e. from periods in which bonus assignment varied between workers of a given gang; the remaining observations are used to reduce residual variance and improve estimation precision.

The empirical analysis in this paper relies on exogenous variation in assignments of workers to bonus schemes *within gangs for a given week*. Since overall variation in the administrative individual level outcomes is large across gangs and over time, differences in aggregate measures of the outcome variables across periods during the experiment and periods outside of the experiment are not necessarily indicative of the effect of the bonus schemes. Thus, for comparability, individual-level statistics reported in Table 1.1 for the period of the experiment are computed only among those workers who in a given week were

Table 1.1: Summary statistics

A. Worker-week level observations

Number of observations	
Total in data	38,295
During bonus experiment	
All worker-weeks	15,603
No-bonus condition worker-weeks only	3,642
Outside of bonus experiment	22,692

Variable	mean	sd	min	p25	p50	p75	max	N
Individual-week level --- during bonus experiment for workers under no-bonus scheme								
Full weekly attendance [0/1]	0.555	0.497						3,642
Weekly attendance [days]	5.29	1.04	0.00	5.00	6.00	6.00	6.00	3,642
Total weekly output [kg]	384.5	196.4	0.0	248.0	360.0	504.0	1387.0	3,642
Average daily productivity if plucking [kg/day]	81.0	29.5	11.0	58.7	76.6	99.2	232.4	3,614
Regular income [MK]	3,360	1,723	0	2,150	3,142	4,415	11,033	3,642
Individual week-level --- outside of bonus experiment								
Full attendance [0/1]	0.640	0.480						22,692
Weekly attendance [days]	5.46	0.89	0.00	5.00	6.00	6.00	6.00	22,692
Total weekly output [kg]	284.3	199.7	0.0	126.0	257.0	404.0	1711.0	22,692
Average daily productivity if plucking [kg/day]	65.0	32.5	2.0	43.3	62.0	84.7	285.2	22,692
Regular income [MK]	2,366	1,652	0	1,046	2,160	3,362	13,610	22,692

B. Gang level observations

Variable	mean	sd	min	p25	p50	p75	max	N
Gang size --- over-time average of daily number of pluckers								
During bonus experiment	60.8	19.9	32.2	45.4	57.2	71.7	106.3	25
Outside of bonus experiment	50.7	17.3	28.4	37.1	46.9	63.2	90.6	25

Notes: Calculations based on administrative data of tea firm for the 2012/2013 main season. Data "during bonus experiment" refers to data points from weeks in which gangs were part of the bonus experiment; "outside of bonus experiment" refers to data points before or after observations during the experiment; see text for details. Exchage rate during the time of the experiment was ca. MK360 per US$1.

in the no-bonus condition and did not receive additional incentives beyond the regular wage (i.e. their regular piece rate pay on most days; and fixed daily wage on days in which they did not work as pluckers).

Table 1.1 describes key outcomes for the analysis. The average number of days per week in attendance is 5.29. The six-day work week therefore implies an average daily attendance of 88.1%. On average nearly half of workers[5] miss at least one day: the share of pluckers who work all days of the week —those with "full attendance"– is 55.5% on average. When plucking tea, pluckers pluck an average of 81.0kg per day and the average total weekly output is 384.5kg, with a standard deviation of 196.4kg (51% of the average). Average weekly regular income is MK3,360– about US$9 (not considering payments for bonus scheme payments and not including the value of at-work benefits like lunch or food subsidies). Output levels and, accordingly, income from tea plucking are on average lower during the period that is outside of the experiment since it includes pre-season weeks in which growth is slower. Attendance during that period is slightly higher both in days per week and share of full attendance.

Table 1.1, panel B shows the size distribution for plucking gangs. For the 25 gangs that were part of this study, gang size —measured as the across-time average per gang of the daily number of pluckers assigned to a gang— ranged from about 32 to about 106 tea pluckers during period of the bonus experiment, with a mean of 60.8 pluckers. The same numbers for the period outside of the bonus experiment are somewhat lower (mean of 50.7).

1.2.3 The Field Experiment

The experiment was designed to generate exogenous variation in the payoff structure of incentives for attendance and weekly output of tea pluckers. Temporary bonus incentives were phased-in over a span of three weeks starting February 2013 in calendar week 5 after a pilot in January in one of the gangs in weeks 3 and 4.[6] Table 1.2 summarizes the timeline of the experiment and shows the phase-in of gangs into the bonus schemes. A week prior to the start of the experiment each gang was visited by research assistants to explain how the bonus scheme worked. Workers were informed that the bonus incentives — which were paid in addition to their usual piece-rate compensation — would end at the end of the main

[5] More precisely, "half of worker-weeks" since the data is pooled across workers and periods; but for simplicity I refer to workers only.

[6] The pilot differed from the main experiment only in two ways: a) bonus group assignment was explicitly randomized, by workers' drawing from a bag (instead of by last digit of workers' employee numbers), for everyone in week 3 and for everyone in week 4 who was in the no-bonus group in week 3; everyone who was in bonus group in week 3 was assigned to the no-bonus condition in week 4; b) the value of the bonus was half of that of the initial value in the main experiment. The two weeks of the one-gang pilot are included in the analysis and are counted as part of the first phase of the experiment. In the empirical analysis I include the pilot observations but results are robust to excluding them and are available on request.

season. The bonus incentives lasted until calendar week 17, so the total time under the bonus ranged from 9 - 13 weeks depending on when the bonus was first introduced in a gang (15 weeks for one gang in which the pilot took place in weeks 3 and 4).

The study was conducted with two adjacent tea estates of the firm. In one of the estates all five divisions participated; in the second estate only two out of four divisions were included.[7] A total of 1,678 pluckers worked in one of the included seven divisions at any point during the experiment.

During each week of the experiment workers were assigned to one of two bonus incentive conditions — either to a "lottery bonus" condition or a "fixed bonus" condition — or a "no-bonus" condition. The assignment of workers to experimental conditions was quasi-random and was done at the *individual* level, not at the gang level. Assignment was determined by the last digit of a worker's employee number. Each of the ten possible digits was assigned to one of the three bonus conditions. Workers originally receive these numbers when they are first employed with the firm. The numbers are given out in the order workers arrive during initial recruitment. Workers know their number well, since they regularly use it for identification on paydays. The assignment of digits to the experimental conditions was randomized in week 5, and in all weeks after that, it was set by a pre-determined 10-week cycle (the estimated length of the remainder of the main season) with different starting points for each digit. In a given week the assignment by digit was the same across all workers; i.e. a "1" in one gang is under the same bonus scheme as a "1" in another gang. The exact sequence for each last digit is shown in Table 1.3. The schedule was generated by imposing a number of restrictions to generate desirable patterns, mainly to avoid extreme constellations, to ensure week-to-week variation for each worker, and to keep a constant ratio of experimental conditions (set initially to 2:2:1 for the lottery bonus, the fixed bonus and the no-bonus conditions respectively, and changed to 4:4:3 in the last two periods).[8]

The actual distribution of experimental conditions in given gang and week varied depending on the employee numbers of the workers assigned to the different gangs. The resulting

[7] The partial selection was done to limit organizational and financial burden. The two included divisions were selected over the non-included because of historically higher rates of absenteeism which the firm considered a problem. Variation of absenteeism across divisions is strongly correlated with geography: division in locations with more diverse economic activity tend to have higher rates of absenteeism across all of the firm's estates.

[8] The following restrictions applied to each cycle for each digit: Neither of the bonus conditions should occur three times in row; the no-bonus condition should not occur twice in a row; there should be at one tuple of (lottery bonus, lottery bonus) for consecutive weeks. Each four consecutive weeks should have at least one lottery bonus. There should be at least one tuple of (lottery, fixed) and (fixed, lottery) for two consecutive weeks during the cycle. There should be at most two lottery bonus weeks in each set of four consecutive weeks. There should be not be two consecutive weeks of bonuses of one type followed by two consecutive weeks of another bonus. The actual sequence was then chosen randomly from the set of all possible sequences.

Table 1.2: Timeline of incentive experiment

Month	Calendar week 2012/2013	Period	Gangs (divisions) included in bonus schemes	Expected value of additional bonus pay at low / high kg threshold [MK]
Oct '12	43		0 (0)	*(beginning of data set)*
	…			
Jan	2		0 (0)	
Jan	3	(pilot)	1 (1)	100 / 100
Jan	4		1 (1)	100 / 100
Jan / Feb	5	1	1 (1)	200 / 200
Feb	6		4 (1)	200 / 200
Feb	7		19 (5)	200 / 200
Feb	8		22 (6)	200 / 200
Feb / Mar	9		25 (7)	200 / 200
Mar	10		25 (7)	200 / 200
Mar	11	2	25 (7)	350 / 350
Mar	12		25 (7)	350 / 350
Mar	13		25 (7)	350 / 350
Apr	14	3	25 (7)	430 / 870
Apr	15		25 (7)	430 / 870
Apr	16		25 (7)	430 / 870
Apr	17		25 (7)	430 / 870
Apr/May	18		0 (0)	
	…			
May	21		0 (0)	*(end of data set)*

Note: Average exchange rate during experiment was about MK 360 per US$1.

Table 1.3: Schedule of assignment of workers to experimental conditions

	Period 1						Period 2			Period 3			
Calendar week:	5	6	7	8	9	10	11	12	13	14	15	16	17
Last digit of emloyee number													
1	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**	**Lot**	0	**Fix**	0	**Fix**
2	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**	**Lot**	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**
3	0	**Lot**	**Lot**	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**	**Lot**
4	**Lot**	**Lot**	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**	**Lot**	0
5	**Fix**	**Fix**	**Fix**	0	**Lot**	**Lot**	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	0
6	**Fix**	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**	**Lot**	0	**Fix**	**Lot**	0	0
7	**Lot**	**Fix**	0	**Lot**	**Lot**	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	**Fix**	0
8	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**	**Lot**	0	**Fix**	**Lot**
9	**Lot**	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**	**Lot**	0	**Fix**
0	**Lot**	0	**Lot**	**Lot**	0	**Fix**	**Lot**	**Fix**	**Lot**	**Fix**	**Fix**	0	**Lot**
Number of digits assinged to													
lottery bonus	4	4	4	4	4	4	4	4	4	4	4	3	3
fixed bonus	4	4	4	4	4	4	4	4	4	4	4	3	3
no-bonus	2	2	2	2	2	2	2	2	2	2	2	4	4

Notes: "Lot" = lottery bonus condition; "Fix" = fixed bonus condition; "0" = no-bonus condition distribution is shown in Table 1.4. Across all gangs the distribution of within-gang shares average out to values close to the ratios of the last digits. The shares vary by gang over time as well as across gangs at any given time.

At the end of each week, either Friday or Saturday, pluckers were informed by their supervisors about their assignment to the following week's experimental conditions (see Figure 1.1). In addition, to remind workers of their groups, starting in calendar week 8 the weighment clerks informed each worker during each of the tea leaf weighing in the first couple of days of the week. At the end of each week the supervisor also announced the details of the requirements to qualify for the upcoming week's bonus, which varied from week to week.

1.2.3.1 The bonus schemes

I cooperated with the tea firm to introduce two bonus schemes for tea pluckers. The firm's goal with the introduction of these bonuses was to increase worker attendance and to increase plucking productivity. Therefore, bonus qualification was based both on attendance and total weekly output (as measured in kg of tea leaves plucked). The firm values high levels of attendance in order to maintain optimal plucking cycles that matter for quality of the final product and for total quantity of plucked leaves. The regular, non-experiment labor contract does not strongly incentivize attendance. Simultaneous incentivizing of both attendance and productivity ensures that workers do not just show up and pluck the daily minimum; it also dis-incentivizes plucking high numbers conditional on attending and not attending on other days in return. In general, higher attendance is more valuable with respect to maintaining optimal plucking cycles if productivity is also high.

I divide the experiment into three separate periods based on the bonus amounts (see Table 1.2 again) and based on eligibility criteria. Over time, the bonus amounts were successively increased. The output eligibility criteria in first two periods were ex-ante announced absolute weekly output thresholds, whereas in the third period output eligibility partially took the form of a relative threshold.

In order to qualify for a bonus in a given week, workers that were assigned to one of the two bonus conditions had to have full attendance whenever work was offered. Usually this meant showing up for work all six working days of the week. In some instances, when leaf growth was unexpectedly low compared to a gang's available manpower, a gang, or a subset of it, could not be offered work on a given day. In some instances the work offered was an activity different from plucking. In those cases, the attendance requirement still applied. Official sick days, however, did not count against the attendance requirement.[9]

[9] When a worker was sick and wanted to register a sick day he or she first needed to stop by the division's office to get a sick-leave form, needed to be seen in the estate's free clinic and needed to take the form,

Table 1.4: Means of per-gang shares of tea pluckers in experimental conditions by week (standard deviations in parentheses)

Calendar week	Lottery bonus	Fixed bonus	No-bonus
5	**0.463**	**0.341**	**0.195**
6	**0.375**	**0.389**	**0.236**
	(0.034)	(0.048)	(0.032)
7	**0.398**	**0.396**	**0.206**
	(0.056)	(0.062)	(0.053)
8	**0.392**	**0.417**	**0.191**
	(0.086)	(0.082)	(0.051)
9	**0.408**	**0.391**	**0.201**
	(0.079)	(0.067)	(0.056)
10	**0.393**	**0.417**	**0.190**
	(0.054)	(0.053)	(0.047)
11	**0.396**	**0.409**	**0.194**
	(0.059)	(0.065)	(0.042)
12	**0.394**	**0.396**	**0.210**
	(0.062)	(0.065)	(0.052)
13	**0.412**	**0.402**	**0.186**
	(0.036)	(0.044)	(0.037)
14	**0.402**	**0.387**	**0.211**
	(0.066)	(0.054)	(0.054)
15	**0.392**	**0.403**	**0.204**
	(0.084)	(0.075)	(0.050)
16	**0.301**	**0.313**	**0.387**
	(0.049)	(0.066)	(0.067)
17	**0.304**	**0.296**	**0.401**
	(0.076)	(0.055)	(0.092)

Notes: Assingment was based on last digit of employee number as shown in Table 3. Calculations based on gang-week level observations. In week 5 only one gang was included in the experiment so there is only one gang level observation.

Figure 1.1: Relative timing of bonus group announcements and bonus receipts

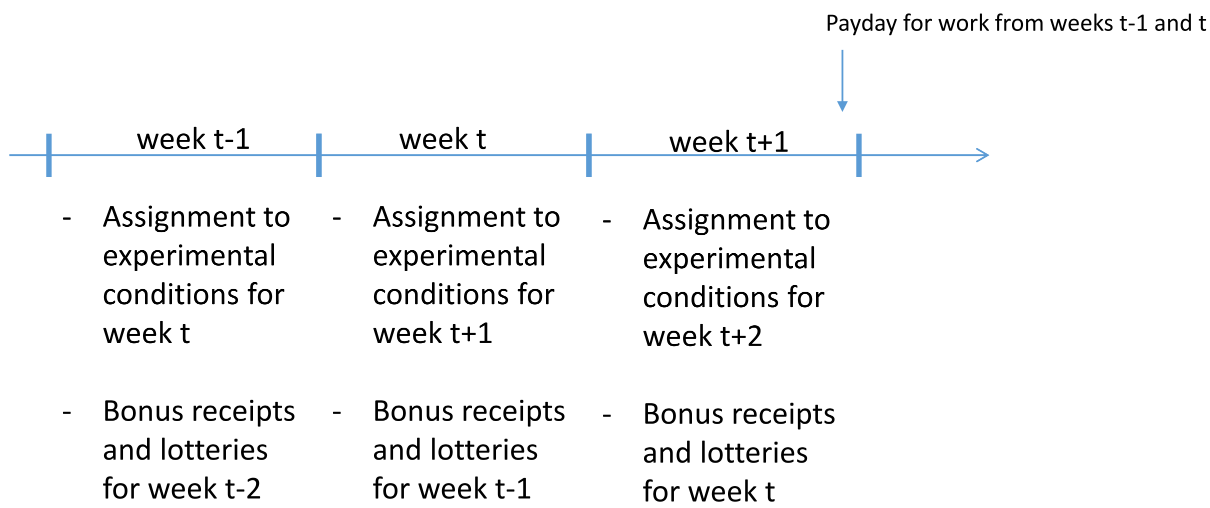

In addition to the attendance requirement, qualifying for a bonus required meeting two weekly output thresholds — a lower threshold resulting in a lower bonus or a higher threshold resulting in a higher bonus (conditional on meeting the attendance requirement and being assigned to the bonus condition in a given week). Columns 1 and 2 of Table 1.5 list the average kg thresholds for each week. For a large part of the experiment, until including calendar week 13, bonus thresholds were absolute weekly kg thresholds that were announced ahead of time: for each week they were announced at the end of the prior week. The thresholds were set at the division level and, thus, were constant across gangs in a given division for a given week. The thresholds were based on a moving average of the kg of tea-plucked per man-day in the latest two weeks for which data was available. In the first weeks, the lower threshold was set to the average and the higher threshold was set to 20% above the lower threshold. Since the thresholds were deemed relatively high, partially due to a slump in plant growth, the low thresholds were set to lower levels in week 13, and the higher thresholds were set 35% above the low threshold for the same week.

Table 1.5 columns 3 and 4 list the per-gang shares of pluckers who were above the respective kg thresholds separately for each week. Note that in a given week, across-gang variation in these measures is substantial. This variation highlights that field conditions vary between gangs, and while in the first two periods of the experiment kg thresholds varied at the level of the division, the implied relative difficulty of achieving the output targets varied a lot.

The fixed bonus worked as follows: If a plucker was under the fixed bonus scheme in a given week and met the attendance requirement, then if his/her weekly total output was above the first threshold he/she received bonus of a known amount (MK 200 –ca. \$.56– until week 10, MK 350 –ca. \$.97– until week 11; after that the system changed as described below). If the worker also exceeded the second threshold he/she received another bonus of the same amount. So during this part of the experiment the total bonus at the high threshold, the "high bonus", was twice the amount of the "low bonus". Each week each gang was visited by a research assistant (RA) to inform workers if they had qualified for the bonus and to hand out paper receipts indicating the amounts and at which payday they would receive the bonus (either the upcoming one or the one after that).

Table 1.5 columns 5 and 6 provide a sense of the magnitude of the bonus amounts relative to regular income. Especially in the first few weeks before the amounts were increased, the bonus amounts were relatively small with the bonus set at around 5% of the regular weekly

signed by the clinic's physician, back to the division's office. Since this procedure usually takes several hours including travel and wait times management was not concerned with fraudulent sick pay claims to meet the bonus incentive requirement.

Table 1.5: Characteristics of bonus output thresholds by week, means (standard deviation in parentheses)

Period	Week	(1) Average output thresholds [kg] low	(2) Average output thresholds [kg] high	(3) Average of per-gang share of workers above output thresholds low	(4) Average of per-gang share of workers above output thresholds high	(4) Average size of expected bonus as share of regular piece-rate pay at respective thresholds low	(5) Average size of expected bonus as share of regular piece-rate pay at respective thresholds high
1	**5**	**222.0**	**266.0**	**0.854**	**0.756**	**0.113**	**0.095**
		(0.0)	(0.0)			(0.000)	(0.000)
	6	**501.0**	**601.0**	**0.416**	**0.281**	**0.050**	**0.042**
		(0.0)	(0.0)	(0.248)	(0.176)	(0.000)	(0.000)
	7	**484.1**	**580.9**	**0.367**	**0.242**	**0.053**	**0.044**
		(67.5)	(80.8)	(0.254)	(0.224)	(0.008)	(0.007)
	8	**490.1**	**588.3**	**0.195**	**0.084**	**0.052**	**0.043**
		(57.1)	(68.6)	(0.133)	(0.075)	(0.007)	(0.005)
	9	**489.6**	**587.7**	**0.303**	**0.173**	**0.052**	**0.044**
		(68.6)	(82.1)	(0.228)	(0.165)	(0.007)	(0.006)
	10	**472.5**	**566.9**	**0.337**	**0.218**	**0.054**	**0.045**
		(58.8)	(70.7)	(0.240)	(0.202)	(0.007)	(0.006)
2	**11**	**465.5**	**558.3**	**0.341**	**0.199**	**0.097**	**0.081**
		(67.9)	(81.5)	(0.222)	(0.173)	(0.015)	(0.012)
	12	**450.4**	**540.3**	**0.306**	**0.168**	**0.100**	**0.083**
		(65.5)	(78.8)	(0.217)	(0.164)	(0.015)	(0.013)
	13	**359.7**	**486.5**	**0.609**	**0.339**	**0.125**	**0.093**
		(49.1)	(81.4)	(0.189)	(0.204)	(0.018)	(0.018)
3*	**14**	**264.0**	**428.1**	**0.794**	**0.347**	**0.167**	**0.208**
		(0.0)	(147.6)	(0.167)	(0.087)	(0.000)	(0.051)
	15	**264.0**	**393.4**	**0.744**	**0.330**	**0.167**	**0.226**
		(0.0)	(156.5)	(0.192)	(0.076)	(0.000)	(0.108)
	16	**264.0**	**452.1**	**0.773**	**0.326**	**0.167**	**0.197**
		(0.0)	(110.1)	(0.192)	(0.058)	(0.000)	(0.081)
	17	**264.0**	**475.3**	**0.827**	**0.327**	**0.167**	**0.187**
		(0.0)	(126.1)	(0.183)	(0.071)	(0.000)	(0.089)

Note: Calculations in columsn 1, 2 and 5, 6 are based on worker-week level observations; in columns 3, 4 on gang-week level observations. *) In weeks 14-17 the low thresholds were not weekly but daily thresholds of 44kg and the high bonus kg threshold was determined by relative within-gang ranking and so thresholds were not announced ex-ante and only computed ex-post for this table.

pay at the low thresholds. The bonus amounts increased substantially in the last period of the experiment, which is described further below.

The lottery bonus worked very similarly to the fixed bonus. But instead of receiving a set bonus amount workers would draw in a lottery during the weekly RA visit to determine the amount of the bonus. During this period of the experiment the worker would either draw once in the lottery or twice (with replacement) depending on whether he/she had exceeded the first output threshold only or both thresholds. The lottery draw was carried out as follows: An RA would hold up a black plastic bag with 100 plastic tokens inside and the eligible worker would draw once. The plastic tokens were labeled with the money prize amounts. Prizes were revealed only privately; but virtually all respondents chose to share the results with fellow workers who were standing by.

The distribution and amounts of the lotteries are shown in Table 1.6. The expected values of the lottery draws are equal to the corresponding bonus amounts under the fixed scheme. The distributions were modeled after common lottery structures but included a non-trivial lower bound instead of a zero in case of not winning a high prize.[10] Importantly, the probabilities of getting each of the 4 prizes in a given lottery draw stay constant, even if the amounts changed over the course of the experiment. This helped people remember the lottery details and also meant that whenever the bonus amounts were changed, *the probabilities* were not new.

The weekly RA visits to notify bonus winners and conduct lotteries took place between Tuesday and Thursday, usually Tuesday and Wednesday, in exceptional cases on Friday, and tended to take place on the same day of the week for a given gang.

Starting in week 14 the output portion of the bonus criteria was changed from an absolute threshold set at the division level to relative, gang-level specific output criteria for the high threshold combined with a low absolute daily threshold equal to the standard daily minimum requirement of 44kg for the low bonus. The bonus at the low end essentially meant that the low bonus was a bonus for attendance only, whereas the high bonus was a bonus for attendance and high output. The relative threshold was introduced to accommodate large differences in output potential both across gangs within the same division and over time. Using a division-level threshold had lead to large variations in bonus qualifications across gangs in the beginning of the experiment. These differences were driven, by and large, by variations in field conditions and not by variations in group effort.

The lottery bonus was also modified in period 3. Instead of drawing in the lottery once

[10] Quiggin (1991) derives some results for the design of an optimal lottery given rank-dependent preferences with overweighting of extreme positive; the results are broadly similar to the structure used in this experiment.

Table 1.6: Prize distributions under lottery bonus scheme

Period 1

Calendar weeks: 6-10
Expected value: MK 200
One draw if exceeded low threshold; two draws if exceeded high threshold

Prize rank	Number of tokens	Prize amount (MK)
4th	89	100
3rd	7	500
2nd	3	1,200
1st	1	4,000
	Σ=100	EV = 200

Period 2

Calendar weeks: 11-13
Expected value: MK 350
One draw if exceeded low threshold; two draws if exceeded high threshold

Prize rank	Number of tokens	Prize amount (MK)
4th	89	200
3rd	7	600
2nd	3	2,000
1st	1	7,000
	Σ=100	EV = 350

Period 3 - LOW BONUS

Calendar weeks: 6-10
Expected value: MK 430
One draw if exceeded low threshold but not high threshold; low threshold is daily (44kg) not weekly

Prize rank	Number of tokens	Prize amount (MK)
4th	89	250
3rd	7	1,000
2nd	3	2,250
1st	1	7,000
	Σ=100	EV = 430

Period 3 - HIGH BONUS

Calendar weeks: 6-10
Expected value: MK 1,300
One draw if exceeded high threshold; high threshold set as being in top 1/3 of gang in given week

Prize rank	Number of tokens	Prize amount (MK)
4th	89	750
3rd	7	2,500
2nd	3	6,750
1st	1	25,500
	Σ=100	EV = 1,300

for each threshold crossed, pluckers would now either draw in a small lottery *or* in a large lottery depending on which of the kg thresholds was crossed. The values are provided in the second half of Table 1.6. The sizes relative to regular income are shown in Table 1.5 columns 5 and 6; they sizes increased from around 5% initially to around 15%-20% of weekly pay at the two kg thresholds.

As for the timing of actual payments all earned bonuses were paid together with the paydays associated with the regular income from work during the week when bonus receipts were given out, i.e. together with the pay for regular income earned one week after being under a bonus incentive.

1.3 Effort under stochastic incentives

This section presents a simple illustration of labor supply under stochastic incentives for an agent whose behavior is described by Prospect Theory instead of Expected Utility.[11] I model a worker's decision to provide binary effort in the face of a simple form of stochastic incentives. The worker's evaluation of his options follows Prospect Theory (Kahneman and Tversky, 1992) and nests the standard Expected Utility (EU) case as a special case. The binary effort matches the experimental setup in that the bonus scheme has discrete attendance and output thresholds that can be either achieved or not. The stochastic incentives mimic the bonus schemes of this study which are either a "lottery bonus" under which, conditional on meeting the attendance and output requirements, a large prize is paid out with known, small probability.

Consider the one-time labor supply of an agent that can decide whether to provide a unit of (one-dimensional) effort or not. The worker has utility cost of effort of c and faces the following incentives: the worker receives compensation of x with probability p, if effort is provided, and 0 otherwise. The probability of receiving compensation is strictly positive and possibly equal to one, $0 < p \leq 1$. In case no effort is provided, the agent receives his reservation utility $\underline{u}$. The experiment is about non-negative bonuses and therefore we restricted to the gains domain (as opposed to both gains and losses). The worker derives $v(x)$ from compensation x. The compensation valuation is normalized to zero: $v(0) = 0$. The value associated with the compensation prospect (p, x) is the weighted sum of the value of the compensation for effort weighted with weights $w(p)$, which are not necessarily equal to the objective probabilities.

The worker's value function is written as:

[11] Models with alternative mechanisms such as excitement, regret or salience would produce identical predictions. But since distinguishing between models or underlying mechanisms is not the goal of this paper I focus on an illustration with a commonly accepted descriptive theory.

$$V(p,x) = \begin{cases} V_1(p,x) \equiv w(p)v(x) - c, & \text{if } e = 1 \\ V_0(p,x) \equiv \underline{u}, & \text{if } e = 0 \end{cases}$$

The worker decides to provide effort *iff* $V_1(p,x) \geq V_0(p,x) \iff w(p)v(x) \geq \underline{u} + c$. The higher the value of the compensation prospect $w(p)v(x)$, the larger the region $\underline{u}+c$ for which effort is provided. In an environment with heterogeneous workers that have different levels of reservation utility and vary in their disutility from effort, the larger compensation x is and the larger the weight $w(p)$ associated with that compensation is, the larger the share of workers that decides to provide effort.

In the case of a deterministic bonus incentive with payoff B, i.e. $p = 1$ and $x = B$, a worker provides effort *iff* $B \geq \underline{u} + c$. In comparison, for a stochastic incentive with the same expected value B, i.e. a compensation prospect of $(p, x = B/p)$, effort is provided *iff* $w(p)v(B/p) \geq \underline{u} + c$.

In the case of standard Expected Utility (EU), the probability weights and objective probabilities coincide, i.e. $w(p) = p$. Assuming decreasing marginal utility and standard risk-aversion under EU, $v(\cdot)$ is concave. Since now $B > pv(B/p)$ for $p < 1$, there is a larger region $\underline{u}+c$ for which workers will choose to provide effort under the deterministic incentive relative to the stochastic incentive.

In contrast, the reverse statement is true if probabilities are over-weighted sufficiently, i.e. if $w(p) >> p$ so that the the overweighting of p dominates the concavity of $v(\cdot)$, then effort provision would occur over a larger range of $\underline{u} + c$.

In this model, to what extent the firm can motivate workers to provide effort under the two incentive schemes —deterministic incentives on the one hand and stochastic incentives with the same expected payoff on the other hand— depends the degree of overweighting in the population and the curvature of the utility for money. The more over-weighting and the less risk aversion, the greater the difference will be between labor supply under the stochastic incentive relative to the deterministic incentive.

1.4 Effects of the bonus schemes on worker behavior

1.4.1 Graphical analysis

I investigate the effect of the bonus schemes on key administrative variables, in particular on the two dimensions of effort that were incentivized by the bonus schemes: attendance and output per worker. I begin with a simple graphical comparison of outcomes by bonus assignment, without any additional controls or structure imposed, before presenting regressions estimates in the next section. The experiment was carried out over several weeks and can be divided into three 'periods' according to the value of the bonus and the output requirements for eligibility (see description in section 1.2.3.1). I first present effects averaged across the entire duration of the experiment before disaggregating results by time.

Figure 1.2 shows the distribution of one set of measures of the two dimensions that were incentivized by the bonus schemes: weekly attendance and total weekly output. The first part of Figure 1.2 shows the PDF of number of days worked per week and suggests that both bonuses did have an effect on attendance. The positive effect of the bonus is concentrated at the top end of the distribution, with the probability mass moving from four and, more so, five to six days which was the was the attendance requirement of the bonus scheme. Effects appear to be stronger for the lottery bonus than for the fixed bonus. The second part of Figure 1.2 shows a CDF of individual workers' total weekly output, averaged across the entire duration of the experiment. The bonus incentives seem to have had no substantial effect on output on average across weeks.

Figures 1.3 and 1.4 disaggregate the distributions of Figure 1.2 by period of the experiment. Figure 1.3 documents that the higher overall full weekly attendance compared to no bonus documented in Figure 1.2, is due to the higher rates of full attendance periods 1 and 3, with effects relatively larger in period 3. In period 2, rates of full attendance are very similar across all groups. 1.4 suggests the absence of large effects of either bonus scheme on output compared to the no-bonus condition, especially for periods 1 and 2. In period 3, the CDFs of total weekly output are weakly separated between the no-bonus condition and the two bonus schemes over some parts of the graph, indicating that the bonuses may have had positive effects at that stage of the experiment.

1.4.2 Empirical model

In order to formally estimate the effects the bonus incentives on labor supply and related outcomes I estimate a linear panel model of the following form. For a worker i in gang j at time t, let the dependent variable of interest y_{ijt} be a linear function of current and past bonus assignments, conditions in a given gang at a given time, individual time-invariant component and an idiosyncratic error:

$$y_{ijt} = (\boldsymbol{bonus_{i,t}}, \boldsymbol{bonus_{i,t-1}}, \dots)\boldsymbol{\beta} + \gamma_{jt} + c_i + u_{ijt}. \tag{1.1}$$

$(\boldsymbol{bonus_{i,t}}, \boldsymbol{bonus_{i,t-1}}, \dots)$ are a set of 1 x 2 *vectors* that indicate assignments to the two different bonus schemes and varies both across i and across t. The vector $\boldsymbol{bonus_{is}}$ indicates bonus assignment in s and includes two indicators: $lottery_{is}$ is equal to 1 if worker

Figure 1.2: Distribution of attendance and output by bonus assignment

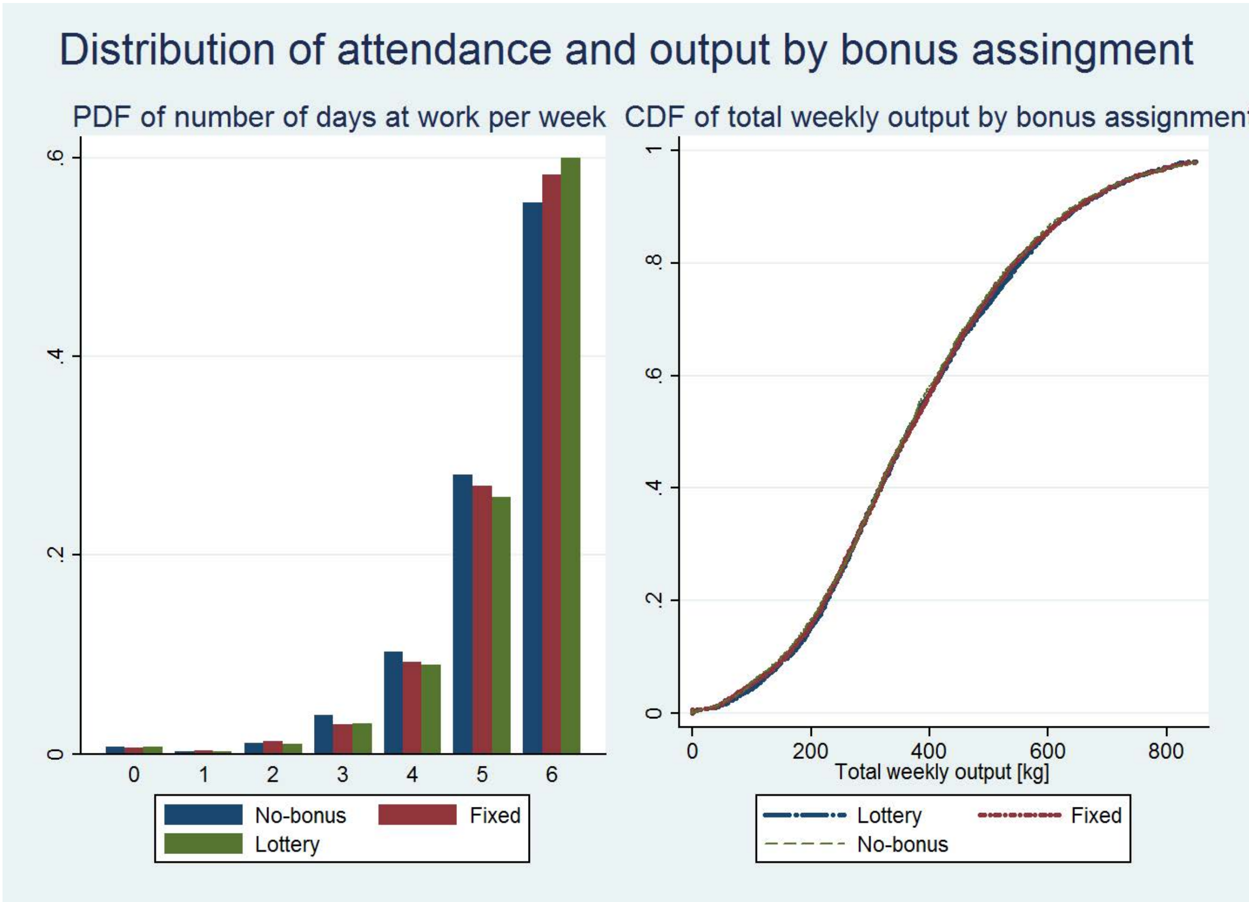

Figure 1.3: PDF of weekly attendance by bonus assignment, by period

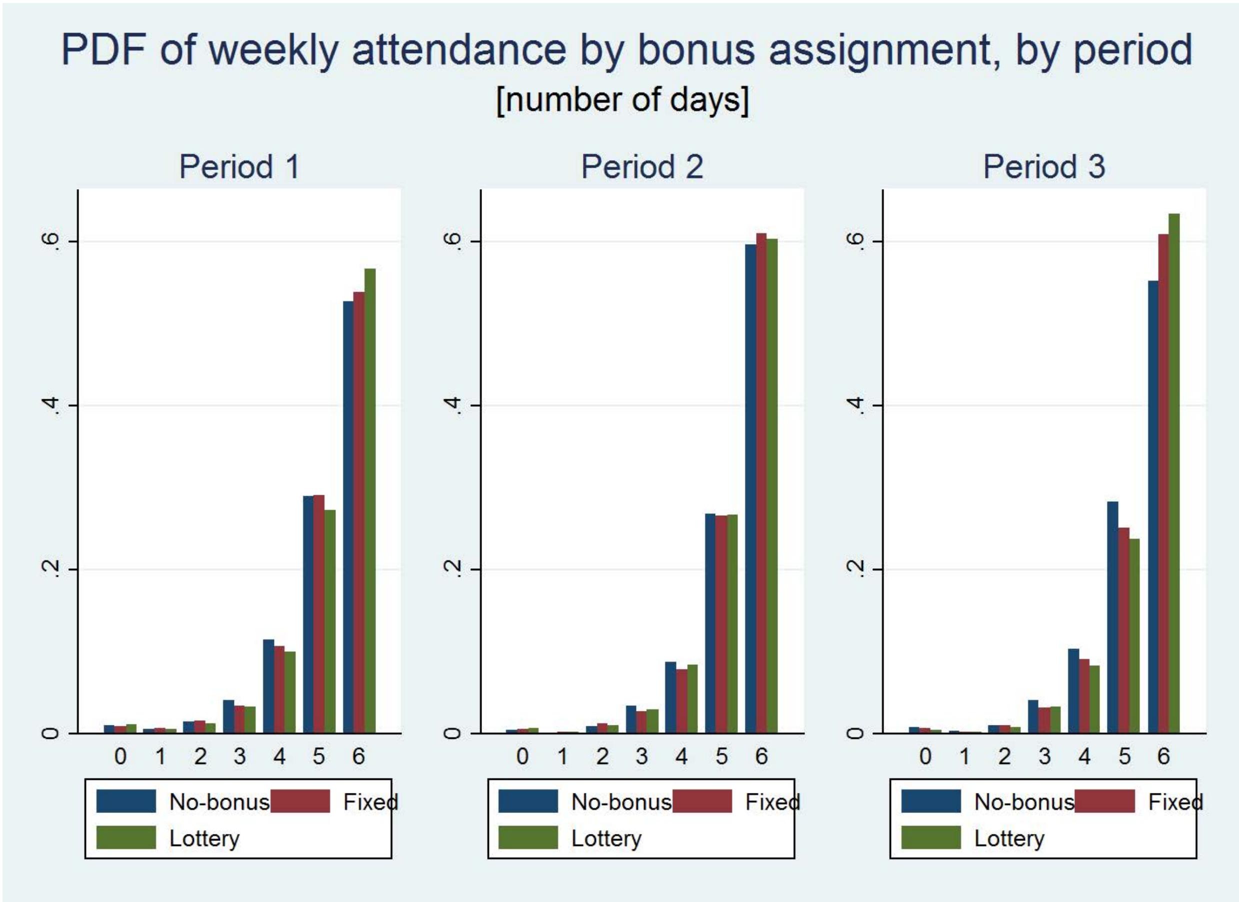

Figure 1.4: CDF of total weekly output by bonus assignment, by period

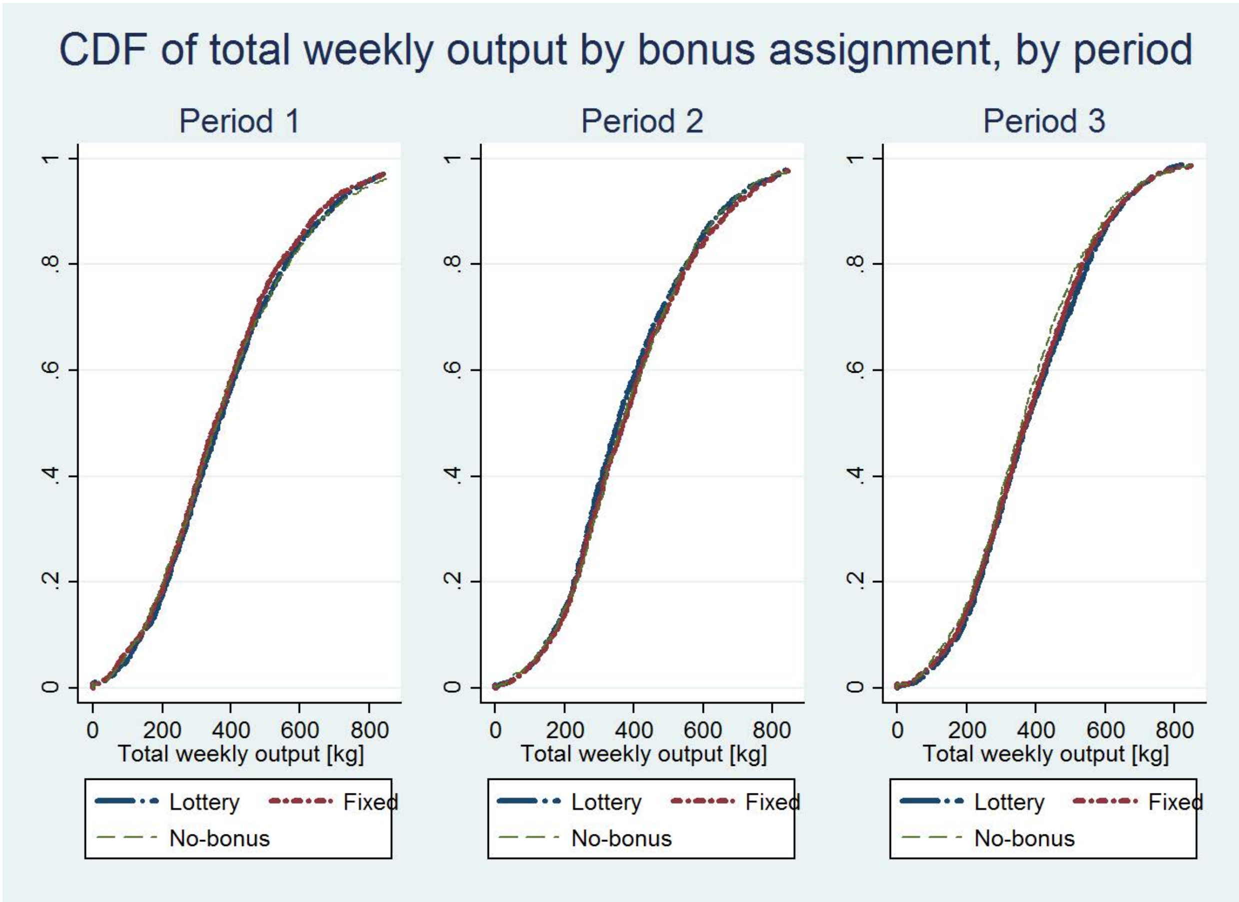

i at time s is assigned to the Lottery Bonus and 0 otherwise; $fixed_{is}$ is equal to 1 if i is assigned to the Fixed Bonus. The omitted category is no-bonus group and so the elements of $\boldsymbol{\beta}$ corresponding to $lottery_{is}$ and $fixed_{is}$ give the differences in outcome y_{ijt} between the respective bonus types and the no-bonus group.

γ_{jt} are a full set a of of coefficients for each gang-week combination. c_i are time-invariant unobserved individual effects. u_{ijt} are idiosyncratic errors with mean zero.

Given the exogenous assignment of bonus schemes, a natural starting point for estimation of the model above would be to ignore the unobserved individual effect in (1.1), subsume c_i under the idiosyncratic error and consistently estimate $\boldsymbol{\beta}$ by running OLS on the data set that pools all weeks for each worker. However, pooled OLS does not make use of between- and within-worker variation over time efficiently. Pooled OLS weights between- and within-worker variation equally, irrespective of the relative precision with with each source of variation is contributing to the estimation. In contrast, a random effects estimator weights the two dimensions optimally and thus potentially increases estimation precision. In many applications, random effects models suffer from unobserved individual heterogeneity that bias the estimator of the parameters of interest, and individual fixed effects models are employed to re-gain consistency. In this case, however, the exogeneity assumptions that are natural for pooled OLS estimation also imply the required exogeneity assumption for random effects.

I include lagged values of the bonus assignment indicators and the reasoning for that is as follows. First, in this experiment's schedule of bonus condition assignments $\boldsymbol{bonus_{i,t}}$ was correlated with $\boldsymbol{bonus_{i,t-1}}$. Secondly, in addition, as I argue below, it is reasonable in the setting of this experiment that workers behavior in t is influenced by their assignment in $t-1$ ($\boldsymbol{bonus_{i,t-1}}$). As a result, *not* including lagged values of bonus assignments on the right hand side of (1.1) would lead to omitted variable bias and the estimator of the coefficients on contemporaneous bonus assignments, $\boldsymbol{bonus_{i,t}}$, would be inconsistent.

Contemporaneous bonus condition assignments were correlated with past assignments by construction of the weekly shifting last-digit schedule. Weekly assignments were not stratified on past assignments as this was technically infeasible in this setting: the 10 different digits were assigned to three bonus conditions, and the no-bonus group was underrepresented to maximize power for tests between lottery and fixed bonus (see Table 1.3 on page 12 for reference). In addition, the schedule was chosen to increase over-time variation in assignments for individual workers (see the discussion in section 1.2.3 on page 9 for details). Table 1.7 shows current weeks' assignment probabilities relative to prior weeks' assignments to illustrate the in-built asymmetries. As an example, if a worker was assigned to the lottery bonus last week then 47% of the time he was assigned to the fixed bonus in the current week and was assigned to the lottery bonus or no bonus only 25% and 28%, respectively, of the time.

The correlations of current and past assignments discussed above are problematic for consistent estimation of contemporaneous effects if past assignments have an effect on the value of contemporaneous outcomes. There are variety of potential mechanisms for such an effect. I outline two examples for illustration: lack of time-separability of utility between one week and the next, and attention effects.

1) *Non-time-separable utility.* Consider a simple model with an agent that faces convex costs of effort and thus wants to smooth effort over time; assume no borrowing or saving. Over the short horizon of a two-week period, it is reasonable to assume that disutility from physical effort is not time-separable and that high effort now leads to both disutility now as well as next week – an "exhaustion" effect. Faced with time varying shocks to his wage schedule, an agent would decide to both work *more* than when the wage is high *and* to work *less* in the following week to smooth effort costs across weeks.

2) *Attention effects.* If being in the lottery group this week leads a worker to pay more attention to the bonus scheme overall in this *and* the following week, for example because lotteries are different and exciting, then next week the same worker might react more strongly to the bonus schemes, either lottery or fixed, and this week's assignment would have a direct affect on next week's effort.

In theory several lags of bonus assignments could included in the specification and it is an empirical question whether lags of bonus assignments do affect current outcomes. Note that no observations are dropped by included lags of the bonus assignment. For data points outside of the experiment, group assignment is by default set to the no-bonus group, and so in all weeks lagged assignments are defined.

In theory all past lags of bonus assignments could be correlated with current outcomes, even after conditioning on all other weeks' bonus assignments. The number of available empirical variation in assignment histories is limited, however (not least because the assignments follow a fixed ten-week schedule). In practice, while I reject that one-week lags have no effect for outcomes for which the bonuses had an effect, I cannot reject that lags beyond the first week have no impact. Therefore, the main specification I estimate includes only current week's bonus assignment and one-week lags of bonus assignment (alternative specifications with further lags produce very similar results). I estimate:

$$y_{ijt} = (\boldsymbol{bonus_{i,t}}, \boldsymbol{bonus_{i,t-1}})\boldsymbol{\beta} + \gamma_{jt} + c_i + u_{ijt} \tag{1.2}$$

$$= (lottery_{it}, fixed_{it}, lottery_{i,t-1}, fixed_{i,t-1})\boldsymbol{\beta} + \gamma_{jt} + c_i + u_{ijt}. \tag{1.3}$$

Table 1.7: Distribution of bonus assignment in current and relative to prior week

Assignment in …	...prior week:			
...current week:	Lottery bonus	Fixed bonus	No bonus	Total
Lottery bonus	25%	47%	43%	38%
Fixed Bonus	47%	23%	47%	38%
No bonus	28%	30%	9%	23%
Total	100%	100%	100%	100%

The exogenously determined bonus assignment schedule justifies a strict exogeneity assumption in the sense of Wooldridge (2001, p. 257, RE.1b): For all i and j

$$\mathbb{E}(c_i|\boldsymbol{bonus_{i1}}, \ldots, \boldsymbol{bonus_{iT}}) = E(c_i) = 0 \tag{1.4}$$

The assumption implies that the individual effect can be treated as an unobserved "random effect" and allows for consistent estimation of $\boldsymbol{\beta}$ in a random effects model. To check if in fact weekly assignment to the different groups across all weeks of the experiment is uncorrelated with c_i, I proxy for c_i with pre-experimental averages of weekly output and weekly attendance, respectively, and run a pooled OLS estimation of these averages on a full set of interactions of week dummies and bonus group assignments. F-tests of the joint null hypothesis that bonus groups indicators are equal to zero in every week of the bonus experiment reveal no overall balance across the experiment. Empirically, I can cannot reject balance for either weekly output or number of days in attendance (p-values of 0.265 and 0.739, respectively).

By design of the bonus experiment there are two sources of variation for the assignment of workers to bonus schemes: assignment varies between workers at a given point in time and "within" workers, i.e. for a given worker across time. Under the exogeneity assumption above (and further standard regularity assumptions), $\boldsymbol{\beta}$ can be estimated as pooled OLS, using "between" and "within" variation in equal proportion; with an individual fixed effect (FE) estimator, using only "within" variation only; or as a "random effects" (RE) model, using Feasible GLS. The RE estimator optimally weights between- and within-individual variation according to the sampling variance in either component and is the most efficient of the three estimators. Therefore I use RE specifications as the main method of estimation. Results from pooled OLS estimation and fixed effects are shown in the appendix and give essentially identical results for the effects of interest albeit with lower precision.

Tea pluckers work together in gangs. This can lead to statistical dependence of observations within the same gang in a given week for two reasons: gang(-week) level shocks and peer effects. Correlations induced by gang-week effects (γ_{jt}) can be controlled for by including the according set of dummies in the regression. However, even conditional on γ_{jt} observations can be expected to be correlated within gangs in the presence of peer effects. To allow for this possibility, standard errors are clustered at the gang level and as such inference is robust to arbitrary correlation within gangs at a given moment in time and across time. The gang level clustered are also valid under the likely individual worker outcomes auto-correlation.

1.4.3 Regression results – overall effects

Table 1.8 presents the overall effects —across all weeks of the experiment— of the two bonus schemes vis-a-vis the no-bonus condition on key outcome variables using individual random effects regressions. Columns 1 and 2 show that the bonus incentives lead to higher attendance compared to the no-bonus condition: the probability of a worker attending all days of the week, which was an explicit target variable of the bonus scheme, and the number of days in attendance increased for both bonus bonus groups (the p-value of a joint test of the two bonus conditions being equal to zero is 0.0032 in column 1 and 0.0052 in column 2). The effect was more pronounced for the lottery bonus scheme. Full weekly attendance increased by 3.73 percentage points when a plucker was under the lottery bonus relative to the no-bonus condition. Full attendance was about 7% higher under the lottery bonus relative to the mean of 55.5% under no bonus. The difference between lottery and no-bonus conditions is more than twice as large as the difference between the fixed bonus and no bonus, and the difference in effects is statistically significant (p-value 0.0437). The effect on the number of days is also larger for the lottery bonus and the difference is weakly significant (p-value 0.0758). The relative effect of the lottery bonus vis-a-vis the no-bonus condition is lower for the number of days in attendance than for full attendance: 0.0654 days over a mean of 5.29 days, or about 1.2%. This is consistent with the graphical analysis in section 1.4.1, which indicates that attendance shifts are taking place at the top of distribution, with some people not missing that one additional day that they might have otherwise missed.

Along the dimension of weekly output, the bonus schemes had two weekly-varying thresholds that qualified workers for a small(er) or a larg(er) bonus. On average across all weeks, the probability of reaching the low or the high threshold were unaffected by either bonus scheme (Table 1.8, columns 3 and 4). Similarly, productivity per day of plucking (column 6) was not statistically significantly different in either group, and the differences were small (but negative according to the point estimates). For total weekly output, the bonus schemes coefficients are not jointly significantly different from the no-bonus condition (p-value 0.210) and the point estimates are small. However, in this specification the difference of lottery bonus vis-a-vis no bonus is marginally statistically significant, and the difference is larger than for the fixed bonus for which the difference vis-a-vis the no-bonus condition is not significant.

In order to demonstrate the robustness of these key results, I discuss results from pooled OLS and individual fixed effects estimation.[12] The discussion focuses on the attendance

[12] Results in analogue to Table 1.8 from a specification that does not include gang-week dummies but only week dummies are very similar (results not shown): the differences of the two bonus schemes compared to the no-bonus condition are marginally larger, but the difference between lottery and fixed bonus are identical

Table 1.8: Effects of bonus incentives on attendance and output of tea pluckers (individual random effects estimations)

	(1) Full attendance [0/1]	(2) Weekly attendance [days]	(3) Met low output threshold for weekly total output [0/1]	(4) Met high output threshold for weekly total output [0/1]	(5) Total weekly output [kg]	(6) Weekly average of daily productivity if plucking [kg/day]
Lottery Bonus	**0.0373*** **	**0.0654*** **	**-0.00559**	**0.00916**	**4.025***	**-0.108**
	(0.0111)	(0.0204)	(0.00959)	(0.00774)	(2.289)	(0.264)
Fixed Bonus	**0.0177***	**0.0324**	**-0.000203**	**-0.000325**	**2.127**	**-0.305**
	(0.0107)	(0.0212)	(0.00754)	(0.00652)	(2.510)	(0.243)
L1.Lottery Bonus	**0.00333**	**-0.0134**	**-0.00225**	**-0.00335**	**-0.168**	**-0.284**
	(0.0122)	(0.0198)	(0.00761)	(0.00784)	(2.317)	(0.328)
L1.Fixed Bonus	**-0.0178**	**-0.0453*** **	**-0.00116**	**-0.00877**	**-2.918**	**-0.510**
	(0.0122)	(0.0168)	(0.00973)	(0.00843)	(2.600)	(0.331)
P.-val. of Wald tests:						
Lottery Bonus = Fixed Bonus	0.0437	0.0758	0.479	0.201	0.472	0.412
Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.00318	0.00520	0.767	0.393	0.210	0.429
Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0	0.0391	0.0223	0.942	0.524	0.363	0.296
Mean of dep. variable in no-bonus condition during experiment	0.555	5.290	0.533	0.250	384.5	81.02
Number of gang-week observations during experiment	15603	15603	15603	15603	15603	15490
Number of gang-week observations in total	38,295	38,295	38,295	38,295	38,295	38,117
Number of workers	1,678	1,678	1,678	1,678	1,678	1,678

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the gang level. Estimated results are from a worker-level random effect model with a full set of gang-week dummies. Wald-tests: “Lottery Bonus = Fixed Bonus” tests the equality of the two coefficients; “Joint: Lottery Bonus = 0 and Fixed Bonus = 0” and “Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0” test the joint equality of the two pairs of coefficients.

results in more detail, but along both dimensions of specification changes the results can be considered to be the same as in the main specification. Appendix Tables A.1 and A.2 present the same coefficients as in Table 1.8 estimated from pooled OLS and individual fixed effects regressions, respectively. The estimates are very similar and especially the difference between lottery and fixed bonus effects is nearly identical for the attendance outcomes (e.g. the differences are 0.0196 and 0.0330 for columns 1 and 2 in Table 1.8 compared to 0.0193 and 0.0339 in Appendix Table A.1; similarly for comparisons of Table 1.8 vs. Appendix Table A.2). The p-values on these differences are generally lower in Table 1.8 as the random effects model combines efficiently both "between" and "within" variation.

Table 1.9 presents the overall effects of the two bonus schemes vis-a-vis the no-bonus condition on meeting the criteria for receiving a low and high bonus, respectively, and on income. The bonus qualification requirements were a combination of full attendance and total weekly output being above the thresholds. Given the lack of substantial overall effects on output, it is the effects on attendance that leads both types of bonus to have an effect on meeting the bonus qualification requirements, compared to the no-bonus group (columns 1 and 2). The difference in effects between lottery and fixed bonus are smaller (and not statistically significant) relative to the difference in effects on full attendance only (column 1 of Table 1.8). Only a fraction of workers did meet the output thresholds even among those that had full attendance and so these effects are expected to be smaller. Note, however, that the smaller differences between lottery and fixed bonus coefficients in columns 1 and 2 of Table 1.9 compared to column 1 and 2 of Table 1.8 are not fully accounted for by the share of workers meeting the kg thresholds (columns 3 and 4 of Table 1.8). For example, the difference between bonus coefficients in column 1 of Table 1.8 is 0.0196; the average share of pluckers meeting the low threshold is about 53%; yet the difference in bonus coefficients in column 1 of Table 1.9 is not 53% of 0.0196 but somewhat lower at about 22%. Ignoring limitations of statistical power, these computations would imply that, compared to workers' reactions to the fixed bonus, there are relatively more workers that reacted to the lottery bonus with higher attendance who did not make the output thresholds.

Averaged across all weeks of the experiment the lottery bonus scheme had a small effect on regular income (without bonus) compared to the no-bonus condition. This is driven by the slightly higher attendance combined with a lack of potentially off-setting differences in per-day productivity. In contrast, the difference of the fixed bonus compared to the no-bonus

to those of Table 1.8 for the attendance results and are also essentially the same for the remaining outcomes. Alternative specifications that consider lags of the bonus assignment variables (results not shown) also give similar results but generally reduce the precision of estimates. In all specifications with additional lags, the one-week lags of bonus assignment are statistically significant for the attendance outcomes, for which the bonuses do have contemporaneous effects, while further lags are not statistically significant.

Table 1.9: Effects of bonus incentives on bonus criteria qualification and income (individual random effects estimation)

	(1)	(2)	(3)	(4)
	Met bonus criteria for low bonus [0/1]	Met bonus criteria for high bonus [0/1]	Regular income without bonus (output x piece rate) [MK]	Total expected income = regular income + expected bonus [MK]
Lottery Bonus	0.0230**	0.0239**	48.78**	355.3***
	(0.0113)	(0.0100)	(19.87)	(28.04)
Fixed Bonus	0.0187*	0.0139*	27.39	324.8***
	(0.0105)	(0.00839)	(23.21)	(27.95)
L1.Lottery Bonus	-0.00764	-0.00345	0.111	-10.50
	(0.00880)	(0.00811)	(21.40)	(28.78)
L1.Fixed Bonus	-0.0139	-0.0107	-27.41	-50.64
	(0.0117)	(0.00894)	(24.11)	(32.54)
P.-val. of Wald tests:				
Lottery Bonus = Fixed Bonus	0.622	0.200	0.358	0.295
Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.107	0.0578	0.0486	0
Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0	0.489	0.313	0.337	0.165
Mean of dep. variable in no-bonus condition during experiment	0.345	0.196	3360	3360
Number of gang-week observations during experiment	15603	15603	15603	15603
Number of gang-week observations in total	38,295	38,295	38,295	38,295
Number of workers	1,678	1,678	1,678	1,678

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the gang level. Estimated results are from a worker-level random effect model with a full set of gang-week dummies. Wald-tests: “Lottery Bonus = Fixed Bonus” tests the equality of the two coefficients; “Joint: Lottery Bonus = 0 and Fixed Bonus = 0” and “Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0” test the joint equality of the two pairs of coefficients.

condition is not statistically significant. However, the difference between lottery and fixed bonus groups is not statistically significant either. The bonus schemes had an effect on total expected income (including bonus). However, this difference is largely mechanical since in the no-bonus condition expected bonus income was zero. Again, there is no difference between lottery and fixed bonus.

1.4.4 Regression results – effects over time

In this section I present results analogous to the preceding section and estimate equation (1.2) with individual random effects. But instead of averaging effects over time, I interact the bonus condition indicators (both contemporaneous and lagged) with period or week dummies to allow the effect to vary by "period" or by week.

Tables 1.10 and 1.11 show the effect of lottery and fixed bonus on the same set of outcomes as in the preceding section, averaged by period instead of across the entire experiment. The effect of either bonus are in general highest in period 3 when the bonus amounts were largest (and eligibility criteria were arguably easier to understand). While there is a significant effect of the lottery bonus on full attendance in *period 1*, and jointly the two bonus schemes are significantly different from zero for full weekly attendance (p-value 0.0638), there is no statistically significant effect of either bonus scheme in *period 2*.

The differences between lottery and fixed bonus with respect to attendance as discussed in the previous section largely holds up for the different periods. The differences between lottery and fixed bonus for full weekly attendance and the number of days in attendance in Table 1.10 columns 1 and 2 are slightly higher in period 1 than in period 3; for column 1 the difference in period 1 is 0.03102 and in period 3 it is 0.0275. Since the level of the effects vis-a-vis the no-bonus condition are higher in period 3, the relative difference between lottery and fixed bonus is smaller in period 3. In period 2 the point estimates are very small and statistically insignificant, and the coefficient differences between lottery and fixed bonus coefficient are close to zero.

The previous section discussed that together the bonus schemes had no or only little effect on output when averaging over time. Decomposing the effect by periods shows that there are no effects in periods 1 and 2. However, in period 3 the situation is slightly different. While the joint null of no effects of either scheme on weekly output, captured by indicators for crossing the output thresholds (Table 1.10 columns 3 and 4) or by total output (column 5), cannot be rejected, this may be due to lack of power. The estimates for the difference of lottery bonus vs. no-bonus condition are individually statistically significant for both meeting the high bonus output threshold and for total weekly output (columns 4 and 5 respectively). The point estimates for period 3 in column 5 would imply a difference to the no-bonus

Table 1.10: Effects of bonus incentives on attendance and output of tea pluckers, by period (individual random effects estimation)

	(1) Full attendance [0/1]	(2) Weekly attendance [days]	(3) Met low output threshold for weekly total output [0/1]	(4) Met high output threshold for weekly total output [0/1]	(5) Total weekly output [kg]	(6) Weekly average of daily productivity if plucking [kg/day]
Period 1: Lottery Bonus	**0.0394****	**0.0696***	**0.00314**	**-0.00312**	**0.451**	**-0.676**
	(0.0188)	(0.0372)	(0.0169)	(0.0124)	(5.142)	(0.602)
Period 1: Fixed Bonus	**0.00838**	**0.0233**	**-0.0112**	**-0.0200****	**-4.579**	**-0.913**
	(0.0133)	(0.0350)	(0.0162)	(0.00970)	(4.592)	(0.570)
Period 2: Lottery Bonus	**0.00960**	**0.00467**	**-0.0123**	**-0.00559**	**-0.946**	**0.208**
	(0.0166)	(0.0318)	(0.0175)	(0.0117)	(3.687)	(0.454)
Period 2: Fixed Bonus	**0.0117**	**0.0169**	**0.000248**	**-0.00367**	**3.651**	**0.0351**
	(0.0209)	(0.0353)	(0.0159)	(0.0109)	(3.582)	(0.525)
Period 3: Lottery Bonus	**0.0550*****	**0.107*****	**-0.00960**	**0.0293****	**10.91*****	**0.190**
	(0.0142)	(0.0319)	(0.0131)	(0.0142)	(3.035)	(0.407)
Period 3: Fixed Bonus	**0.0275**	**0.0456**	**0.0114**	**0.0197**	**6.396**	**-0.0378**
	(0.0197)	(0.0452)	(0.0134)	(0.0158)	(4.502)	(0.474)
P.-val. of Wald tests:						
Period 1: Lottery Bonus = Fixed Bonus	0.0260	0.107	0.270	0.123	0.165	0.448
Period 1: Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.0638	0.121	0.516	0.0722	0.301	0.246
Period 2: Lottery Bonus = Fixed Bonus	0.912	0.718	0.339	0.827	0.317	0.770
Period 2: Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.809	0.887	0.610	0.893	0.532	0.899
Period 3: Lottery Bonus = Fixed Bonus	0.0799	0.0334	0.0680	0.480	0.226	0.561
Period 3: Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.000262	0.000155	0.189	0.121	0.00116	0.806
Mean of dep. variable in no-bonus condition during period 1	0.527	5.220	0.289	0.182	390.1	80.55
Mean of dep. variable in no-bonus condition during period 2	0.597	5.392	0.403	0.226	394.5	80.84
Mean of dep. variable in no-bonus condition during period 3	0.552	5.283	0.777	0.312	374.7	81.42
Number of gang-week observations during period 1	1611	1611	1611	1611	1611	1598
Number of gang-week observations during period 2	1141	1141	1141	1141	1141	1130
Number of gang-week observations during period 3	890	890	890	890	890	886
Number of gang-week observations in total	38,295	38,295	38,295	38,295	38,295	38,117
Number of workers	1,678	1,678	1,678	1,678	1,678	1,678

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the gang level. Estimated results are from a worker-level random effect model with a full set of gang-week dummies. Model includes interactions of periods with lagged bonus assignments as in models of tables 6 and 7 but the coefficient estimates are omitted here for clarify of presentation and are available on request. Wald-tests: For each period separately, "Lottery Bonus = Fixed Bonus" tests the equality of the two coefficients; "Joint: Lottery Bonus = 0 and Fixed Bonus = 0" test the joint equality of the two coefficients.

Table 1.11: Effects of bonus incentives on bonus criteria qualification and income, by period (individual random effects estimation)

	(1) Met bonus criteria for low bonus [0/1]	(2) Met bonus criteria for high bonus [0/1]	(3) Regular income without bonus (output x piece rate) [MK]	(4) Total expected income = regular income + expected bonus [MK]
Period 1: Lottery Bonus	**0.0165**	**0.00452**	**8.524**	**88.81***
	(0.0156)	(0.0128)	(42.17)	(45.79)
Period 1: Fixed Bonus	**0.000269**	**-0.00845**	**-35.27**	**38.63**
	(0.0141)	(0.0105)	(37.04)	(37.86)
Period 2: Lottery Bonus	**-0.00906**	**0.00850**	**-3.744**	**182.8*****
	(0.0188)	(0.0160)	(30.79)	(35.17)
Period 2: Fixed Bonus	**0.00498**	**0.0100**	**32.85**	**225.0*****
	(0.0197)	(0.0150)	(29.73)	(38.87)
Period 3: Lottery Bonus	**0.0496*****	**0.0506*****	**123.1*****	**702.8*****
	(0.0165)	(0.0159)	(32.89)	(46.42)
Period 3: Fixed Bonus	**0.0421****	**0.0353****	**72.44**	**629.8*****
	(0.0206)	(0.0146)	(46.73)	(58.28)
P.-val. of Wald tests:				
Period 1: Lottery Bonus = Fixed Bonus	0.224	0.229	0.140	0.114
Period 1: Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.426	0.430	0.273	0.129
Period 2: Lottery Bonus = Fixed Bonus	0.299	0.887	0.328	0.298
Period 2: Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.572	0.799	0.502	4.01e-10
Period 3: Lottery Bonus = Fixed Bonus	0.647	0.278	0.184	0.192
Period 3: Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.0110	0.00502	0.000655	0
Mean of dep. variable in no-bonus condition during period 1	0.225	0.152	3408	3408
Mean of dep. variable in no-bonus condition during period 2	0.330	0.191	3452	3452
Mean of dep. variable in no-bonus condition during period 3	0.438	0.230	4041	4041
Number of gang-week observations during period 1	1611	1611	1611	1611
Number of gang-week observations during period 2	1141	1141	1141	1141
Number of gang-week observations during period 3	890	890	890	890
Number of gang-week observations in total	38,295	38,295	38,295	38,295
Number of workers	1,678	1,678	1,678	1,678

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the gang level. Estimated results are from a worker-level random effect model with a full set of gang-week dummies. Model includes interactions of periods with lagged bonus assignments as in models of tables 6 and 7 but the coefficient estimates are omitted here for clarify of presentation and are available on request. Wald-tests: For each period separately, “Lottery Bonus = Fixed Bonus” tests the equality of the two coefficients; “Joint: Lottery Bonus = 0 and Fixed Bonus = 0” test the joint equality of the two coefficients.

condition of about 2.9% and 1.7% of the average weekly output in the no-bonus condition during that period (374.7 kg) for the lottery bonus and the fixed bonus, respectively. The point estimates would imply the lottery bonus effect is larger than that of the fixed bonus with a relative difference vis-a-vis the no-bonus that is broadly similar to that for attendance (about 2 times); however, the difference is not statistically significant (p-value of 0.226 for the usual two-sided test).

To further disaggregate results over time Figure 1.5 shows the difference in effects between lottery and fixed bonus schemes on attendance and output from regressions with all weeks of the experiment interacted with bonus scheme assignments. In the beginning the confidence bands are wide because of the gradual phase-in of gangs into the experiment: in week 5 only one gang was part of the experiment, in week 6 there were 4, and only in week 8 did all 25 gangs of the study participate (see Table 1.2 on page 11). Horizontal bars in weeks 11 and 14 mark the change to the experiment's periods 2 and 3, respectively, for which bonus conditions changed. The left panel of the figure plots the estimated coefficient differences for number of days in attendance; the right panel plots the same for total weekly output. While the right panel is characterized by moderate-sized swings of the coefficient difference around zero, the left panel shows that the difference is positive for most weeks and non-decreasing over time.

In conclusion, results from effects averaged over time and disaggregated by time show that both bonus schemes lead to higher attendance compared to the no-bonus group. The bonus schemes at most weakly increased output and the effects are concentrated in the last phase of the experiment when bonus amounts were largest and for which the lottery bonus is marginally statistically significantly different from the no-bonus condition.

Averaged across all weeks the lottery bonus increased attendance by a little more than twice as much as the fixed bonus. While there is substantial variation over time the results do not suggest that the effects waned over time. A caveat to that interpretation is that the amounts of the lottery were changed and so while the probabilities of winning the different bonus amounts did not change (see Table 1.6 on page 19) temporary probability weighting biases may have "reset" to some extent as the bonus amounts changed. However, there is no obvious pattern of declining effects within each of the three separate periods of the experiment.[13]

[13] In addition, bonus assignment effects did not vary significantly with the history of individual lottery draws and winnings of the lottery or with that of worker's peers (results now shown); these tests, however, are likely to suffer from low power because of limited overall effects and limited variation in the number of winners or individuals winning the large amounts.

Figure 1.5: Effect of lottery bonus vs. fixed bonus over time

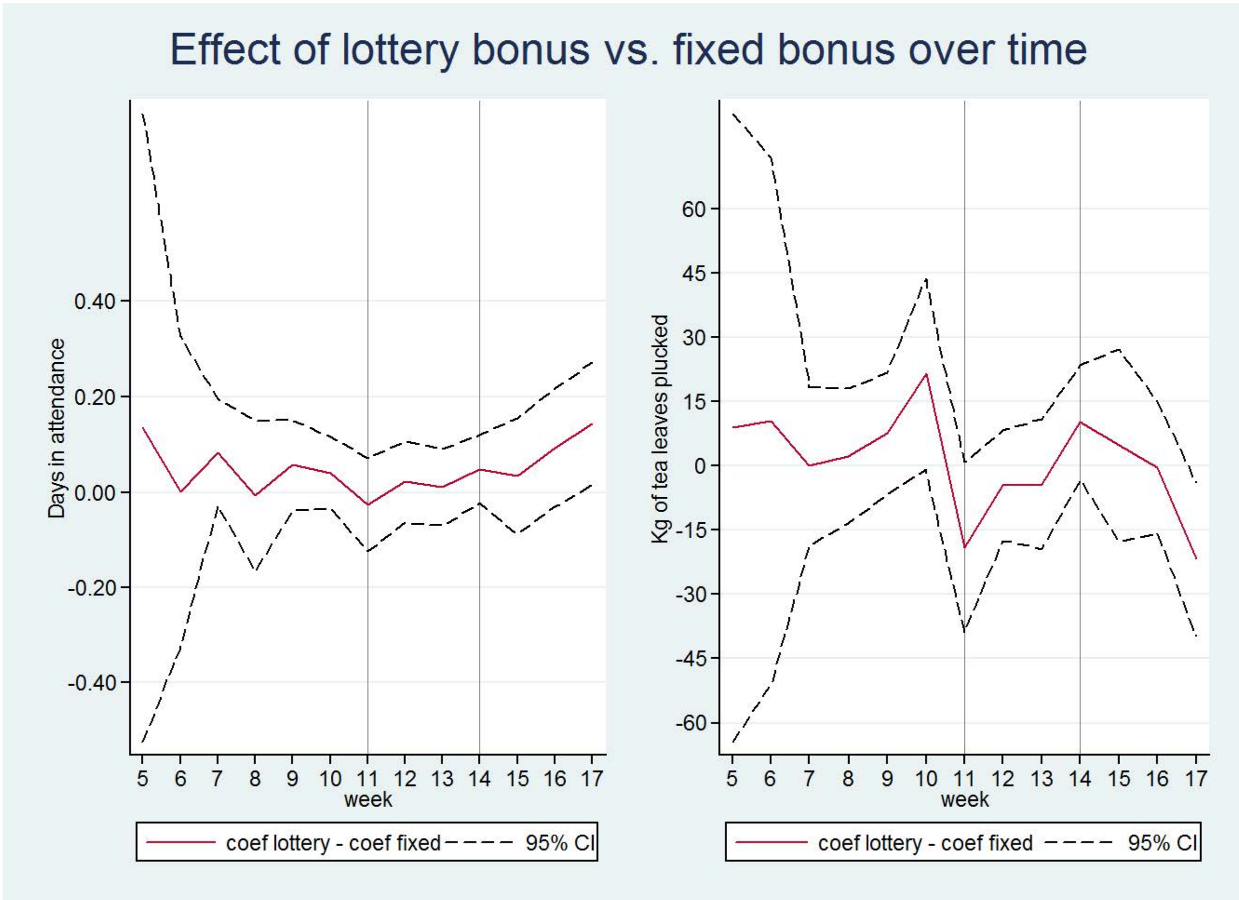

Notes: The figure shows coefficients from an estimation following equation (4) and analogous to results of Table 1.8 but with bonus assignments (both contemporaneous and one-week lags) interacted with a dummy for every week of the experiment. Standard errors from weeks 5 – 7 are clustered only at the individual level, not the gang level, due to the gradual phase-in of the experiment — with only one gang in the experiment in week 5 gang-level clustering is not possible; until weeks 8 a low number of gangs was enrolled so that the asymptotic approximation of the clustered variance estimator is poor and likely to be anti-conservative; for those weeks individually clustered standard errors are larger.

1.4.5 Further discussion of results: attendance vs. output impacts

The analysis above documents that the bonus schemes had a sizable impact on attendance, had no effects on output until period 3, and had only moderate-sized effects on output in period 3. Since both attendance and output were incentivized with the bonus one could have expected an increase in outcomes in both dimensions. In addition, even if workers only react along the dimension of attendance one might have expected total weekly output to increase mechanically, if workers continue to do be at least as productive per day as before. However, viewing attendance and productivity at work as two dimensions of effort, the effect of a bonus that discretely changes marginal benefits, as in this experiment, is theoretically ambiguous. This is especially true if the two dimensions of effort are complements in the worker's cost of effort function, which is likely to be the case, e.g. if the marginal cost of increasing output per day is increasing in the total number of days worked per week.

Under the standard piece rate attendance was not explicitly incentivized. That changed with the introduction of the bonus which required full weekly attendance. This changed discretely the benefits of coming to work every day of the week; but conditional on coming to work every day it might not be optimal for pluckers to keep up the same average daily productivity as before; it might be optimal to decrease output per day, keeping total weekly output constant while increasing weekly attendance.

Thus, the attendance-output results fall well within standard theory as applied to this setting. In addition, however, several other factors specific to the context of this experiment may have also led workers to optimally choose higher attendance and not increase total weekly output in period 1.

First, in general the marginal cost of attendance at work is subject to shocks, e.g. own illness or illness of family members; however, weekly output is subject to the same shocks *plus* additional uncertainty stemming from field conditions: a day of absence will also reduce total weekly output, and a day off work automatically disqualified a worker for the bonus scheme, and any extra productivity will go uncompensated (besides the usual piece rate). Furthermore, field conditions vary substantially over time even within gangs and daily productivity conditional on own effort is partially unpredictable. Secondly, pluckers have very precise knowledge of daily output —this is announced at the end of the day and also is committed to memory by pluckers in order to check for errors in pay calculations (payday slips list daily output); however, many pluckers do not have have good knowledge of cumulative weekly output. Computing cumulative output requires the skills of addition and also attention and effort. But education among workers is low. In addition, while adding and keeping track of output figures require additional effort, the expected return in the form of expected bonus payments was initially quite low.

Greater uncertainty about the marginal cost and marginal benefits of effort decreases pluckers expected payoff of increasing effort and with uncertainty arguably higher for weekly output, this may have led to a larger response in attendance than in weekly output. This mechanism could also explain why there are signs of an output response in the last period of the experiment. In period 3, the low output threshold was set to meeting only the daily minimum amount that is expected even outside the bonus scheme, and the high output threshold was a relative threshold. Since within-gang ranks are relatively more stable over time as varying field conditions no longer play a role, this may have reduced uncertainty associated with the output targets, and thus maybe have increased effort in that dimension, relative to attendance.[14]

In the following I evaluate if the above mechanisms related to uncertainty over output due to over-time variablility in field conditions is empirically plausible. First, to provide a sense of the relative variation in attendance and output between workers and over-time for a given worker, I decompose the overall standard deviation variation. For the number of days worked per week the ratio of within- to between standard deviation is 1.03; for total weekly kg plucked the same ratio is .83. Thus, for both outcome variables both between-worker and over-time variation are relevant. In particular for output, note that besides cross-sectional variation, output of a given worker over time varies substantially. Furthermore, this over-time variability in output for a given worker is driven to a substantial degree by time-varying field conditions and only to some degree by varying degrees of supply of effort at work or attendance. During the weeks of the experiment, of the over-time variation in weekly output remaining after netting out worker fixed effects, over a third (37%) can be explained by gang-week fixed effects alone. In conclusion, over-time variation in output is substantially, and a large fraction of it can be attributed directly to changing field conditions rather than time-varying factors such as variation in worker opportunity cost of coming to work (own sickness, family sickness, other tasks that need attention) or variation in worker effort.

1.5 Conclusion

I study the introduction of a bonus scheme with lottery payoffs at a large agricultural firm in Malawi. The "lottery bonus" was motivated by evidence from lab studies and by real world behavior that suggests that people often become risk-seeking when faced with small chances of relatively large gains. This result is borne out by lab results with hypothetical

[14] In addition, larger bonus amounts in period 3 could have contributed to the differential in effects along the two dimensions of effort. Suppose costs of attention and calculation for total weekly output are non-negligible; then since these are fixed costs with respect to the amount of the bonus, the larger the bonus amount the more likely it is that such calculations are worth it in expectation.

and real choices over gambles with small and high stakes and have helped formulate Prospect Theory.

In this paper I test if ‘excess’ risk-seeking by otherwise risk-averse agents holds for real world labor supply decisions. I collaborated with the management of a tea estate in the south of Malawi to set up a field experiment that exogenously assigned piece-rate workers to two bonus incentives that varied over time and between workers. In general, both workers’ weekly attendance and their productivity enter directly in this firm’s profit function. Attendance- and output-based bonuses were introduced for over 1,600 piece-rate workers for up to 11 weeks. Workers were assigned weekly to one of two bonus schemes or a no-bonus group. Weekly bonuses, conditional on qualifying, were paid either as a lottery or as a fixed amount of the same expected value. The (expected) values of the bonuses increased over time and ranged from between about 5% to around 15%-20% of total weekly pay at the relevant margins for qualifying.

The outcome of this experiment confirms anecdotes and one strand of previous empirical results that long-odd prospects of large gains are relatively more motivating than expected-value equivalent certain prospects. The results suggest that this relative difference holds for actual labor supply decisions with variation in actual monetary compensation. I find that the effect of the lottery bonus on the probability of full weekly attendance —a requirement for bonus eligibility and an explicit target variable of the incentive schemes— was about twice as large as the effect of the fixed bonus. When output was affected at all by the bonus scheme, namely towards the end of the experiment, the point estimate for the lottery bonus effect was larger than the fixed bonus effect though the differences was not statistically significantly different.

The differential effect of the lottery bonus on attendance is statistically significant towards the end, in the last weeks of the experiment, suggesting that despite experience with the prize distribution the lottery effect persisted, at least over the horizon of this experiment. In contrast, recent findings from lab experiments suggest that lottery preference disappears over time. One reason for the difference in results could be that lab experiments that were carried out over many repeated choices over different gambles involved only small-stakes. This experiment, in contrast, featured substantial stakes, decisions that were spaced out over time, and interactions with as well as observation of peers. On the other hand, given the limited statistical power of this design I would not be able to reject steady declines in effect sizes over time. Additionally, some details of the bonuses were changed in the last phase of the experiment —higher bonus amounts and lottery prizes, changes in output requirements— which could have led to a “reset” of overweighting of low the probabilities of winning a large prize. Future work would be able to determine the longer-run stability of the real-stakes labor supply effects measured in this experiment.

Besides the larger stakes, one obvious difference of this experiments' setting to prior studies cited above was that this study took place in a developing country. In general, take-up of lottery products is relatively higher for poorer individuals and so the incentive effect of right-skewed payoffs may be especially important in developing countries. In fact, relative subjective wealth maybe one moderator of lottery take-up[15] that could keep up the interest in the lotteries as people feel they have a chance to make more than a trivial difference in their lives and this chance is valued more than proportionally.

[15] See Haisley, Mostafa and Loewenstein (2008).

CHAPTER II

Income Timing, Temptation and Expenditures: A Field Experiment in Malawi

From a work with Jason Kerwin.

2.1 Introduction

Savings rates in developing countries appear to be very low. People save very little, whether in cash or other liquid assets. Moreover, despite evidently high returns to investment in domains ranging from health (Jones et al., 2003) to agriculture and small business (de Mel, McKenzie and Woodruff, 2008, 2012), people do not seem to be making those investments. In theory, even in the face of borrowing constraints, if returns are high enough households should be able to save up and invest. However, households appear to have trouble saving: households in developing countries act as if they are "savings constrained", meaning that transforming liquid wealth across time is costly.

In developing countries in particular, households face a range of explicit and implicit "external" costs to savings, e.g. risk of theft, high transaction costs, or lack of access to formal savings. In addition, savings constraints can be "internal" – people might be present-biased, causing them to save less than they would like. Present-biased preferences have been documented extensively in laboratory studies, and recent field research has confirmed that some people do exhibit present-biased preferences in the context of real-life choices (Giné et al., 2012). A number of papers have studied the potential of commitment savings accounts to manage this kind of internal savings constraint (Dupas and Robinson, 2013*b*; Ashraf, Karlan and Yin, 2006; Bune et al., 2014). However, the cause of present-biased preferences, and the best way to mitigate their impact on the poor's ability to save, remains unclear: in their review of the constraints that hinder savings among the poor, Karlan, Ratan and Zinman (2014) conclude that "remarkably little is known about which behavioral biases actually drive savings behavior." The canonical model of present bias is the Laibson (1997) model of quasi-hyperbolic discounting, but this sheds little light on why some people are present-biased and others are not. A possible explanation for variations in present bias comes from Banerjee and Mullainathan (2013, henceforth BM), who point out that one potential cause of variation in present bias is temptation: people may be biased toward present consumption because they are tempted to spend on goods and services that they later regret spending on, such as alcohol, tobacco, or fatty foods. Savings constraints could prevent people from saving up for large, discrete purchases (such as certain investments or durable goods), and could prevent people from having access to savings in the case of emergencies.

In light of this documented inability or unwillingness to save, the time structure of income streams is likely to be important. People in developing countries invest considerable effort and expenditure into aggregating streams of small installments of income into lump sums, in order to make purchases that cannot be broken up into small pieces (Collins et al., 2009). As a result, larger income installments may lead to more saving by easing this process. Lump-sum payments could also help savings under a BM-style temptation-based model of time inconsistency: BM show that having a larger sum of money on hand can help people overcome the fear that, if they do save, their future self will simply "waste" all of the money on temptation goods. However, converting smooth income streams into larger, deferred sums may also lead to increased temptation and potentially poor choices. Fudenberg and Levine (2006) note that ATMs are frequently placed in locations where lottery tickets are sold, or in nightclubs, in order to induce impulse purchases by myopic consumers. This proverbial effect of "money burning a hole in your pocket" is a potential concern in the Microfinance industry, where recent research has studied whether access to microcredit can induce temptation spending due to the generation of large lump sums (Angelucci, Karlan and Zinman, 2013). In addition, this phenomenon is consistent with both theoretical and empirical work in developed countries (Ozdenoren, Salant and Silverman, 2012; Stephens Jr., 2003; Shapiro, 2005) as well as with anecdotal reports of behavior around payday in developing countries.

In this paper we report results from a field experiment in Malawi designed to examine the role of timing of income for spending and savings decisions and its interaction with issues of self-control. We vary the time structure of wage payments for 363 casual laborers, with workers paid either in four weekly installments or a single lump sum at the end of the month. We also vary the day of the week on which workers are paid, with half of the sample being paid on Fridays, and half on Saturdays. All payments take place at same location, on the site of the local weekend market, which takes place on Saturdays and is reported to be an extremely tempting environment. The travel and time costs of purchasing goods at the market are held constant across study arms by attendance at the payday site by all participants on all potential paydays, even when they do not receive money. Workers who are paid on Saturdays are therefore exposed to a much more tempting environment at the time when they receive their pay (relative to members of the Friday group), with all other factors being held constant. Friday was chosen as the comparison group, rather than Sunday or Monday, in order to eliminate the possibility that people in the Saturday group save less of their income simply because the time frame is longer.

Workers in all study arms receive the same total amount of money: about MK3000, or around 30% of their total cash income over the work period; they are employed in collaboration with a local NGO in two separate rounds of work that are followed by payments with re-randomization of experimental conditions after round 1.

The experiment has both a practical and a conceptual dimension: it was designed to evaluate the role of internal savings constraints in a practically relevant context – temptations to overspend on paydays and at weekend markets and local trading centers in particular – and to test conceptually the role of temptation in mediating the differential effects on spending of income stream frequency.

Additionally, research using exogenous variation in the frequency of income streams is rare. Two very recent papers study the effect of lump sum payments relative to smoother streams of income. Studying the Malawi Social Action Fund's Public Works Project, Beegle, Galasso and Goldberg (2014) compare outcomes for workers who receive their wages in a single lump sum, instead of 5 installments over the course of 15 days (results pending). Haushofer and Shapiro (2013) find a decrease in measured cortisol levels among people who receive lump-sum, as opposed to monthly, transfers, suggesting lower levels of stress.

This paper provides novel empirical evidence in three ways. First, we provide evidence that lump sum payments have an effect on purchases of an actual investment: a high-return, short-term "bond" offered by the project to all respondents. Second, we study the effect of the timing of payments within a week, which has not been examined in the previous literature. Third, we exploit the effect of the timing of payments within a week to explore the role of temptation in driving internal savings constraints.

Qualitative evidence from the area targeted by the study found that people reported market days as tempting environments. These anecdotes were confirmed by survey responses from the respondents in our experimental sample. We also use respondents' own perceptions of regretted or mistaken expenditure, as reported on the surveys, as one of our measures of spending on temptation goods. While goods the respondents self-reported as regretted purchases included alcohol, tobacco, and sweets, the most common category was clothing. This is consistent with anecdotal reports from the local area: clothing is a major expenditure at the markets, with people making expensive purchases and then later regretting them.

Responses from survey data we collected with respondents from this study also inform us about how potential savings constraints interact with payment frequency. Before the start of the experiment we described to workers a non-incentivized, purely hypothetical situation in which they have two choices of wage payments: weekly payments or a lump sum payment in the end. Workers were informed that they would be required to come to the same location the same number of times (just as in the experiment we conducted; the hypothetical wage amounts were also nearly identical to the actual ones). 72% of workers said they preferred a lump sum payment. Of those 72%, a great majority (83%) stated, in an open ended question with at most one answer, that the reason for this preference is that enables people to a "make better plan" for the money, and an additional 13% outright gave avoiding wasteful spending as the reason. These answers imply either a commitment problem as the reason for the lump sum preference or at the least an expected inability to save – be it internal such as self-control problems or external reasons such as fear of theft.

The potential of temptation-driven waste due to market days, the frequency of payments and their interaction are not merely theoretical concerns. Many organizations in Malawi are presently moving to direct-deposit based payment schemes on an infrequent schedule that bring their employees to major cities on focal dates, potentially triggering the sorts of temptation issues discussed above. One example is Malawi's Ministry of Education; teachers now receive their pay via direct deposits into their bank accounts, as opposed to cash payments. This in turn induces a large fraction to travel to urban areas once a month to withdraw all their pay in a lump sum. A similar pattern holds for unconditional cash transfers like GiveDirectly: what makes that the program logistically feasible is that the payments are sent through the M-Pesa mobile payments service. Haushofer and Shapiro (2013) state that GiveDirectly recipients "typically withdraw the entire balance of the transfer upon receipt." Since withdrawals must be done at a participating M-Pesa agent, this will tend to draw recipients to potentially-tempting trading centers at the same time as they receive their pay. This study evaluates how infrequent payments and payments on market days in particular influence spending decisions, for a highly-relevant category of income for people in rural Africa. Prior to the beginning of our study, 77% of our sample reported having done informal agricultural work; it is a more common source of cash income than any other activity except for selling one's own crops for cash. Our intervention also involves a smaller proportion of income these other contexts: GiveDirectly provided income worth more than two months of expenditures, and the direct deposit program covers all of a teacher's income. Our respondents received additional income worth approximately 50% of their existing cash income. This limits our ability to draw conclusions about the effect of changing the timing of larger proportions of income, but also means that our study more closely resembles realistic cash transfer programs for people in rural Africa, who are likely to have existing sources of cash income as well.

We present two sets of findings from the experiment. First, in the context of this experiment, despite strong motivation from anecdotes and suggestive survey data, being paid at the site of the local market on the market day – Saturday – compared to the day before the market day – Friday – does not strongly matter for expenditure decisions in the context of our experiment. Drawing on a range of outcomes we document that neither the level nor composition of expenditures varies significantly by the day of the week that people were paid, and that the frequency of payments does not affect this result. We focus on a set of outcomes related to spending at the market on each Friday and Saturday of the study, for which we can reject even moderate-sized effects of the being paid on Saturdays relative to Fridays. However, some of our alternate outcome measures are noisy enough that we cannot that the day of week of income receipt has moderate-sized effects. This result does not conclusively rule out important payday effects in settings other those of our specific experiment – we discuss external validity in the conclusion – and it does not necessarily imply that self-control more broadly is not a binding constraint for savings. The result should, however, lower our priors about the empirical relevance of the market payday effect, certainly in contexts that are similar to the ones of this study.

Our second set of results relates to the difference in spending patterns by payment frequency. We do not find evidence that the composition of expenditures (including in particular self-reported wasteful consumption) varies with payment frequency.[1] However, we find strong evidence that the mode of payment frequency matters for workers' ability to benefit from high-return investment opportunities with a large minimum investment size. Workers in the monthly group have more cash left in the week after the last payday when the lump sum payment was made and, moreover, they are 9.5 percentage points more likely than the weekly payment group (a relative increase of 151% over the weekly mean of 6.3%) to invest in a risk-free short-term "bond" that required a large minimum installment size payment and that was offered by the project in the week after the last payday. The investment was returned to the respondent together with 33% interest after exactly two weeks. Workers knew about this opportunity before the beginning of round two of the experiment and had gained experience with the product in a pilot offer at the end of the first round. In total, lump sum group workers spent about twice as much as weekly payment group workers on the investment opportunity. We cannot entirely rule out borrowing constraints as an explanation for this result. However, based on other data, we argue that the result is driven by

[1] We elaborate on the specific features of this experiment that maybe have mitigated potential effects in the discussion of the empirical results.

savings constraints.

These results, using a novel outcome measure for investments with a large minimum installment size, also make an important contribution to existing research on the relationship between savings constraints and high returns to investment. Previous research has found that the return to investment is high, but that people do not appear to make those investments – implying that people are constrained in their ability to save up for these investments. However, prior studies either have not measured objective returns (relying on e.g. purchases of health products), or have observed high average returns in a cross-section (e.g. cash drop experiments). Research that uses investments in health products as an outcome relies on the assumption that the return to health investment is actually high, and also that respondents understand these high returns. Cash drop experiments also do not necessarily show that people are failing to pursue high-return investments. Under heterogeneous returns and borrowing constraints it is possible to observe high average returns without a binding savings constraint. Those with access to high-return investments might be limited in how much they invest at any given time because they face either a) borrowing constraints or b) they prefer to not decrease present consumption too much. As a result, people do not take advantage of all their high-return investment opportunities, allowing high returns to persist over time. Our experiment resolves both of these concerns. First, we use an actual investment with high returns and zero risk as an outcome. Second, we ensure that returns are homogeneous. In our experiment everyone has access to the same high-return investment offer, but, compared to the lump sum group, the weekly group – who are otherwise identical due to randomization – need to save to be able to invest. We observe that they do invest, but to a much lesser extent. Thus this paper provides novel evidence for savings constraints being a relevant driver of the persistence of observationally high returns to capital.

2.2 Study Design and Data

We designed a randomized experiment with informal agricultural workers from the Mulanje District of Southern Malawi. These workers took part in an expansion of an existing income-generation program that operates in Mulanje District. The subjects in the study received identical nominal[2] wages for their work, but were randomly assigned to receive the pay with different timing.

We worked with the Mulanje District Executive Council to expand a previously-existing

[2]The official inflation rate in Malawi was about 23% per annum during the study period (`https://www.rbm.mw/inflation_rates_detailed.aspx`), so prices would have risen just 1.7% per month. We therefore ignore the distinction between nominal and real wages for the purposes of our analysis.

income-generation program to an additional 365 workers[3], who worked for a total of up to 15 days in two separate rounds of work and payments. This program was part of the Sustainable Livelihoods program run by Mulanje Mountain Conservation Trust (MMCT), an NGO based in Mulanje District that focused on environmental protection and promoting sustainability in the Mulanje Mountain Forest Reserve and adjoining areas. MMCT provided detailed guidance on how to mirror their existing practices; as with the majority of MMCT's other projects, work oversight was conducted by officials from partnering government departments of the Mulanje District.

The experiment was organized into two rounds that occurred over a period of three months from November 2013 to January 2014, with subjects randomized into treatment conditions separately by round. During each round, subjects worked for two weeks and then received their pay either a) in weekly installments beginning at the end of the second week of work; or b) in a single lump sum, about three weeks after the last day of work. Figure 2.1 shows the timing of experiment with respect to work, payments and data collection. In addition to variation in payment frequency, workers received their pay either c) on Fridays or d) on Saturdays. The two variations on the timing of pay – weekly vs. monthly and Friday vs. Saturday – were cross-randomized, creating four study arms in each round. The distribution of workers into experimental groups is shown in Table 2.1a (pooled) and Table 2.1b (separate by round); details of the randomization follow further below. The payments were made at the site of a major local market that occurs on Saturdays, with the intention of inducing variation in people's temptation to overspend. During the week after the last payday in each round, all workers were visited for a detailed survey about their expenditure and income.

2.2.1 Recruitment of Workers

We worked with MMCT to locate a set of villages that were potential targets for expanding their Sustainable Livelihoods program. The key criteria for a village to be eligible were:

1. *Location.* Villages had to lie within walking distance of the Forest Reserve, because the work activities supported by the program are centered around natural resource management and conservation.

2. *No previous Sustainable Livelihoods program participation.* Because this was an expansion of the program, we excluded areas that were already actively participating in the

[3]The original recruitment included 350 workers two of which dropped early (one never showed up for work; one never showed up to receive his wage); 15 workers were added for round 2 to replace workers who dropped out after the round 1.

Figure 2.1: Timing of work, payments and data collection

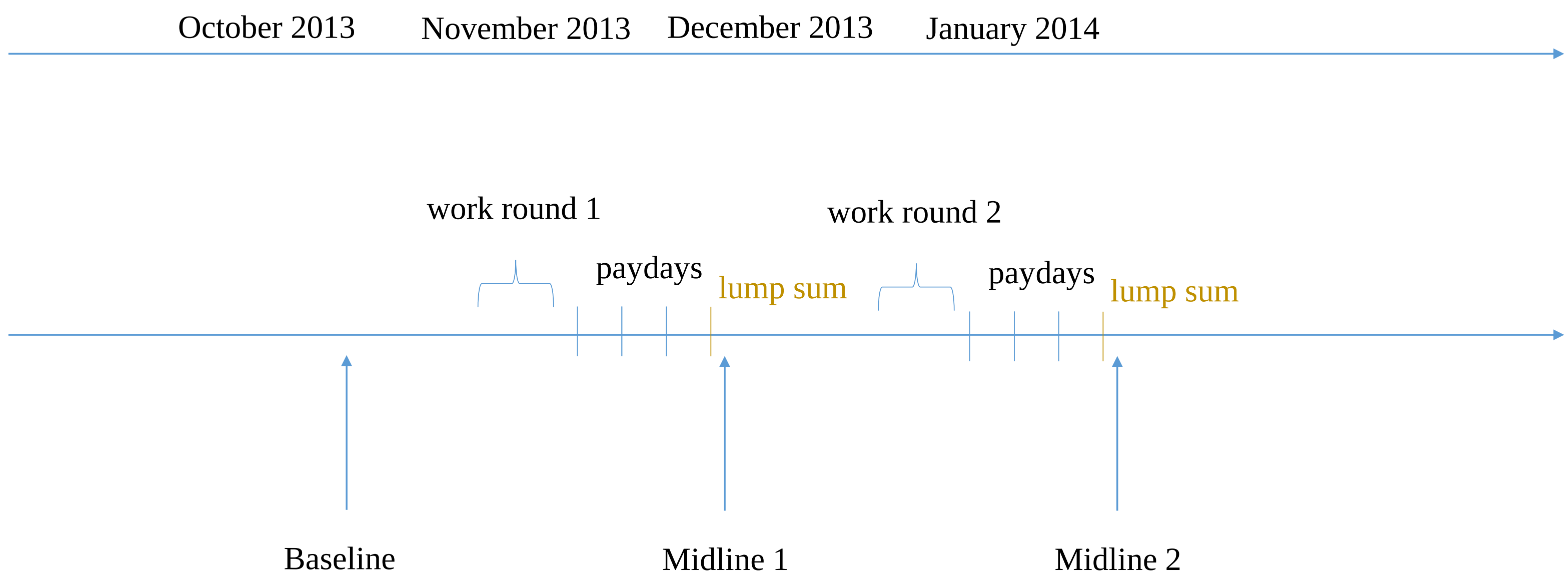

Table 2.1: Distribution of worker-round observations into experimental groups, (a) pooled across round 1 and 2 and (b) separately for round 1 and round 2

a)

Frequency \ Payday	Friday	Saturday	
Weekly	172	177	349
Lump sum	178	172	350
	350	349	699

b)

Experimental group	Round 1	Round 2	Total
Wkly, Fri	86	86	172
Wkly, Sat	89	88	177
Lump, Fri	87	91	178
Lump, Sat	86	86	172
Total	348	351	699

program, or which had been included in the past.

3. *Not included in any other recent income-generation programs.* The expansion was targeted toward underserved communities to maximize the benefits brought to the neediest people.

4. *Limited geographic range.* The villages for the study had to be physically close enough to each other to allow work and payroll to be organized across all of them together.

Given the criteria above, we settled on a region of Traditional Authority (TA) Nkanda near the Forest Reserve as the target location for the project; this area had not previously been included in the Sustainable Livelihoods program, nor recently participated in other major income-generating programs such as the Malawi government's Public Works Programme (PWP). Within that region, we picked seven villages that all lie within the catchment area of Mwanamulanje trading centre, one of the largest markets in TA Nkanda.

The selection of workers was handled by the standard operating procedure employed by the Sustainable Livelihoods program. The nature of the program, including the kind of work, the pay rate, and the expected length of employment, was explained at a meeting with the village head and the village development committee (VDC). Each VDC was then tasked with selecting a set of 50 participants and 15 substitutes. They were told to use the same criteria they generally use for deciding who should benefit from social programs. Discussions with MMCT and the VDCs revealed that the main criterion used was generally poverty, with some tendency to favor women as being more likely to be disadvantaged. The VDCs were asked to list the workers in order of preference from 1 to 65, and told we would replace workers who dropped out of the program by moving in order from position 51 to position 65 on the list of workers from their own village. This was done for a total of 15 workers at the end of the first round of the study.

This process generated an initial sample of 350 workers all of whom were interviewed in a baseline survey. One person dropped out before the work started and one person never showed up at payday (only an additional nine people missed any day of work). After all payments of round 1 were done, 343 workers were successfully interviewed in the Midline 1 survey. Before the start of round 2 of the program, 13 workers left the study, and a total of 15 replacement workers were added.[4] A total of 352 workers participated in round 2 of the study, of which all but 3 workers had full attendance and 346 were surveyed at Midline 2. The sample is similar to the broader population of the local region in most respects, differing chiefly in ways that are consistent with the selection criteria; for example, we recruited more women (69% compared 55% in the district) and our sample is slightly worse of

[4] Only the 13 drop-outs should have been added to replace the dropped workers according to the protocol.

socio-economically than the rest of Mulanje District.[5] We consider the sample to be representative of the type of person likely to be involved in government- or non-government-provided income-generation programs in Mulanje district.

2.2.2 Randomization

Our study exploits exogenous variation in the timing of individual's pay. We designed this to vary in two ways. First, the payments are either in weekly installments for four weeks, on a single lump sum during the last week. Second, the payments are made either on Fridays or Saturdays.

The effect of monthly lump sum payments, as opposed to weekly installments, is theoretically ambiguous. In a context where people have problems aggregating streams of income, receiving one's pay in a lump sum at the end of the payment period would increase take-up of profitable investments that are available after the end of the fourth week. However, if people's temptation to overspend is an increasing function of their potential immediate consumption, lump-sum payments could reduce savings instead. This would be the case if the lump sum were received concurrently with opportunities to purchase temptation goods, in which case the money could "burn a hole in people's pockets", causing them to spend money on things that *ex ante* they would prefer not to purchase. If these were the only two potential mechanisms, the variation in the frequency of pay would allow us to see which one dominates in our sample. However, the lump-sum payment could also increase savings through borrowing constraints, if people would prefer a smoother stream of income and would ideally borrowing against the future lump sum payment.

The variation in the day of the week of the payment is designed to shed light on the mechanisms behind the savings constraints people face. If money is received in a tempting environment, like the local market day, then arguably costs to resisting that temptation increase and workers would decide to spend and consume more right at the market when receiving their pay.

We picked Saturdays at the local trading center – so that payroll for this group happened during the major market in the local area – as a tempting context for the receipt of income. This choice was based on extensive qualitative and descriptive work with people in the local area. Anecdotally, people in Mulanje District often describe market days as tempting situations, in which excitement can cause them to purchase things they would rather not. Our survey data confirms this: for a free-response question about situations that are tempting or in which respondents may "waste" money, 37% of all respondents volunteered Market

[5] See Appendix B for detailed summary statistics on demographic characteristics.

Days as a tempting situation, by far the most common among those being ever tempted.[6,7] Multiple-choice questions confirmed this pattern: 69% of people said that market days are more tempting than the day before market days, and 65% of people said having a lot of cash on hand at the trading center was more tempting than having it on hand elsewhere. Based on these answers, payments during market days could exacerbate temptation-based psychological savings constraints, by inducing people to spend money on tempting goods that they would prefer to save. The alternate day – Friday – should not have the same effect on temptation spending, because the market does not take place on that day.

We chose Friday as the alternate day for several reasons. First, it was logistically simpler to manage payments on two consecutive days than on different ones; Sunday was not an option because the vast majority of our sample goes to church on Sunday mornings. Second, using the day before the market ensured that all respondents had the liquid cash needed to make purchases at the market – if we had paid the comparison group on a later day, then for the first week they would not have had any money to spend at the market on Saturday. Third, and most important, if the comparison group was paid after the Saturday group, then any differences in savings could simply be a function of having to hang on to the money for a shorter period. By choosing Friday as the comparison group, we ensured that any such effects worked against the expected direction of the results.

There are also a number of reasons why the Saturday payday might not increase temptation, as well as mechanisms that might mute the effects. First, as noted above many respondents report that having cash at the trading center is more tempting than having it elsewhere. While this is likely due to the market day itself, part of it could be independent of market days: people might just be more tempted to spend at the trading center even if the weekend market is not currently active; the selection of goods is always greater than at the village. Second, while Saturday is the major market day for the local region, there are other markets nearby that operate on Friday. Third, on an open-ended question about reasons they waste money (where the options were not read aloud), only 42% of people report spending on temptation as reason they spend money they later regret spending. This is an appreciable fraction, but if it represents all the people who could possibly be affected by the Saturday treatment, any measured effects will tend to be muted.

We employed a within-person cross-randomized design in order to maximize statistical

[6] Since 39% of respondents said they were never tempted, this constituted 58% of people who believe they ever waste money. The next-most frequent answer was "Going to the Trading Centre in general (not just market days)" with 4% mentioning it.

[7] The exact phrasing of the question in English was "In general, what are situations in which you waste money or are tempted to spend money that you would rather not spend?" where "waste" is the commonly used translation used in the local dialect of English and has a clear but less judgmental-sounding translation in the local language.

power. Individuals were randomly assigned to one study arm in the first round of the study and then to another study arm (potentially the same one) for the second round. The randomization for both rounds of the study was done prior to the baseline survey, but the group assignments were not revealed to the workers until the beginning of each round of work. For each round of the study, all workers were randomly assigned to one of four study arms: Weekly payments on Fridays, Weekly Payments on Saturdays, Monthly Payments on Fridays, or Monthly Payments on Saturdays. For the first round, the randomization assignment was stratified by village and gender. The randomization for round 2 was then stratified on the round 1 assignment and village.

2.2.3 Work Activities

Each subject worked for two weeks during each round of the project, for about four days per week, at a daily wage rate of MK400. There were 7 work days during the first round of the project and 8 days during round two. Workers were employed in conservation-oriented activities that promoted the sustainable use of natural resources. At the beginning of each round of work, representatives from the project met with the workers from each village to help them decide on the specific activities to pursue for that round, based on guidance from MMCT's sustainable livelihoods program. The two kinds of work done by the subjects during the study fell under the categories of *Tree Planting* and *Milambala*.

Tree Planting had two separate aspects. During the first round of the project, workers prepared pits for trees to be planted in, and nurseries to house the seedlings for later planting; the seedlings were provided by the Department of Forestry as part of a reforestation program in the area. During round two, which happened once the rainy season had begun, workers did the actual planting of trees. Milambala is a land conservation activity that focuses on building small bund walls to prevent the inundation of fields and limit environmentally harmful erosion of the topsoil. The principal tools needed for the work were hoes, which all the workers already owned. Milambala also required line levels and ropes, which were provided by the project.

Workers were trained in the tasks for each work activity by officials from Mulanje's District Forestry and District Agricultural Offices for Tree Planting and Milambala respectively. Progress on the work was also overseen by officials from the two departments, who set targets for the work to be done on each day and checked in to make sure it was accomplished.

2.2.4 Payroll

Payroll for the project was organized at Mwanamulanje Trading Centre, a major local market in TA Nkanda that was within 4 kilometers of all the villages included in the study. Subjects were informed about how they would be receiving their pay (weekly or monthly, Fridays or Saturdays) at the beginning of each round of work; the procedure was explained verbally, and they were also given a simple handout explaining their group assignment.

To ensure that transit and time costs were held equal across the four study arms, all subjects were required to come to the payroll site on all eight days during each round – even when they were not being paid. This also allowed us to collect high-frequency data on people's cash holdings and spending behavior, via questions that we asked during the payroll administration. In order to encourage attendance and defray some of people's time costs, all subjects received an MK100 show-up fee for each day, on top of any money they were slated to receive as part of their pay for the project. For example, a person who was paid monthly on Fridays was required to come to the market on all the preceding Fridays and Saturdays, and received MK100; on the day she received her pay, she received MK100 plus her entire wages for the project. The payment schedules in each round across the four payday weekends resulting from the show-up fees and payment of wages according to treatment group and number of work days is overviewed in Table 2.2.

MMCT ordinarily manages payroll for its activities using experienced cashiers that work for the organization. For this project, the cashiers were instead employees from the Mulanje District council.

The location and timing of the payroll was specifically chosen to maximize the likelihood that people would be exposed to temptation goods. In pilot testing and qualitative work, people commonly reported market days as periods when they were tempted to spend against their ex ante plans, or tended to waste money. The market at Mwanamulanje happens only on Wednesdays and Saturdays (with Saturdays having the larger market out of the two days), and principally in the morning, which is when people were paid. Shops are still open on Fridays, and there are some mobile vendors, but the majority of market activity happens on Saturdays.

While the purpose of the show-up fee on non-payday days was to equalize transaction costs across treatment groups and make spending patterns comparable, the fact that some amount of money was paid each time may have reduced the potential to observe differences across groups: possibly workers satisfied most of their temptation consumption needs with the MK 100 they received each time they showed up at the market.

Table 2.2: Payment schedules by payday group and round (all values in MK)

	Payday weekends							
	#1		#2		#3		#4	
Round 1	Fri	Sat	Fri	Sat	Fri	Sat	Fri	Sat
Payment group								
Wkly, Fri	**800**	100	**800**	100	**800**	100	**800**	100
Wkly, Sat	100	**800**	100	**800**	100	**800**	100	**800**
Lump, Fri	100	100	100	100	100	100	**2,900**	100
Lump, Fri	100	100	100	100	100	100	100	**2,900**
Round 2								
Payment group								
Wkly, Fri	**900**	100	**900**	100	**900**	100	**900**	100
Wkly, Sat	100	**900**	100	**900**	100	**900**	100	**900**
Lump, Fri	100	100	100	100	100	100	**3,300**	100
Lump, Fri	100	100	100	100	100	100	100	**3,300**

2.2.5 Data

Our data comes from three distinct sources. A detailed survey, focused on expenditures in the past week; several single-item recall questions administered during the payroll; and, as an objective measure of savings behaviors, respondents' choices about purchasing a short-term, high-return, zero-risk investment offered by the project at the end of the second round of the study.

The survey data was collected three times: once at baseline, and once after each round of the study. Subjects were interviewed at their homes, and answered questions about income, assets, savings, and financial transfers, as well as a detailed module about their expenditures since the previous Friday. This module went through a list of goods and asked respondents if they had bought the good since the previous Friday. If they said "yes" to a good, they were asked about how much they bought on each of Friday, Saturday, and Sunday up to now.

Also part of the survey data were a set of questions on wasting money and being tempted to buy things one should not. Respondents were asked about goods that respondents found particularly tempting, or that they thought they wasted money on, as well as situations in which they felt they wasted money. They were also asked ex post about whether they felt they had wasted money in the period since they received their pay; this question was only included on the survey after the second round.

Our second data source is a set of questions asked during the payroll process. On each of the eight paydays, all respondents were required to come to the payroll site as described above. Prior to receiving their pay, they were asked simple aggregate questions about the money they had on them at the time (not including their pay, which they had yet to receive) and the amount of money they spent at the market on the previous payday. Hence on Fridays, people were asked about the money they spent on the Friday of the previous week, and on Saturdays, they were asked about the money they spent yesterday. During the second round of the study, we also asked two additional questions as sensitivity checks: first, we asked people to recall their spending from the Friday of the previous week, to look at the influence of recall bias. Second, we asked people about money they spent outside of the market, in case there were differential patterns in non-market spending.

A third source of data comes from an investment opportunity offered to respondents at the end of the second round of the study. This opportunity was announced before the first payday for the second round, so all respondents had a chance to take part irrespective of their treatment status. Respondents were offered the chance to buy the investment good only once, immediately after we visited them for the midline survey. The investment took the form of a "bond", with shares that cost MK1500 to purchase and that paid back the principal plus MK500 interest after exactly two weeks. Each respondent could buy a maximum of two shares, and no fractional shares were allowed. All respondents who purchased the bond were paid back on time according to the terms of the investment.

The investment good was intentionally offered only once per round, in the week after the final payment was made. This allows us to use it to test for the existence of savings constraints, since members of the weekly group must save their pay in order to use it for the investment good. An alternative design would have offered the investment opportunity each week. This would have lowered the amount of time that the weekly group needed to save in order to purchase it, thus relaxing the savings constraint somewhat. We chose this design in order maximize our statistical power to detect differences across the two groups.

Summary statistics from these data sources for all variables used in the regression analysis are presented in Table 2.3, separately for pre-experiment baseline and for outcome variables. At baseline, the households' total spending considering all expenditures from the last Friday prior to being interviewed up to the day of the survey averages MK2,257 (about US$5.6 or PPP$ 14). Respondents report having an average of MK670 (about US$1.7 or PPP$4.2) left out of the money they had received since the Friday prior to interviewing. Households spend about 69% of their total expenditures on food for preparation at home, another about 6% on immediate consumption away from home and about 28% on non-food items.[8] About a third of food expenditure was on maize, which is the principal staple crop in the region.

Randomization led to a sample with no notable differences in pre-program characteristics across study arms. See discussion in Appendix B.

2.3 Empirical Specification

We study the effects of the experimentally-induced variation in payment timing on several sets of outcomes: expenditure at the market when payment was received; total expenditure levels and composition over the last weekend of each round, including self-reported wasteful expenditures; asset accumulation; and take-up of the large installment-size, risk-free, high-return investment opportunity.

We present three regression specifications reported as separate panels in the main results tables. The first tests the effect of being randomly assigned to a Saturday payday relative to being assigned to a Friday payday. In Panel A of the subsequent tables, we run regressions of the form

$$Y_{ir} = \alpha Saturday_{ir} + \boldsymbol{\beta}'\mathbf{X}_{ir} + \varepsilon_{ir} \tag{2.1}$$

[8]The shares do not add to 1 exactly due to Winsorizing.

Table 2.3: Summary statistics

	Mean	Std. dev.	10th percentile	Median	90th percentile	Obs.
Baseline variables						
Index of asset ownership	-0.02	2.695	-2.489	-0.713	3.061	342
Total spending since incl. last Fri [MK]	2257	3763	200	1000	4600	321
Remaining cash out of received since incl. last Fri [MK]	670	2623	0	20	1400	321
Expenditure shares based on itemized elicitation						
Food for consumption at home	0.69	0.214	0.361	0.742	0.937	341
Maize only	0.234	0.26	0	0.17	0.605	341
Food for consumption out of home	0.061	0.069	0	0.038	0.144	341
Non-Food	0.279	0.235	0.04	0.189	0.655	341
Outcome variables						
Market spending on paydays						
Amount spent on day of wage receipt	1645	1151	200	1500	3200	683
Amount spent at market on Fridays 1, 2, & 3	651	685	200	300	1895	690
Amount spent at market on Saturdays 1, 2, & 3	829	759	200	480	2300	691
Amount spent at market on Friday 4	524	761	50	120	1500	675
Amount spent at market on Saturday 4	823	939	60	500	2300	689
Follow-up survey measures						
Total spending since incl. last Fri [MK]	2509	2395	800	2300	4000	689
Remaining cash out of received since incl. last Fri [MK]	529	996	0	0	2000	689
Expenditure shares based on itemized elicitation						
Food for consumption at home	0.698	0.212	0.371	0.751	0.930	689
Maize only	0.359	0.266	0.000	0.371	0.709	689
Food for consumption out of home	0.051	0.056	0.000	0.034	0.125	689
Non-Food	0.251	0.206	0.043	0.188	0.572	689
Value of net asset purchases since last interview	2154	7486	0	0	5300	689
Self-reported wasteful spending on weekend 4 of round 2						
Total since incl. last Fri [MK]	306	685	0	25	800	346
Friday [MK]	164	462	0	0	400	346
Saturday [MK]	73	256	0	0	150	346
Sunday and after [MK]	66	281	0	0	90	346
Round 2 investment opportunity take-up						
Bought any shares [0/1]	0.108	0.311				351
Total spent on shares [MK]	265	798	0	0	1500	351

Notes: Sample includes 359 respondents who participated in at least one round of the work program and have data from at least one data source for that round (either the payday data, the survey, or both). All money amounts are in Malawian Kwacha (MK); during the study period the market exchange rate was approximately MK400 to the US dollar, and the PPP exchange rate was approximately MK160 to the US dollar. See Appendix B for variable definitions.

Y_{ir} is the dependent variable of interest for worker i in round r. $Saturday_{ir}$ is an indicator variable for individual-level assignment to a Saturday payday in round r. The coefficient α measures the effect of receiving wages on Saturday (either in four weekly installments or in a lump sum at the end). $\mathbf{X}_{ir}$ is a vector that includes stratification cell dummies; two household financial variables measured at baseline prior to the randomized assignment;[9] and a linear function of the weekday of the exogenously-assigned (first attempted) interview date. The available baseline controls are summarized in Table 3. ε_{ir} is a mean-zero error term.

Whenever data from both rounds are used (so r=1, 2 in the equation above) standard errors are clustered at the worker level to account for statistical dependence of outcome measures for the same individual across the two rounds. The stratification cell dummies are separate by round, so these implicitly control for round fixed effects when multiple rounds are used.

In Panel B, we compare the impact of the payday assignment separately for weekly and lump sum payment frequency. Regressions are of the form

$$Y_{ir} = \gamma_1 Sat\_wkly_{ir} + \gamma_2 Lump\_sum_{ir} + \gamma_3 Sat\_lump_{ir} + \boldsymbol{\beta}'\mathbf{X}_i + \varepsilon_{ir} \tag{2.2}$$

where Y_{ir} and $\mathbf{X}_{ir}$ are defined as above. Sat_wkly_{ir} is an indicator for assignment to the Saturday payday group and weekly payments, $Lump\_sum_{ir}$ is an indicator for assignment to the lump sum payment group (either payday), and Sat_lump_{ir} is an indicator for assignment to the Saturday payday group and lump sum payments. The coefficient γ_1 represents the effect of assignment to the Saturday payday group relative to the Friday payday group among those who are assigned to weekly payments. γ_2 represents the effect of assignment to the lump sum payment condition relative to the weekly payment condition among the Friday payment group. γ_3 captures the effect analogous to γ_1 for the Saturday payday effect but among those who are assigned to the lump sum condition. The difference between γ_3 and γ_1, then, captures how the Saturday payday effect varies across payment frequency. The estimated coefficient difference and the p-value for the test of the null hypothesis of no difference (between γ_3 and γ_1) is reported at the bottom of each Panel B.

Finally, panel C – and all specifications shown in Table 7 – is analogous to panel A, except the included experimental group indicator captures the effect of being assigned to the lump sum payment condition relative to the weekly payment condition and gives the according effect averaged across payday assignments (i.e. across Friday and Saturday assignments).

[9] Our baseline financial controls are an index of asset and livestock ownership (using principal component analysis) and the total amount of money the respondent spent out of their income received since the Friday prior to the baseline survey. Results are not sensitive to the specific choice of baseline financial controls.

In general, workers in this project interact with each other and so in theory we cannot exclude that workers assigned to one experimental group had an impact on workers in another. Our design does not allows us to address potential contamination. In the context of our design, this should bias results against finding differences between treatment groups: for example, if monthly payment group members gave loans to weekly payment group members, this should differences in expenditures between the two groups. Additionally, empirically, we find no evidence of increased cash or in-kind transfers for any of the experimental groups (results not shown).

2.4 Empirical Results

2.4.1 Saturday vs. Friday Payday

We examine the effects of the experimentally-induced variation of being paid on Saturday compared to Friday on expenditures. We begin by looking at reported expenditures at the market on all four paydays. Next, we consider total spending over the course of the last payday weekend and into the following week, and the composition of spending. We consider estimates of equations (2.1) and (2.2) in Panels A and B, respectively, of Tables 4 through 6. We return to estimates shown in Panels C of the same tables further below.

We first examine how the specific day on which people were paid affected spending at the market over the course of the eight paydays during each round. Table 2.4 presents estimates for outcomes from the panel of data collected during paydays.

Table 2.4 column 1 presents the effect of the treatment on the amount of money people spend at the market on the day that they receive their wages. This variable measures expenditure on Fridays for the Friday condition and on Saturdays for the Saturday condition; it includes spending on all four paydays for the weekly condition, but only on the fourth week of paydays for the lump sum condition. Panel A indicates that the day of receipt did not matter for same-day market expenditures. If receiving pay in the environment of Saturday's weekend market was tempting for workers then we should expect to see workers in the Saturday group spending more at the market on the day they were paid. The point estimate is close to zero and relatively tightly bounded: the mean of the dependent variable in the Friday group is MK 1656, the point estimate for the Saturday effect is MK -12.5 with a standard error of MK 90. Panel B shows the Saturday effect separately for weekly payment condition and the lump sum condition. The difference in coefficients and p-values of the test of no difference is given below Panel B. There is no differential effect by payment frequency.

Table 2.4 columns 2 through 5 reveal that those workers with payments on Friday spend more money at the market on Fridays – the estimate of the Saturday coefficient is negative for

Friday expenditures – and those with payments on Saturday spend more on Saturdays. The negative coefficient on the Saturday dummy is larger in absolute value for Friday outcomes than for Saturday outcomes, suggesting that Friday wage receivers spend some of their money on Saturday, while Saturday wage receivers do not have extra funds to spend on Friday – the day before their pay receipt.

The natural follow-up question is to ask whether total expenditures over the whole weekend and in the days following the payday weekend were different by Saturday vs. Friday payment. Thus, we turn to Table 2.5 column 1 which presents the effects for days of and after the fourth payday weekend for each of the two rounds, including also non-market expenditures. In Panel A, the point estimate for the Saturday effect is positive but far from statistically significant; taken at face value, the point estimate of MK 59.27 would imply a relative effect of ca. 2.4% of the Saturday assignment on total expenditures compared to the Friday assignment (mean of MK 2500). Compared to the market data of Table 2.4 that was available for all payday weekends, the standard errors are higher, and so moderate Saturday effects cannot be rejected with high confidence for this outcome variable. Panel B shows that the Saturday effect point estimates are higher for the lump sum condition but neither is the sub-group difference individually statistically significant nor is the difference across the two payment frequency conditions significant.

Column 2 shows a statistically-insignificant but negative estimated effect of the Saturday condition on the amount of cash respondents had received since the Friday before the interview but had not yet spent. The difference of about MK 114 is large relative to the Friday payday condition mean of MK 579, and so we cannot reject moderate-sized effects on this outcome.

We have established that there is no detectable Saturday effect on the level of expenditures on market day and beyond. However, if Saturdays are tempting, being paid on Saturdays could also affect the composition of expenditures. To explore this we look at the two sets of outcome variables: self-reported wasteful expenditures and the composition of spending in broad expenditure categories. Table 2.5 columns 3 through 6 show effects on self-reported wasteful spending ("How much did you spend on items that you later thought you should not have spent money on?"), both in total for the last payday weekend as well as separately for Friday, Saturday and after. Table 2.6 columns 1 through 4 show expenditure shares in broad categories. These data are constructed from detailed, itemized listings. Again, we find no significant Saturday effects on average or in interactions with payday frequency. The confidence bands around these sets of point estimates implied by the standard errors are, however, not very narrow and so relatively large differences – relative to the mean in the comparison group – cannot be rejected with confidence.

Table 2.4: Effects of treatment assignment on market spending

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	Amount spent on same day as income receipt	Money spent at market on Fridays 1, 2, 3	Money spent at market on Saturdays 1, 2, 3	Money spent at market on Friday 4	Money spent at market on Saturday 4
Panel A - Saturday vs. Friday					
Saturday payday	-12.50	-558.6***	282.3***	-750.6***	203.1***
	(90.15)	(48.27)	(58.43)	(50.35)	(67.31)
Mean dep var Friday payday	1656	930.6	686.1	912.6	722.7
Panel B -Saturday vs. Friday x Frequency					
(a) *Saturday* payday, *weekly* payment	18.53	-1,151***	542.1***	-375.3***	202.2***
	(114.3)	(64.22)	(89.18)	(41.02)	(76.56)
(b) Any payday, lump sum	-828.2***	-1,189***	-408.1***	671.3***	489.8***
	(107.2)	(61.47)	(62.08)	(87.04)	(95.94)
(c) *Saturday* payday, *lump* sum	-64.11	3.892	1.954	-1,106***	224.7**
	(119.4)	(32.23)	(50.43)	(84.52)	(109.5)
Mean dep var excluded category (Friday payday, weekly payment):	2068	1540	892.8	557.1	474.2
Difference in coefficients and p-values of tests of no difference; H0: (Sat, lump - Fri, lump) - (Sat, wkly - Fri, wkly) = 0					
Coefficient difference: row (c) - row (a)	-82.65	1155	-540.1	-730.8	22.54
P-value of H0: coef difference = 0	0.620	0	0	0	0.869
Panel C - Lump sum vs. weekly					
Lump sum payment	-869.8***	-598.3***	-683.4***	309.4***	496.6***
	(83.55)	(49.26)	(53.83)	(56.15)	(73.40)
Mean dep var weekly payment	2067	950.6	1170	362.3	576.3
Number of observations	683	690	691	675	689

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Regressions are run on pooled data from round 1 and round 2 (see Empirical Strategy for details). Standard errors are clustered at the individual level in parentheses. USD 1 is ca. MK 400 for study period. All regressions include stratification cell fixed effects and an index of baseline asset ownership based on first principal components, difference in days between date of interview and the preceding weekend, baseline total spending. For complete variable definitions, see Appendix B, and Table 2 for summary statistics.

Table 2.5: Effects of treatment assignment on total spending and cash saving and wasteful spending

	(1)	(2)	(3)	(4)	(5)	(6)
			Self-reported wasteful spending on weekend 4 of *round 2*			
Dependent variable:	Total spending since last Fri, inclusive [MK]	Remaining cash out of received since last Fri, inclusive [MK]	Total since last Fri, inclusive [MK]	Friday [MK]	Saturday [MK]	Sunday and after [MK]
Panel A - Saturday vs. Friday						
Saturday payday	59.27	-113.6	-59.74	-123.8**	52.14*	-3.199
	(165.4)	(76.36)	(71.68)	(47.85)	(27.58)	(32.45)
Mean dep var Friday payday	2500	579.2	324.3	220.7	43.49	64.11
Panel B -Saturday vs. Friday x Frequency						
(a) *Saturday* payday, *weekly* payment	-87.93	-51.66	-126.2	-116.6*	12.61	-17.59
	(232.8)	(109.3)	(88.62)	(60.60)	(30.30)	(48.82)
(b) Any payday, lump sum	1,275***	196.3*	26.22	71.81	-6.284	-21.24
	(247.5)	(106.8)	(107.6)	(83.14)	(24.94)	(44.61)
(c) *Saturday* payday, *lump* sum	265.8	-170.0	6.301	-129.5*	91.02**	10.60
	(214.9)	(109.2)	(112.7)	(76.15)	(43.85)	(42.22)
Mean dep var excluded category (Friday payday, weekly payment):	1881	483.5	322.3	189.1	52.59	73.53
Difference in coefficients and p-values of tests of no difference; H0: (Sat, lump - Fri, lump) - (Sat, wkly - Fri, wkly) = 0						
Coefficient difference: row (c) - row (a)	353.7	-118.3	132.5	-12.90	78.41	28.19
P-value of H0: coef difference = 0	0.266	0.448	0.358	0.897	0.129	0.659
Panel C - Lump sum vs. weekly						
Lump sum payment	1,451***	139.2*	92.77	66.54	32.28	-7.165
	(159.1)	(71.40)	(70.53)	(47.43)	(27.63)	(31.06)
Mean dep var weekly payment	1833	468.5	261.8	132.3	58.36	67.60
Number of observations	689	689	346	346	346	346

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Regressions of columns 1 and 2 are run on pooled data from round 1 and round 2 for which standard errors are clustered at the individual level; remaining columns use only round 2 data since outcomes are not available in round 1. All regressions include stratification cell fixed effects and an index of baseline asset ownership based on first principal components, difference in days between date of interview and the preceding weekend, baseline total spending and -if available- the baseline value of the outcome variable. For complete variable definitions, see Appendix B, and Table 2 for summary statistics.

Table 2.6: Effects of treatment assignment on expenditure composition and asset accumulation

	(1)	(2)	(3)	(4)	(5)
	Expenditure shares based on itemized elicitation				
Dependent variable:	Food for consumption at home	Maize only	Food for consumption out of home	Non-Food	Value of net asset purchases since last interview
Panel A - Saturday vs. Friday					
Saturday payday	-0.00366	0.00327	0.00443	-0.00125	-805.3
	(0.0154)	(0.0183)	(0.00413)	(0.0155)	(534.4)
Mean dep var Friday payday	0.700	0.359	0.0491	0.250	2581
Panel B -Saturday vs. Friday x Frequency					
(a) *Saturday* payday, *weekly* payment	0.0159	0.0181	0.00902	-0.0257	-520.6
	(0.0217)	(0.0270)	(0.00570)	(0.0213)	(842.3)
(b) Any payday, lump sum	0.00443	0.0328	0.000595	-0.00579	287.1
	(0.0233)	(0.0282)	(0.00591)	(0.0228)	(865.7)
(c) *Saturday* payday, *lump* sum	-0.0237	-0.0107	-0.000288	0.0238	-1,088
	(0.0221)	(0.0246)	(0.00618)	(0.0219)	(725.0)
Mean dep var excluded category (Friday payday, weekly payment):	0.702	0.348	0.0473	0.250	2604
Difference in coefficients and p-values of tests of no difference; H0: (Sat, lump - Fri, lump) - (Sat, wkly - Fri, wkly) = 0					
Coefficient difference: row (c) - row (a)	-0.0396	-0.0289	-0.00931	0.0495	-567.4
P-value of H0: coef difference = 0	0.207	0.430	0.276	0.103	0.623
Panel C - Lump sum vs. weekly					
Lump sum payment	-0.0153	0.0182	-0.00416	0.0190	19.61
	(0.0162)	(0.0192)	(0.00449)	(0.0161)	(525.7)
Mean dep var weekly payment	0.707	0.352	0.0523	0.240	2271
Number of observations	689	689	689	689	689

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Regressions are run on pooled data from round 1 and round 2 (see Empirical Strategy for details). Standard errors clustered at the individual level in parentheses. USD 1 is ca. MK 400 for study period. All regressions include stratification cell fixed effects and an index of baseline asset ownership based on first principal components, difference in days between date of interview and the preceding weekend, baseline total spending and -if available- the baseline value of the outcome variable. For complete variable definitions, see Appendix B, and Table 2 for summary statistics.

The wasteful spending variables in Table 2.5 are only available for round 2; we choose to shows this set of outcomes as it most unambiguously reflects temptation spending and avoids constructing outcomes with researcher-imposed ideas of what are temptation expenditures. There are multiple ways of constructing outcomes with the same intention. One variation that we have explored is based on reports of unplanned purchases of items: we have considered both items that are commonly unplanned purchases across the whole sample, as well as individual self-reports that a specific purchase was not planned. Neither of these variations affects the pattern of no significant treatment effects, and so we omit these alternative specifications for brevity. Lastly, column 5 of Table 2.6 shows that over the course of the entire payment period, Saturday payments did not differentially affect asset accumulation compared to Friday payments.

We also explore the possibility that the day of the week on which people are paid might interact with the frequency of their pay. The estimates and p-values in Panel B of the Tables 4 to 7 show that there is no significant interaction of Saturday payments with payment frequency for the remaining cash outcome variable. We consistently see no significant interactions between the day of the week on which respondents are paid and the frequency of their payments.

2.4.2 Lump Sum Payment vs. Weekly Payments

In the preceding section, we showed that Saturday payments did not affect expenditures compared to Friday payments. In this section we focus on the effect of receiving a lump-sum payment relative to receiving weekly installments. Workers were randomized into one of the two payment frequency conditions; the lump sum group received wage payments on the last of four weekends at which the weekly payment condition received their wages. However, all workers were required to come to the site where payroll was administered every Friday and Saturday on all four payday weekends, even if no wages were received. Workers received a small “show-up fee" of MK 100 and were also asked the payday questions described in Section 2 above.

We briefly return to the outcomes of Tables 4 through 6 discussed above. Panel C shows the effect of lump sum payments vis-à-vis the weekly payment condition, all of which are strongly statistically significant. Table 2.4 column 1 shows market expenditures on the day of receipt of pay across the entire payment period (e.g. across all four Fridays of reach round for those paid weekly on Friday; or spending on Saturday of the last payday only for those in the Saturday lump sum group). On average, workers in the lump sum condition spent MK 870 less of their total pay at the market on the same day that they received it. In payday weekends 1 through 3, when the lump sum condition was not receiving any wages, market expenditure on paydays was lower in the lump sum condition: on Fridays 1, 2 and 3 in total workers only spend about 37% of the average in the weekly payment condition (Column 2) and the same rate is about 42% for Saturdays (Column 3). On the last payday weekend, when those in the lump sum group receive their wages, expenditures are higher by MK 309 and MK 497, respectively. The increase in the monthly group's expenditures during the fourth weekend is smaller than their decline in expenditures in weekends 1, 2 and 3.

Table 2.4 concerned expenditures at the market; Table 2.5 (Panel C) looks at survey measures of total expenditures during the fourth payday weekend. Consistent with the payday data about market expenditure, total expenditures over the weekend and into the following week are higher for the lump sum group (by MK 1,451, column 1). Despite the higher spending, cash remaining on hand out of the money received since the Friday prior to the follow-up interview are marginally statistically significantly higher with a point estimate of ca. MK 139. Wasteful spending, however, was not significantly different for the lump sum group (columns 3 through 6), suggesting that the higher receipt of cash in one chunk does not lead recipients to overspend on goods they later regret – at least in this context. While the standard errors are large enough that we cannot reject a doubling of wasteful expenditure, the results from Panel C of Table 2.6 are consistent with the idea that the composition of expenditure did not change in the monthly group (Columns 1 to 4). The shares of expenditure in different broad item categories were not significantly different.

Column 5 of the same table examines if higher expenditures lead to differential asset purchases. The estimates show that net asset accumulation over the course of the all payday weekends does not appear to be different between lump sum and weekly payment conditions. However, standard errors are large and so economically significant effects cannot be ruled out by these estimates.

Lastly, in Table 2.7 we examine the effect of lump sum payments on take up of a large minimum-installment, high-return, risk-free "investment opportunity" that was offered to respondents right after the follow-up interview.[10] Workers were able to buy either 1 or 2 "shares" from the project that had a risk-free return of 33% and were repaid after exactly two weeks. This investment opportunity was offered to test if the timing of payments affects respondents' ability to take up profitable but lumpy investment opportunities. The main advantages of this novel outcome variable is that it provides a controlled investment instrument with known features, and, moreover, that it makes a high-return investment opportunity, that requires a large minimum investment, homogeneously available to every respondent at

[10] There is no effect of Saturday vs. Friday payments on these outcomes, consistent with the lack of difference in remaining cash after weekend 4. For clarity of presentation we omit the specifications of Panels A and B and focus only on the regressions analogous to Panel C in the preceding results tables.

the time of surveying. In real life respondents' opportunities vary widely cross-sectionally and, importantly, over time – e.g. farming investments are largely only available during a limited period of the year.

In round 1 the opportunity to invest was only announced in the week preceding the final payday. This limits the usefulness of the round 1 results, because workers already knew their treatment status but did not know about the investment opportunity until a week before it was made available to them. This could bias any estimated effects either upwards or downwards. An upward bias ccould occur because the weekly payment group members did not know about this opportunity until they had received three quarters of their wage. The wage amount remaining to paid in the last payday weekend was smaller than minimum requirement amount for the investment opportunity (the remaining payment was MK 800 but one unit of the investment offer was priced at MK 1500); this would eliminate the subset of weekly workers who had less than MK700 in weekly income from being able to purchase the investment good. A downward bias could occur because lump sum payment group members may have already committed their pay to other expenditures. This would limit their ability to purchase the investment good, thus understating any measured effects.

In contrast, in round 2 the investment opportunity was announced before the start of the round, so all respondents across both groups knew they would have the opportunity prior to learning which payment group they were in. Workers therefore had advance notice of the prospect of this opportunity before any wage payments began, and before they could potentially commit any of their wages to other expenditures in a way that depended on their study arm assignment. Because of these differences in setup across rounds, we show results both from regressions on pooled data from both rounds and then specifically for round 2.

Table 2.7 columns 1 and 2 repeat outcome variables from Table 2.5 columns 1 and 2 (cf. Panel C) to be able to track differences due to changing sets of observations across Panels I through III (pooled, round 2 only, round 1 only, respectively). Columns 3 and 4 of Table 2.7 show effects on take-up of the investment opportunity. In Panel I, pooling observations across the two rounds, lump sum payment group members had a 4.8 percentage points higher probability of buying any share (significant at the 10% level) and the total amount spend on the investment opportunity was about MK 122 higher (significant at the 5% level). The comparison to Panels II and III show that this effect is concentrated in round 2 where the effect of lump sum payments on probability of taking up 9.5 percentage points, relative to a base of only 6.3% among the weekly payment group.[11] Total spending – number of shares

[11] Takeup actually remains the same across rounds for the monthly group and declines from round 1 to round 2 for the weekly group. However, we cannot draw any strong conclusions from this pattern because of general seasonal variations in behavior - for example, spending levels are generally higher in round 1 before the start of the lean season in round 2.

Table 2.7: Effects of treatment assignment on post-interview risk-free, high-return investment offer

	(1)	(2)	(3)	(4)
Dependent variable:	Total spending since last Fri, inclusive [MK]	Remaining cash out of received since last Fri, inclusive [MK]	Bought any shares [0/1]	Total spent on shares [MK]
Panel I - Round 1 and 2 pooled				
Lump sum payment	1,451***	139.2*	0.0484*	121.7**
	(159.1)	(71.40)	(0.0247)	(58.81)
Mean dep var weekly payment	1836	468.5	0.106	223.5
Number of observations	689	689	699	699
Panel II - Round 2 only				
Lump sum payment	1,658***	274.0***	0.0949***	196.2**
	(190.6)	(96.82)	(0.0327)	(84.80)
Mean dep var weekly payment	1634	393.1	0.0632	172.4
Number of observations	346	346	351	351
Panel III - Round 1 only				
Lump sum payment	1,252***	-4.320	0.00396	52.51
	(245.2)	(109.6)	(0.0381)	(79.20)
Mean dep var weekly payment	2036	543.0	0.149	274.3
Number of observations	343	343	348	348

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Regressions in Panel A are run on pooled data from round 1 and round 2 (standard errors clustered at the individual level in parentheses); Panels B & C are run separately on round 1 and round 2, respectively (robust standard errors in parentheses) . USD 1 is ca. MK 400 for study period. All regressions include stratification cell fixed effects and an index of baseline asset ownership based on first principal components, difference in days between date of interview and the preceding weekend, baseline total spending and -if available- the baseline value of the outcome variable. For complete variable definitions, see Appendix B, and Table 2 for summary statistics.

times the price per share – was about MK 196 higher in the lump sum group, relative to a base of MK 172 in the weekly payment group. Both differences are statistically significant, at the 1 percent and 5 percent levels, respectively.

The results from Table 2.7 suggests that paying workers in a lump sum enabled them to hold enough cash to make use of a high-return lumpy investment opportunity, while the weekly group did not have sufficient extra cash holdings at the time opportunity was offered – despite experience with the product (from round 1) and sufficient advance notice.

In theory, the higher investment by the lump sum payment group could be driven by credit constraints alone, as opposed to savings constraints. Consider the case in which workers assigned to the lump sum payment group really wanted to smooth their consumption in the way the weekly payment group was able to, but could not due to a borrowing constraint. In that case, lump sum workers would "involuntarily" end up with more cash at the time the investment opportunity was offered and so they make use of it. While borrowing constraints are likely binding for many in the economic environment of this study, several arguments make this model an unlikely driver of our result: 72% of workers at baseline report preferring to be paid in a lump sum after four weeks as opposed to receiving four weekly installments (with the same twice-weekly attendance requirements in the hypothetical scenario that respondents were asked about as were imposed in this experiment). Of those 72%, a great majority (83%) state, in an open ended question with at most one answer, that the reason for this preference is that it enables them to a "make better plan" for the money. 13% outright list avoiding wasteful spending as the reason. These answers imply either a commitment problem as the reason for the lump sum preference, or at the least an expected inability to save – either due to internal constraints, such as self-control problems, or external constraints, such as fear of theft. Lastly, if lump sum payment group members truly preferred to smooth consumption in the way the weekly group was able to then they should not prefer to invest in the shares offered in this project as it locks up half (if they bought one share) or all (if they bought two shares) of total received wage payments for two weeks without any opportunity to access it. While in theory workers could have potentially borrowed against the future income receipt to access the money in the investment, this would also have held for the receipt of their wages, implying borrowing constraints could not be driving the results.

If lump sum condition households were limited in their ability to smooth consumption in the face of shocks then we would also expect that lump sum condition households would – relative to the weekly payment condition – receive more transfers from their social network over the course of the four payday weekends or request more loans – two of the most common risk coping mechanisms for workers of this study. However, we do not find statistically-significant effects on either of these outcomes; the point estimates are small, but the standard errors are large and so even sizeable effects along these dimensions cannot be ruled out (results not shown).

2.5 Discussion and Conclusion

Markets for financial intermediations in developing countries are imperfect. Besides the "external" constraints this creates for households, these market imperfections may exacerbate "internal" constraints such as time-inconsistent preferences and limited attention. In such a setting the exact timing of income streams can matter for spending and savings decisions. Spending may be higher, or skewed towards unplanned or wasteful expenditures in environments that are tempting, and spending may be different depending on the frequency of payments. If the timing of income receipt matters, this may have implications for the payment policies of employers and cash transfer programs, who may be interested in structuring payments to maximize benefits to income recipients.

In the specific context of this study, and in developing countries in general, there are two concerns about how wage payments are structured across time. First, when income is received in tempting environments, recipients may end up spending more, or may spend more on different items than they had planned ex ante, or than deemed prudent ex post. Second, when income is received in small installments, people may find it harder to generate meaningful sums that can be used for large-installment expenditures such as durable goods purchases, buying in bulk to receive quantity discounts, or high-return investments. In order to determine if these concerns are empirically relevant we designed a field experiment that varied the degree of temptation people faced when receiving payments, as well as whether payments were received in small installments or as a lump sum.

Based on ample qualitative evidence suggesting that spending – in particular frivolous spending – might be higher if income is received on market days, our experiment used the day of the week that workers were paid to vary the level of temptation workers faced when receiving income. Half of our sample received their income during the major local market day, which happened on Saturdays; the other half received their income at the same site on Fridays. However, we do not find evidence, for the sample of casual workers in Malawi that were part of our study, that the specific day of the receipt of income is an important driver of expenditures. Observed spending and savings behavior had no statistically-significant differences between those paid on Fridays and those paid on Saturdays, and we can rule out moderate-sized effects. This pattern does not depend on whether people are paid in a single lump sum or in small installments.

These findings do not reject the general idea that the environment in which people are paid matters. We worked in seven villages around one particular trading center in Malawi. In this setting, other trading centers with complementary market days – e.g. ones that take place on Fridays, when the payday trading center's market was not occurring – are within 30 minutes' travel. In other settings in which there are no complementary nearby market days, the day of payment may matter more. However, the setting of our study is fairly typical for many rural areas in Malawi and other countries in the region, where there are very often trading centers with a market day covering most days of the week, located within distances that can be traveled in reasonable times. Thus, the findings of our study should imply that the specific day of income receipt is not a major driver of spending decisions in a broad range of settings in rural Africa.

We also investigate the impact of paying workers in one lump sum compared to weekly payments. Our findings suggest that organizations can help income recipients overcome savings constraints by providing income in larger installments rather than smaller ones. Workers in the lump sum payment group spend relatively less of it immediately on receipt. Since they also receive more money on the last payday weekend – the full amount of wages compared to the weekly group that is receiving only the fourth of four equal installments – lump sum payment group members remain with more cash in the week after the last payday. In general, receiving income in a lump sum does not appear to affect the composition of expenditure, only the level. This mitigates concerns that lump sums "burn a hole in workers' pockets". Moreover, we find evidence that lump-sum income receipt promotes saving: people in the lump sum payment group show higher propensity to save in a high interest, relatively short-term asset that was offered to all respondents and required a large minimum investment. We argue that the differential investment is largely a function of the weekly payment group workers' inability to have cash available at the time of the investment offer (the timing of which was known to all workers before any payments were made).

The findings suggest that it is preferable for recipients that organizations pay at least part of wages or cash transfers in lump sums as a form of pre-committed savings. There is a trade-off between a desire to smooth consumption and the ability to generate lump sums; and so in an environment with borrowing constraints and generally high costs of risk coping, receiving all household income infrequently is unlikely to be desirable for households. In the context of this study, however, almost all households had some other means – besides the income from this project – of securing basic levels of consumption. Furthermore, a majority of households reports that they prefer to receive this additional income as a lump sum. This supports the idea that projects designed to generate income for people in developing countries, such as GiveDirectly, should provide income in strategically-timed lump sums (or at least offer this option) in order to maximize benefits to recipients.

The investment opportunity was artificially provided to study participants as part of this project in order to improve measurement of investment behavior in a small sample observed over a small time horizon and in a context where absolute income differentials across treatment groups were small. In addition, overall take-up of the investment opportunity was low. As such, the observed effects mainly support the overall conceptual point. However, the implied magnitudes are also interesting: we provided both the weekly and the lump group households with identical total additional income of MK 4000 (MK 3200 wages + 8 x MK 100 show up fees) over the course of the second round of this project. The point estimates imply that on average each member of the lump sum group was able to increase household income by an additional MK 65[12] – about 1.6% of income from the project's employment – within two weeks of the last payday via the investment opportunity, solely because of the changed timing of payments.

Practically speaking, the effect of changing the payments from small installments to lump sums will depend on the return to the relevant investment. We can get a sense of this by considering an example of an investment that is very similar to the one we offered: secondary school fees, which are approximately MK3000 ($7.50) per year in Malawi, and which generally must be paid in total at the beginning of the school year, rather than in installments. If people do think about education as an investment, we would expect that a project that pays respondents' total wages of MK3000 in a single lump sum timed for the beginning of the school year, rather than in small installments, to increase school fees payments by as much as 9 percentage points. This could have significant social benefits: if school fees are the only barrier to attending secondary school (and they are commonly cited as a reason teenagers do not go to school in Malawi) then that shift would have similar effects on the rate of school attendance. To get a sense of the total social benefit of this change in timing, note that Malawi has a GNI per capita of $320, and that research on the returns to education generally estimates figures of at least 10% per year in developing countries. Thus the additional 9% of children who are able to attend school would earn an additional $32 per year. Over a 40-year working life, starting 4 years after the investment, and at a social discount rate of 10%, this would raise a child's income by $213, for a net benefit of $206 per person. This is a substantial payoff for a relatively minor change.

School fees also highlight the external validity of our results for the investment good: they are time-sensitive, as are many other investment opportunities in the developing world, such as farm input purchases, which must be timed for the planting season.[13] This exacerbates

[12] 33% of 196.2, from Table 7, column 4.

[13] While some farm inputs can be bought and stored, others cannot for various reasons. For example, Malawi's government subsidizes fertilizer purchases immediately before the planting season, so farmers must have the cash to purchase the subsidized fertilizer within a fairly tight window.

the savings constraints that people face: it is easier to save up for an investment if you can make the purchase whenever you have the money, as opposed to needing to bring the money on a specific day. There are other important investments that do not have this same time-sensitive feature: for example, metal roofing has a large minimum installment size, but can be purchased whenever people have the money for it. Due to the design of the investment option used in this study, we cannot be sure that our results hold for alternative, less time-sensitive goods.

These benefits would come at relatively little cost, and organizing payroll just once a month, could even be cheaper for the paying organization. We also see no significant downsides to partial lump sums payments, even when they are received during one of the most tempting environments that people typically experience in rural Africa. However, further research is needed in order to better-establish whether lump-sum payments can potentially backfire in developing countries.

Our results provide several lessons for future research on lump sum payments as well as on the role of self-control problems in driving savings constraints. First, people are aware of the self-control problems they face, and thus survey questions that directly ask people about temptation and wasteful spending are a useful way to measure people's self-control issues. Second, offering study participants a meaningful investment opportunity that bears actual interest can be a helpful way to isolate an intervention's effects on savings constraints. Other outcomes have two kinds of limitations: non-financial investments such as health and education may not be perceived as investments by respondents, and heterogeneity in returns may generate misleading inferences about the extent of savings constraints. Third, to the extent that self-control problems are generating internal savings constraints in rural Africa, they may not be particularly amenable to policy interventions. Receiving one's pay during the market – a location commonly listed as being tempting by the respondents in our study – generated only small variations in their level of self-reported wasteful spending, possibly because people continue to select into other tempting situations. This suggests that other causes of savings constraints may merit further research.

CHAPTER III

Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi

From a work with Xavier Giné, Jessica Goldberg, and Dean Yang.

3.1 Introduction

Agriculture in Sub-Saharan Africa employs two-thirds of the labor force and generates about one-third of GDP growth. According to the 2008 World Development Report, GDP growth originating in agriculture is about four times more effective in reducing poverty than GDP growth originating outside agriculture. For this reason, policies that foster agricultural productivity can have a substantial impact on food security and poverty reduction.

In recent decades, there has been substantial interest among policy-makers, donors, and international development institutions in microfinance (financial services for the poor) as an anti-poverty intervention. Provision of microcredit has perhaps attracted the most attention. In 2009, the Microcredit Summit estimated that there were more than 3,500 microfinance institutions around the world with 150 million clients (Daley-Harris, 2009). While these outreach numbers are impressive, microcredit today is largely devoted to non-agricultural activities (Morduch, 1999; Armendáriz and Morduch, 2005) due to the substantial challenges inherent in agricultural lending.[1] Given the limited supply of credit for agriculture, many donors and academics (for example, Deaton, 1990; Robinson, 2001, and more recently the Bill and Melinda Gates Foundation) have emphasized the potential for increasing access to formal savings.[2]

[1] Giné, Goldberg and Yang (2012) find that imperfect personal identification leads to asymmetric information problems (both adverse selection and moral hazard) in the rural Malawian credit market.

[2] Aportela (1999) finds that a post-office savings expansion in Mexico raised savings by 3-5 percentage points. Burgess and Pande (2005) find that a policy-driven expansion of rural banking reduced poverty in India, and provide suggestive evidence that deposit mobilization and credit access were intermediating

The motivating question of this study is whether facilitating formal savings can promote agricultural development. To this end, we collaborated with a bank and private sector firms to implement a randomized controlled trial of a program facilitating formal savings for Malawian cash crop (tobacco) farmers. To our knowledge, this is the first randomized study of the agricultural impacts of an intervention facilitating savings in a formal banking institution.

In advance of the May-July 2009 harvest season, farmers were randomized into a control group or one of several treatment groups. Formal savings were facilitated for farmers in the treatment group by offering them the opportunity to have their cash-crop proceeds from the upcoming harvest channeled into bank accounts that would be opened for them, in their own names. Two main varieties of this treatment were implemented: 1) an "ordinary" savings treatment, where the bank accounts offered had no special features, and 2) a "commitment" savings treatment, in which farmers had the option of saving in special accounts that disallowed withdrawals until a set date (chosen by the account owner).[3]

Treated farmers were encouraged to use these accounts to save for future agricultural input purchases. Farmers in the control group, on the other hand, received no such facilitation of formal savings accounts, and were simply paid their crop sale proceeds in cash (which was the status quo).[4] We examine treatment impacts on savings at the partner bank (observed in administrative data) as well as on agricultural and other household outcomes (via a household survey).

The first key finding is that there are positive and statistically significant treatment effects on a range of outcomes. Facilitating formal savings leads to higher deposits into formal savings accounts at the partner bank, higher savings at the partner bank immediately prior to the next planting season (November-December 2009), higher agricultural input expenditures in that season, higher output in the subsequent harvest (May-July 2010), and higher per capita consumption in the household after that harvest. Impacts on agricultural input expenditures and on output are substantial, amounting to increases over the control group mean of 13.3% and 21.4% respectively.

The second key finding is somewhat unexpected, and has to do with the mechanism

channels. Bruhn and McKenzie (2009) find that bank branch openings by consumer durable stores in Mexico leads to increases in the number of informal business owners, in total employment, and in average income.

[3]In addition, these treatments were cross-randomized with another treatment intended to create variation in the public observability of savings balances (details are explained in Section 3.2). In total, there were six different randomly-assigned treatment types. Differences in impacts across treatments are typically not statistically significantly different from one another, so we place little emphasis on differentiating impacts across treatment types in this paper.

[4]Control group farmers also received a generic encouragement to save for future agricultural input purchases, so as to distinguish this effect from the effect of formal savings facilitation.

through which treatment translates into agricultural outcomes. *Ex ante*, the leading candidate mechanism was the alleviation of savings constraints. In the status quo, farmers have imperfect means of preserving funds between harvest and the subsequent planting season. Depletion of funds not held in bank accounts over this period could be due to self-control problems, demands for sharing with one's social network, and losses due to other factors (e.g., theft, fire). Improving access to formal savings would therefore give farmers a better means of preserving funds between harvest and the subsequent planting, leading to increases in agricultural input expenditures (and then to improvements on other subsequent related outcomes).

Our results indicate, however, that only a minority of the treatment effect on agricultural input expenditures is likely to be attributable to alleviating formal savings constraints. While amounts initially deposited into the accounts would have been sufficient to pay for the increase in agricultural input expenditures that we observe, administrative data from the bank reveals that the majority of these funds were withdrawn almost immediately after being deposited. Three months later, just prior to the end-of-2009 planting season treated farmers still had 1,863 Malawi kwacha (US$12.85) higher savings than did control-group farmers, but the treatment effect on agricultural input expenditures is higher by a factor of four: MK 8,023 (US$55.33). Therefore, only about a quarter of the effect of the treatment on agricultural input expenditures can be attributed to alleviation of savings constraints *per se*.[5]

We discuss a variety of mechanisms for which we are able to provide incomplete evidence as well other mechanisms that can be ruled out. In the end, with the design implemented and data available we are not able to identify the precise mechanisms through which our treatment effects operated. For example, the funds held in accounts may have served as a buffer stock, allowing farmers to self-insure and take on more risk (by investing more in agricultural inputs.) Alternately, the existence of the accounts could have helped study participants resist demands to share resources with their social network. Behavioral phenomena such as mental accounting or reference-dependence also provide possible explanations. We must leave exploration of these and other possible mechanisms to future work.

The remainder of this paper is organized as follows. The next section describes the experimental design and data sources. Section 3.3 describes our empirical specification. Section 3.4 presents the treatment effect estimates. Section 3.5 then considers evidence on the mechanisms through which the treatment effects may have operated. Section 3.6 concludes.

[5]Dupas and Robinson (2013*a*) conduct a randomized experiment of a savings intervention in a sample of Kenyan microentrepreneurs, and interpret impacts as due to this mechanism.

3.2 Experimental design and survey data

The experiment was a collaborative effort between Opportunity Bank of Malawi (OBM),[6] Alliance One, Limbe Leaf, the University of Michigan and the World Bank. Opportunity International is a private microfinance institution operating in 24 countries that offers savings and credit products; in Malawi, it has a full banking license that allows it to collect deposits and on-lend funds. Alliance One and Limbe Leaf are two large private agri-business companies that offer extension services and high-quality inputs to smallholder farmers via an out-grower tobacco scheme.[7] These two companies work with smallholder out-growers by organizing them geographically into clubs of 10-20 members who obtain tobacco production loans under group liability from OBM.[8] Tobacco clubs meet regularly and sell their crop output collectively to the tobacco auction floor. In the central Malawi region we study, tobacco farmers have similar poverty and income levels to those of non-tobacco-producing households.[9]

While all farmers in the study were loan customers of OBM at the start of the project, the loans provided a fixed input package that for the majority of farmers fell short of optimal levels of fertilizer use on their tobacco plots.[10] This is important because it suggests that there is room for savings to increase input utilization. In addition, while a minority of farmers was using optimal levels of fertilizer for the amount of land they were cultivating at baseline, even those farmers could use savings generated by the intervention to obtain additional inputs and expand land under tobacco cultivation, or shift land from other crops

[6] At the time of th e study, our bank partner went by the company name Opportunity International Bank of Malawi (OBM), but has since changed its name to Opportunity Bank of Malawi (OBM).

[7] Tobacco is central to the Malawian economy, as it is the country's main cash crop. About 70% of the country's foreign exchange earnings come from tobacco sales, and a large share of the labor force works in tobacco and related industries.

[8] The cost of an input loan includes an interest rate of 28% percent per year and a one-time 2.5% processing fee.

[9] Based on authors' calculations from the 2004 Malawi Integrated Household Survey (IHS), individuals in tobacco farming rural households in central Malawi live on PPP$1.46/day on average, while the corresponding average for non-tobacco farmers is PPP$1.51/day. That said, the two groups are different in other ways. Tobacco farmers have somewhat larger households (6.68 persons compared to 4.94 persons for households not farming tobacco), higher levels of education of the household head (5.61 years compared to 4.63 years) and a higher share of school age kids (6-17 years) currently in school conditional on having school age children (88.1% compared to 77.9%).

[10] The input package was designed for a smaller cultivated area. As a result, 60.4% of farmers were applying less than the recommended amount of nitrogen on their tobacco plots at baseline. The figures for the two other key nutrients for tobacco are even more striking: 83.2% and 84.7% of farmers used less than the recommended amount of phosphorus and potassium, respectively. For each of the three nutrients, among farmers using less than recommended levels, the mean ratio of actual use to optimal use was about 0.7. Optimal use levels were determined by Alliance One and Limbe Leaf in collaboration with Malawi's Agricultural Research and Extension Trust (ARET), and are similar to nutrient level recommendations in the United States (Pearce and Denton, 2011).

towards tobacco. Finally, the savings intervention could also affect use of fertilizer and other inputs on maize (the main staple crop in Malawi) and other crops.[11]

The experiment was designed to test the impact of facilitating savings in formal bank accounts. In addition, we sought to test whether offering accounts with "commitment" features would have a greater impact than offering "ordinary" bank accounts without such features.[12] Farmer clubs were randomly assigned to either a control group offered no savings facilitation, an "ordinary savings" treatment group that was offered assistance setting up direct deposit into individual, liquid savings accounts, and a "commitment savings" treatment group that was offered assistance setting up direct deposit into individual ordinary savings accounts *and* additional accounts with commitment features.

The design of the experiment also aimed to explore the role of savings accounts in helping farmers resist pressure to share resources with others in their social network. Farmer clubs in the ordinary and commitment savings treatment groups were further cross-randomized into sub-groups that were or were not entered into a raffle wherein they could win prizes based on their account balances (described further below).

In sum, the two cross-cutting interventions result in seven treatment conditions: a pure control condition without savings account offers or raffles; ordinary savings accounts with no raffles, with private distribution of raffle tickets, and with public distribution of raffle tickets; and commitment savings accounts with no raffles, with private distribution of raffle tickets, and with public distribution of raffle tickets (see Table 3.2).

Figure 3.1 presents the timing of the experiment with reference to the Malawian agricultural season. The baseline survey and interventions were administered in April and May 2009, immediately before the 2009 harvest. As a result, farmers in the commitment treatment group made allocation decisions into the commitment and ordinary accounts in the "cold state" prior to receiving the net proceeds from tobacco sales.[13] Planting starts between

[11] At baseline, 89.5% and 99.9% of farmers were applying less than the recommended amount of nitrogen and phosphorus, respectively, on their maize plots and 44.1% and 98.6% of farmers applied less than half the recommended amounts for the two nutrients. Among farmers applying less than the recommended amount of nitrogen (phosphorus) on maize, the ratio of actual use to optimal use was 0.48 (0.14). Potassium is not recommended for maize cultivated in central Malawi. Nutrient recommendations are from Benson (1999).

[12] Research on savings accounts with features that self-aware individuals can use to limit their options in anticipation of future self-control problems includes Ashraf, Karlan and Yin (2006), who investigate demand for and impacts of a commitment savings device in the Philippines and find that demand for such commitment devices is concentrated among women exhibiting present-biased time preferences. Duflo, Kremer and Robinson (2011) find that offering a small, time-limited discount on fertilizer immediately after harvest has an effect on fertilizer use that is comparable to that of much larger discounts offered later, around planting time. Giné, Goldberg and Yang (2012) find that Malawian farmers with present-biased preferences are more likely to revise a plan about how to use future income, a result that supports the potential of commitment accounts to improve welfare for those with self-control problems.

[13] If decisions had been made the day that tobacco sales were transferred to OBM then the allocations into the commitment accounts by present-biased individuals would have been lower.

November and December depending on the arrival of the rains. We will therefore refer to the time from harvest until end October as the pre-planting period.

Randomization of the savings and raffle treatments was conducted at the club level in order to minimize cross-treatment contamination.[14] The sample consists of 299 clubs with 3,150 farmers surveyed at baseline (February-April 2009), for whom we can track savings deposits, withdrawals, and balances in our partner bank's administrative data. In addition, we have data from an endline survey administered in July-September 2010, after the 2010 harvest, for 2,835 farmers from 298 clubs. Attrition from the baseline to the endline survey was 10.0% and is not statistically significantly different across different treatment groups (as shown in Online Appendix Table 1). The endline survey will be used to examine impacts on outcomes such as farm inputs, production, and household per capita expenditures.

3.2.1 Financial education

Members of all clubs attended a financial education session immediately after the baseline survey was administered. The session reviewed basic elements of budgeting and explained the benefits of formal savings accounts, with an emphasis on how such accounts could be used to set aside funds for future consumption and investment. The full script of the financial education session can be found in Appendix A.

The same financial education session was deliberately provided to all clubs - including those subsequently assigned to the control group - so that treatment effects could be attributed solely to the provision of the financial products, abstracting from the effects of financial education that are implicitly provided during the product offer (for example, strategies for improved budgeting). For this reason, we can estimate neither the impact of the ordinary and commitment treatments without such financial education, nor the impact of the financial education alone.

3.2.2 Savings treatments

Implementation of the savings treatments took advantage of the existing system of depositing crop sale proceeds into OBM bank accounts. At harvest, farmers sold their tobacco to the company at the price prevailing on the nearest tobacco auction floor.[15] For farmers in the control group, the proceeds from the sale were then electronically transferred to

[14] Prior to randomization, treatment clubs were stratified by location, tobacco type (burley, flue-cured or dark-fire) and week of scheduled interview. The stratification of treatment assignment resulted in 19 distinct location/tobacco-type/week stratification cells.

[15] The tobacco growing regions are divided among the two tobacco buyer companies. In their coverage area each buyer company organizes farmers into clubs and provides them with basic extension services.

Figure 3.1: Project Timing

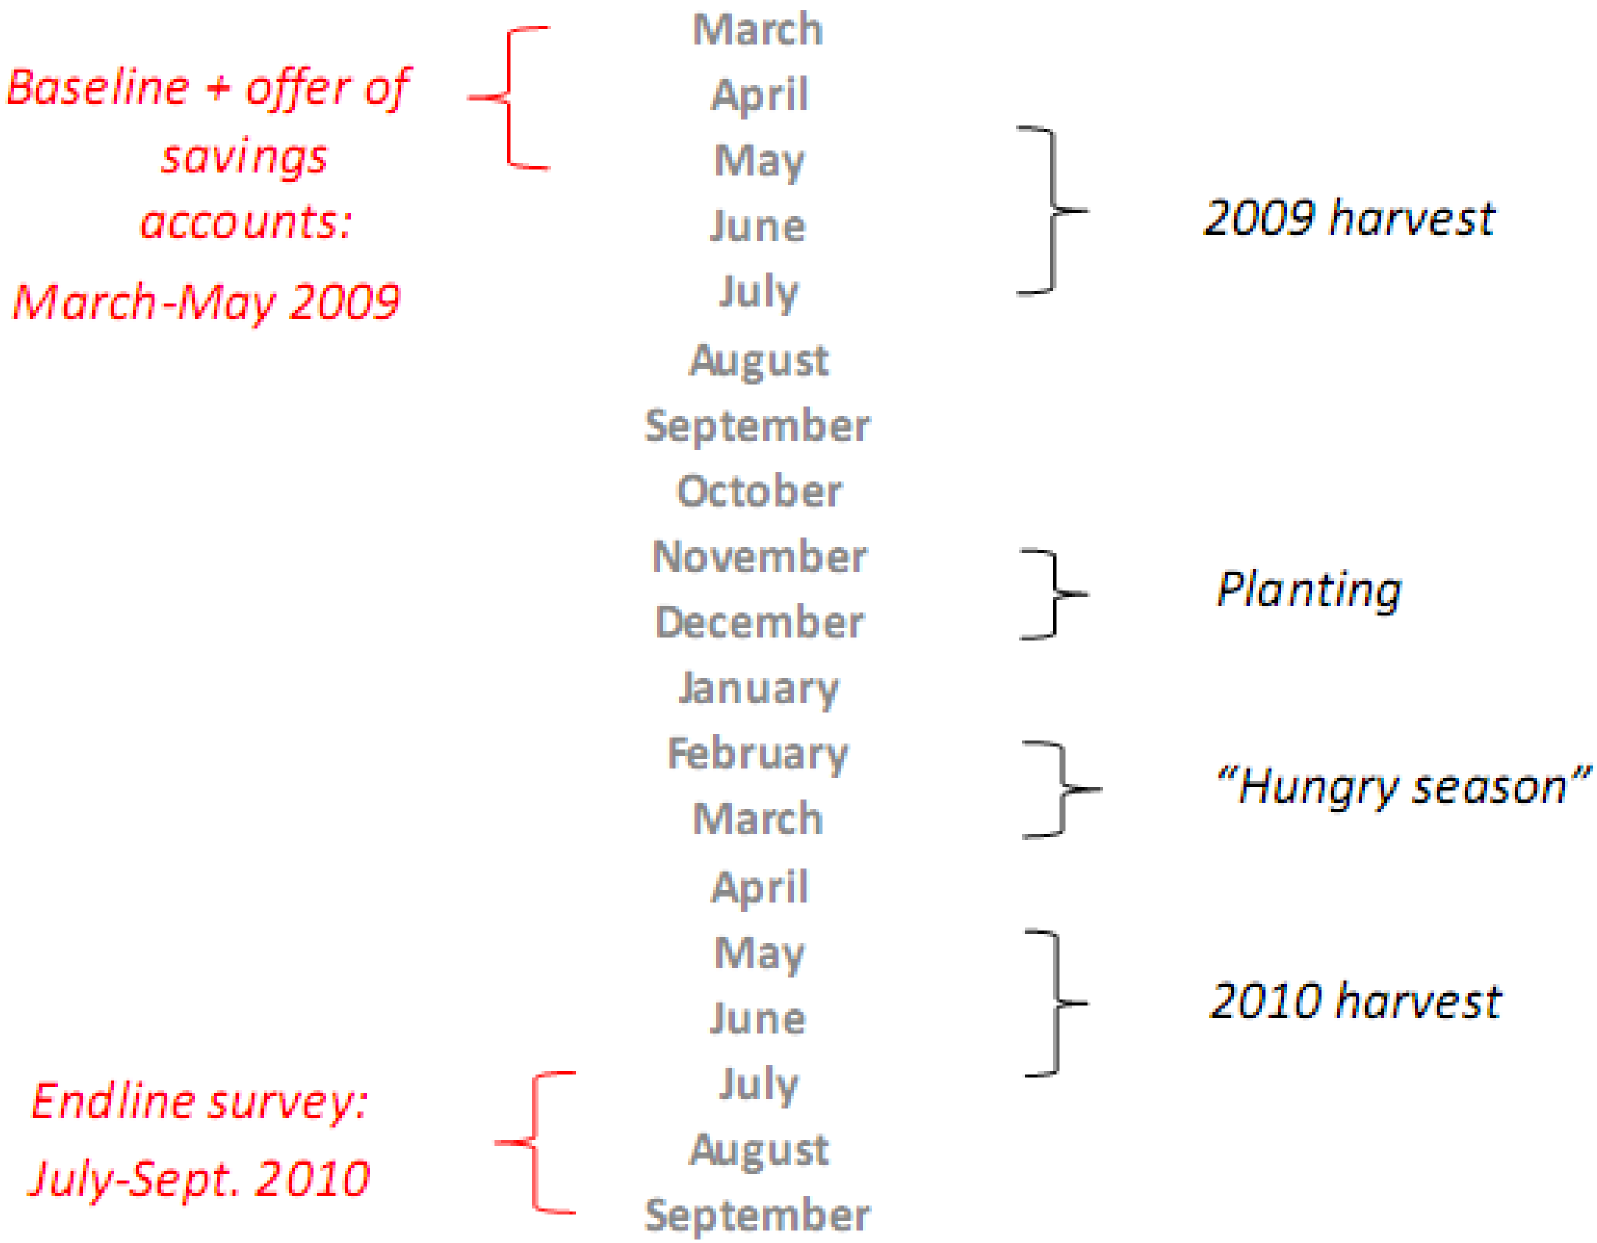

OBM, which deducted the loan repayment (plus fees and surcharges) of all borrowers in the club, and then credited the remaining balance to a club account at OBM. Club members authorized to access the club account (usually the chairman or the treasurer) came to OBM branches and withdrew the funds in cash.

Farmers in the ordinary savings treatment were offered account opening assistance and the opportunity to have their harvest proceeds (net of loan repayment) directly deposited into individual accounts in their own individual names (see Figure 3.2 for a schematic illustration of the money flows). These ordinary savings accounts are regular OBM savings accounts with an annual interest rate of 2.5%. After their crop was sold, farmers traveled to the closest OBM branch to confirm that funds were available at the club level, i.e. that club proceeds exceeded the club's loan obligation. Authorized members of the clubs (often accompanied by other club members) then filled out a sheet specifying the division of the balance of the club account between farmers. Funds were transferred into the individual accounts of club members who had opted to open them. Other club members received their share of the money in cash.

Farmers in clubs assigned to the ordinary savings treatment were offered only one (ordinary) savings account. Farmers assigned to the commitment treatment had the option of opening an additional account with commitment features. The commitment savings account had the same interest rate as the ordinary account, but allowed farmers to specify an amount to be transferred to this illiquid account, and a "release date" when the bank would allow access to the funds.[16] During the account opening process, farmers stated how much they wanted deposited in the ordinary and commitment savings accounts after the sale of their tobacco crops. For example, if a farmer stated that that he wanted MK 40,000 in an ordinary account and MK 25,000 in a commitment savings account, funds would first be deposited into the ordinary account until MK 40,000 had been deposited, then into the commitment savings account for up to MK 25,000, with any remainder being deposited back into the ordinary account. The choice of a "trigger amount" that had to flow into the ordinary account before any money would be deposited into the commitment account turns out to be important, because many farmers chose triggers higher than their eventual crop sale revenue, and therefore ended up without deposits into their commitment accounts. Opening only a commitment account was not an option, though farmers who wanted to deposit all of their proceeds into the commitment account and none into their ordinary account were free to do so. No fees were charged for the initial post-crop-sale deposits into the ordinary or com-

[16] By design, funds in the commitment account could not be accessed before the release date. In a small number of cases OBM staff allowed early withdrawals of funds when clients presented evidence of emergency needs, e.g. health or funeral expenditures.

Figure 3.2: Tobacco Sales and Bank Transactions

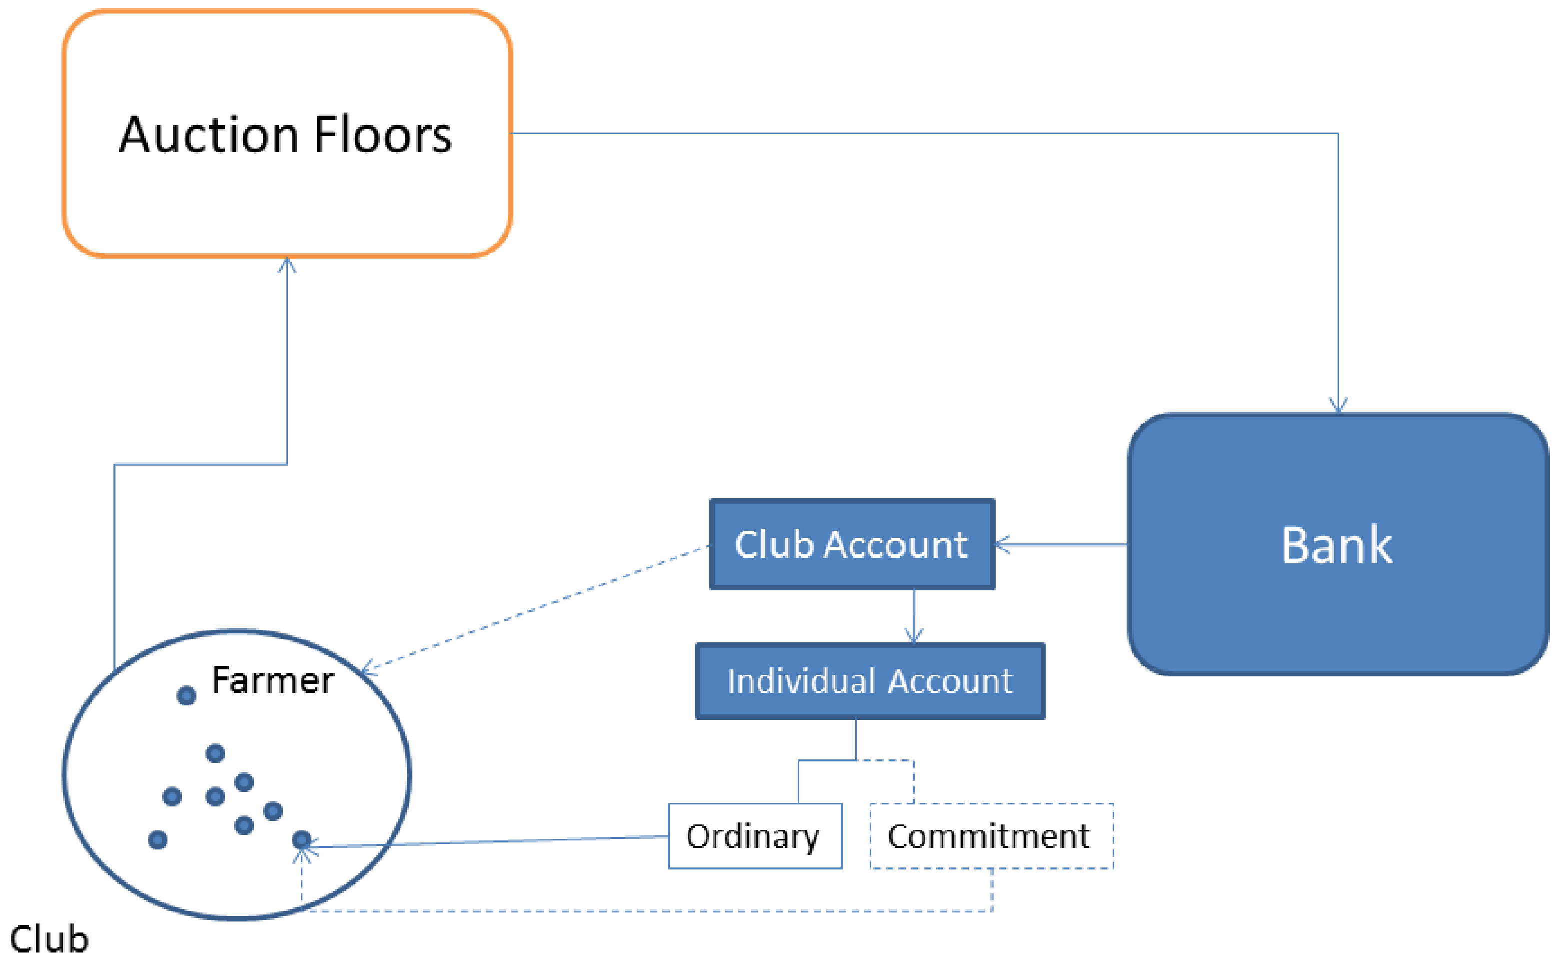

mitment accounts. Further details on account features and fees can be found in Appendix A.

Farmers who were not offered a particular account type due to their treatment status (e.g., control group farmers who were not offered either type of account, or ordinary treatment group farmers who were not offered the commitment account) but learned about and requested them were not denied those accounts, but they were not given information about or assistance in opening them.[17] In other words, the savings treatments were implemented as an encouragement design.

3.2.3 Raffle Treatments

To study the impact of public information on savings and investment behavior, we implemented a cross-cutting randomization of a savings-linked raffle. Participants in each of the two savings treatments were randomly assigned to one of three raffle conditions (members of the control group were not eligible for raffle tickets, because the tickets were based on savings account balances).

We distributed tickets for a raffle to win a bicycle or a bag of fertilizer (one of each per participating branch), where the number of tickets each participant received was determined by his or her savings balance as of pre-announced dates that fell before large expenditures (like fertilizer purchases) were likely to deplete savings balances. Every MK 1,000 in an OBM account (in total across ordinary and commitment savings accounts) entitled a participant to one raffle ticket. Ticket allocations would be on the basis of average balances from July 1 to August 1 (first distribution) and from September 1 to October 1 (second distribution). By varying the way in which tickets were distributed, we sought to exogenously vary the information that club members had about each other's savings balances.

Because the raffle itself could provide an incentive to save or could serve as a reminder to save (Karlan, McConnell, Mullainathan, Zinman, 2012; Kast, Meier and Pomeranz, 2012), one third of clubs assigned to either ordinary or commitment savings accounts was randomly determined to be ineligible to receive raffle tickets (and was not told about the raffle). Another one third of clubs with savings accounts was randomly selected to have raffle tickets distributed privately. Study participants were called to a meeting for raffle ticket distribution but were handed their tickets out of view of other study participants. The final third of clubs

[17] During the baseline interaction with study participants, no farmers in the control group expressed to our survey staff a desire for either ordinary or commitment accounts, and none in the ordinary treatment group requested commitment accounts. According to OBM administrative records, seven individuals in the control group (1.7%) and 52 farmers in the ordinary treatment group (3.7%) had commitment accounts by the end of October 2009 (these were opened without our assistance or encouragement). None of these farmers had any transactions in the accounts.

with savings accounts was randomly selected for public distribution of raffle tickets. In these clubs, each participant's name and the number of tickets received was announced verbally to everyone that attended the raffle meeting.

A feature of the simple formula for determining the number of tickets was that farmers in clubs where tickets were distributed publicly could easily estimate other members' savings balances. Private distribution of tickets, though, did not reveal information about individuals' account balances. The raffle scheme was explained to participants at the time of the baseline survey with a participatory demonstration. Members were first given hypothetical balances, and then given raffle tickets in a manner that corresponded to the distribution mechanism for the treatment condition to which the club was assigned. In clubs assigned to private distribution, members were called up one by one and given tickets in private (out of sight of other club members). In clubs assigned to public distribution, members were called up and their number of tickets was announced to the group. Since real tickets based on actual account balances were distributed twice during the experiment, the first distribution also functioned as an additional demonstration. (As reported in Section 3.4 below, however, substantial withdrawals from both the ordinary and commitment accounts occurred soon after funds were deposited, and as a result, this public revelation treatment was likely to have had little effect.)

3.2.4 Sample

Table 3.1 presents summary statistics of baseline household and farmer club characteristics. All variables expressed in money terms are in Malawi Kwacha (MK145/USD during the study period). Baseline survey respondents own an average of 4.7 acres of land and are mostly male (only six percent were female). Respondents are on average 45 years old. They have an average of 5.5 years of formal education, and have low levels of financial literacy.[18] Sixty three percent of farmers at baseline had an account with a formal bank (mostly with OBM).[19] The average reported savings balance in bank accounts at the time of the baseline was MK 2,083 (USD 14), with an additional MK 1,244 (USD 9) saved in the form of cash at home.

[18] In particular, 42% of respondents were able to compute 10% of 10,000, 63% were able to divide MK 20,000 by five and only 27% could apply a yearly interest rate of 10% to an initial balance to compute the total savings balance after a year.

[19] This number includes a number of "payroll" accounts opened in a previous season by OBM and one of the tobacco buyer companies as a payment system for crop proceeds, and which do not actually allow for savings accumulation. Our baseline survey unfortunately did not properly distinguish between these two types of accounts.

Table 3.1: Summary Statistics

	Mean	Standard Deviation	10th Percentile	Median	90th Percentile	Observations
Treatment conditions						
Control group	0.135	0.341				3,150
Panel A						
Savings	0.865	0.341				3,150
Panel B						
Commitment Savings	0.417	0.493				3,150
Ordinary Savings	0.448	0.497				3,150
Panel C						
Commitment, no raffle	0.136	0.342				3,150
Commitment, priv. raffle	0.142	0.349				3,150
Commitment, pub. raffle	0.139	0.346				3,150
Ordinary, no raffle	0.146	0.354				3,150
Ordinary, priv. raffle	0.149	0.356				3,150
Ordinary, pub. raffle	0.153	0.360				3,150
Baseline Characteristics						
Number of members per club	13.88	6.44	9.00	11.00	23.00	299
Female	0.063	0.243				3,150
Married	0.955	0.208				3,150
Age [years]	45.02	13.61	28.00	44.00	64.00	3,150
Years of education	5.45	3.53	0.00	6.00	10.00	3,150
Household size	5.79	1.99	3.00	6.00	9.00	3,150
Asset index	-0.02	1.86	-1.59	-0.67	2.46	3,150
Livestock index	-0.03	1.15	-1.00	-0.36	1.37	3,150
Land under cultivation [acres]	4.67	2.14	2.50	4.03	7.50	3,150
Cash spent on inputs [MK]	25,169	41,228	0	10,000	64,500	3,150
Proceeds from crop sales [MK]	125,657	174,977	7,000	67,000	300,000	3,150
Has bank account	0.634	0.482				3,150
Savings in cash at home [MK]	1,244	3,895	0	0	3,000	3,150
Savings in bank accounts [MK]	2,083	8,265	0	0	3,000	2,949
Hyperbolic	0.102	0.303				3,117
Patient now, impatient later	0.304	0.460				3,117
Net transfers made in past 12m [MK]	1,753	7,645	-2,990	500	8,100	3,150
Missing value for formal savings and cash	0.064	0.244				3,150
Missing value for time preferences	0.010	0.102				3,150
Transactions with Partner Institution						
Any transfer via direct deposit	0.154	0.361				3,150
Deposits into ordinary accounts, pre-planting [MK]	18,472	82,396	0	0	38,907	3,150
Deposits into commitment accounts, pre-planting [MK]	615	5,367	0	0	0	3,150
Deposits into other accounts, pre-planting [MK]	296	3,804	0	0	0	3,150
Total deposits into accounts, pre-planting [MK]	19,383	84,483	0	0	40,694	3,150
Total withdrawals from accounts, pre-planting [MK]	18,600	82,744	38,600	0	0	3,150
Net of all transactions, pre-planting [MK]	762	13,857	0	0	649	3,150
Net of all transactions, Nov-Dec [MK]	-848	6,870	0	0	2	3,150
Net of all transactions, Jan-Apr [MK]	-269	4,032	0	0	4	3,150
Endline Survey Outcomes						
Land under cultivation [acres]	4.52	2.66	2.00	4.00	8.00	2,835
Cash spent on inputs [MK]	21,632	32,853	500	11,000	51,500	2,835
Total value of inputs [MK]	68,046	84,014	1,500	43,750	157,272	2,835
Proceeds from crop sales [MK]	109,604	162,580	0	56,000	270,000	2,835
Value of crop output (sold & not sold) [MK]	177,747	201,131	27,480	115,582	387,203	2,835
Farm profit (output-intput) [MK]	110,703	156,747	0	70,372	264,953	2,835
Total expenditure in last 30 days [MK]	11,905	13,219	2,250	7,500	26,000	2,835
Household size	5.80	2.15	3.00	6.00	9.00	2,835
Total transfers made [MK]	3,152	5,099	0	1,300	8,000	2,835
Total transfers received [MK]	2,204	4,377	0	500	6,050	2,835
Total net transfers made [MK]	939	5,896	-3,000	350	5,750	2,835
Tobacco loan amount [MK]	40,787	77,962	0	0	130,000	2,835
Has fixed deposit account	0.067	0.250				2,835
Not interviewed in follow-up	0.100	0.300				3,150

Data based on two surveys conducted in February to April 2009 (baseline) and July to September 2010 (endline), and on administrative records of our partner institution. "MK" is Malawi kwacha (MK145 = US$1 during study period). Withdrawals presented as negative numbers. See Appendix B for variable definitions.

Table 3.2: Assignment of clubs to treatment conditions

	No savings intervention	Savings intervention: ordinary accounts offered	Savings intervention: ordinary and commitment accounts offered
No raffle	Group 0: 42 clubs	Group 1: 43 clubs	Group 4: 42 clubs
Public distribution of raffle tickets	N/A	Group 2: 44 clubs	Group 5: 43 clubs
Private distribution of raffle tickets	N/A	Group 3: 43 clubs	Group 6: 42 clubs

3.2.5 Balance of baseline characteristics across treatment conditions

To examine whether randomization across treatments achieved balance in pre-treatment characteristics, Table 3.3 presents the differences in means of 17 baseline variables in the same format as used for the subsequent analysis. Panel A checks for balance between the control group and the treatment group, the latter pooled across all of the savings and raffle treatments. Panel B looks for differences between the control group, the ordinary savings group, and the commitment savings group, with each of the savings treatments pooled across their respective raffle sub-treatments.

With a few exceptions, the sample is well balanced. We test balance for 17 baseline variables. In Panel A, respondents assigned to the savings treatment are four percentage points more likely to be female and two percentage points less likely to be married than those assigned to the control group. At baseline, they report spending nearly MK 4,000 more in cash on agricultural inputs, a difference that is statistically significant at the 90 percent confidence level.

Panel B reveals that respondents in both the commitment and ordinary treatment groups are more likely to be female and less likely to be married. The treatment-related imbalance with respect to cash spent on inputs found in Panel A appears to be driven by imbalance in the ordinary treatment group, which is different from the control group at the 5% level (the difference between the commitment treatment group and the control group for that variable is not statistically significant at conventional levels.) Those in the commitment treatment group are also less likely to be patient now and impatient later, compared to the control group (significant at the 5% level).

To alleviate any concerns that baseline imbalance may be driving our results, we follow Bruhn and McKenzie (2009) and include the full set of baseline characteristics in Table 3.3 as controls in the main regressions, in addition to the stratification cell fixed effects.

3.3 Empirical specification

We study the effects of our experimental interventions on several sets of outcomes: deposits into and withdrawals from savings accounts, savings balances, agricultural outcomes from the next year's growing season and household expenditure following that season, households' financial interactions with others in their network, and future use of financial products. These data come from the endline survey administered after the 2010 harvest, and from administrative data on bank transactions and account balances collected throughout the project.

We present two regression specifications reported as separate panels in the main results

Table 3.3: Test of Balance in Baseline Characteristics (ordinary least-squares regressions)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)
<u>Dependent variable:</u>	Female	Married	Age [years]	Years of edu-cation	House-hold size	Asset index	Live-stock index	Land under culti-vation [acres]	Pro-ceeds from crop sales [MK]	Cash spent on inputs [MK]	Has bank account	Savings in accounts and cash [MK]	Hyper-bolic	Patient now, im-patient later	Net transfers made in past 12m [MK]	Missing val.: formal savings and cash	Missing val.: time prefe-rences
PANEL A																	
Any treatment	0.044***	-0.018**	-1.42	0.14	-0.03	0.08	-0.07	-0.01	6,997	3,918*	-0.021	371	0.012	-0.054	72	-0.002	0.001
	(0.012)	(0.009)	(0.93)	(0.20)	(0.13)	(0.11)	(0.09)	(0.14)	(8,891)	(2,027)	(0.029)	(550)	(0.017)	(0.034)	(452)	(0.013)	(0.005)
<u>P-values of F-tests for joint significance of baseline variables:</u>					0.1481												
PANEL B																	
Commitment treatment	0.045***	-0.019*	-1.39	0.09	-0.04	0.07	-0.06	-0.05	5,604	3,337	-0.039	376	0.024	-0.076**	-195	-0.004	0.003
	(0.013)	(0.010)	(0.97)	(0.22)	(0.13)	(0.12)	(0.09)	(0.15)	(9,779)	(2,357)	(0.032)	(612)	(0.019)	(0.036)	(476)	(0.014)	(0.005)
Ordinary treatment	0.042***	-0.018*	-1.45	0.19	-0.02	0.09	-0.07	0.02	8,294	4,459**	-0.005	367	0.000	-0.034	320	0.000	0.000
	(0.013)	(0.010)	(0.98)	(0.22)	(0.13)	(0.12)	(0.09)	(0.15)	(9,639)	(2,209)	(0.031)	(588)	(0.018)	(0.037)	(475)	(0.015)	(0.005)
<u>P-val. of F-test</u>: Coefficients on commitment and ordinary treatments are equal	0.790	0.912	0.924	0.557	0.857	0.825	0.936	0.549	0.731	0.592	0.219	0.985	0.083	0.110	0.094	0.730	0.661
<u>P-values of F-tests for joint significance of baseline variables:</u>																	
Commitment savings					0.6168												
Ordinary savings					0.8851												
<u>Mean dep. var. in Control group</u>	0.024	0.972	46.23	5.31	5.81	-0.11	0.03	4.67	117,495	21,798	0.658	3,235	0.095	0.352	1,655	0.066	0.009
<u>Number of observations</u>	3,150	3,150	3,150	3,150	3,150	3,150	3,150	3,150	3,150	3,150	3,150	2,949	3,117	3,117	3,150	3,150	3,150

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the club level. USD 1 is ca. MK 145. All regressions include stratification-cell fixed effects. For variable definitions, see Appendix B. F-test of Panel B: "Commitment savings = Ordinary savings" tests the equality of means in commitment and ordinary treatment groups. F-tests of joint significance: test of joint significance in regression of respective treatment dummies on all 17 baseline variables.

tables. The first tests the effect of being randomly assigned to any of the savings facilitation treatments, relative to being assigned to the control group. In Panel A of the subsequent tables, we run regressions of the form

$$Y_{ij} = \delta + \alpha Savings_j + \beta' X_{ij} + \varepsilon_{ij}. \tag{3.1}$$

Y_{ij} is the dependent variable of interest for farmer i in club j. $Savings_j$ is an indicator variable for club-level assignment to either of the two savings treatment groups. The coefficient α measures the effect of being offered direct deposit into an individual savings account (either ordinary savings accounts only or ordinary plus commitment accounts). X_{ij} is a vector that includes stratification cell dummies and the 17 household characteristics measured in the baseline survey prior to treatment, and summarized in Table 3.3, and ε_{ij} is a mean-zero error term. Because the unit of randomization is the club, standard errors are clustered at this level (Moulton, 1986).

In Panel B, we compare the impact of assignment to the ordinary savings treatment to the impact of assignment to the commitment savings treatment. Regressions are of the form

$$Y_{ij} = \delta + \gamma_1 Ordinary_j + \gamma_1 Commitment_j + \beta' X_{ij} + \varepsilon_{ij}, \tag{3.2}$$

where Y_{ij} and X_{ij} are defined as above. $Ordinary_j$ is an indicator for club-level assignment to the ordinary savings treatment, and $Commitment_j$ is an indicator for assignment to the commitment savings treatment. The coefficient $\gamma_1$represents the effect of eligibility for direct deposit into ordinary accounts only, relative to the control group. γ_2 captures the analogous effect for eligibility for direct deposit into ordinary accounts *and* automatic transfers into commitment savings accounts. The difference between those two coefficients, then, captures the marginal effect of the commitment savings account relative to direct deposit into the ordinary account. The p-value for the test of the null hypothesis that $\gamma_1 = \gamma_2$ is reported at the bottom of each Panel B.

Both regression equations (3.1) and (3.2) measure treatment effects that pool the raffle sub-treatments. Results with full detail on the raffle sub-treatments (six treatments in all) are presented in Online Appendix Tables 3-6.

Throughout the analysis, we focus on intent-to-treat (ITT) estimates because not every club member offered account opening assistance decided to open an account. We do not report average treatment on the treated (TOT) estimates because it is plausible that members without accounts are influenced by the training script itself or by members who do open accounts in the same club, either of which would violate the stable unit treatment value assumption (Angrist, Imbens and Rubin, 1996, SUTVA).

3.4 Empirical results

We first examine the effects of our experimental interventions on formal savings: the flow of funds into and out of accounts, and savings account balances. In and of themselves, however, these are not measures of wellbeing. Therefore, we also analyze impacts on agricultural input use, farm output, household expenditures, and other household behaviors.

3.4.1 Take-up and impacts on savings transactions

The first question of interest is whether the experimental treatments changed use of individual savings accounts. Table 3.4 presents estimates of equations (3.1) and (3.2) (in Panels A and B, respectively) for outcomes from administrative data on account transactions.

Column 1 presents treatment effects on "take-up" of the offered financial services: opening of individual bank accounts coupled with direct deposit of tobacco crop proceeds.[20] Panel A indicates that take-up was 19.4% among respondents offered any treatment (this dependent variable is zero by design in the control group). Take-up is very similar across the commitment and ordinary treatments (Panel B), and statistically indistinguishable across them (the p-value of the difference in take-up across the two groups is 0.432.)

Owing to the study's aim to promote agricultural input investments in the Nov-Dec 2009 planting season, for the remaining dependent variables in Table 3.4, we examine transactions over the months preceding that period, March through October 2009. In column 2, the dependent variable is total deposits into all accounts at the partner bank (these are direct deposits from the tobacco companies as well as other deposits made by account holders). The mean of this variable in the control group is MK 3,281 (USD 21.72). Compared to this amount, the impact of being assigned to any treatment group shown in Panel A is large (MK 17,609, or USD 121.44) and statistically significantly different from zero at the 1% level. Given that take-up was very similar across the two treatment groups, and that take-up by design meant that all crop proceeds were deposited with the partner bank, it should not be surprising that the treatment effect is very similar across commitment and ordinary treatment groups (Panel B). Each separate treatment effect is statistically significantly different from zero at the 1% level, but the treatment effects are not statistically significantly different from one another (p-value 0.642).

The next three columns provide more detail on the types of account into which deposits were destined, examining treatment effects on deposits into ordinary accounts, commitment accounts, and "other" accounts that study participants might have held at the partner bank

[20] The time period over which this dependent variable is calculated is intentionally very broad (Mar 2009 to Apr 2010), so as to capture any direct deposit from the tobacco purchase companies into the study respondent accounts. In practice the vast majority of direct deposits took place in the May-July 2009 harvest season.

Table 3.4: Impact of Treatments on Deposits and Withdrawals (ordinary least-squares regressions)

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Any transfer via direct deposit (take-up)	Total deposits into accounts [MK]	Deposits into ordinary accounts [MK]	Deposits into commitment accounts [MK]	Deposits into other accounts [MK]	Total withdrawals from accounts [MK]
Time period:	Mar 2009 - Apr 2010	Mar-Oct 2009	Mar-Oct 2009	Mar-Oct 2009	Mar-Oct 2009	Mar-Oct 2009
PANEL A						
Any treatment	0.194***	17,609***	16,807***	668***	134	-16,761***
	(0.036)	(3,910)	(3,773)	(224)	(163)	(3,819)
PANEL B						
Commitment treatment	0.207***	18,801***	17,021***	1,490***	290	-17,511***
	(0.039)	(4,360)	(4,137)	(358)	(202)	(4,235)
Ordinary treatment	0.181***	16,513***	16,611***	-88	-9	-16,071***
	(0.040)	(4,840)	(4,743)	(181)	(163)	(4,745)
P-val. of F-test: Coefficients on commitment and ordinary treatments are equal	0.432	0.642	0.931	0.000	0.074	0.764
Mean dep. var. in Control group	0.000	3,281	3,107	0	174	-3,256
Number of observations	3,150	3,150	3,150	3,150	3,150	3,150

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the club level. USD 1 is ca. MK 145. All regressions include stratification cell fixed effects and the following baseline variables: Dummy for male respondent; dummy for married; age in years; years of completed education; number of household members; asset index; livestock index; land under cultivation; proceeds from tobacco and maize sales during the 2008 season; cash spent on inputs for the 2009 season; dummy for ownership of any formal bank account; amount of savings in bank or cash (with missing values replaced with zeros); dummy for hyperbolic (missing values replaced with zeros); dummy for "patient now, impatient later" (missing values replaced with zeros); net transfers made to social network over 12 months; dummy for missing value in savings amount; dummy for missing value in hyperbolic and "patient now, impatient later". For complete variable definitions, see Appendix B. F-test of Panel B tests the equality of means in commitment and ordinary treatment groups. Planting season is Nov-Apr. Fertilizer application occurs in Nov-Dec. Fertilizer purchases occur in both pre-planting period (Oct and before) and start of planting season (Nov-Dec). Net deposits are deposits minus withdrawals.

(which we did not assist in opening). The vast majority of deposits were into ordinary savings accounts. Treatment effects on that outcome (Panels A and B of column 3) are very similar in magnitude and statistical significance levels to those for total deposits in column 2.

In contrast, treatment effects on deposits into commitment accounts were much smaller (column 4). Panel A reveals that respondents assigned to any treatment group deposited less than MK 700 into a commitment account (significant at the 1% level), but that figure pools across individuals offered the commitment savings accounts and those offered ordinary accounts only. In Panel B, as we might expect, the impact of the ordinary treatment is very close to zero (and not statistically significant), while the impact of the commitment treatment is MK 1,490 (USD 10.28) and statistically significant at the 1% level. Results in column 4 reveal that the encouragement design had the intended effect of increasing use of illiquid savings instruments in the commitment treatment group. While impacts on commitment savings balances are positive and statistically significant, it is clear commitment savings deposits are substantially lower than deposits into ordinary accounts, even among those offered the commitment treatment.

Column 5 indicates that there were no large or statistically significant treatment effects on deposits into other partner bank accounts that were not offered by the project.

Treatment effects on withdrawals in the pre-planting period (column 6) are nearly as large in magnitude as effects on deposits. The "any treatment" coefficient in Panel A as well as the separate commitment and ordinary treatment coefficients in Panel B are all statistically significantly different from zero at the 1% level.

3.4.2 Time patterns of deposits and withdrawals

A key aim of this project was to promote savings for agricultural input investments, by facilitating individual bank account opening and channeling substantial resources (respondents' own crop proceeds) into those accounts. The results in Table 3.4 are therefore sobering, in that both deposits into and withdrawals from OBM accounts in the 2009 pre-planting period were substantial for both the commitment and ordinary treatments.

A question of interest is whether funds remained deposited in the accounts until the following planting period (November-December 2009), when agricultural inputs are typically applied. As it turns out, in many cases funds in ordinary accounts were withdrawn relatively quickly after the initial deposit of crop proceeds was made. About 22 percent of the initial deposits into ordinary accounts were followed by withdrawals on the same day of nearly equal amounts.[21] On average, only 26 percent of the original balance remained in an ordinary

[21] See Appendix B for details about the construction of deposit spells underlying these calculations.

savings account two weeks after it was initially deposited.

Figure 3.3 presents average deposits into and withdrawals from ordinary and other (non-commitment) accounts, by month, from March 2009 to April 2010.[22] The sample in Figure 3.3.a is individuals in the commitment treatment, while the sample for Figure 3.3.b is individuals in the ordinary treatment. For comparison, the sample used in Figure 3.3.c is individuals in the control group.

The figures indicate that peak deposits occurred in June, July, and August 2009, coinciding with the peak tobacco sales months. Average deposits in every month for individuals in both the commitment and ordinary treatments are quite similar in magnitude to average withdrawals, indicating that the majority of deposited funds were withdrawn soon thereafter. As a result, savings balances during the pre-planting period were much lower than deposited amounts, explaining why most farmers did not participate in the raffle.[23]

One likely reason funds in the ordinary accounts were withdrawn soon after they had been deposited has to do with transaction costs. Farmers lived on average 20 kilometers away from the bank branch and would typically travel there by foot, bus, or bicycle.[24] In addition to travel time, farmers report a median waiting time at the branch to withdraw money of one hour.

In contrast to the time pattern of the ordinary accounts, funds into commitment accounts do stay in accounts for longer periods of time. Figure 3.3.d displays average deposits into and withdrawals from commitment accounts, by month, for individuals in the commitment treatment. For deposits, the peak months are June, July, and August, coinciding with the peak deposit months for the ordinary accounts. But withdrawals from the commitment accounts are delayed substantially, occurring in October, November, and December, coinciding with the key months when agricultural inputs must be purchased and applied on fields. Of course, as revealed in Table 3.4, the amounts of money involved in these transactions are much lower than those in ordinary accounts.

3.4.3 Impacts on savings balances

Notwithstanding the fact that substantial amounts were withdrawn from accounts very soon after the direct deposits occurred, it is still possible that enough funds remained in total across both types of accounts to be able to detect statistically significant effects on savings balances. Due to our interest in facilitating savings for agricultural input utilization in the

[22] The data presented are the sum of the dependent variables in columns 4 and 6 of Table 3.4.

[23] The pattern is similar for individuals in the control group, but levels are much lower owing to the fact that direct deposit from the tobacco auction floor into farmer accounts was not enabled for that group.

[24] The median round-trip bus fare is MK 400 and takes two hours each way.

Figure 3.3: Deposits into and withdrawals from ordinary accounts

a. Commitment treatment group deposits into ordinary accounts

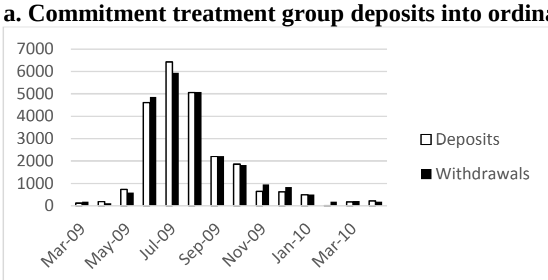

b. Ordinary treatment group deposits into ordinary accounts

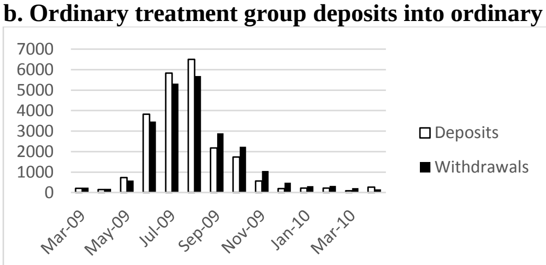

c. Control group deposits into ordinary accounts

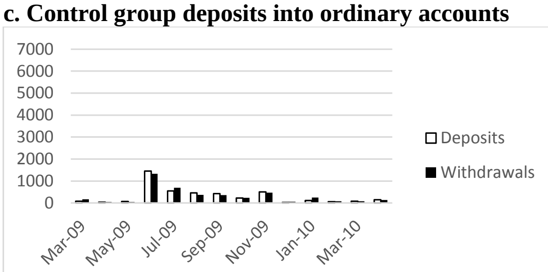

d. Commitment treatment group deposits into commitment accounts

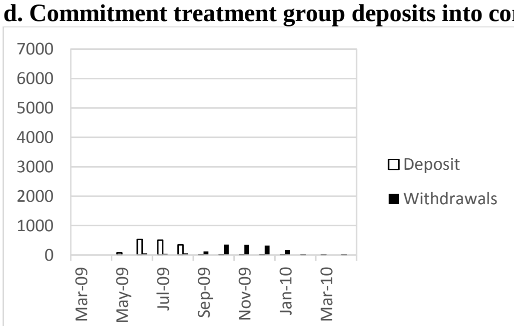

Notes: Deposits and withdrawals are denominated in Malawi kwacha (MK). Figures a, b, and c include transactions in ordinary accounts opened as part of the intervention as well as other non-commitment accounts owned by study participants (sum of dependent variables in columns 3 and 5 of Table 4). Figure d shows deposits into and withdrawals from commitment accounts, for individuals in commitment treatment group.

November-December 2009 planting season, we now examine treatment effects on savings balances immediately prior to that period.

Table 3.5 reports coefficient estimates from estimation of equations (3.1) and (3.2) for savings balances in the different types of OBM accounts, on October 22, 2009. In Panel A, which presents the impact of "any treatment," we find that the treatment effect is positive and statistically significantly different from zero at the 1% level for total savings balances (column 1), ordinary savings balances (column 2), and commitment savings balances (column 3). In addition, the coefficient in the regression for savings balances in other accounts (column 4) is also positive and statistically significantly different from zero at the 5% level.

In Panel B, which estimates separate effects for the commitment and ordinary treatments, we find that the effects of each treatment on total savings balances (column 1) are positive and statistically significantly different from zero at the 1% level. That said, the effect of the commitment treatment is larger than that of the ordinary treatment, and this difference is statistically significant at the 5% level. Effects of the treatments are very similar on savings in ordinary accounts and on savings in other accounts (columns 2 and 4); we cannot reject equality of the ordinary and commitment treatment effects for these outcomes at conventional significance levels. By contrast, the two treatments (unsurprisingly) differ in their impact on savings balances in commitment savings accounts: the commitment treatment effect is positive and statistically significantly different from zero at the 1% level, while the ordinary treatment effect is very close to zero and is not statistically significant. Equality of these two coefficients is rejected at the 1% level. It is therefore clear that the difference in the impacts of the commitment and ordinary treatments on total savings (shown in column 1) is being driven by the differing impacts on savings in commitment accounts (column 3).

These results reveal that both types of savings accounts have positive impact on savings preservation between the May-July 2009 harvest and the November-December 2009 planting season, with the commitment treatment providing an additional boost to savings on top of the impact of the ordinary account. The magnitudes of these effects are not negligible, in absolute terms for rural Malawian households as well as in comparison to control group savings of MK 364 (USD 2.36). The impact of "any treatment" on savings from Panel A is MK 1,863 (USD 12.85). From Panel B, the impact of the commitment savings treatment is MK 2,475 (USD 17.07) and the impact of the ordinary treatment is MK 1,301 (USD 8.97).

3.4.4 Impacts on agricultural outcomes and household expenditure

In Table 3.6, we turn to the impacts of the treatments on agricultural outcomes in the 2009-10 season (land cultivation, input use, crop output) and on household expenditures

Table 3.5: Impact of Treatments on Savings Balances (ordinary least-squares regressions)

	(1)	(2)	(3)	(4)
Dependent variable:	Savings balance immediately prior to planting period (on Oct 22, 2009)			
Account type:	All accounts, in total	Ordinary	Commitment	Other
PANEL A				
Any treatment	1,863***	1,167***	435***	262**
	(412)	(302)	(154)	(124)
PANEL B				
Commitment treatment	2,475***	1,167***	935***	372**
	(524)	(364)	(238)	(187)
Ordinary treatment	1,301***	1,167***	-26	160
	(442)	(349)	(129)	(129)
P-val. of F-test: Coefficients on commitment and ordinary treatments are equal	0.019	0.999	0.000	0.290
Mean dep. var. in Control group	364	302	0	62
Number of observations	3,150	3,150	3,150	3,150

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Dependent variable is savings balance on Oct 22, 2009, just prior to November-December 2009 planting season. Standard errors are clustered at the club level. USD 1 is ca. MK 145. All regressions include stratification cell fixed effects and the following baseline variables: Dummy for male respondent; dummy for married; age in years; years of completed education; number of household members; asset index; livestock index; land under cultivation; proceeds from tobacco and maize sales during the 2008 season; cash spent on inputs for the 2009 season; dummy for ownership of any formal bank account; amount of savings in bank or cash (with missing values replaced with zeros); dummy for hyperbolic (missing values replaced with zeros); dummy for "patient now, impatient later" (missing values replaced with zeros); net transfers made to social network over 12 months; dummy for missing value in savings amount; dummy for missing value in hyperbolic and "patient now, impatient later". For complete variable definitions, see Appendix B. F-test of Panel B: "Commitment savings = Ordinary savings" tests the equality of means in commitment and ordinary treatment groups.

Table 3.6: Impact of Treatments on Agricultural Outcomes in 2009-2010 Season and Household Expenditure after 2010 Harvest

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Land under cultivation [acres]	Total value of inputs [MK]	Proceeds from crop sales [MK]	Value of crop output (sold and not sold) [MK]	Farm profit (output-input) [MK]	Total expenditure in 30 days prior to survey [MK]
PANEL A						
Any treatment	0.30**	8,023*	19,595**	23,921**	16,927*	1,151*
	(0.15)	(4,131)	(8,996)	(11,529)	(9,117)	(601)
PANEL B						
Commitment treatment	0.33**	10,297**	26,427***	31,259**	21,369**	1,442**
	(0.16)	(4,563)	(9,979)	(12,510)	(10,064)	(656)
Ordinary treatment	0.27*	5,946	13,358	17,223	12,872	885
	(0.16)	(4,504)	(9,518)	(12,204)	(9,577)	(650)
P-val. of F-test: Coefficients on commitment and ordinary treatments are equal	0.614	0.246	0.086	0.117	0.246	0.283
Mean dep. var. in Control group	4.28	60,372	91,747	155,685	95,210	10,678
Number of observations	2,835	2,835	2,835	2,835	2,835	2,835

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the club level. USD 1 is ca. MK 145. All regressions include stratification cell fixed effects and the following baseline variables: Dummy for male respondent; dummy for married; age in years; years of completed education; number of household members; asset index; livestock index; land under cultivation; proceeds from tobacco and maize sales during the 2008 season; cash spent on inputs for the 2009 season; dummy for ownership of any formal bank account; amount of savings in bank or cash (with missing values replaced with zeros); dummy for hyperbolic (missing values replaced with zeros); dummy for "patient now, impatient later" (missing values replaced with zeros); net transfers made to social network over 12 months; dummy for missing value in savings amount; dummy for missing value in hyperbolic and "patient now, impatient later". For complete variable definitions, see Appendix B. F-test of Panel B tests the equality of means in commitment and ordinary treatment groups.

Table 3.7: Impact of treatments on household size, transfers and fixed deposit demand

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Household size	Tobacco loan amount [MK]	Total transfers made [MK]	Total transfers received [MK]	Total net transfers made [MK]	Has fixed deposit account
PANEL A						
Any treatment	0.14	3,158	215	-301	477	0.032***
	(0.09)	(4,583)	(249)	(248)	(322)	(0.012)
PANEL B						
Commitment treatment	-0.004	3,418	304	-316	568	0.050***
	(0.019)	(4,897)	(275)	(258)	(347)	(0.014)
Ordinary treatment	-0.010	2,920	134	-288	394	0.016
	(0.019)	(5,068)	(267)	(262)	(342)	(0.012)
P-val. of F-test: Coefficients on commitment and ordinary treatments are equal	0.748	0.899	0.431	0.856	0.483	0.008
Mean dep. var. in Control grou	5.72	40,147	2,872	2,492	418	0.039
Number of observations	2,835	2,835	2,835	2,835	2,835	2,835

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the club level. USD 1 is ca. MK 145. All regressions include stratification cell fixed effects and the following baseline variables: Dummy for male respondent; dummy for married; age in years; years of completed education; number of household members; asset index; livestock index; land under cultivation; proceeds from tobacco and maize sales during the 2008 season; cash spent on inputs for the 2009 season; dummy for ownership of any formal bank account; amount of savings in bank or cash (with missing values replaced with zeros); dummy for hyperbolic (missing values replaced with zeros); dummy for "patient now, impatient later" (missing values replaced with zeros); net transfers made to social network over 12 months; dummy for missing value in savings amount; dummy for missing value in hyperbolic and "patient now, impatient later". For complete variable definitions, see Appendix B. F-test of Panel B: "Commitment savings = Ordinary savings" tests the equality of means in commitment and ordinary treatment groups.

after the 2010 harvest.[25]

Column 1 presents treatment effects on land under cultivation in acres. Panel A indicates that land cultivated was higher by 0.30 acres among respondents offered any treatment (statistically significant at the 5% level), compared to 4.28 acres in the control group. Treatment effects are very similar when estimated for the commitment and ordinary treatments separately (Panel B), and the difference between the two is not statistically significantly different from zero.[26]

Results in column 2, Panel A show that the treatment had a positive impact on the total monetary value of agricultural inputs used in the 2009-10 planting season, which is statistically significant at the 10% level. Estimating the effects separately for the commitment and ordinary treatments reveals that both effects are positively signed, and the effect of the commitment treatment is statistically significantly different from zero at the 5% level. While the commitment treatment coefficient is larger in magnitude than the ordinary treatment coefficient, we cannot reject at conventional statistical significance levels that the two treatment coefficients are equal to one another.

The increase in agricultural input utilization caused by the treatment appears to have, in turn, caused increases in agricultural output. Columns 3-5 show treatment effects on, respectively, crop sale proceeds, value of crop output (both sold and unsold), and farm profit (value of output minus value of inputs). For each of these outcomes, the "any treatment" coefficient in Panel A is positive and statistically significant at the 5% or 10% level. In Panel B, the commitment treatment coefficient is positive and statistically significant in each of the regressions at the 1% or 5% level, and is larger in magnitude in each case than the corresponding ordinary treatment coefficient. Only in column 3 (proceeds from crop sales) can we reject at conventional levels (10% in this case) the hypothesis that the commitment and ordinary treatment coefficients are equal.[27, 28]

[25] All outcomes in Table 3.6 are for the total household, not per capita. We show in Table 3.7, column 1 that the treatments have no effect on household size, so interpretation of impacts in Table 3.6 is not clouded by concurrent changes in household size.

[26] We investigated whether the treatment effects on land are due to increased land rentals, and found no large or statistically significant effect (for "any treatment" and for the commitment and ordinary treatments separately). Results available from authors on request.

[27] The increase in farm profit in column 5 and in the value of the inputs in column 2 suggests a high rate of return to inputs. Most of the increases in expenditures were on firewood to cure tobacco and on fertilizer. Among the different varieties of tobacco grown, the highest value one needs more curing, so the increased profits could be due to a shift in the crop mix towards higher value tobacco as well as the increased inputs. In addition, historical production and weather data suggest that 2010 was a good production year with average crop prices.

[28] In results available upon request, we find that increases in production caused by the treatments are relatively concentrated in tobacco production. In the control group, tobacco accounts for 66.5% of the kwacha value of production, but increases in tobacco production account for 81.4% of the treatment effect (MK19,477 of the MK23,921 increase in the value of crop output).

Given the positive treatment effects on agricultural production, it is of interest to examine effects on household expenditures, in column 6. The effect of any treatment is positive and statistically significant at the 10% level (Panel A). Results in Panel B show that both commitment and ordinary treatment effects are positive in magnitude, and the commitment treatment effect is statistically significantly different from zero at the 5% level. We cannot reject at conventional significance levels that the commitment and ordinary treatment effects are equal.

The treatment effects identified in Table 3.6 are economically significant. In Panel A, the treatment effect on total value of inputs is MK 8,023 (USD 55.33), amounting to an increase of 13.3 percent over the control group mean, while the treatment effect on value of crop (sold and unsold) is MK 23,921 (USD 164.97), an increment of 15.4 percent over the control group mean. The increase in household expenditure is 10.8 percent vis-a-vis the mean in the control group. These results show large, consistent effects of "any treatment" on outcomes that are likely connected to household well-being.

Consistent with these findings, column 6 of Table 3.7 shows that being assigned to a savings treatment group increased the probability of owning a fixed-deposit account over a year later by 0.03 percentage points, a statistically significant increase of 75 percent relative to the control group mean of 0.039.[29] In addition, study participants continue to use the offered ordinary accounts. Using the bank's administrative data we find that treatment effects on deposits, withdrawals, and net deposits persist during the May to July 2010 period, more than a year after the initial intervention, particularly in the ordinary treatment group.

The continued usage of ordinary accounts and the increased take-up of fixed deposit accounts one year after the intervention suggest that farmers in the treatment group found something of value in the savings products offered.

3.5 Mechanisms

We now turn to considering the mechanisms through which our treatment effects may have operated. Studies of the impact of savings account access typically posit (implicitly or explicitly) that effects would operate via alleviation of savings constraints (Dupas and Robinson, 2013*a*; Prina, 2013, for example). A study population is typically thought to have imperfect methods for preserving funds, which can be depleted for a variety of reasons such as self-control problems, demands for sharing with one's social network, or theft. In our study

[29] Fixed deposit accounts (also known as time deposit or term deposit accounts in other countries) are accounts in which the depositor accepts lower liquidity (an agreement not to withdraw) for a certain specified time period, in return for a higher interest rate. Such accounts could be seen as providing a similar commitment function to the commitment savings accounts offered in the experiment.

population, alleviation of savings constraints via provision of formal savings accounts could help farmers preserve funds between harvest and the subsequent planting season, leading to positive impacts on agricultural input expenditures (and then on other subsequent related outcomes).

While we do find positive treatment effects on both savings balances and on subsequent agricultural input utilization, the relative magnitudes of the effects are inconsistent with alleviation of savings constraints being the only mechanism at work. Consider the impact of "any treatment" on the value of agricultural inputs used (Table 6, column 2), MK 8,023. While the treatment did cause an increase in deposits exceeding that amount (MK 17,609, from Table 4, column 2), withdrawals happened quite soon after deposits, so that very little remained in the accounts some months later once the time came for the November-December input purchases: the treatment effect on savings balances at the end of October is just MK 1,863 (Table 5, column 1), which is just 23% of the increase in the value of inputs.[30] Therefore, no more than about a quarter of the effect of the treatment on agricultural input expenditures can be attributed to alleviation of savings constraints *per se*.

In Table 7, we estimate treatment effects on other outcomes, to test for other operative mechanisms behind our main results. One possible explanation for the increase in total expenditure on inputs for the savings treatment group could be that increased savings at the bank led to increased eligibility for loans, and it is these loans that funded the increased purchases of inputs.[31] Column 2 examines the size of loans provided by a lender in the subsequent season. While coefficients in Panels A and B are positive, none are statistically significantly different from zero.[32] It should be said, however, that the point estimates are relatively imprecise, and 95% confidence intervals do include the estimated treatment effects on the value of agricultural inputs.

Other alternate explanations have in common the hypothesis that while most funds deposited in the accounts at harvest time were withdrawn fairly soon thereafter, they may have nonetheless been spent on agricultural inputs. They could have been spent on inputs sometime between harvest and the November-December planting (making immaterial our finding of low savings balances in late October.) Or they could have been preserved outside the bank (say in cash held at home or with "money guards") and used for input purchases during

[30] A one-sided test that the "any treatment" effect on the value of agricultural inputs (8,023) is larger than the treatment effect on end-of-October savings balances (1,863) is statistically significant at the 10% level (p-value 0.061). Corresponding tests for the ordinary treatment and commitment treatment have p-values of 0.143 and 0.038 respectively.

[31] Loans from informal lenders and friends and family account for a small fraction of total borrowing. At any rate, conducting this analysis for total credit instead of just tobacco credit yields very similar results.

[32] Similarly, we find no difference across treatment and control groups in the probability of accessing a loan (results not shown).

the planting season. In either case alleviation of savings constraints via provision of formal accounts *per se* cannot be the operative mechanism, so we search for other mechanisms.

One hypothesis is that the existence of the accounts allowed households to resist social network demands for resources (what one might call "other-control" problems) in the period between the harvest and planting seasons. While the data from our partner bank show relatively low savings overall, with only a minority in the restricted-access commitment accounts, neither total balances nor the share in commitment accounts were public knowledge to the community. The existence of formal accounts may have provided an excuse to turn down requests for assistance from the social network by claiming that savings were inaccessible.[33] We test this hypothesis in Table 7 by regressing three direct measures of transfers between households (transfers made, transfers received, and net transfers) on the treatment variables. We find no effect of either intervention in any of these outcomes, however. All coefficients (in both Panels A and B) are relatively small in magnitude and none are statistically significantly different from zero at conventional levels. That said, these measures span the pre-planting to post-harvest period, and are thus consistent with lower transfers during the pre-planting season, when commitment accounts were active and therefore could serve as a valid excuse for reducing transfers, followed by higher transfers after the harvest, when farmers with commitment accounts realized larger revenues. Unfortunately, we lack the data needed to examine the timing of transfers. In addition, it is still possible that the commitment treatment allowed study participants to keep funds from others within the household, or to refrain from consuming resources early in anticipation of future requests from others (as in Goldberg, 2011).

Another possibility is that the ability to hold a buffer stock in formal savings accounts made farmers willing to take on the risk of making higher input investments (Angeletos and Calvet, 2006; Kazianga and Udry, 2006).

Alternatively, treatment may also have affected agricultural production decisions via one or more of several mechanisms suggested by research in psychology and behavioral economics. Because the savings accounts were framed by the experiment as vehicles for accumulating funds for agricultural inputs, the very act of signing up for deposits into savings accounts could have been viewed by farmers as a commitment to raise expenditures of this type. This mere elicitation of farmers' intentions may have influenced their later behavior (Feldman and Lynch, 1988; Webb and Sheeran, 2006; Zwane et al., 2011). Relatedly, the act of signing

[33] To be sure, one of the "raffle" arms involved public distribution of raffle tickets based on savings balances. We do not find that these effects are distinguishable from the effects of treatments with no distribution of tickets. Also, the distribution of funds across ordinary and commitment accounts was not public knowledge because the cross-randomized raffle treatments awarded raffle tickets on the basis of total funds across all accounts, so even the public raffle did not reveal how little was saved in commitment accounts.

up for direct deposits into savings accounts may have created an "agricultural input" mental account for the deposited funds (Thaler, 1990), even if most funds were withdrawn soon after being deposited and relatively small amounts remained in the accounts. Finally, signing up for direct deposit into accounts could have altered study participants' reference points about future input use, farm output, and consumption. In this context, prospect theory (Kahneman and Tversky, 1979*b*) would predict that farmers offered savings accounts could have become more willing to invest in agricultural inputs, so as to avoid losses in the form of failing to achieve their (experimentally-induced) higher reference points for input use, output, and consumption. Unfortunately, we can offer no direct evidence to support or contradict that such psychological channels may have been at work.

3.6 Conclusion

Viewed as a policy intervention for increasing the use of agricultural inputs by households in developing countries, savings accounts have appealing features. Unlike subsidies, they do not require major government budget commitments. While the supply of credit for agricultural inputs is often constrained, banks are eager to attract new savings customers. The results of our field experiment among cash crop farm households in Malawi show that offering access to individual savings accounts not only increases banking transactions, but also has statistically significant and economically meaningful effects on measures of household wellbeing, such as investments in inputs and subsequent agricultural yields, profits, and household expenditure. Ours is one of the first randomized studies of the economic impact of savings accounts, and the first (to our knowledge) to measure impacts on important agricultural outcomes (input use and farm output) and household consumption levels.

An important direction for future research would be to provide evidence on the mechanisms underlying the effects we found, since our treatment effects on input utilization are larger than can be explained by alleviation of savings constraints alone. Other mechanisms that might be explored might be the role of savings as a buffer stock for self-insurance, increases in credit access, reductions in demands from others in the social network ("other-control" problems), as well as mechanisms suggested by behavioral economics (e.g., mental accounting and reference dependence).

APPENDIX A

Robustness checks

A.1 Robustness checks

This section includes alternative specifications for the main regression results. See text for description.

Table A.1: Effects of bonus incentives on attendance and output of tea pluckers (pooled OLS estimation)

	(1) Full attendance [0/1]	(2) Weekly attendance [days]	(3) Met output low threshold for weekly total output [0/1]	(4) Met output high threshold for weekly total output [0/1]	(5) Total weekly output [kg]	(6) Weekly average of daily productivity if plucking [kg/day]
Lottery Bonus	**0.0357*** **	**0.0607*** **	**-0.00521**	**0.00947**	**3.380**	**-0.0849**
	(0.0112)	(0.0208)	(0.00985)	(0.00839)	(2.726)	(0.407)
Fixed Bonus	**0.016**	**0.03**	**4.99e-05**	**0.000396**	**1.473**	**-0.210**
	(0.0107)	(0.0218)	(0.00748)	(0.00668)	(2.569)	(0.297)
L1.Lottery Bonus	**0.00176**	**-0.0192**	**-0.00144**	**-0.00326**	**-1.123**	**-0.316**
	(0.0125)	(0.0200)	(0.00768)	(0.00789)	(2.684)	(0.408)
L1.Fixed Bonus	**-0.0196**	**-0.0523*** **	**-0.000343**	**-0.00825**	**-3.971**	**-0.465**
	(0.0127)	(0.0176)	(0.00985)	(0.00869)	(2.756)	(0.392)
P.-val. of Wald tests:						
Lottery Bonus = Fixed Bonus	0.0611	0.0747	0.506	0.238	0.515	0.663
Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.0134	0.0205	0.793	0.447	0.473	0.714
Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0	0.0602	0.0227	0.965	0.590	0.284	0.498
Mean of dep. variable in no-bonus condition during experiment	0.555	5.290	0.533	0.250	384.4	81.01
Number of gang-week observations during experiment	15603	15603	15603	15603	15603	15490
Number of gang-week observations in total	38,295	38,295	38,295	38,295	38,295	38,117
Number of workers	1,678	1,678	1,678	1,678	1,678	1,678

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the gang level. Estimated results are from a worker-level random effect model with a full set of gang-week dummies. Wald-tests: “Lottery Bonus = Fixed Bonus” tests the equality of the two coefficients. “Joint: Lottery Bonus = 0 and Fixed Bonus = 0” and “Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0” test the joint equality of the two respective pairs of coefficients.

Table A.2: Effects of bonus incentives on attendance and output of tea pluckers (individual fixed effects estimation)

	(1)	(2)	(3)	(4)	(5)	(6)
	Full attendance [0/1]	Weekly attendance [days]	Met output low threshold for weekly total output [0/1]	Met output high threshold for weekly total output [0/1]	Total weekly output [kg]	Weekly average of daily productivity if plucking [kg/day]
Lottery Bonus	**0.038*** **	**0.07*** **	**-0.006**	**0.009**	**4.1***	**-0.1**
	(0.011)	(0.02)	(0.010)	(0.007)	(2.2)	(0.3)
Fixed Bonus	**0.018**	**0.04***	**-0.001**	**0.000**	**2.4**	**-0.3**
	(0.011)	(0.02)	(0.008)	(0.006)	(2.6)	(0.2)
L1.Lottery Bonus	**0.005**	**-0.01**	**-0.003**	**-0.004**	**-0.1**	**-0.3**
	(0.012)	(0.02)	(0.008)	(0.008)	(2.3)	(0.3)
L1.Fixed Bonus	**-0.016**	**-0.04****	**-0.001**	**-0.009**	**-2.7**	**-0.5**
	(0.012)	(0.02)	(0.010)	(0.008)	(2.7)	(0.3)
P.-val. of Wald tests:						
Lottery Bonus = Fixed Bonus	0.051	0.113	0.479	0.213	0.536	0.407
Joint: Lottery Bonus = 0 and Fixed Bonus = 0	0.003	0.004	0.754	0.387	0.181	0.438
Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0	0.047	0.043	0.934	0.531	0.407	0.335
Mean of dep. variable in no-bonus condition during experiment	.55491488	5.29	0.533	0.250	384.4	81.0
Number of gang-week observations during experiment	15603	15,603	15,603	15,603	15,603	15,490
Number of gang-week observations in total	38295	38,295	38,295	38,295	38,295	38,117
Number of workers	1,678	1,678	1,678	1,678	1,678	1,678

Notes: Stars indicate significance at 10% (*), 5% (**), and 1% (***) levels. Standard errors are clustered at the gang level. Estimated results are from a worker-level random effect model with a full set of gang-week dummies. Wald-tests: “Lottery Bonus = Fixed Bonus” tests the equality of the two coefficients. “Joint: Lottery Bonus = 0 and Fixed Bonus = 0” and “Joint: L1.Lottery Bonus = 0 and L1.Fixed Bonus = 0” test the joint equality of the two respective sets of coefficients.

APPENDIX B

Balance and comparison demographic characteristics of sample to census data

B.1 Balance and comparison demographic characteristics of sample to census data

Appendix Table B.1 shows summary statistics for important baseline characteristics: all available baseline measures corresponding to outcome variables used in the results tables as well as an index of asset holdings (used as controls in the main results tables). Columns 4 and 5 present formal statistical tests of the null hypothesis that pre-program characteristics have equal means across all four study arms. For each covariate, the test is conducted by running two linear regressions as seemingly-unrelated regressions (SURs) of the variable on a saturated set of categorical indicator variables for study arm, one regression for each round. We then run a joint test of the null hypothesis that all the coefficients are zero. Column 5 shows the p-values, which are uniformly above 0.3. The last row shows the test statistic and p-value for a joint test of the hypothesis that all the coefficients equal zero across all 26 regressions. We fail to reject the null of no differences (p-value of 0.81). The sample is similarly balanced on demographic covariates; see analogue test statistics in Appendix Table B.2.

Table B.2, columns 1 to 3, presents summary statistics of demographic characteristics for the 350 workers from our sample for whom baseline data is available. As a basis for comparison, we also present statistics for Mulanje District as a whole, taken from the IPUMS-International 10% sample of the 2008 Malawi Population and Housing Census. A comparison of our sample with the rest of the district suggests that it is generally representative of the

Table B.1: Balance of baseline variables

	Worker Sample Summary Statistics			Test for Difference Across Study Arms	
	(1)	(2)	(3)	(4)	(5)
Variable	Mean	SD	N	Chi^2	p-value
Income and Spending					
Total spending since last Friday, inclusive [MK]	2271.04	3728.39	329	3.84	0.70
Cash remaining out of total received since last Friday, inclusive [MK]	683.81	2618.13	329	7.00	0.32
Expenditure Composition					
Food for consumption at home	0.66	0.23	349	3.05	0.80
Maize only	0.23	0.26	349	2.24	0.90
Food for consumption out of home	0.06	0.07	349	1.77	0.94
Non-Food	0.28	0.23	349	3.83	0.70
Assets					
Baseline Asset Ownership Index (First Principal Component)	0.00	2.68	350	5.53	0.48
Combined Test Across All Variables				28.50	0.81

Notes: Sample includes 359 respondents who participated in at least one round of the work program and have data from at least one data source for that round (either the payday data, the survey, or both). All money amounts are in Malawian Kwacha (MK); during the study period the market exchange rate was approximately MK400 to the US dollar, and the PPP exchange rate was approximately MK160 to the US dollar. Tests for any difference in means across study arms use seemingly-unrelated regressions of a variable on a full set of categorical indicator variables for study arm, clustered by respondent, to do pooled tests of the null hypothesis that all study arms have equal means in both rounds; the test statistics are chi-square distributed with 6 degrees of freedom. The Combined Test Across All Variables is a combined SUR of all 8 covariates in both rounds; its chi-square test statistic has 36 degrees of freedom due to some constraints being dropped.

local area, with differences that are likely due to the criteria used by the Village Development Committee (VDCs) to select workers for the program. Our sample is 69% female, which is substantially higher than the district average of 55%. It also has a larger share of people from the Lomwe ethnic group, at 90% compared with 75%. It is otherwise quite similar to the district as a whole, with similar rates of marriage (70%) and Christian religion (90%). The differences in the other variables are fairly small, and consistent with the VDCs selecting people of lower socioeconomic status for the program. For example, our sample averages 3.5 years of completed schooling, compared with 4.4 years for the district as a whole, and has a mean age of 40 compared with 37 for Mulanje District. Our workers are also more likely to be divorced and less likely to be single.

Table B.2: Demographic characteristics of sample - balance and comparison to census

	Worker Sample Summary Statistics			Test for Difference Across Study Arms		Mulanje District 2008 Census Summary Statistics	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Variable	Mean	SD	N	Chi-square	p-value	Mean	SD
Male	0.31	0.46	344	3.79	0.70	0.45	0.50
Religion							
Christian	0.90	0.30	341	5.93	0.43	0.91	0.28
Muslim	0.10	0.30	341	5.93	0.43	0.05	0.22
Marital Status							
Married	0.69	0.46	338	5.46	0.49	0.71	0.45
Divorced/Widowed	0.25	0.44	338	5.44	0.49	0.17	0.37
Single	0.05	0.21	338	8.74	0.19	0.12	0.33
Ethnic Group							
Lomwe	0.89	0.31	344	1.66	0.95	0.75	0.43
Yao	0.07	0.26	344	2.29	0.89	0.05	0.22
Mang'anja	0.02	0.13	344	6.15	0.41	†	
Other	0.02	0.15	344	8.27	0.22	0.20	0.40
Years of Education Completed	3.54	3.15	341	3.47	0.75	4.45	3.91
Age (Years)	40.03	15.40	344	4.24	0.64	37.35	17.27
Combined Test Across All Variables				61.42	0.73		

Notes: Pooled Sample includes 359 respondents who participated in at least one round of the work program. Tests for any difference in means across study arms use seemingly-unrelated regressions of a variable on a full set of categorical indicator variables for study arm, clustered by respondent, to do pooled tests of the null hypothesis that all study arms have equal means in both rounds; the test statistics are chi-square distributed with 6 degrees of freedom. The Combined Test Across All Variables is a combined SUR of all 13 covariates in both rounds; its chi-square test statistic has 69 degrees of freedom due to some constraints being dropped. † The 2008 Malawi Census does not report Mang'anja ethnicity as a separate category, si it is included in "other".

APPENDIX C

Variable definitions

C.1 Variable definitions

Data used in this paper come from three rounds of "full length" surveys (a baseline and two follow-up interviews), from two- to four-question surveys during paydays as well as from administrative records of the project. We conducted a baseline survey from 4 Oct 2013 to 19 Oct 2013 and two follow-up surveys after the last payday weekend of each round, once from 2 Dec 2013 to 7 Dec 2013 and once from 27 Jan 2014 to 31 Jan 2014. All variables that are created from survey data are Winsorized at the 1st and 99th percentile. All figures in money terms are in local currency units, Malawi Kwacha (MK).

C.1.1 Variables from payday surveys

Amount spent on same day as income receipt is total market spending on all days that workers received their wages (sum of all four payday Fridays or Saturdays for the weekly payment group; the fourth payday Friday or Saturday for the lump sum payment group).

Money spent at market on Fridays 1, 2, 3 is the sum of total market spending on the first three payday Fridays.

Money spent at market on Saturdays 1, 2, 3 is the sum of total market spending on the first three payday Saturdays.

Money spent at market on Friday 4 is the total market spending on the fourth payday Friday.

Money spent at market on Saturday 4 is the total market spending on the fourth payday Saturday.

C.1.2 Variables from follow-up surveys

Total spending since last Friday, inclusive [MK] is the total household spending starting from the fourth payday Friday until the day of interview in the week after the fourth payday. The variable is derived from the difference of the answers to the questions "Since last Friday, how much cash have you received?" and "How much of that cash do you have left?", respectively.

Remaining cash out of received since last Friday, inclusive [MK] is household's remaining cash holdings out of money received starting from the fourth payday Friday until the day of interview.

Self-reported wasteful spending on weekend 4 of round 2 variables ask for money that respondents report as "wasted" or spending which the respondent was tempted into spending that he/she should not have spent?

- *Total since last Friday, inclusive [MK]* is the sum of total wasteful spending starting from the fourth payday Friday until the day of interview in the week after the fourth payday.
- *Friday [MK]* is total wasteful spending on the fourth payday Friday.
- *Saturday [MK]* is total wasteful spending on the fourth payday Saturday.
- *Sunday and after [MK]* is the sum of total wasteful spending starting from the fourth payday Sunday until the day of interview in the week after the forth payday.

Expenditure shares based on itemized elicitation is the sum of itemized expenditures grouped different categories as a share of total expenditures across all items based on an large listing of possible items (with items derived from Malawi's Integrated Household Survey; a select number of items was consolidated or omitted but including "other" items in each category; total number of 105 items in 12 categories).

- *Food for consumption at home* includes eight categories of food items typically used for home consumption.
- *Maize only* includes only maize flour and maize grain.
- *Food for consumption out of home* includes all items from categories "cooked foods from vendor" and "Beverages" which are typically consumed away from home.
- *Non-Food* includes all non-food items.

Value of net asset purchases since last interview is the sum of difference between value of assets bought and assets sold from an itemized list of common assets as well as an "other" category considering purchases and sales since the last interview, i.e. since baseline interview for follow-up 1 and since follow-up 1 for follow-up 2.

C.1.3 Variables from baseline surveys

Assets index is an index based on the first principal component of the number of items owned of 64 common non-financial, non-livestock assets and the number of animals owned of 9 common types of livestock.

Total spending is defined similarly to "Total spending since last Friday, inclusive" described under follow-up variables above, covering the last Friday prior to the interview until the day of the interview.

C.1.4 Variables from project records

Bought any shares is an indicator for whether respondent bought at least one "share" of the investment opportunity offered after the follow-up interviews (see details in main text in Data Collection section).

Total spent on shares is the total amount spent on the investment opportunity offered and equal the number shares bought times the price of one share (MK 1,500).

APPENDIX D

Account details and full text of training script

D.1 Savings account details

We offered farmers training and account opening assistance for two types of accounts depending on treatment status (control, ordinary savings or commitment savings). The " ordinary" account referred to in the main text is OBM's Kasupe account. Kasupe accounts had an account opening fee of MK500, no monthly fee, three free withdrawals transactions via ATM per month, and a MK25 fee per ATM withdrawal thereafter (all withdrawals at the teller were free). The minimum balance for Kasupe accounts was MK500 and there was an account closing fee of MK1,000. Kasupe accounts paid an interest rate of 2.5% p.a. with interest accruing quarterly. Deposit transactions into Kasupe accounts were free.

Farmers were given the option to have their proceeds directly deposited into an existing account if they already had a savings account with OBM. Another type of savings account not actively marketed in this experiment but part of OBM's product portfolio was standard savings accounts with the following fee structure: an opening fee of MK500; a monthly fee of MK75; no withdrawal fees; minimum balance of MK1,000; a closing fee of MK1,000; an interest rate of 6.5% p.a. with quarterly accrual. This less common account type is included in the category " ordinary" accounts together with Kasupe accounts.

The " commitment" account referred to in the main text was an account newly developed for the project called " SavePlan." SavePlan accounts paid the same interest rate as Kasupe accounts, but had no minimum balance requirement. SavePlan accounts also had no account opening or closing fees. Deposit transactions into SavePlan accounts were free. The only withdrawals permitted for SavePlan accounts were transfers to ordinary (Kasupe or other) savings accounts, for which no fee was charged.

D.2 Scripts for savings training, account offers, and raffle training

(Scripts were administered in club meeting immediately following administration of baseline survey. Malawian research project staff played the roles of Persons 1 and 2.)

D.2.1 Section 1: Savings Accounts (All Clubs)

Person 1: Saving money in an individual bank account is a very smart way to protect your money and improve your wellbeing. As you know, OBM has Kasupe accounts that are easy and affordable to use.

Person 2: But I already have a savings account with my club. What is better about this Kasupe account?

First ask the group to list things that are good about the Kasupe account. When the group has come up with several suggestions, move on to the next line:

Person 1: The Kasupe account is yours alone. You don't share it with the rest of your club members. You are the only one who can take money out of the account and the only one who knows how much money you have saved in the account.

Person 2: What are the details of the account? How much does it cost, and what is the interest?

Person 1: MK 500 for smartcard, MK 500 for initial deposit, no monthly charge, MK 25 transaction charge (ATM fee, withdrawal fee).

Person 2: But I can just keep money at home. What are some of the benefits of saving my money in a Kasupe account instead of at home?

Let the group make suggestions. After several things have been suggested, agree with the group and then move on to the next line.

Person 1: Money is safer in a bank account than at home. If you keep your money at home, it could be stolen or lost in a fire. If you keep it at the bank, it is protected. Also, if you keep money at home, you may feel obligated to give money to your family or friends if they ask for it. If your money is in the bank, you can say that you don't have any money to give.

Person 2: That is interesting, but I think my money is safe at home.

Ask the group: "Do you think money is safe at home?" Let the group come up with answers, then move on.

Person 1: There are other reasons to keep money in the bank, too. Keeping money in a bank account can help you save for the future. If you have money at home, it is easy to be tempted to spend it on food or drinks or household items. If you have money in the bank, you will think twice about taking it out to spend. Instead, you can leave it in the bank to save for important purchases like school fees or buying fertilizer or accumulating the deposit for a new loan. Also, you can be sure to put away money in case you have an emergency in the future, like someone gets sick and needs to go to the hospital.

D.2.2 Section 2: Saving for the future (All Clubs)

Person 2: It would be good to save for the future, but I have many needs now. How can I afford to save?

Person 1: It is important to make a plan for how to spend your money. One way to do this is to divide the money you will have after selling your tobacco and paying your loans into two amounts. One amount is to use now, and the other amount is to use in the future. Then, you can commit to keeping the future amount safe, and not touching it now.

Person 2: How can I do that?

Person 1: Think about how much money you will have after you sell your tobacco and repay your loan to OBM. Then, think about expenses you have immediately.

Have the group list things they need to spend money on immediately. Get a list of 5-6 things, then move on.

Person 2: Yes, I will have to pay someone who has done weeding for me. Also, I need to buy some soap and other household goods. My children need new clothes, too.

Person 1: Yes, these are the kinds of things you need to spend money on right away, when you get paid. But now think of things you will need to spend money on in the future. What do you want to be absolutely sure you can afford?

Ask the group to list things they want to save for in the future. Make sure they are thinking of long-term things or expenses that will happen in a few months. Get the group to list 5-6 things, then move on.

Person 2: I can think of many things. I will need to pay school fees. Also, I want to make sure I can buy fertilizer for my maize. And I want to have money for food next year during the hungry season.

Person 1: These are important expenses. You should plan to protect some of your money so that it is available for those expenses. You can do that by committing to locking it away until a date in the future, when you will need it. What is a date that makes sense? Choose a time that is close to when you will need the money for the reasons you just described, so that you aren't tempted to spend it on other things.

Ask the group: "When do you think you want to access money you would save for the future?" Let the group discuss several dates. Make sure they consider purchasing inputs, and also food during the hunger season.

Person 2: Hmmm. November 1 is probably a good time. That will be in time for me to buy fertilizer and pay my loan deposit.

Person 1: Now that you have chosen a date, you have to decide how to divide your money between things you will buy before that date, and the things you are saving for in the future. This is an important choice. You have to make sure that you have enough money for your immediate needs and things you will have to buy before the date you have chosen. You also have to estimate how much money you will need for the things you want to buy in the future. Start with money you need soon. Of the money you will have after you sell your tobacco and repay your loan, how much do you need to have available for spending before November 1, which is the date you have chosen?

Have the group suggest amounts of money they will spend on immediate expenses.

Person 2: Well, I need to pay someone for ganyu. And I need to buy clothes, and some household items right away. I will also need to spend some money after the harvest season on small things like soap. I will need to spend MK 25,000 between when I get money and November 1.

Person 1: Ok. How much do you want to make sure to have for the future, after that date you have chosen?

Person 2: I will need MK 4,500 for fertilizer, and MK 3,000 for a deposit on a new loan. Also, I want to keep MK 2,000 for food in the hungry season. That is MK 9,500 total.

Person 1: So in total, your plan is to spend at least MK 25,000 now, and MK 9,500 in the future. That is MK 34,500. Do you think you will have at least that much profit after selling your tobacco and repaying your loan?

Person 2: Yes, I think I will have about MK 40,000.

Person 1: Good. If you earn that much, then the extra money can be available immediately. Then you can commit to saving MK 9,500 for the future, and keep your other money available to spend sooner. You don't have to spend it all before your date of November 1, of course, but it will be available while you are committing to lock away MK 9,500 until then. You made three decisions: You decided how much money you needed immediately, you decided how much money to lock away for the future, and you decided when you needed to access that locked away money.

Person 2: Yes. Those weren't hard decisions. But let's demonstrate how it would work if I had chosen different options.

D.2.3 Section 3: Account Allocation Demonstration (All Clubs)

In this section, the two enumerators will work together to do a demonstration with bottle caps. You will need 12 bottle caps for this demonstration. Draw two big circles in the dirt, *and make sure everyone can see them.*

These circles represent money available for use immediately *(point at one circle)* and money committed to be saved for the future *(point at the other circle).* These bottle caps represent money. Think of each cap as MK 1,000. So, the 12 caps I have here represent MK 12,000 that someone has after selling his crop and repaying his loan.

Now, if I need MK 3,000 now and commit to saving MK 5,000 for the future, then the first MK 3,000 I earn goes in this circle, for use immediately *(put 3 bottle caps in the immediate use circle).* Then, the next MK 5,000 I earn gets locked away for the future (*put 5 bottle caps in the future circle).* Any extra money is available for use in the future, even though I don't have to spend it immediately it is not locked away (*put the remaining 8 bottle caps in the immediate use circle).*

(*Collect all of the bottle caps).* Think of this like a debt. I owe the ordinary account 3 bottle caps, and I owe the commitment account 5 bottle caps. I must pay the ordinary account first, before I pay the commitment account. Suppose I get 10 bottle caps after I sell my tobacco and repay my loans. *(Hold up 10 caps).*

First, I put 3 for immediate use. (*Put 3 caps in the immediate use circle.*) Next, I lock 5 away for use in the future. (*Put 5 caps in the future use circle.*) Then, since I've met the targets for immediate use and future use, I put all the other caps in the immediate use circle. *(Put the remaining 2 caps in the immediate use circle.*)

What if I only get 3 caps? (*Have someone come up to demonstrate. Give the person 3 caps. See where he puts them. All 3 should go in the immediate circle, and none in the future circle. If he gets this wrong, ask if anyone has a different idea. Explain if necessary.)*

(Enumerator, if farmers don't understand the demonstration you just performed, please skip back to the start of the demonstration and explain the bottle caps idea again.)

What if I get 6 caps? (*Have a volunteer come up and give him 6 caps. Correct answer: 3 in immediate, 3 in future.)*

What if you get 12 caps? (*Have another volunteer come up, etc. Correct answer: first put 3 in immediate, then 5 in future, then 4 more in immediate. Total is 7 immediate, 5 in future.*)

Dividing the bottle caps between the two circles is just like the spending plan you made before. You decide how much money you need to have available for immediate use. When you get money, it is first made available for immediate use, up to the goal you set. (*Point at the immediate use circle).* Then, you decide how much to save for the future. After making sure you have money for immediate use, you protect money for the future. (*Point at the future use circle).* Then, if there is money left after you meet both your immediate and future goals, that extra money remains available for use whenever you choose. (*Point at the* *immediate use circle).* This way, you can make a plan for how to divide your money between money you need now, and money you can commit to saving for the future, even when you don't know exactly how much you will earn.

D.2.4 Section 4: Offer of Kasupe (Ordinary) Accounts (All Clubs Except Group 0)

Person 1: We have talked a lot about how to make a budget that gives you enough money for immediate needs and commits you to saving money for the future. Also, we've discussed why saving at the bank is useful.

Person 2: Yes. I can make a plan about the amount of money I need for the short term, an amount I want to be sure to save for the future, and a date in the future when I will want that money. But how am I to use the bank?

Person 1: Usually, when you are paid for your tobacco, money is put into your group account. Then, the club officers give you your share of the cash. You leave it in the group account if you want. Or, you can save it at the bank, but to do that, you have to take your cash to the bank and deposit it into your individual account.

Person 2: Yes. It is inconvenient to have to take the money back to the bank, and often, I am tempted to spend the money as soon as I receive it.

Person 1: This season, we are offering you a new option. You can sign up to have your money transferred directly into your own Kasupe account. That means that when your bales of tobacco clear the auction floor, OBM would automatically put the money you have earned after repaying your loan into your own Kasupe account.

Person 2: How would OBM know which money was mine and which money belongs to others in my club?

Person 1: You would have to agree that OBM could get a copy of your seller sheet from Auction Holdings. OBM would use the information on the seller sheet to figure out how much money should go into your account.

Person 2: So if I agree to this, what do I have to do?

Person 1: The first thing to do is to open a Kasupe account, if you don't already have one. We can help with filling out the forms. The next thing to do is to sign a form authorizing the direct deposit. You can do both of those things today.

Person 2: That's all I have to do?

Person 1: Yes. It is very easy. If you open an account or already have one, and fill out the form for direct deposit, then your money will be put into your individual account automatically when your tobacco is sold and your loan has been recovered.

Ask the group if there are any questions about how to sign up for direct deposit.

Person 2: What if I decide I don't want to try this system and I would rather have my money go into the club account?

Person 1: You can still open a Kasupe account. Just don't fill out the [BLUE] form. Then, you will continue to get your money from the club officers, who will withdraw it from the club account for you. But if you do choose to have the money sent directly to your individual account, then ALL of your money for tobacco this season will go to the individual account. You can't change your mind part way through the season.

Person 2: Ok. I think I want the direct deposit. If I sign up for that, how do I get my cash?

Person 1: You can withdraw cash from the bank. You can either use your smart-card, or make the withdrawal by talking to a teller. You can do this at the branch or kiosk, or when the mobile bank comes to town. The closest place to make a withdrawal is ______________.

Person 2: So I can take money out whenever I want?

Person 1: Yes, you can, but you should remember the commitment you thought about to save money for a date in the future.

D.2.5 Section 5: Offer of SavePlan (Commitment) Accounts (Commitment Clubs Only)

Person 2: Is there a way that OBM can help me keep that commitment?

Person 1: Yes. You can open a special "SavePlan" account in addition to your Kasupe account.

Person 2: How would that work?

Person 1: Opening a SavePlan just tells the bank to follow the plan you made before. You will fill out a form with the three decisions you made earlier: how much money you need to have available for immediate use, the amount of money you want to lock away for the future, and the date you want that money released.

Person 2: That is easy. It's just writing down decisions I've already thought about. What happens after I fill out the form?

Person 1: Once you fill out the form, OBM will use it to put the money you are saving for the future in a special, individual, commitment account. You won't be able to take money out of that account until the date you have chosen, and you can't change your mind about the date or the amount of money.

Person 2: Do I earn interest on money in this special account?

Person 1: Yes. You earn the same interest on money in the commitment account as in the ordinary Kasupe account. The only difference is that the money in the commitment account is locked away until the date you have chosen.

Person 2: What if I earn more or less money than I thought I would have?

Person 1: It works just like the bottle caps. After the loan is recovered, money first goes into your ordinary Kasupe account, up to the amount you said you needed to have available immediately. Then, money goes to the SavePlan to be locked away for the future. When you have reached your target for saving for the future, extra money earned after that amount goes back to the ordinary Kasupe account.

Person 2: So if I don't earn as much as I thought, I will still have money available immediately?

Person 1: Yes. Money goes to the Kasupe account first, and you can withdraw from that whenever you want. It only goes to the special commitment account when you have reached your target for immediate spending.

Person 2: So this form just tells the bank to stick to the commitment I made to myself about how much to save for the future, and when I can use that money.

Person 1: That's right. You can choose any amount and date you want, and OBM will hold it for you so that you stick to the plan. We can help you fill out the form if you would like to use this special account in addition to the regular Kasupe account.

D.2.6 Section 6: Raffle (All Raffle Clubs)

As an extra incentive to save money, there will be a raffle draw where some farmers in this project may have a chance to win a prize. You have to save to have a chance to win, and the more you save, the better your chance to win. There will be two prizes in each district. The first prize will be a new bicycle, and the second prize will be a 50 kg bag of D-compound.

The raffle tickets will be based on the amount of money you save in your bank account. The prizes will be awarded in November. The raffle tickets will be given out at two times before then. The first time will be in August when we will come back and give you tickets based on the money you have saved between July 1 and August 1. OBM will calculate the average balance in your savings account for those 30 days and the number of tickets you will get will be based on this amount. The second time we hand out tickets will be in October. OBM will calculate your average balance from September 1 to October 1, and give you additional tickets based on that balance. Each person will get individual tickets based on their account balance. The prize is for individuals and not for the club.

You can increase your chance of winning by saving more money and saving it for a longer time. You will get one ticket for every MK 1000 in your average balance. If you put MK 10000 in your account by July 1 and keep it there until at least August 1, then you will get

10 tickets. If you don't have any money in your account from July 1 to July 14, and then put MK 10000 into your account on July 15 and keep it there until at least August 1, you will only get five tickets. If anyone here has two accounts with OBM, we will add up the balance in both accounts. Money saved with other banks will not count for the raffle, though.

D.2.7 Section 7A: Public Raffle (Public Raffle Clubs Only)

We will hand out the raffle tickets in August and October during group meetings like the one we are having today. We will give out the tickets in front of others, so your friends will know how many tickets you are getting.

I will demonstrate how tickets will be handed out. I am going to hand you a piece of paper with a number on it. Pretend that is your average account balance from July 1 to August 1. No one but you and OBM knows this number, so don't tell anyone!

(Distribute the papers with fake account balances to 5 volunteers)

Now, I will give you the number of raffle tickets you get for that balance. Come up one at a time and show me your piece of paper, so I can give you your tickets.

(Have the farmers come up one at a time. Look at the paper and hand out tickets. Make sure to say out loud for every farmer how many tickets he gets. Make sure that the other farmers are paying attention to this.)

When we hand out tickets in August and October, it will work the same way. You will each be called up one at a time to receive tickets based on the amount you have saved, and your club will see how many tickets you receive.

D.2.8 Section 7B: Private Raffle (Private Raffle Clubs Only)

We will hand out the raffle tickets in August and October during group meetings like the one we are having today. We will give out the tickets one at a time, so no one will know how many tickets you are getting.

I will demonstrate how tickets will be handed out. I am going to hand you a piece of paper with a number on it. Pretend that is your average account balance from July 1 to August 1. No one but you and OBM knows this number, so don't tell anyone!

(Distribute the papers with fake account balances to 5 volunteers)

Now, I will give you the number of raffle tickets you get for that balance. Come up one at a time and show me your piece of paper, so I can give you your tickets.

(Have the farmers come up one at a time. Look at the paper and hand out tickets. Make sure no one sees how many tickets you hand to each person.)

When we hand out tickets in August and October, it will work the same way. You will each be called up one at a time to receive tickets based on the amount you have saved, and no one will know how many tickets you have received.

APPENDIX E

Variable definitions

E.1 Variable definitions

Data used in this paper come from two surveys as well as from administrative records of our partner financial institution (OBM). We conducted a baseline survey from March to April 2009 and an endline survey from July to September 2010.

All variables that are created from survey data are top coded at the 99th percentile for variables with a positive range and bottom and top coded at the 1st and 99th percentile respectively for variables with a range that spans both negative and positive values. All figures in money terms are in Malawi Kwacha (MK).

Baseline characteristics (from baseline survey):

Number of members per club is the number of listed club members per information provided by the buyer companies (Alliance One and Limbe Leaf). Not all club members were interviewed.

Female equals 1 for female respondents and 0 for male respondents.

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

Age is respondent's age in years.

Years of education is the respondent's years of completed schooling.

Household size is the number of people counted as members of the respondent's household at the time of the baseline survey.

Asset index is an index based on the first principal component of the number of items owned of 14 common non-financial, non-livestock assets and indicators of presence of 4 major types of housing characteristics (iron sheet roof, glass windows, concrete floor, electricity connection).

Livestock index is an index based on the first principal component of the number of animals owned of 7 common types of livestock.

Land under cultivation is the total of area of land under cultivation, measured in acres, for the late-2008 planting season.

Proceeds from crop sales is the sum of sales from the two main cash crops, maize and tobacco, in the 2008 harvest.

Cash spent on inputs is the total amount of cash spent - excluding the value of input packages that are part of a loan – on seeds, fertilizer, pesticides, and hired labor for the 2008-2009 planting season

Has bank account is 1 if a household member has an account with a formal financial institution, and 0 if not.

Savings in accounts and cash is the sum of current savings with formal institutions and in cash at home.

Hyperbolic is 1 if the respondent exhibited strictly more patience in one month, hypothetical monetary trade-offs set 12 months in the future than in the same trade-offs set in the present, and 0 otherwise. See section 5 above for more details.

Patient now, impatient later is 1 if the respondent exhibited strictly *less* patience in one month, hypothetical monetary trade-offs set 12 months in the future than in the same trade-offs set in the presence and 0 otherwise.

Net transfers made in past 12m is the total of transfers made to the social network minus the sum of transfers received from the social network, summed across six categories (social events, health shocks, education of children, agricultural inputs, hired labor and 'other').

Missing value for formal savings and cash is 1 if the variable "*Savings in accounts and cash*" is missing and 0 if it has valid values.

Missing value for time preferences is 1 if the respondent has missing values for the time preferences variables ("*Hyperbolic*" and "*Patient now, impatient later*") is missing, and 0 if these variables have valid values.

Transactions with Partner Institution (from internal records of OBM):

Any transfer via direct deposit is 1 if the respondent receives any deposit from his or her tobacco club's account to his or her individual savings account, and 0 if not.

Deposits into ordinary accounts, pre-planting is the sum of (positive) transactions into the respondent's OBM ordinary savings accounts during the period of March to October 2009.

Deposits into commitment accounts, pre-planting is the sum of (positive) transactions into the respondent's OBM commitment savings accounts during the period of March to October 2009.

Deposits into other accounts, pre-planting is the sum of (positive) transactions into the respondent's OBM non-ordinary, non-commitment savings accounts during the period of March to October 2009.

Total deposits into accounts, pre-planting is the sum of transactions into the respondent's OBM accounts (sum across all accounts) during the period of March to October 2009.

Total withdrawals from accounts, pre-planting is the sum of transactions out of the respondent's OBM accounts (sum across all accounts) during the period of March to October 2009.

Net deposits, pre-planting is the difference between all deposits and withdrawals in the respondent's OBM accounts during the period of March to October 2009.

Net deposits, Nov-Dec is the difference between all deposits and withdrawals in the respondent's OBM accounts during the period of November and December 2009.

Net deposits, Jan-Apr is the difference between all deposits and withdrawals in the respondent's OBM accounts during the period of January through April 2010.

Construction of deposit spells used to calculate the share of deposits withdrawn on same day and the share of initial deposit amount remaining in account after two weeks

For these calculations we only consider deposits into ordinary accounts that are greater than MK 500 to avoid small positive transactions like interest payments to count as deposits. A deposit is considered fully withdrawn when the cumulative net transactions are within MK 400 of the initial deposit or 99% has been withdrawn, whichever is greater. This is to avoid considering deposits not withdrawn for a long time when respondents left a very small amount in the account (absolute or relative to the initial deposit). However, the calculations of the share of deposits withdrawn on same day and the average share of initial deposit amount remaining in accounts after two weeks are robust to decreasing these " buffer" amounts. The deposit and withdrawal " spells" are coded as non-overlapping: as long as the initial deposit is not withdrawn the spell is considered active. That means when another deposit is made before the initial deposit was fully withdrawn the second deposits is added to the cumulative net transactions, i.e. reduces the amount considered withdrawn. Only spells with initial deposits after March 1, 2009 are considered. Spells with initial deposits that are not counted as fully withdrawn by August 31, 2010 are set to end on that date.

Agricultural outcomes, household expenditure, and other variables, from endline survey (all planting and harvest variables refer to the 2009-2010 planting season):

Land under cultivation is the total area of land under cultivation, measured in acres.

Cash spent on inputs is the total amount of cash spent - excluding the value of input packages that are part of a loan - on seeds, fertilizer, pesticides, and hired labor for the 2009-2010 planting season.

Total value of inputs is the sum of cash spent on agricultural inputs plus the value of inputs included in-kind in loan packages for the 2009-2010 planting season. Input categories include seeds, pesticides, fertilizer, hired labor, transport and firewood (for curing tobacco).

Proceeds from crop sales is the sum of sales from the two main crops, maize and tobacco for the 2009-10 planting season.

Value of crop output (sold & not sold) is the sum of revenue from crop sales and the value of the unsold crop for seven main crops (maize, burley tobacco, dark fire tobacco, flue-cured tobacco, ground nuts, beans, soya). Value of harvest not sold equals the kilograms of crops not sold multiplied by the price/kilogram, summed across the seven main crops. Price/kilogram for each crop is obtained by calculating crop-specific revenue/kilogram for each observation in the sample and then taking the sample average.

Farm profit (output - input) is the difference between "Value of crop output" and "Total value of inputs" defined above.

Total expenditure in last 30 days is the sum of three categories household expenditures (food, non-food household items and transport) over the last 30 days prior to the endline survey.

Household size is the number of people counted as members of the respondent's household at the time of the endline survey.

Total transfers made is the total of transfers made to the social network over the 12 months prior to the endline interview, summed across six categories (social events, health shocks, education of children, agricultural inputs, hired labor and 'other').

Total transfers received is the total of transfers received from the social network over the 12 months prior to the endline interview, summed across six categories (social events, health shocks, education of children, agricultural inputs, hired labor and 'other').

Total net transfers made is the difference between "Total transfers made" and "Total transfers received" defined above.

Tobacco club loan is the total amount owed as part of a tobacco club loan for the 2009-2010 planting season.

Not interviewed in endline is 1 if the respondent was not interviewed and is 0 if the respondent was interviewed during the endline.

BIBLIOGRAPHY

Angeletos, George-Marios, and Laurent-Emmanuel Calvet. 2006. "Idiosyncratic production risk, growth and the business cycle." *Journal of Monetary Economics*, 53(6): 1095–1115.

Angelucci, Manuela, Dean S. Karlan, and Jonathan Zinman. 2013. "Win Some Lose Some? Evidence from a Randomized Microcredit Program Placement Experiment by Compartamos Banco." Center for Global Development Working Papers 330.

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association*, 91(434): pp. 444–455.

Aportela, F. 1999. "Effects of Financial Access on Savings by Low-Income People." Banco de Mexico mimeo.

Armendáriz, Beatriz, and Jonathan Morduch. 2005. *The economics of microfinance.* MIT Press, Cambridge, Mass.

Ashraf, N., D. Karlan, and W. Yin. 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines*." *Quarterly Journal of Economics*, 121(2): 635–672.

Banerjee, A. V, and S. Mullainathan. 2013. "The shape of temptation: Implications for the economic lives of the poor." *NBER Working Paper*.

Barberis, Nicholas, and Ming Huang. 2008. "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices." *American Economic Review*, 98(5): 2066–2100.

Barberis, Nicholas C. 2013. "Thirty Years of Prospect Theory in Economics: A Review and Assessment." *Journal of Economic Perspectives*, 27(1): 173–96.

Beegle, Kathleen, Emanuela Galasso, and Jessica Goldberg. 2014. "The design of public works and the competing goals of investment and food security." Extended Abstract prepared for Population Association of America Annual Meetings.

Benson, T. 1999. "Area-specific fertilizer recommendations for hybrid maize grown by Malawian smallholders: A manual for field assistants." Maize Commodity Team, Chitedze Agricultural Research Station, Malawi.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice Under Risk." *The Quarterly Journal of Economics*, 127(3): 1243–1285.

Bruhn, Miriam, and David McKenzie. 2009. "In Pursuit of Balance: Randomization in Practice in Development Field Experiments." *American Economic Journal: Applied Economics*, 1(4): 200–232.

Bune, Lasse, Xavier Giné, Jessica Goldberg, and Dean Yang. 2014. "Facilitating Savings for Agriculture." University of Michigan mimeo.

Burgess, Robin, and Rohini Pande. 2005. "Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment." *The American Economic Review*, 95(3): pp. 780–795.

Collins, D., J. Morduch, S. Rutherford, and O. Ruthven. 2009. *Portfolios of the Poor, How the World's Poor Live on $2 a Day.* New Jersey: Princeton University Press.

Conlisk, John. 1993. "The utility of gambling." *Journal of Risk and Uncertainty*, 6(3): 255–275.

Daley-Harris, S. 2009. "State of the Microcredit Summit Campaign Report 2009." Washington, DC.

Deaton, A. 1990. "Saving in Developing Countries: Theory and Review."

de Mel, Suresh, David McKenzie, and Christopher M. Woodruff. 2012. "Business Training and Female Enterprise Start-Up, Growth, and Dynamics: Experimental Evidence from Sri Lanka." Social Science Research Network SSRN Scholarly Paper ID 2161233, Rochester, NY.

de Mel, Suresh, David McKenzie, and Christopher Woodruff. 2008. "Returns to Capital in Microenterprises: Evidence from a Field Experiment." *The Quarterly Journal of Economics*, 123(4): 1329–1372.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review*, 101(6): 2350–90.

Dupas, Pascaline, and Jonathan Robinson. 2013*a*. "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya." *American Economic Journal: Applied Economics*, 5(1): 163–92.

Dupas, Pascaline, and Jonathan Robinson. 2013*b*. "Why Don't the Poor Save More? Evidence from Health Savings Experiments." *American Economic Review*, 103(4): 1138–1171.

Erev, Ido, Eyal Ert, Alvin E. Roth, Ernan Haruvy, Stefan M. Herzog, Robin Hau, Ralph Hertwig, Terrence Stewart, Robert West, and Christian Lebiere. 2010. "A choice prediction competition: Choices from experience and from description." *Journal of Behavioral Decision Making*, 23(1): 15–47.

Feldman, Jack M., and John G. Lynch. 1988. "Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior." *Journal of Applied Psychology*, 73(3): 421 – 435.

Fudenberg, Drew, and David K. Levine. 2006. "A dual-self model of impulse control." *The American Economic Review*, 96(5): 1449–1476.

Gajic, Aleksandra, David Cameron, and Jeremiah Hurley. 2012. "The cost-effectiveness of cash versus lottery incentives for a web-based, stated-preference community survey." *The European Journal of Health Economics*, 13(6): 789–799.

Giné, Xavier, Jessica Goldberg, and Dean Yang. 2012. "Credit Market Consequences of Improved Personal Identification: Field Experimental Evidence from Malawi." *American Economic Review*, 102(6): 2923–54.

Giné, Xavier, Jessica Goldberg, Dan Silverman, and Dean Yang. 2012. "Revising commitments: Field evidence on the adjustment of prior choices." National Bureau of Economic Research Working Paper.

Goldberg, Jessica. 2011. "The Lesser of Two Evils: The Roles of Impatience and Selfishness in Consumption Decisions." University of Maryland mimeo.

Guillen, Mauro, and Adrian Tschoegl. 2002. "Banking on Gambling: Banks and Lottery-Linked Deposit Accounts." *Journal of Financial Services Research*, 21(3): 219–231.

Haisley, Emily, Romel Mostafa, and George Loewenstein. 2008. "Subjective relative income and lottery ticket purchases." *Journal of Behavioral Decision Making*, 21(3): 283–295.

Halpern, Scott D., Rachel Kohn, Aaron Dornbrand-Lo, Thomas Metkus, David A. Asch, and Kevin G. Volpp. 2011. "Lottery-Based versus Fixed Incentives to Increase Clinicians' Response to Surveys." *Health Services Research*, 46(5): 1663–1674.

Haushofer, Johannes, and Jeremy Shapiro. 2013. "Household Response to Income Changes: Evidence from an Unconditional Cash Transfer Program in Kenya." Massachussetts Institute of Technology Working Paper, Cambridge, MA.

Hertwig, Ralph, and Ido Erev. 2009. "The description–experience gap in risky choice." *Trends in Cognitive Sciences*, 13(12): 517 – 523.

Jones, Gareth, Richard W. Steketee, Robert E. Black, Zulfiqar A. Bhutta, and Saul S. Morris. 2003. "How many child deaths can we prevent this year?" *The lancet*, 362(9377): 65–71.

Kachelmeier, Steven J, and Mohamed Shehata. 1992. "Examining Risk Preferences under High Monetary Incentives: Experimental Evidence from the People's Republic of China." *American Economic Review*, 82(5): 1120–41.

Kahneman, Daniel, and Amos Tversky. 1979*a*. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2): 263–91.

Kahneman, Daniel, and Amos Tversky. 1979*b*. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2): pp. 263–292.

Kahneman, Daniel, and Amos Tversky. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5(4): 297–323.

Karlan, Dean, Aishwarya Ratan, and Jonathan Zinman. 2014. "Savings by and for the Poor: A Research Review and Agenda." *Review of Income and Wealth*, 60(1): 36–78.

Kazianga, Harounan, and Christopher Udry. 2006. "Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso." *Journal of Development Economics*, 79(2): 413–446.

Kwang, Ng Yew. 1965. "Why do People Buy Lottery Tickets? Choices Involving Risk and the Indivisibility of Expenditure." *Journal of Political Economy*, 73: 530.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *The Quarterly Journal of Economics*, 112(2): 443–478.

Loomes, Graham, and Robert Sugden. 1982. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty." *Economic Journal*, 92(368): 805–24.

Morduch, Jonathan. 1999. "The Microfinance Promise." *Journal of Economic Literature*, 37(4): 1569–1614.

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

Ozdenoren, Emre, Stephen W. Salant, and Dan Silverman. 2012. "Willpower and the Optimal Control of Visceral Urges." *Journal of the European Economic Association*, 10(2): 342–368.

Pearce, B., and P. Denton. 2011. "2011-2012 Kentucky & Tennessee Tobacco Production Guide." , ed. G. Schwab, K. Seebold and B. Pearce, Chapter Field Selection, Tillage, and Fertilization, 24–29. University of Tennessee and University of Kentucky.

Prina, Silvia. 2013. "Banking the Poor Via Savings Accounts: Evidence from a Field Experiment." Case Western Reserve University mimeo.

Quiggin, John. 1991. "On the Optimal Design of Lotteries." *Economica*, 58(229): 1–16.

Robinson, M. 2001. *The Microfinance Revolution: Sustainable Finance for the Poor.* World Bank Publications.

Shapiro, Jesse M. 2005. "Is there a daily discount rate? Evidence from the food stamp nutrition cycle." *Journal of Public Economics*, 89(2–3): 303–325.

Singer, Eleanor, and Cong Ye. 2013. "The Use and Effects of Incentives in Surveys." *The ANNALS of the American Academy of Political and Social Science*, 645(1): 112–141.

Snowberg, Erik, and Justin Wolfers. 2010. "Explaining the Favorite-Long Shot Bias: Is it Risk-Love or Misperceptions?" *Journal of Political Economy*, 118(4): 723–746.

Stephens Jr., M. 2003. "'3rd of tha Month': Do Social Security Recipients Smooth Consumption Between Checks?" *American Economic Review*, 93(1).

Thaler, Richard H. 1990. "Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives*, 4(1): 193–205.

Volpp, Kevin, George Loewenstein, Andrea Troxel, Jalpa Doshi, Maureen Price, Mitchell Laskin, and Stephen Kimmel. 2008*a*. "A test of financial incentives to improve warfarin adherence." *BMC Health Services Research*, 8(1): 272.

Volpp, Kevin, LK John, Andrea Troxel, L Norton, J Fassbender, and George Loewenstein. 2008*b*. "Financial incentive–based approaches for weight loss: A randomized trial." *JAMA*, 300(22): 2631–2637.

Webb, Thomas L., and Paschal Sheeran. 2006. "Does changing behavioral intentions engender behavior change? A meta-analysis of the experimental evidence." *Psychological Bulletin*, 132(2): 249 – 268.

Wooldridge, Jeffrey M. 2001. *Econometric Analysis of Cross Section and Panel Data.* Vol. 1 of *MIT Press Books*, The MIT Press.

World Bank. 2008. "World Development Report: Agriculture for Development."

Zabojnik, Jan. 2002. "The Employment Contract as a Lottery."

Zwane, Alix Peterson, Jonathan Zinman, Eric Van Dusen, William Pariente, Clair Null, Edward Miguel, Michael Kremer, Dean S. Karlan, Richard Hornbeck, Xavier Giné, Esther Duflo, Florencia Devoto, Bruno Crepon, and Abhijit Banerjee. 2011. "Being surveyed can change later behavior and related parameter estimates." *Proceedings of the National Academy of Sciences.*