

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

Jessica Goldberg*

December 2010

Abstract

Individuals in poor agrarian economies sometimes exhibit high marginal propensities to consume that are suggestive of high discount rates. I test whether informal insurance networks or social norms about sharing income can explain the 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 support 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 35 percent more than private winners in the period immediately following the lotteries. This spending pattern is consistent with a seven percent tax on surplus income in a simple model where a fraction of money that is not spent immediately must be shared with others in the social network.

JEL: O1, O21

Keywords: development, social norms, sharing, informal insurance, Malawi

*Department of Economics and Ford School of Public Policy, University of Michigan. E-mail jegoldbe@umich.edu. This project was supported with research grants from the Rackham Graduate School at the University of Michigan. I thank Brian Jacob, David Lam, Jeff Smith, and Dean Yang for their extremely helpful comments, and Santhosh Srinivasan and Lutamayo Mwamlima for their assistance with field work. All errors and omissions are my own.

1 Introduction

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

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

public condition spend 35 percent more windfall income within the first week of receiving it than those who receive money under the private condition.

The notion that social norms constrain consumption and savings decisions is not new. Townsend develops a model in which informal transfers within a village smooth consumption against idiosyncratic shocks.[7] Ligon et al. show that with unenforceable commitments to village networks, informal insurance schemes do not fully insure against these shocks.[5] Enforcement, then, constrains the private first-best allocation of income under some circumstances. Social anthropologists expand upon this notion, explaining that the pressure to share cash with family or neighbors is pervasive, and that people respond to these pressures by spending income quickly: “Consequently, on those infrequent occasions when they were able to earn money, they often made wasteful or ill-considered expenditures just to keep friends from borrowing it.”[6] (p. 18). Baland et al. document similar behavior among micro-finance customers in Cameroon who borrow money despite having savings of equal or greater value than their loans, thus incurring unnecessary interest costs for projects that could be self-financed. They report, “Excess borrowing is purposefully used by some members to signal financial difficulties to their relatives in search of financial help. Reimbursement obligations are then used as an argument to discourage such demands.”[2] (p. 9) Comola and Fafchamps demonstrate that gift-exchange between rural Tanzanian households is sometimes involuntary rather than an optional reciprocal strategy for coping with idiosyncratic shocks.[3]

I extend this literature by using experimental variation in whether windfall income is received in a public or private setting to formally test the effect of social pressure on consumption decisions. My paper proceeds as follows. I describe the model in Section 2. I explain the experiment in Section 3, and the data in Section 4. I present results in Section 5. Section 6 concludes.

2 A Simple Model

Individuals with high discount rates consume income rapidly and have low marginal propensities to save. Hyperbolic discounters apply an additional penalty to consumption in future periods, which also leads to high marginal propensities to consume. In a two-period model with windfall income, consuming the windfall income in the initial period rather than spread-

ing it between the two periods would be consistent with a high discount rate or hyperbolic discounting. Here, I develop a model in which the same behavior is explained by pressure to share, which is an economic constraint, rather than the discount rate, which is a behavioral parameter. In this model, exponential discounters and even people who are perfectly patient will consume income rapidly in order to avoid sharing their income when others observe their income. The combination of rapid consumption and obligatory transfers to people outside of the household results in high expenditures in the period in which income is received.

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

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

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

Spending income immediately demonstrates that a household has legitimate consumption needs. Conceptually, it is a strategy similar to non-cash-constrained families taking out loans, as a way to signal inability to give money to relatives.[2] In contrast, holding cash indicates a relatively low marginal utility of consumption, such that it might be socially optimal to reallocate money to someone else in the network. Individuals can choose to signal that they have high needs and should therefore be exempt from contributions to the network by

spending their income immediately. Even though this rapid spending reduces contributions to the group, it escapes sanction because negative shocks are not perfectly observed, and thus others in the network are not sure whether the rapid spending represented genuine or strategic need to consume. Social sanctions can only be enforced for income that was known to others, so holding cash from w^u is subject to sharing but holding cash from w^v is not.

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

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

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

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

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

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.

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.

The solution to this model has four cases, which I describe in detail in the appendix. The sharing norm binds in two cases. In one, the sharing norm causes individuals to consume more of public income in period 0 than they would have otherwise, driving the marginal propensity to consume from public income above the marginal propensity to consume out of private income. At the extreme, individuals are driven to the corner solution. At the corner solution, they consume all of public income and none of private income in period 0.¹

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

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

3 Experimental Design

Enforcing sharing norms on savings or assets requires public information about the level of savings or assets, even when the “tax” is a contribution to an informal insurance network. When information about an individual’s income is public, that individual faces constraints that lead to a higher marginal propensity to consume immediately after receiving income and more money given to others in the social network. To test those implications of the simple optimization model above and determine the causal effect of public information about individuals’ choices on their financial decision-making, I study the allocation of windfall

¹These cases are described as cases 1 and 3, respectively, in the appendix.

income under different information conditions. In principle, my experiment is similar to Ashraf's.[1] While Ashraf focuses on information conditions that affect within-household bargaining, I am interested in community-level dynamics.

Maranz's description of social pressure to share turns on the assumption that "if someone receives money, those people who are socially close will know it." [6] It is exactly that assumption that I manipulate in a simple field experiment involving 1,553 farmers in central Malawi. I run 316 lotteries in 158 agricultural clubs in central Malawi in May 2008. These clubs of approximately 10 members each were formed in late 2007 for the purpose of receiving extension services and borrowing through group liability schemes, and the lotteries are conducted when the clubs are assembled for meetings and surveys related to their loans for the coming agricultural season. The lotteries are facilitated by trained, Chichewa-speaking enumerators.

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

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

increase in private income, w_0^v .

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

4 Data

My final sample is of 315² 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 where 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 revisited³. Some 627 members of those clubs were present and were administered a follow-up questionnaire about assets, savings, and the recent harvest. At least one of the lottery winners was present in 77 of those clubs; in total, 114 lottery winners were administered a supplemental survey asking how they had actually used their prize money.

Table 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

²One individual who won a private lottery declined the prize.

³Budget constraints precluded revisiting all clubs.

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

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

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

My analysis of how prize money is used divides spending into five categories: consumption by the winner, consumption by others in the winner’s household, consumption by persons not in the winner’s household, investment or purchase of durable goods for the household, and savings. These categories are mutually exclusive and exhaustive. I include purchase of agricultural inputs such as fertilizer and pesticides, purchase of livestock, and purchase of building materials in the “investment” category. Results for analysis of these categories are not sensitive to alternative definitions of investment, such as removing livestock.

5 Results

5.1 Timing of expenditures

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

Table 3 shows the difference in expenditures within one week of the lottery for public and private winners. Columns (1) through (4) use data about anticipated expenditures collected at the time of the lotteries for all winners. Note first that all lottery winners anticipated using a large amount of their prize money very quickly. In the full sample, winners of private lotteries anticipate spending MK 1,815, or 72 percent of their prize money, within one week of the lottery. Winners of public lotteries anticipate spending MK 2,051 (MK 2,041 controlling for baseline characteristics including gender, age, household size, education, land owned, physical attributes of the house, and the household’s durable assets) in the same time period. The difference of MK 236 (MK 226 controlling for baseline characteristics) is statistically significant and represents a 13 percent increase relative to the immediate expenditures of winners of private lotteries. Limiting the analysis to the subsample of respondents who were resurveyed three months after the lotteries, the magnitude of the difference between public and private lottery winners’ expenditures persists, though the standard errors are larger.

I use data from the follow-up survey about actual expenditures in columns (5) and (6).

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

5.2 Strategic spending and sharing

Table 3 shows that gross expenditures immediately after the lotteries were higher for those who won in public settings than in private settings.⁴ This is strong evidence that social pressure affects consumption decisions. In Table 4, I look more directly at strategic spending and sharing. Strategic spending is the difference between immediate consumption of public lottery winners compared to private lottery winners. It has three components: expenditures for the winner him/herself, expenditures for the winner's household, and expenditures on durable (relatively illiquid) goods for the winner's household. The model predicts that public winners spend strategically to reduce the amount of money they are obligated to give to others.

As shown in columns (1) to (3), private winners spend somewhat more than public winners on themselves, their households, and durable goods. However, these differences are not statistically significant. Column (4) is the sum of the three previous columns; the coefficient on the indicator for winning a public lottery in column (4) is the total amount of strategic spending. Receiving money in public causes lottery winners to consume MK 246 more of their prize in the first week than they would have if they had received money privately. This difference is not statistically significant, but the magnitude is large and in the predicted

⁴Estimates in this and subsequent tables include baseline covariates. Including the covariates does not affect the results.

direction. MK 246 is ten percent of the total prize money, and a 25 percent increase in immediate spending relative to winners of public lotteries. It is an amount equivalent to about two-and-a-half day’s wages, and enough to purchase a live chicken or pay for a course of medicine to treat malaria.

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

$$\tau = \frac{\text{transfers}}{w_0^u - c_0} \tag{3}$$

In this case, $\tau = 0.07$. Winners were obligated to share with others seven percent of income they did not spend on their own immediate consumption.⁵

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

Most of these predictions are not borne out in the data from the follow-up survey. The categories in Table 5 correspond to those in Table 4, but Table 5 covers spending in three months following the lotteries. Public winners spend approximately the same amount as private winners on themselves. They spend MK 186 less than private winners on consumption goods for their households, but MK 196 more on durables. Neither difference is statistically significant, though the reallocation is consistent with an effort to smooth consumption by winners of public lotteries.

More surprisingly, over the three month horizon, there is no meaningful difference in the

⁵Since this experiment studies use of windfall income, results may not generalize to earned income. Jakeila finds a higher tax rate imposed on windfall than “earned” income in a laboratory experiment in Kenya.[4]

amount of money given to people outside the household. Winners of private lotteries give away MK 190 over three months, and winners of public lotteries give away MK 186. Winners of public lotteries give away more money immediately than winners of private lotteries, but less in subsequent months. Receiving money in public affects the timing of consumption and sharing, but not the total level. This result is not explained by the model described in Section 2. One possible explanation is that social obligations are absolute instead of relative to income, and that efficiency concerns dictate collecting money from each person as soon as practical. Public winners therefore paid their “tax” immediately, while private winners paid small installments over time. Alternatively, it may imply that information about winning the prize in the private condition became public over time, and private winners were subject to taxation as information was revealed. Either of these alternative explanations still have welfare implications for receiving money in public compared to in private, since both imply constraints on individuals’ flexibility in smoothing consumption.

5.3 Ability to predict spending

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.

I analyze errors in prediction of immediate expenditures in Table 6. Column 1 estimates the equation $Y_{i, August} = \alpha + \beta Y_{i, May} + \epsilon_i$. If individuals perfectly predict their spending, then α equals zero and β equals one. Alternatively, if the prediction contains no information about actual spending, then β equals zero. The coefficient is not statistically different from zero, consistent with predictions about spending within one week containing no information about actual spending in that time period. Column 2 adds an indicator for whether the individual won a public or private lottery, and an interaction between winning a public lottery and the predicted level of spending in the first week. The statistically significant coefficient on the interaction term suggests that public and private winners differ in their ability to predict spending in the week following the lottery. Indeed, private winners actually spend 0.32 kwacha for every one kwacha they predicted spending within one week of the lottery. For winners of public lotteries, though, the correlation between predicted and actual spending

within one week is zero. Public winners faced apparently unexpected constraints in allocating their prize money relative to private winners. This might indicate uncertainty surrounding the social norms or tax rate governing sharing of windfall income from an unusual source.

In Tables 7 and 8, I examine the ability to predict spending in each of the five categories. The estimates in Table 7 are category-by-category OLS regressions of realized spending on anticipated spending and a constant. As in column (1) of table 5, there are two interesting tests of each coefficient in table 8. A coefficient of zero indicates that the predicted level of spending in a given category contained no information about actual spending in that category. A coefficient of one means that the prediction was perfect. For each category of spending, I reject the hypothesis that lottery winners perfectly anticipated their spending with their May survey responses. Anticipations do have some predictive value for all categories except spending on one's self, however. The strongest correlation between predicted and actual spending is for money shared with people outside of the household. For every kwacha lottery winners anticipated sharing with others, they actually shared MK 0.40.

Table 8 asks whether winners of public lotteries differed from winners of private lotteries in their ability to predict spending across the five categories. The regressions in this table include an indicator for whether the individual won a public lottery and an interaction term. As before, interaction terms that are significantly different from zero indicate that public winners differed from private winners in their ability to predict spending. The coefficient on the predicted level of spending is the marginal spending by August for each kwacha predicted in May for private lottery winners. The sum of the coefficients on the predicted level and the interaction term captures the same concept for public lottery winners.

Winners of public and private lotteries are not significantly different in their ability to predict spending in any of the categories, but there are interesting differences in the categories for which each group of winners is able to make meaningful predictions. Public lottery winners could predict about three times more of the variance in their spending on durable goods than private lottery winners. For each kwacha public lottery winners anticipated spending on durables, they actually spent MK 0.37, compared to MK 0.13 for private lottery winners. Though the difference is not statistically different from zero, public lottery winners' predictions explained a statistically significant fraction of the variance in their spending on durables and private lottery winners' predictions did not. In contrast, private lottery winners

were three-and-a-half times more accurate at predicting the amount of money they would give to others. For each kwacha they anticipated sharing, private lottery winners actually gave MK 0.67 to others. Public lottery winners, though, actually gave away only 0.19 kwacha for each kwacha they anticipated sharing.

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

6 Conclusion

I use a simple experiment of allocating money to members of agricultural clubs in public and private lotteries to measure the impact of public information on farmers' anticipated use of their prizes. While all winners spend a large share of their prize money in the one-week period immediately following the lottery, those who won money in a public setting have 35 percent higher expenditures in that short window. The tendency to spend quickly could suggest very high discount rates, but the difference between the rapid spending by individuals who win the money publicly and that of those who win privately requires additional explanation. Strong sharing norms may constitute a tax on income that others know about. Then, spending such income quickly, before others lay claim to it, is the rational optimizing behavior of an individual with a standard discount rate facing a budget constraint with parameters such that a tax on public income is binding. The short-run spending patterns in my data suggest a tax of seven percent on surplus income. Over a longer horizon, however, the use of income

received in public is statistically indistinguishable from income received in private. The welfare implications of such a finding are ambiguous. Sharing norms may be a constraint that force individuals to accept a second-best solution, where spending quickly limits the ability to shop for better prices or leads to hasty decisions that are regretted in the future. Even though identical fractions of public and private income are given to people outside the household over a longer horizon, the accelerated giving and spending when income is received in public reduces consumption smoothing and the ability to adjust sharing to unanticipated negative shocks. However, sharing norms may also provide an important means of insuring against idiosyncratic shocks. Public information about income may increase the enforceability of these informal insurance networks, therefore increasing the consumption-smoothing benefits they provide. Nonetheless, it is clear from my data that it is important to model this additional constraint when studying consumption decisions of individuals when income is easily observed and such norms are likely to be present. These individuals may exhibit high marginal propensities to consume despite not having high or hyperbolic discount rates.

Figures

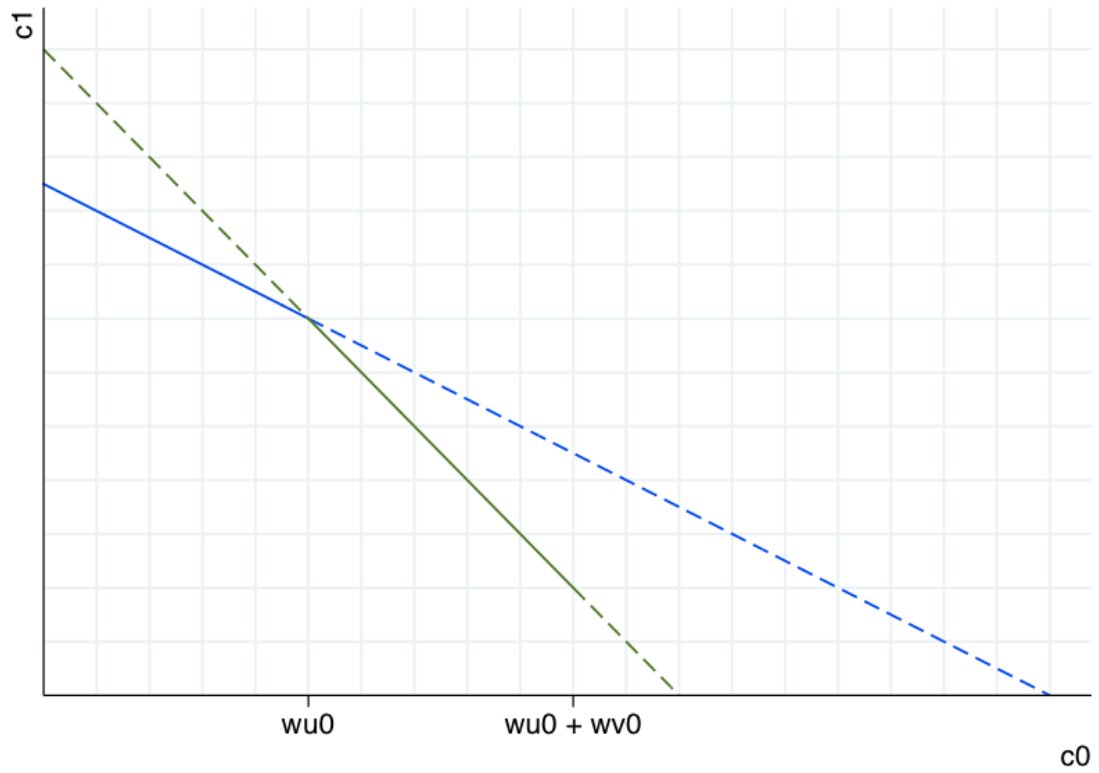


Figure 1: Possible combinations of c_0 and c_1

Tables

Table 1: Sample means, public and private winners

	(1)	(2)	(3)	(4)	(5)
	Male	Age	Years of Education	Land Owned	In August Sample
Public	0.947 (0.018)	43.49 (1.133)	5.933 (0.297)	4.520 (0.205)	0.414 (0.040)
Private	0.923 (0.022)	44.65 (1.090)	6.130 (0.290)	5.067 (0.238)	0.323 (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)	(2)	(3)	(4)
	Male	Age	Years of Education	Land Owned
In August Sample	0.947 (0.024)	44.83 (1.273)	5.491 (0.365)	4.948 (0.251)
Not in August Sample	0.927 (0.020)	43.62 (0.931)	6.349 (0.282)	4.701 (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)	(2)	(3)	(4)	(5)	(6)
		Predicted			Realized	
Public	236.097** (113.878)	225.642* (117.319)	293.754 (189.479)	351.604* (197.181)	348.711* (186.395)	342.715* (191.255)
Covariates		x		x		x
Observations	294	294	114	114	114	114
Mean for private winners	1815.07	1815.07	1715.29	1715.29	985.10	985.10
R^2	0.01	0.07	0.02	0.11	0.03	0.11

OLS estimates. All standard errors are clustered at the club level.

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

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

Table 4: Immediate spending and sharing

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

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.

See footnote for Table 3 for a list of covariates. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$

Table 5: Eventual spending and sharing

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

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.

See footnote for Table 3 for a list of covariates. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$

Table 6: Relationship between actual and predicted spending within one week of lottery

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

OLS estimates. All standard errors are clustered at the club level.

Sample includes winners interviewed at follow up.

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

Table 7: Actual versus predicted spending and sharing

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

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.

See footnote for Table 3 for a list of covariates. * p<0.10, ** p<0.05, *** p<0.001

Table 8: Actual versus predicted spending and sharing, by lottery type

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

OLS estimates. All standard errors are clustered at the club level. Sample is all lottery winners who were interviewed at follow-up.

See footnote for Table 3 for a list of covariates. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$

References

- [1] N Ashraf. Spousal control and intra-household decision making: An experimental study in the philippines. *American Economic Review*, 99(4):1245–1277, Aug 2009.
- [2] J Baland, C Guirkinger, and C Mali. Pretending to be poor: Borrowing to escape forced solidarity in credit cooperatives in cameroon. *Working Paper, Center for Research in Economic Development, University of Namur*, Jul 2007.
- [3] M Comola and M Fafchamps. Are gifts and loans between households voluntary? *Centre for the Study of African Economies Working Paper Series*, (20), Jan 2010.
- [4] P Jakiela. How fair shares compare: Experimental evidence from two cultures. 2009.
- [5] E Ligon, J P Thomas, and T Worrall. Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of Economic Studies*, 69(1):209–244, Dec 2002.
- [6] D Maranz. *African Friends and Money Matters: Observations from Africa*, volume 37 of *Publications in Ethnography*. SIL International, 2001.
- [7] R M Townsend. Risk and insurance in village india. *Econometrica*, 62(3):539–591, Dec 1994.

A Mathematical Appendix

Recall that

$$w_t^u = y_t^u + (1+r)s_{t-1}^u$$

$$w_t^v = y_t^v + (1+r)s_{t-1}^v$$

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

There are four different cases as illustrated in figure 2.

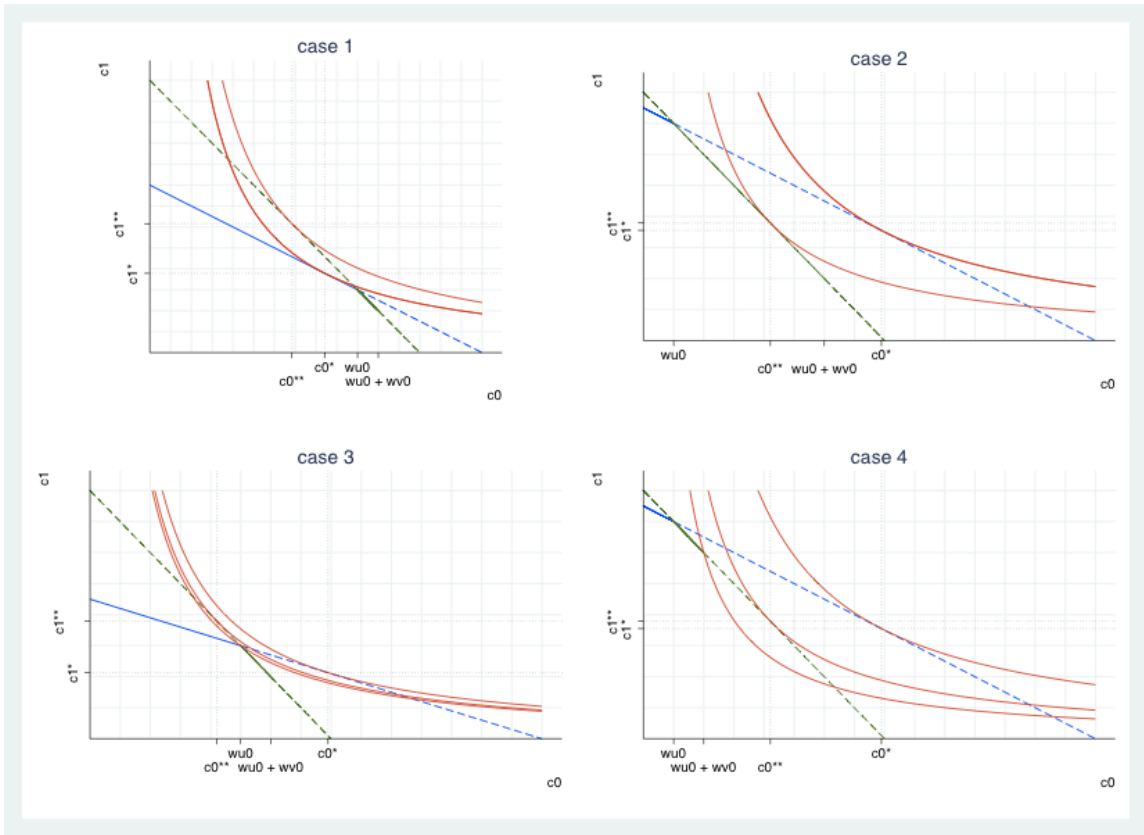


Figure 2: Different cases of the model

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

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

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

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

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

First, I maximize utility with respect to the sharing budget constraint, such that equation 1 holds, to find the utility-maximizing level of consumption c_0^* .

$$\begin{aligned} \max_{c_0} \ln c_0 + \beta \ln(y^u + y^v + \theta w_0^u + w_0^v - \theta c_0) \\ \frac{\partial F}{\partial c_0} = \frac{1}{c_0} + \beta \frac{-\theta}{y^u + y^v + \theta w_0^u + w_0^v - \theta c_0} = 0 \\ \beta \theta c_0 = y^u + y^v + \theta w_0^u + w_0^v - \theta c_0 \\ c_0^* = \frac{1}{\theta(1 + \beta)} (y^u + y^v + \theta w_0^u + w_0^v) \end{aligned}$$

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

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

Next, I maximize utility with respect to the no-sharing budget constraint, such that equation 2 holds, and call the utility-maximizing level of consumption under this constraint c_0^{**} .

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

$$\begin{aligned}\frac{\partial F}{\partial c_0} &= \frac{1}{c_0} + \beta \frac{-1}{y^u + y^v + w_0^u + w_0^v - c_0} = 0 \\ \beta c_0 &= y^u + y^v + w_0^u + w_0^v - c_0 \\ c_0^{**} &= \frac{1}{1 + \beta} (y^u + y^v + w_0^u + w_0^v)\end{aligned}$$

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

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

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

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

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