Saving Behavior and Social Norms in Malawi: Evidence from a Field Experiment

Jessica Goldberg *

March 2009

Preliminary and Incomplete. Please do not quote.

Abstract

Reasons for under-saving in poor agrarian communities are manifold, and individuals sometimes exhibit high marginal propensities to consume that are suggestive of hyperbolic discounting. I test whether informal insurance networks or social norms about savings can explain low savings and a tendency to spend income rapidly. The mechanism through which informal insurance networks and social pressure may affect consumption and savings is through public information about individuals' assets. To test the extent to which informal insurance networks or the social norms that develop from them affect the timing of consumption and level of savings, I employ a field experiment to distinguish between the use of windfall money when receipt of the money is known to others in the community versus when it is private information. I run two lotteries in each of 158 agriculture clubs in central Malawi. In each club, one lottery and its winner are publicly announced to the whole group. The other lottery is private, and only its winner knows that a second lottery was held. I measure differences in expected use of the windfall income between public and private lottery winners. I find that public winners spend more money in the period immediately following the lotteries than private winners spend in the same time frame. In communities where income is easily observed and subject to social sharing norms, this pattern could lead to behavior that appears to result from very high discount rates or hyperbolic discounting but is instead the rational response of individuals with standard discount rates to an unobserved constraint in consumption.

*Jessica Goldberg, Department of Economics and Ford School of Public Policy, University of Michigan, jegoldbe@umich.edu.

1 Introduction

For smallholder farmers in Malawi, income is not distributed evenly throughout the year, but rather concentrated at harvest time (May-August, depending on the crop). Expenditures on some household necessities obviously take place throughout the year, though luxuries are purchased following the harvest. For many households, the largest annual expenditures are those for farming inputs. These expenditures typically occur in October or November, several months after harvest.

Despite the lag between receipt of income and major purchases, farmers appear to save little and spend their income quickly after receiving it. In a baseline survey of participants in an agricultural micro-lending program, only 18 percent of respondents reported saving money to use for the subsequent season's farming inputs, though larger fractions save for precautionary reasons or for the purchase of household goods throughout the year (BFIRM baseline survey, October 2007). Anecdotally, many farmers rely on loans to finance their use of fertilizer, pesticide, and other inputs. Interest rates on loans are between 30 and 33 percent per year, so the opportunity cost of relying on loans to finance inputs year after year is substantial. Additionally, in a classification scheme initially devised by Keynes (1936), savings are used to develop private business enterprise. Individual under-saving, then, can lead to underdevelopment and can suppress individual income in the long run.

Reasons for under-saving in poor agrarian communities are manifold. They include binding sustenance constraints, frequent negative shocks, lack of access to financial services, and high discount factors. Here, I focus on two related reasons for reduced savings rates: informal insurance networks, and social pressure to share.

In their most extreme form, informal insurance networks involve collecting from every individual in each period any income in excess of the individual's permanent income, and reallocating that surplus to others in the network. There is no room for private savings in this structure. More moderate informal insurance networks may allow for some private savings, but will lead to private savings at lower rates than without the insurance networks.

Social pressure to share, a phenomena documented by anthropologists and frequently cited by Malawian farmers in focus groups, may simply be the institutionalization of informal insurance networks. Nonetheless, it constitutes a high marginal tax rate on income for

individuals who are unable to deny a relative's or neighbor's request for money. Consequentially, social norms may contribute to under-savings and lead to behavior that suggests a high discount rate, when in fact individuals are responding rationally to standard discount rates but face constraints not typically modeled by economists.

While there is substantial research about the extent of informal insurance networks in developing countries, there is limited evidence of how these networks affect consumption and savings behavior. To test the extent to which these networks or the social norms that develop from them affect savings, I employ a field experiment to distinguish between the use of windfall money when receipt of the money is known to others in the community versus when it is private information. I run two lotteries in 158 agriculture clubs in central Malawi. One lottery, and its winner, is publicly announced to the whole group. The other lottery is private, and only its winner knows that a second lottery was held. This experimental variation in information about cash assets allows me to distinguish between the private, first-best use of unexpected income, and the second-best use of such income when constrained by obligations to other members of the community.

The paper proceeds as follows. In Section 2, I briefly discuss the related literature. In Section 3, I derive predictions about savings from a simple optimization model. In Section 4, I describe my field experiment in detail. In Section 5, I discuss the data. In Section 6, I present results. I discuss specification checks in Section 7, and conclude in Section 8.

2 Literature Review

Four categories of literature are relevant to my question of how informal insurance networks or social norms affect savings. The first is the literature about informal insurance networks themselves. Townsend (1994) tests a full-insurance model in rural India. His focus is on consumption smoothing rather than savings or investment; indeed, he does not even report summary statistics for savings or agricultural inputs. However, his finding that informal insurance networks contribute to consumption smoothing suggests that at least some fraction of money that might otherwise be saved is diverted to these networks. Ligon et al. (2002) return to Townsend's model in a setting where commitments to informal networks cannot be fully enforced. They find that limited commitment explains much of the deviation from

the full-insurance model. An important conclusion in its own right, this also suggests that obligations to the informal insurance network constrain private allocations of income.

The second category of literature is research by anthropologists on social norms or pressure to share income. Social norms may simply be the institutionalization of obligations to informal insurance networks. The pressure to share cash with family or neighbors is considered pervasive throughout Africa:

> Of course people do not always spend all their money as soon as they receive it. However, a major factor in people's use of money is the expectation that friends and relatives will ask to "borrow" from them. This is doubly true as people are keenly aware of each other's business. Almost always if someone receives money, those people who are socially close will know it. Yet certain monies must be reserved for essential expenditures, like food and shelter. Short-term reserves for essentials are allowed within the system, but even then the expenditures need to be made without delay. If this is not done, the cash will appear to be available for borrowing (Maranz, p. 18).

In response to the social pressure to share cash, Maranz claims that people over-consume and under-save: "Consequently, on those infrequent occasions when they were able to earn money, they often made wasteful or ill-considered expenditures just to keep friends from borrowing it" (Maranz, p. 18).

Economists have documented the impact of social pressure on financial decision-making. Baland et al. (2007) interview borrowers in Cameroon who borrow money despite having savings of equal or greater value than their loans, thus incurring unnecessary interest costs for projects that could be self-financed. They report, "Excess borrowing is purposefully used by some members to signal financial difficulties to their relatives in search of financial help. Reimbursement obligations are then used as an argument to discourage such demands" (Baland et al., p. 9).

The third category of literature I consider is the growing work on behavioral or experimental economics. DellaVigna (2007) sorts behavioral economics research into three categories of deviations from standard economic assumptions: non-standard preferences, incorrect beliefs, and systematic biases in decision making. Research about social norms, which falls under the

category of systemic biases or non-standard decision making, has mostly focused on workplace productivity (Falk and Ichino 2006, Mas and Moretti 2006). That line of research demonstrates that people conform to the expectations of those around them even when conforming may have economic costs. Ashraf (2006) uses a field experiment to show that when their financial decisions are observable to their spouses, married Filipino men are more likely to share income with their wives than to keep it for themselves.

The literature on informal networks with limited commitment and on social norms suggests that whether or not others know about an individual's income may be a key factor in how the individual uses the income. Specifically, it may be harder or less profitable to save for future agricultural inputs if others in your community know that you have cash and feel entitled to a portion of it. Ashraf (2006) demonstrates a simple experimental design where changing the individual's choice from a private to a public action reveals information about financial decision-making.

The last group of papers I reference are those about time preferences for consumption. Laibson (1997) develops a model of hyperbolic discounting, where individuals treat the present differently from all future periods. This model rationalizes immediately consuming a higher fraction of income than would be chosen in a model with standard discount rates. I will present an alternative model that leads to observationally equivalent consumption patterns but includes standard discount rates with a tax on income that can be evaded through rapid consumption.

3 A Simple Model

Laibson's (1997) model of hyperbolic discounting explains a high marginal propensity to consume present income by applying an additional discount factor to future periods. In a two-period model, consuming windfall income in the initial period, rather than spreading the windfall income between the two periods, would be behavior consistent with hyperbolic discounting. Here, I develop a model in which the same behavior is explained by pressure to share, which is an economic constraint, rather than the additional discount factor, which is a behavioral parameter.

In each period, a individual's wealth consists of both savings from the previous period

and income earned that period. Both savings and income can be either public or private – that is, known either to the whole community or known only to the individual. Let w denote wealth, s savings, c consumption and y income. The superscript u will denote that it is public, and v that it is private. The subscript t indexes time periods. We have thus

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

In addition, if individuals are known to have extra cash – that is, if $c_t^u < w_t^u$ – they are taxed at rate τ on this surplus cash. This is consistent with Maranz's notion that cash is seen as "available" to others in the social network, while goods are not. Note that individuals can evade taxes by consuming all of their publicly-observed income in the period in which it is received. This immediate consumption converts income from cash, which is taxable, to goods, which are not.

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

First if $c_0 < w_0^u$ then individuals are taxed on the portion of first period income that is not consumed in the first period:

$$\begin{aligned} w_1^u &= \theta(w_0^u - c_0) + y^u \\ w_1^v &= w_0^v + y^v \\ c_1 &= \theta(w_0^u - c_0) + y^u + w_0^u + y^v \\ &= y^u + y^v + \theta w_0^u + w_0^v - \theta c_0 \end{aligned} \tag{1}$$

Second if $c_0 \geq w_0^u$ then individuals evade taxes:

$$\begin{aligned} w_1^u &= y^u \\ w_1^v &= (w_0^v + w_0^v - c_0) + y^v \\ c_1 &= y^u + (w_0^u + w_0^v - c_0) + y^v \\ &= y^u + y^v + w_0^u + w_0^v - c_0 \end{aligned} \tag{2}$$

In addition, I assume that there is no borrowing so $c_0 \leq w_0^u + w_0^v$. The set of possible combinations of c_0 and c_1 is shown in figure 1. The blue line corresponds to equation 1 and

the green to equation 2. The solid portion of the line represents combinations of c_0 and c_1 that are actually possible.

Figure 1: Possible combinations of c_0 and c_1

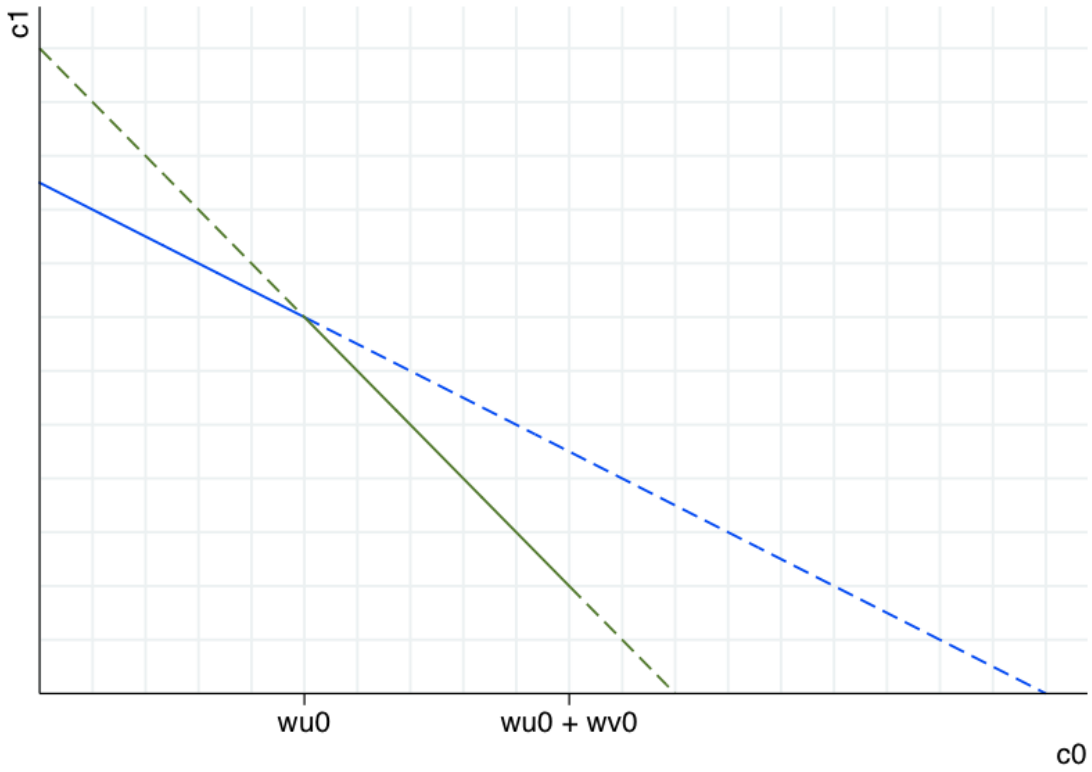

An individual wants to maximize utility, $u(c_0, c_1) = \ln c_0 + \beta \ln c_1$. Maximizing utility such that equation 1 holds, I find the point (call it c_0^*) where an indifference curve of the utility function is tangent to the blue line. If I maximize such that equation 2 holds, I find the point (c_0^{**}) where an indifference curve is tangent to the green line.

There are essentially four different cases as illustrated in figure 2.

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

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

Figure 2: Different Cases

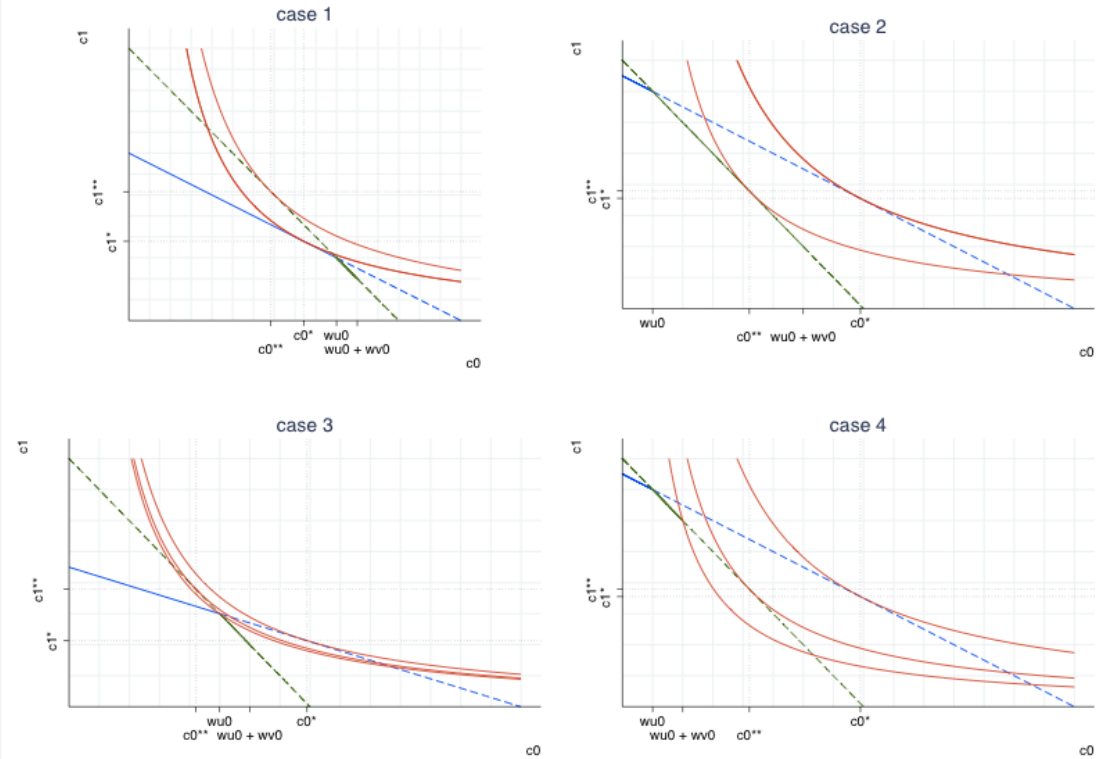

Case 3 $c_0^* > w_0^u$ and $c_0^{**} < w_0^u$. Thus $c_0 = w_0^u$ is the optimal point allowed.

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

Next, I maximize utility in the tax-paying case, such that equation 1 holds to find an expression for c_0^*.

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

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

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

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

I sustitute into the budget constraint (equation 1 to find c_1^*:

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

Next, I maximize utility in the tax-evading scenario, such that equation 2 holds to find an expression for c_0^{**}.

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

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

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

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

Again, I substitute this solution into the budget constraint (equation 2) to find c_1^{**}:

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

I am interested in $\partial c_0/\partial w_0^u$ and $\partial c_0/\partial w_0^v$.

	Case 1	Case 2	Case 3	Case 4
$c_0 =$	$\frac{1}{\theta(1+\beta)}\big(y^u + y^v + \theta w_0^u + w_0^v\big)$	$\frac{1}{1+\beta}\big(y^u + y^v + w_0^u + w_0^v\big)$	w_0^u	$w_0^u + w_0^v$
$\frac{\partial c_0}{\partial w_0^u} =$	$\frac{1}{1+\beta}$	$\frac{1}{1+\beta}$	1	1
$\frac{\partial c_0}{\partial w_0^v} =$	$\frac{1}{\theta(1+\beta)}$	$\frac{1}{1+\beta}$	0	1
$c_1 =$	$\frac{\beta}{(1+\beta)}\big(y^u + y^v + \theta w_0^u + w_0^v\big)$	$\frac{\beta}{(1+\beta)}\big(y^u + y^v + w_0^u + w_0^v\big)$	$y^u + y^v + w_0^v$	$y^u + y^v$
$\frac{\partial c_1}{\partial w_0^u} =$	$\frac{\beta\theta}{1+\beta}$	$\frac{\beta}{1+\beta}$	0	0
$\frac{\partial c_1}{\partial w_0^v} =$	$\frac{\beta}{1+\beta}$	$\frac{\beta}{1+\beta}$	1	0

In case 1, the income effect dominates the substitution effect, and a shock to first period private income increases consumption by more than a shock to first period public income. In case 2, the tax on unconsumed public income is not binding because individuals want to consume more than their public income in the first period. The effects of public and private

shocks on first period consumption are symmetric. In case 3, the tax binds: individuals would like to save some of their public income, but the substitution effect dominates the income effect and forces them to the corner solution. Here, first period consumption increases one-for-one with first period public income, while first period private income has no effect on first period consumption. In case 4, the no borrowing constraint binds and individuals can consume only as much as they receive in the first period.

Case 3 is plausible if a large share of income is publicly observed and subject to a high tax rate. It would lead to high marginal propensities to consume, behavior that is observationally equivalent to hyperbolic discounting. However, the consumption patterns would be explained by a tax on holding income that others know about rather than a non-standard discount rate. In other words, case 3 this model provides an alternate explanation for behavior that may appear to be hyperbolic discounting, for individuals facing a high tax on holding cash.

I conduct lotteries in clubs of farmers in Malawi to generate exogenous shocks to public and private income and to test whether consumption of public income is consistent with the patterns predicted by case 3 of this model. If the model holds, I expect to see that a shock to public income is consumed immediately while a shock to private income is not.

4 Experimental Design

Enforcing a tax on savings or assets requires public information about the level of savings or assets, even when the "tax" is a contribution to an informal insurance network. To test that implication of the simple optimization model above and determine the causal effect of public information about individuals choices on their financial decision-making, I study the allocation of windfall income under different information conditions. In principle, my experiment is similar to Ashraf (2006). While Ashraf focuses on within-household bargaining, I am interested in community-level dynamics.

Maranz's description of social pressure to share turns on the assumption that "if someone receives money, those people who are socially close will know it." It is exactly that assumption that I manipulate in a simple field experiment involving 1,553 farmers in central Malawi. I run 316 lotteries in 158 agricultural clubs in central Malawi in May 2008. These clubs of approximately 10 members each were formed in late 2007 for the purpose of receiving exten-

sion services and borrowing through group liability schemes, and the lotteries are conducted when the clubs are assembled for meetings and surveys related to their loans for the coming agricultural season. The lotteries are facilitated by trained, Chichewa-speaking enumerators.

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

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

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

5 Data

My final sample is of 315[1] lottery winners, half of whom won in "public" settings and the other half of whom won under "private" conditions. I have data from four surveys. Surveys were administered on two occasions, in May and August 2008, at a central meeting location were all group members were asked to gather. Chichewa-speaking enumerators conducted one-on-one interviews with respondents.

In each club, the May surveys were conducted on the same day as the lotteries. The baseline survey was administered to all 1,553 members of groups where lotteries were conducted. The supplemental lottery survey asked the 315 winners to list the ways in which they would use their prize money, and then indicate when each transaction would take place and who would be the beneficiary. In August, a subset of 81 participating clubs were revisited[2]. Some 627 members of those clubs were present and were administered a follow-up questionnaire about assets, savings, and the recent harvest. At least one of the lottery winners was present in 77 of those clubs; in total, 116 lottery winners were administered a supplemental survey asking how they had actually used their prize money.

Table 1 presents summary statistics for baseline characteristics of the public and private lottery winners. Public and private winners do not differ significantly in their gender, age, or years of education. The apparent difference in land ownership is due to outliers and becomes insignificant when trimming the top one percent of land holdings. Including or excluding these baseline characteristics does not affect the sign or significance of subsequent results. However, public and private winners do differ substantially in their likelihood of being resurveyed in August. I examine this apparently selective attrition in Table 2.

Not all clubs were resurveyed in August, but since each club has one private and one public winner, equal response rates for the two types of lottery winners were desirable. Public lottery winners (41.4%) were about as likely as all respondents (40.4%) to appear in the August sample. Only 32.3 percent of private lottery winners were resurveyed, however. Among lottery winners, those who did and did not respond to the August survey were about equal in gender, age, and land owned. However, those who did not respond have significantly fewer years of education than those who did. It is possible that the more educated face a higher

[1] One individual who won a private lottery declined the prize.
[2] Budget constraints precluded revisiting all clubs.

opportunity cost of time, though unlikely because the sample consists entirely of farmers who do not do regular wage labor and the August survey took place in the lull period between the harvest and the next season's planting.

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

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

My analysis of how prize money is used divides spending into five categories: consumption by the winner, consumption by others in the winner's household, consumption by persons not in the winner's household, investment, and savings. These categories are mutually exclusive and exhaustive. I include purchase of agricultural inputs such as fertilizer and pesticides, purchase of livestock, and purchase of building materials in the "investment" category. Results for analysis of these categories are not sensitive to alternative definitions of investment, such as removing livestock or adding transfers to other people (which in principle might be repaid in the future).

I have supplemental data about savings in the bank and at home. These data come from the main survey administered to all club members in August. The savings data pertain to

savings from all sources, not just the lottery prize.

6 Results

Timing

The most basic question addressed by the experiment is whether winners of public and private lotteries spend their prize money differently. Because public and private winners are randomly selected, any differences in their behavior can be attributed to the type of lottery they won. In all of my results, including covariates from the background survey to soak up residual variation does not change the results. Furthermore, having two winners per club allows me to include club fixed effects; however, including club fixed effects does not change any of my results. For simplicity, I discuss results without fixed effects or additional covariates except as noted; all standard errors are clustered at the club level.

The model predicts that if there is a binding tax on public income – that is, that people would prefer to consume less than the full value of their prize immediately – then winners of public lotteries will spend their prize immediately, while winners of public lotteries will not. In my data, this is equivalent to public winners spending all MK 2,500 of their prize immediately, and private winners spending none of their prize immediately. While this strong prediction is not supported by the results, public winners are marginally more likely to spend all of their prize immediately than private winners, with 42 percent of public winners compared to 31 percent of private winners spending all of the prize immediately. Differences in the fraction spending none of the prize immediately are much smaller: 31 percent of public winners and 33 percent of private winners do not spend any of their prize immediately.

Despite not finding exactly $\frac{\partial c_0}{\partial w_0^u} = 1$ and $\frac{\partial c_0}{\partial w_0^v} = 0$, there are reasons to believe that my results are consistent with case 3 of the model. First, different individuals are almost certainly in different cases of the model, since individuals with different endowments y^u and y^v may find themselves in different cases even if they have the same discount rates. Testing whether all public winners spend their entire prize immediately, and all private winners don't spend any of the prize immediately, tests whether *all* individuals are behaving consistently with case 3 of the model. One way to interpret the results in table 3 is that 41 percent of public winners and 33 percent of private winners behave precisely as predicted by case 3 of the model. Second,

while the budget constraint in the model is continuous, it is not continuously differentiable. This precludes finding an intermediate solution where $0 \leq \frac{\partial c_0}{\partial w_0^v} \leq \frac{\partial c_0}{\partial w_0^u}$. However, case 3 is logically consistent with the notion that winners of public lotteries will spend money more quickly in order to evade the social tax. As illustrated in Figure 3, this is exactly the behavior observed in the sample. Public and private winners spend approximately the same amount on the day of the lottery (probably because the exercise itself took most of the day and left few opportunities to spend money that day), but spending for public winners out paces that for private winners for more several weeks following the lotteries before eventually converging.

Table 3 shows the timing of use of the prize money for public and private winners. The first thing to note is that both public and private winners anticipate using and actually use a large fraction of their prize very quickly after receiving the money. Column 1 uses data on anticipations for the full sample of 315 lottery winners, and shows that public winners anticipated using MK 2053 within one week of the lottery, while private winners anticipated using MK 1786 in the same time period. These figures represent 82 percent and 71 percent of the prize, respectively. There are many possible explanations for spending a high fraction of income very quickly. One possibility is that people are cash constrained, especially as the experiment was conducted before the harvest. Another possibility is that people are impatient. Other explanations have to do with the experiment itself: for example, winners might believe that their prize could be taken away from them. However, none of these explanations affect individuals who won public lotteries differently than those who won private lotteries, and thus, they cannot explain the significant difference in the expected spending of individuals who won public and private lotteries.

The difference in anticipated timing of spending persists when limiting the sample to only the 116 individuals who were administered the follow up lottery survey in August. The magnitudes of expected spending for this limited sample, shown in column 2, are very similar to those in the full sample in column 1, though the standard errors are almost two times as large. The stability of this result is reassuring in light of the differential attrition between public and private lottery winners.

Column 3 shows results about realized spending, from the August survey. Winners of public lotteries report that they spent MK 1,414 of their prize within a week of receiving it. Winners of private lotteries spent MK 1,073 in the same period. The difference is statistically

and economically significant. The gap between immediate spending by public and private lottery winners is MK 341, or nearly 14 percent of the total prize. MK 341 is about three days' wages for casual labor in Malawi. There is clear evidence that participants in this experiment spent money more quickly when others knew they had won it than when others did not know.

Table 4 breaks down immediate spending into five categories discussed in the previous section: spending on own consumption, consumption by others in the household, consumption by others outside of the household, investment, and savings. The only statistically significant difference in spending by pubic and private winners is on consumption by others outside of the household, shown in column 3: public winners spend MK 105 more than private winners on non-household members in the week immediately following the lottery. The difference in spending on one's self (column 1) is of comparable magnitude though not significantly different from zero. Interestingly, public winners invest about MK 150 more than private winners in the week immediately after receiving the prize. The higher level of investment, though only on the margin of significance, provides some evidence that at least part of the rapid spending by public lottery winners is an extra effort relative to private winners to preserve assets for future use.

Recall that in Table 1, I show that public and private winners differed in their *predicted* as well as *realized* spending within one week of the lottery. While winners of public lotteries reported more immediate spending in both surveys, both public and private lottery winners actually spent less within one week than they anticipated. Figure 4 plots differences in realized and predicted spending; negative numbers indicate that people actually spent less than they anticipated spending in the week following the lottery. Public and private winners both have high variance in their prediction error. The spike at zero is not driven by people who correctly predicted spending nothing within the week after the lottery.

Table 5 provides a more systematic analysis of prediction errors in immediate spending. Column 1 estimates the equation $Y_{i,August} = \alpha + \beta Y_{i,May} + \epsilon_i$. If individuals perfectly predict their spending, then α equals zero and β equals one. Alternatively, if the prediction contains no information about actual spending, then β equals zero. The coefficient is a precisely estimated zero, consistent with predictions about spending within one week containing no information about actual spending in that time period. Column 2 adds an indicator for

whether the individual won a public or private lottery, and an interaction between winning a public lottery and the predicted level of spending in the first week. The statistically significant coefficient on the interaction term suggests that public and private winners differ in their ability to predict spending in the week following the lottery. Indeed, private winners actually spend 0.28 kwacha for every one kwacha they predicted spending within one week of the lottery. Public winners, though, spent 0.21 kwacha less for each kwacha they predicted spending. The correlations between predicted and actual spending for each category of winner are statistically different from zero.

Categories of Spending

I turn now to analysis of how lottery winners spend their prizes. The model predicts that the marginal propensity to save from public income is lower than the marginal propensity to save from private income. In my data, three types of results would be consistent with this prediction:

1. Public winners save less of the prize than private winners.
2. Public winners invest more of the prize than private winners.
3. Public winners spend more on people outside of their households than private winners.

The first item is a direct prediction of case 3 of the model. The second embodies the idea that public winners try to protect assets for the future by converting them into non-taxable goods instead of saving cash, which is more liquid but subject to taxation. The third directly captures the tax, in the form of transfers to others.

Table 6 summarizes anticipated spending for public and private lottery winners. Differences in anticipated spending on themselves, their households, individuals outside the household, and investment are not significant. Indeed, the difference for anticipated investment is a precisely estimated zero. However, public winners anticipate saving about MK 150 less than private winners, which is significant at the 95 percent confidence interval and consistent with the first type of supporting evidence mentioned above.

Table 7 repeats the exercise, using data about spending realized as of August. Differences in spending on themselves and on others are neither economically nor statistically significant.

Differences in spending on the household and in investing, though, are substantial in monetary value and at the cusp of statistical significance. Public winners spent about MK 236 less on their own households, and invested about MK 225 more than private winners. Public winners also saved MK 46 less than private winners, though that difference is not significant. There is no evidence that money received in public was taxed away, since spending on people outside the household was effectively the same for public and private winners.

Tables 8 and 9 examine the ability to predict spending in each of the five categories. Table 8 regresses realized spending on anticipated spending and a constant, one category at a time. As in column (1) of table 5, there are two interesting tests of each coefficient in table 8. A coefficient of zero indicates that the predicted level of spending in a given category contained no information about actual spending in that category. A coefficient of one means that the prediction was perfect. For each category of spending, we reject the hypothesis that lottery winners perfectly anticipated their spending with their May survey responses. Anticipations about spending on the household and about savings may have some predictive value. For each kwacha lottery winners anticipated spending on individuals outside of the household their actual spending increased by 0.39 kwacha, and for each kwacha individuals predicted saving, they actually saved an additional 0.28 kwacha. These coefficients on anticipated spending, shown in columns (2) and (5), are on the margin of being significantly different from zero and have the largest magnitudes of all five categories. The relationship between predicted and realized investment is statistically different from zero.

Table 9 asks whether winners of public lotteries differed from winners of private lotteries in their ability to predict spending across the five categories. The regressions in this table include indicator for whether the individual won a public lottery and an interaction term; they correspond to table 5, column (2). As before, interaction terms that are significantly different from zero indicate that public winners differed from private winners in their ability to predict spending. The coefficient on the predicted level of spending is the marginal spending in by August for each kwacha predicted in May for private lottery winners. The sum of the coefficients on the predicted level and the interaction term captures the same concept for public lottery winners. Ability to predict spending on oneself is not statistically different from zero, though the magnitude of the coefficient is large. Private lottery winners spent MK 0.22 less on themselves for each kwacha they predicted spending on themselves, while

public lottery winners spent 0.06 kwacha more. Difference in ability to predict spending on the household is a precisely estimated zero. Spending on others is an interesting case: private winners predicted a statistically significant share of spending for others and each kwacha of predicted spending translated into MK 0.68 of realized spending. Public winners, though, were marginally significantly less able to predict spending on others, and spending on others increased only MK 0.07 for every MK 1 of predicted spending. Public winners were marginally more able to predict their investment levels, and nearly identically to private winners in ability to predict savings.

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

7 Specification Checks

The results I presented in Section 6 are from OLS regressions. I also test specifications where I include controls for gender, amount of land owned, years of education, and club fixed effects. Including these covariates in an attempt to absorb residual variation does not change the results.

[THIS SECTION IS INCOMPLETE.]

8 Conclusions

I use a simple experiment of allocating money to members of agricultural clubs in public and private lotteries to measure the impact of public information on farmers' anticipated use of their prizes. While all winners spend a large share of their prize money in the one-week period immediately following the lottery, those who won money in a public setting spent

a significantly larger amount of money in that short window. The tendency to consume quickly could suggest hyperbolic or very high discount rates, but the difference between rapid spending by individuals who win the money publicly and privately requires an additional explanation. Strong sharing norms, suggested by anthropologists and by Malawian farmers in focus groups, may constitute a tax on income that others know about. Then, spending such income quickly, before others lay claim to it, is the rational optimizing behavior of an individual with a standard discount rate facing budget constraint with parameters such that a tax on public income is binding. Note that the welfare implications of such a finding are ambiguous. Sharing norms may be a constraint that force individuals to accept a second-best solution, where spending quickly limits the ability to shop for better prices or leads to hasty decisions that are regretted in the future. However, sharing norms may also provide an important means of insuring against idiosyncratic shocks. Public information about income may increase the enforceability of these informal insurance networks, therefore increasing the consumption-smoothing benefits they provide. Nonetheless, it is clear from my data that it is important to model this additional constraint when studying consumption decisions of individuals when income is easily observed and such norms are likely to be present.

Tables

Table 1: Sample Means, Public and Private Winners

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

Means for winners of public and private lotteries.

Table 2: Sample Means, Attriters and Non-attriters

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

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

Table 3: Spending within One Week of Lottery

	(1) Predicted	(2) Predicted	(3) Realized
Public	2,053	1,983	1,414
	(86.13)	(137.8)	(136.2)
Private	1,786	1,715	1,073
	(78.05)	(141.2)	(151.6)
Observations	315	116	116
p-value: public=private	0.015	0.148	0.064

Sample in column (1) includes all winners.
Samples in columns (2) and (3) includes winners interviewed at follow up.
Standard errors clustered at the club level.
p-values for test that spending was equal for public and private winners.

Table 4: Immediate Spending on Consumption, Investment and Savings

	(1)	(2)	(3)	(4)	(5)
	Self	HH	Non-HH	Invest	Save
Public	338.3	456.7	155.4	405.46	38.46
	(82.32)	(88.47)	(54.42)	(96.87)	(38.58)
Private	213.7	476.7	47.06	247.6	88.23
	(81.39)	(105.4)	(23.68)	(88.17)	(62.22)
Observations	116	116	116	116	116
p-value: public=private	0.255	0.882	0.047	0.171	0.505

Sample includes winners interviewed at follow up.
Standard errors clustered at the club level.
p-values for test that spending was equal for public and private winners.

Table 5: Relationship between Realized and Predicted Spending within One Week of Lottery

	(1)	(2)
	Realized	Realized
Predicted	0.004	0.281
	(0.089)	(0.134)*
Public		1237.6
		(327.4)**
Public × Predicted		-0.490
		(0.160)**
Constant	1256.7	590.5
Observations	116	116
R-squared	0.00	0.08
p-value: public+interaction = 0		0.00

Sample includes winners interviewed at follow up.
Standard errors clustered at the club level.

Table 6: Anticipated Use of Prize Money

	(1)	(2)	(3)	(4)	(5)
	Self	HH	Non-HH	Invest	Save
Public	219.3	793.9	100.3	572.6	248.4
	(58.29)	(76.24)	(25.06)	(72.13)	(44.66)
Private	182.5	827.1	73.86	571.5	397.3
	(41.83)	(76.21)	(23.41)	(74.97)	(60.95)
Observations	315	315	315	315	315
p-value: public=private	0.602	0.747	0.443	0.991	0.042

Sample includes all winners.
Standard errors clustered at the club level.
p-values for test that spending was equal for public and private winners.

Table 7: Realized Use of Prize Money

	(1)	(2)	(3)	(4)	(5)
	Self	Household	Non-HH	Invest	Save
Public	524.8	747.2	209.2	784.8	153.8
	(104.6)	(109.3)	(65.02)	(125.1)	(72.05)
Private	515.7	983.5	186.3	559.0	200.0
	(119.5)	(141.1)	(65.45)	(117.9)	(89.31)
Observations	116	116	116	116	116
p-value: public=private	0.949	0.179	0.789	0.159	0.701

Sample includes winners interviewed at follow up.
Standard errors clustered at the club level.
p-values for test that spending was equal for public and private winners.

Table 8: Realized versus Predicted Spending

	(1) Self	(2) Household	(3) Non-HH	(4) Invest	(5) Save
Predicted Self	-0.099				
	(0.113)				
Predicted HH		0.198			
		(0.105)			
Predicted Non-HH			0.346		
			(0.245)		
Predicted Invest				0.257	
				(0.111)*	
Predicted Save					0.297
					(0.153)
Constant	540.4	686.5	169.6	544.9	70.64
Observations	116	116	116	116	116
R-squared	0.00	0.04	0.04	0.06	0.10
p-value: $\beta_{pred} = 1$	0.000	0.000	0.009	0.000	0.000

Sample includes winners interviewed at follow up.
Standard errors clustered at the club level.
p-values for test that spending was equal for public and private winners.

Table 9: Realized versus Predicted Spending for Public and Private Winners

	(1) Self	(2) Household	(3) Non-HH	(4) Invest	(5) Save
Predicted	-0.220	0.230	0.679	0.115	0.356
	(0.110)*	(0.156)	(0.297)*	(0.130)	(0.192)
Public	-55.96	-245.4	69.54	83.20	62.13
	(160.4)	(228.7)	(91.68)	(194.3)	(94.81)
Predicted × Public	0.270	-0.031	-0.605	0.294	-0.121
	(0.219)	(0.204)	(0.355)	(0.182)	(0.288)
Constant	571.4	813.6	132.6	491.3	35.69
Observations	116	116	116	116	116
R-squared	0.01	0.06	0.07	0.09	0.11
p-val: public+interaction=0	0.760	0.162	0.722	0.008	0.303

Sample includes winners interviewed at follow up.
Standard errors clustered at the club level.

Table 10: Savings of Lottery Winners, August

	(1)	(2)	(3)
	Amount Saved at Home	Amount Saved in Bank	Total Savings
Public	9910.9	3259.4	13170.3
	(3428.5)	(1169.4)	(3486.1)
Private	5101.6	8428.6	13443.7
	(1213.8)	(3566.6)	(3725.5)
Observations	122	123	123
p-value: public=private	0.200	0.169	0.959

Sample includes winners interviewed at follow up.
p-values for test that savings were equal for public and private winners.

References

Ashraf, N. (2006). "Spousal Control and Intra-Household Decision Making: An Experimental Study in the Phlippines," Working Paper.

Baland, J., Guirkinger, C., and Mali, C. (2007). "Pretending to be poor: borrowing to escape forced solidarity in credit cooperatives in Cameroon," Working Paper.

DellaVigna, S. (2007). "Psychology and Economics: Evidence from the Field," NBER Working Paper #13420.

Falk, A. and Ichino, A. (2006). "Clean Evidence on Peer Effects," *Journal of Labor Economics*, Vol. 24, No. 1, pp. 39-58.

Keynes, J. M. (1936). The General Theory Of Employment, Interest, and Money.

Ligon, E., Thomas, J. P., and Worrall, T. (2002). "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies*, Vol. 69, No. 1, pp. 209-244.

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

Mas, A. and Moretti, E. (2006). "Peers at Work," NBER Working Paper #12508.

Townsend, R. (1994). "Risk and Insurance in Village India," *Econometrica*, Vol. 62, pp. 539-592.