

Sharing Norms and Saving Behaviour

30 01 2017

Hanna Fromell*, Daniele Nosenzo*, Trudy Owens* and Fabio Tufano*

PRELIMINARY AND INCOMPLETE

PLEASE DO NOT CITE

Abstract:

We run a field experiment in Kenya and provide empirical evidence of the existence of sharing norms in rural poor communities and their impact on individual savings. Using the Krupka-Weber (2013) norm-elicitation method we show that these communities live by social norms that prescribe extensive sharing of income at the expense of own savings. For women, these norms are weaker if information about income is private rather than publicly known by the community. Using information on subjects' bank account usage we show that these norms, and their interaction with public information, systematically affect actual saving behaviour.

* School of Economics, University of Nottingham.

1. INTRODUCTION

In poor communities in the developing world, social networks serve as an important tool to substitute for the lack of legal institutions such as contract enforcement, credit provision, financial insurance and resource redistribution (Coate and Ravallion, 1993; Townsend, 1994; Fafchamps and Lund, 2003; Angelucci and De Giorgi, 2009). Social networks have also been argued to foster entrepreneurship by encouraging individuals to make risky investments and enhancing social capital and trust (Fafchamps and Minten, 1999; Fafchamps and Minten, 2002).

Despite their informal nature, social networks come with strong obligations. In particular, they often impose an informal “tax” on the relatively wealthier members to redistribute income to public services and to the less wealthy (Platteau, 2000; Olken and Singhal, 2011). This practice of income sharing in poor communities has been argued to be supported by so called “sharing norms” that represent a commonly accepted set of rules and understandings that sharing income with the relatively less wealthy members of the social network is socially appropriate (Platteau, 2000; Hoff and Sen, 2011). Failure to adhere to these norms is often associated with sanctions from the community in the form of economic retaliation and, in some cases, ostracism (Platteau, 2000; Barr and Stein, 2008; Hoff and Sen, 2011). Moreover, Platteau (2000) argues that sharing norms create an anticipation of being faced with claims from others on one’s personal income in the future and therefore may have a negative effect on economic development.¹ This implies that sharing norms may constitute an obstacle to wealth accumulation and reduce incentives to save among the poor.

However, to date there is only indirect evidence of the link between sharing norms and saving behaviour. For example, Dupas and Robinson (2013) and Brune et al. (2015) show that the individuals who are financially “taxed” by their community are those who benefit the most from using a savings technology with a commitment component. This suggests that commitment devices, such as restrictions of withdrawals or penalties of failure in making regular deposits, help individuals to save partly because commitment devices may give them means to resist demands from others to share. There is also evidence that individuals are willing to bear costs in

¹ Personal wealth accumulation may even be perceived as “anti-social behaviour” in some communities as it can be interpreted as an attempt to break free from the social network itself (Platteau, 2000).

order to keep information about income private (Baland et al., 2011; Beekman et al., 2015; Boltz et al., 2015; Jakiela and Ozier, 2015; Di Falco et al., 2016), which again may suggest the existence of a willingness to pay to avoid the pressure of sharing norms. These studies show that financial behaviour in poor communities is sensitive to features that affect others' ability to make claims on one's income, which in turn suggests that sharing norms may play an important role in these communities. However, none of the existing literature has, to the best of our knowledge, produced direct evidence of the existence of these norms or their impact on saving decisions.

In this paper we aim to fill this important gap. Our study is based on a field experiment conducted in poor communities in rural Kenya. We test for the existence of sharing norms that prescribe individuals to share income at the expense of personal savings. We do this by eliciting these norms directly using a modified version of the norm-elicitation method developed by Krupka and Weber (2013). This involves an incentivized task that elicits social norms by using coordination games. The task asks subjects to consider a situation described in a vignette and to give their perceptions of how socially appropriate or inappropriate different behaviours would be in the specific situation. In our experiment, subjects are read a vignette that describes an individual in their village who recently earned extra income and wishes to save it. This individual is thereafter asked to share their income with a friend and must thus decide how much to save and how much to share. We elicit the norm in this situation by asking subjects to rate how socially appropriate or inappropriate different alternatives of saving and sharing are. We use two different versions of this vignette and vary whether the income of the individual in the vignette is private or publicly known by everyone else in the community. This allows us to test whether the norm is sensitive to privacy about earnings.

Our results confirm the existence of a sharing norm that clearly prescribes individuals to share income at the expense of savings. On average, subjects perceived behaviour that involved higher levels of sharing but lower savings, significantly more socially appropriate compared to behaviour that involved less sharing and more saving. In fact, sharing all the extra income (and saving nothing) was perceived as the most socially appropriate action in our sample. We also find that, in the case of women, the extent to which higher sharing is considered more socially appropriate is significantly reduced when information about income is private. This finding is in line with results by for example Baland et al. (2011) and Jakiela and Ozier (2015) who show that

individuals in poor communities are willing to pay costs to keep information about earnings private. Our results provide a microfoundation for the behaviour that they observed: when information about earnings is private, the sharing norm is weaker compared to when information about earnings is public, which may in turn motivate individuals with a preference for norm compliance to make more self-interested choices.

We next test whether privacy in the information about earnings, and its impact on norms which we identified in the experiment, systematically affects subjects' actual saving behaviour. We do this by varying whether individuals have their experimental earnings communicated to them in private or have this information publicly announced to everyone in the experimental session. By paying subjects their experimental earnings into bank accounts we study the effect of private versus public information about experimental earnings on individual savings one month after the experiment. While a large share of our subjects save all of their experimental earnings in both treatments, we detect important effects among subjects at the lower end of the savings distribution. Using quantile regressions, we find that, for women with relatively low savings, keeping the experimental earnings private has a positive and significant effect on savings. This is thus in line with our finding that the sharing norm is weaker when information is kept private.

As we have shown that private information about earnings both weakens the social norm that prescribes sharing at the expense of individual savings and positively affects women's savings, we confirm Platteau's (2000) claim of that sharing norms may have negative impact on individual savings. Our results further suggest that there is scope for policy improvements. Formal means of savings, such as an individual savings account, commonly ensure that the individual has private information about the amount saved, in contrast to what is largely the case for various informal means of savings, such as saving groups. Increasing access to and the affordability of such formal saving devices among the poor may thus help individuals to save more by providing an environment where the sharing norm is less strong.

2 EXPERIMENTAL DESIGN

The field experiment collected data via three instruments: a household survey conducted at subjects' homes, an experiment run one day after the household survey and data on bank-account usage provided by our partnering bank, Kenya Commercial Bank, for the one month following

the experiment. The data from the survey and the experiment was collected by a team of 10 local enumerators.

2.1 Location, Sampling and Timeframe

The field experiment took place in Uasin Gishu County, in six villages outside Eldoret town (reachable within 30 minutes). Our bank partner, Kenya Commercial Bank, is the largest bank in the country and is also well represented in the local area, both directly through its network of bank branches and agents and indirectly through the network of M-Pesa agents, which is a mobile money transfer service offered by the bank's partner, Safaricom.

The six villages were randomly selected conditioned on that any two selected villages did not lie in direct proximity to one another, that they were located at a given distance from Eldoret town and contained around 100 households. 227 household representatives took part in the household survey, the experiment and the account opening. The study ran between April and June in 2016.

2.2 Household survey

The household survey was preceded by a survey with the village elders, which gathered basic information about the village and identified the location of all households and the names of the respective heads of households. The household survey was conducted with a subsample of the households in the village, interviewing either the female or the male head of household.² The survey collected data on individual demographics, individual preferences, perceptions about financial demands from others, and financial transfers to and from their social network.³

2.3 Experiment

The experiment was designed to study the effect of sharing norms on savings by studying both the prevailing social norm as well as actual saving behaviour under two different contexts: one in which information about income is held privately by the individual and one in which this is

² The gender was randomly determined beforehand but if the sought gender was unable or unwilling to participate in the survey, the subsequent experiment and account opening, the other household-head of the opposite gender would be selected instead.

³ The survey also collected data on household characteristics, basic demographics on all household members, assets, savings, loans, income, expenditure, and social network characteristics.

publicly known by other members of the village. Social norms are elicited in the experiment using a norm-elicitation task and saving behaviour is studied outside the experiment by observing savings in a bank account. In our two treatments, PRIVATE and PUBLIC, which are run using two different subject pools in each village, we manipulate information about income in both the norm-eliciting task and in the “experimental earnings stage” (i.e. the part of the experiment when earnings are communicated to subjects). This means that a subject allocated to the PRIVATE treatment takes part in the PRIVATE versions of both the norm-elicitation task and the experimental earnings stage.

The experiment took place on the day after the household survey and was run in community buildings inside the villages. The two treatments were included in each village and constituted two separate sessions, ran one after the other, each lasting for approximately 4.5 hours. Each participant was randomly allocated to one of the two treatments.⁴ The experiment contained four incentivised games where one of the games comprised of the norm-elicitation task that elicited sharing norms.⁵

The norm-elicitation task consisted of a modified version of the norm-elicitation method developed by Krupka and Weber (2013) to elicit the norm of sharing at the expense of individual savings in the village. Subjects were first read a vignette that described a situation faced by an individual in their village. In the vignette, the individual has recently earned an extra income of 900 Ksh and wishes to save this.⁶ A friend of the individual asks for the money since they need it to pay school fees. The individual in the vignette faces six options: save 100% and give nothing to their friend, save 80% and give the rest to their friend, save 60% and give the rest to their friend, and so on, proceeding to the action of saving nothing and give everything to their friend. After having listened to the vignette, our participants were asked to rate each of the actions

⁴ Due to variation in the number of people that showed up to the experimental session, one treatment included slightly more participants than the other (there are 105 participants in PRIVATE and 122 in PUBLIC). The morning session lasted between 9.30 am and 2.00 pm and the afternoon session between 2.00 pm and 5.30 pm. Participants in the afternoon session were strictly prevented from interacting with participants in the morning session before the start of the afternoon session. The order of the treatments was randomized across villages.

⁵ The other three games were a one-shot public goods game, a financial literacy task and a cognitive ability task. The games were played in the listed order with the norm-elicitation task at the end in all sessions. The preceding three games were identical across treatments.

⁶ 900 Ksh translate to about £7.20, corresponding to more than two weeks per capita income among rural Kenyans (Kenya, 2014).

according to how socially appropriate it would be for the individual in the vignette to take that action. Participants were incentivised to rate the actions in accordance with the prevailing norm in their village since they could earn money only by giving the exact same rating for a given action as rated by a randomly selected other participant in the experimental session.⁷

The norm-elicitation task varied the version of the vignette that was told to the participants in the two treatments, PRIVATE and PUBLIC. In PRIVATE, information about the extra income earned by the individual in the vignette was held privately by the individual only. The individual had to make a private decision on how much to save for themselves and how much to give to their friend. In PUBLIC, information about the extra income was publicly known by everyone in the village. The individual's decision on how much to save for themselves and how much to give to their friend would become publicly known by everyone in their village. The norm-elicitation task thus measures the sharing norm in two different contexts: when an individual's earnings and saving decision is public and when it is private.

After all games had been played, participants were informed of their earnings. This stage was manipulated across treatments to make information about experimental earnings either privately held by each participant or publicly known by everyone in the session. To make this manipulation meaningful, we designed the way that earnings were made such that they varied extensively across individuals.⁸ Total experimental earnings ranged from 300 Ksh to 1500 Ksh averaging 900 Ksh.⁹

In the PRIVATE treatment, participants were informed privately about their earnings and it was not possible for others to find out about their earnings without the individual disclosing this information themselves. In the PUBLIC treatment, participants' earnings were announced

⁷ Only one of the six actions would be selected for payment in this game. The participant earned 900 Ksh if they matched their rating with the rating of their partner and earned 0 Ksh if they failed to do so.

⁸ Earnings ranged from 450 KSh-1050 Ksh; 0 Ksh -900 Ksh; 0 Ksh -1000 Ksh; and 0 Ksh -900 Ksh in the four different games. The calculation of earnings also involved a random component in that only one of the games was selected for the calculation of earnings. Each subject made a random selection by drawing a piece of paper from a bag. If the participant selected the norm-elicitation task, they made a second draw to determine which of the six rated actions that would be used for the calculation. 300 Ksh was the lowest possible earnings since a show up fee of 300 Ksh was added to any earnings made from one of the games.

⁹ As we compare the elicited sharing norm with actual savings, we aimed to set the average amount that was earned by the individual in the vignette equal to the expected average earnings received by participants in the experiment.

publicly among all participants in the session.¹⁰ Thus a participant in the PRIVATE treatment had the ability to credibly deny the size of their experimental earnings to participants in the session while this was not possible in the PUBLIC treatment.

After the experiment had ended, subjects were assisted by bank officers of Kenya Commercial Bank in opening an individual bank account. On the following day, their bank account was activated and their experimental earnings were paid into their bank account.

2.4 Account savings

Savings are measured as the percentage of experimental earnings left in subjects' bank account one month after the experiment. By comparing savings between individuals in the PUBLIC treatment and the PRIVATE treatment, we can test whether private information about earnings positively affects savings. To capture the impact on savings rather than a lack of outside opportunities, the account paid a competitive interest on savings.¹¹ Information about the fees associated with account activity such as depositing, withdrawing and transferring money via M-Pesa was communicated to participants verbally and via a handed out information sheet.¹²

2.5 Attrition

Out of 302 individuals who were initially surveyed and invited to the experiment on the following day, 227 participated in the experiment.¹³ While women are found more likely to participate in the experiment than men, individuals who participated were overall not

¹⁰ To ensure that participants in the PRIVATE treatment trusted that earnings were private information, this information was communicated to them one-by-one in a separate room by a single experimenter. They had been told that the experimenter would keep this information strictly to themselves and that the bank would not be informed about their earnings. Participants in the PUBLIC treatment were given a proof of participation that contained information about their earnings from the experiment. To make treatments as identical as possible, participants in the PRIVATE treatment also received a proof of participation that however contained no information about earnings.

¹¹ To ensure that the interest paid was competitive both in the informal and the formal market, the account paid an interest of 50 % on the amount saved one month after the experiment. Apart from our intervention of paying an interest after the first one month, the account was a fully liquid standard bank account.

¹² We did not deduct any charged bank fees from the calculation of the amount left in the account when paying the interest rates. This was done to make the interest easy to understand. Interest was not paid on deposits made by the participant. With the exception of fees associated with viewing account statements using the phone, bank fees are however considered as expenses in our measure of savings in the analysis since subjects were aware of these fees. To further ensure that all participants understood how the interest would be paid and to maintain their trust in the procedure, participants received two text messages: one directly after earnings had been paid into their account and another 15 days later. These informed of when and how the interest would be paid out.

¹³ We intentionally invited more participants than needed since some level of attrition was expected.

significantly different from those who did not participate (see section 3 for detailed analysis). Out of the 227 participants, 13 had to be excluded from the analysis of savings due to either delay in the activation of their bank account (three subjects) or because they had loan arrears from previous financial activities which were immediately and automatically deducted from their balance after the opening of their account (ten subjects).

3 DATA AND BALANCE CHECK

3.1 Data

To measure individual preferences we employed both incentivized and non-incentivised techniques. Risk aversion, time discounting, trust, altruism and positive and negative reciprocity are measured by survey questions that have been confirmed to capture relevant behaviour and especially developed to be suitable in Swahili-speaking rural communities (Falk et al., 2016). Social capital, cognitive ability and financial literacy were measured with financial incentives in the experiment.¹⁴

We include a series of measures in the survey to capture the extent to which people perceive and give in to demands over their income from others. Similarly to John (2016), we measure the amount of claims that would be made by others in the event that the person would have cash savings at home. We measure the amount of gifts and loans that have been transferred to others and produce a variable of “net giving” by subtracting the amount of gifts and loans that the individual themselves have received from others. This measure is used to capture the extent to which the individual is “taxed” by their social network. Following Dupas and Robinson (2013) we ask the individual, using a Likert scale from 0-10, whether they feel obligated to give something to someone who would come to ask them for money.

3.2 Sample Characteristics and Balance Check

Table 1 presents baseline summary statistics. Column 1 tests whether the individuals who participated in our experiment are different from individuals who were surveyed and invited to

¹⁴ Social capital was measured using a one-shot public goods game (Gächter et al., 2004). Cognitive ability was measured using test scores on solving Raven’s matrices (Raven, 2000). We followed the method by Mani et al. (2013). Financial literacy was measured using four questions developed by Cole et al. (2011).

the experiment but did not end up participating. Coefficients and robust standard errors clustered at the village level are presented from the regressions of the relevant variable on whether the surveyed individual participated in the experiment. As the joint test of the hypothesis that all coefficients being equal to zero cannot be rejected, as is clear from the p-value in the last row, we conclude that the sample of surveyed individuals who participated in our experiment are not systematically different from those who did not participate. From the balance checks done for each variable separately, we find that three are significantly different. Women were more likely to participate in the experiment compared to men (57 percentage points). Individuals with fewer years of education were also more likely to take part in the experiment. Part of the reason for this simultaneous bias in education and gender lies in that women are less educated than men. Finally the coefficient on time discounting is significant, indicating that our participants are slightly more patient than the other surveyed individuals.

In the savings analysis, 13 subjects are excluded since we do not have accurate information about savings due to problems with the account or with loan arrears. Column 2 performs the same test as in Column 1 but testing whether the 214 participants included in the savings analysis are different from the rest of our surveyed sample, i.e. the individuals who were surveyed and invited to the experiment but did not end up participating as well as the 13 participants for whom we do not have accurate savings data. Also for these 214 subjects, the joint test of the hypothesis that all coefficients being equal to zero cannot be rejected.

Column 4 reports sample mean and standard deviations for the sample of experimental participants of the variables of interest. 69% of our participants are female, average age is 41 and the share of literate individuals is a fairly high 92%. A balance check is performed in column 5 by presenting the coefficient and robust standard errors clustered at village level from regressions of the relevant variable on the PRIVATE treatment. None of the individual characteristics were significantly different across treatments. However, participants in PRIVATE earned 77 Ksh more on average and this difference is significant at the 10 percent level. However, the joint test confirms that we succeeded with the randomization procedure.

Table 1 – Sample Characteristics and Balance Check

	Participate	Savings	N	Variable Mean	PRIVATE	N
<i>Demographic characteristics</i>						
Female	.57 (.19)***	.49 (.23)**	302	.69 [.46]	.45 (.38)	227
Married	-.09 (.44)	.15 (.36)	299	.78 [.41]	.07 (.36)	225
Age	2.82 (2.20)	2.94 (1.97)	302	41 [12.89]	-.50 (1.57)	227
Years of education	-.87 (.30)**	-.97 (.29)**	299	8.65 [3.35]	-.14 (.38)	224
Can write simple Swahili	-.42 (.47)	-.63 (.44)	287	.92 [.28]	.35 (.63)	216
<i>Incentivized behavioural measures</i>						
Experimental earnings				889.01 [339.42]	76.89 (33.29)*	227
Public good contribution				6.36 [4.82]	-.21 (.54)	227
Cognitive ability				1.85 [.83]	.06 (.10)	227
Financial literacy				2.71 [1.06]	-.11 (.12)	227
<i>Preferences</i>						
Risk Taking	.07 (.25)	.22 (.27)	302	6.76 [2.11]	.11 (.20)	227
Time Discounting	.29 (.08)**	.29 (.08)**	302	6.93 [1.94]	.16 (.25)	227
Altruism	-.28 (.22)	-.12 (.23)	302	7.04 [2.27]	.15 (.26)	227
Positive Reciprocity	.17 (.36)	.24 (.29)	302	7.51 [1.92]	.12 (.20)	227
Negative Reciprocity	.21 (.37)	.42 (.20)*	300	4.33 [3.20]	-.60 (.54)	225
Trust	.13 (.33)	.22 (.29)	301	6.47 [1.93]	-.03 (.27)	226
<i>Income sharing</i>						
Net giving	-294.72 (262.79)	-226.75 (272.00)	297	247.85 [1767.62]	-68.71 (232.97)	223
Feel obligated to give	.04 (.15)	.28 (.09)**	301	7.61 [1.56]	.12 (.11)	227
P-value – Joint Test	.39	.18			.69	

Notes: standard deviations in brackets. Robust standard errors clustered at the village level in parentheses.
 ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

4. RESULTS

The results section is divided into two parts. We start by presenting our findings on sharing norms measured in the norm-elicitation task. We then present findings on saving behaviour from subjects' usage of their bank accounts.

4.1 Sharing Norms

We use subjects' ratings of social appropriateness in the norm-elicitation task to produce a numerical score of the sharing norm. In line with Krupka and Weber (2013) we assign “very socially inappropriate” a value of -1, “socially inappropriate” a value of $-1/3$, “socially appropriate” a value of $1/3$ and “very socially appropriate” a value of 1. We measure the sharing norm under two different contexts: when earnings and the decision of how much to save and how much to give is privately held by the individual, and when this information is publicly known by everyone in the individual's village. This is done by eliciting the norm under two different treatments, PRIVATE and PUBLIC, using two different subject pools. In line with Krupka and Weber (2013) we use the average norms reported for each action available to the individual in the vignette and produce a separate average for each treatment.

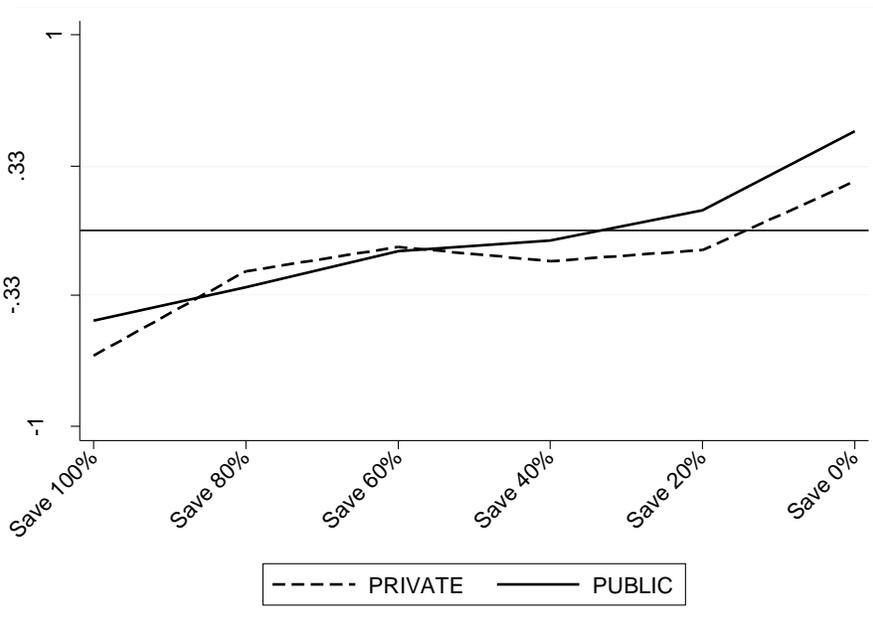
Graph 1 illustrates the elicited sharing norm in the PRIVATE and the PUBLIC treatment resulting from the computed averages. The graph shows the average appropriateness rating on the y-axis, where -1 is the possible rating of “very socially inappropriate” and 1 the possible rating of “very socially appropriate”. The x-axis represents the different available actions to the individual in the vignette, where “Save 100%” is the action of saving 100% and sharing nothing with the friend and “Save 0%” is the action of saving 0% and sharing everything with the friend.

It has been argued that individuals in poor communities face sharing norms that prescribe them to share their income with others in their community and that this may discourage individual savings (Platteau, 2000). Some suggestive evidence exists indicating that such norms may reduce savings (Dupas and Robinson, 2013; Brune et al., 2015). This leads us to predict that the social norm elicited in our norm-elicitation task will be one that prescribes positive levels of sharing at the expense of personal savings. For our elicited social norm, this would imply that an action other than “Save 100%” is rated as the most socially appropriate action.

Hypothesis 1. *There exist sharing norms in poor communities that prescribe people to share their income with others at the expense of own savings.*

We start by confirming that there is a distinct sharing norm among our subjects and that this is true in both treatments. In both treatments the action “Save 0%” (and share all with the friend) receives the highest rating of social appropriateness compared to any other action.¹⁵ Although only strictly true in the PUBLIC treatment, as we look across the spectrum of the rated actions, higher levels of sharing seems overall to be associated with a higher level of social appropriateness. We confirm that this relationship is significant as the regression of the social appropriateness ratings on the percentage of sharing implied by the different actions report a positive and significant coefficient in both treatments ($p < 0.001$).

Figure 1 – Social appropriateness ratings by treatment



Field evidence show that individuals are able to act more in their self-interest when their income can be kept private (Boltz et al., 2015).¹⁶ Furthermore, Krupka and Weber (2013) elicited

¹⁵ In each treatment separately, we compare the rating for “Save 0%” with the rating for each of the other saving decisions. Using a Wilcoxon Rank-sum test, the result is significant at a 1 percent level for all comparisons apart from one (Save 0% vs. Save 40% in PRIVATE), which is significant at a 5 percent level.

¹⁶ Boltz et al., 2015 who study the impact of being able to hide gains made in the lab on resource allocation outside the lab among urban dwellers in Senegal. They find that individuals with a preference for income privacy decrease the share that they transfer to kin when they are able to hide gains.

social norms to study a related finding from dictator-game behaviour in the lab. They find that the social norm is more forgiving of behaviour that involves no sharing if information about income and the possibility to share is kept private.¹⁷ Together this evidence point towards a possible explanation for why individuals in poor communities may be willing to pay a cost to keep information about income private: privacy in income may reduce the sharing obligations prescribed by the social norm in their community which in turn may enable the individual to keep a greater share of their income for themselves. We therefore expect that higher levels of sharing and lower levels of saving are not prescribed as strongly in the PRIVATE treatment compared to the PUBLIC treatment.

Hypothesis 2. *There are different sharing norms that apply to cases where an individual's saving decision is private or publicly observed. Sharing norms emphasize sharing behaviour less if the saving decision is private.*

From the ratings displayed in Figure 1 we have noted that higher levels of sharing overall are considered more socially appropriate and that the most socially appropriate action is to share everything. Comparing the lines of the two different treatments we see that the rating of the most socially appropriate action “Save 0%” (i.e. share everything) is lower in the PRIVATE treatment compared to PUBLIC. This difference is statistically significant at the 5 percent level using a Wilcoxon Rank-sum test. This suggests that the sharing norm prescribes highest possible sharing at a lesser extent when information about income and the saving decision is private compared to public.¹⁸

Table 2 formally tests our second hypothesis by testing whether the slope between social appropriateness and savings is less steep in PRIVATE compared to PUBLIC. From the first regression including all subjects we confirm the existence of a sharing norm that assigns lower levels of social appropriateness to actions of higher levels of savings as the coefficient “Savings”

¹⁷ Lazear et al. (2012) and Dana et al. (2006) both use lab experiments with students and demonstrate that they have preferences for keeping information about their ability to share income private. Krupka and Weber (2013) complement this literature by directly eliciting the social norms that apply in these situations. They show that choosing to keep all money in a dictator game is considered more socially appropriate if the recipient thereby is made unaware of that a dictator game is taking place.

¹⁸ There is no significant difference in ratings between PUBLIC and PRIVATE for the action representing the strongest violation of the norm “Save 100%” (i.e. share nothing).

is negative and significant.¹⁹ The interaction term “PRIVATE × Savings” measures the effect of savings in the PRIVATE treatment compared to the effect of savings in the PUBLIC treatment. As the coefficient is insignificant we do not have evidence that an increase in savings implies a greater change in social appropriateness in one treatment compared to the other. Yet, as predicted, the most socially appropriate action “Save 0%” is assigned a significantly lower rating of social appropriateness in PRIVATE compared to PUBLIC. This is confirmed from the negative and significant treatment dummy “PRIVATE”.

The second and third columns run the same regression as the first column for women and men separately. For women, higher levels of savings are strongly negatively associated with lower levels of social appropriateness when the saving decision is public information. Since the interaction term is positive and significant, this confirms that, for women, an increase in savings in the PRIVATE treatment results in a smaller decrease in social appropriateness than in the PUBLIC treatment. In contrast, for men, an increase in savings are not as strongly associated with a decrease in social appropriateness. Furthermore, for men, there is no evidence of a difference in norm ratings across treatments.²⁰

Taken together, the results from the norm-elicitation task suggest that there is some evidence for our second hypothesis that the sharing norm emphasizes sharing behaviour less if income and the decision of how much to save and to share is private. However, these results are mainly driven by our female subjects.

¹⁹ Note that savings here refers to the level of savings associated with the different actions in the vignette. Actual savings measured for each subject is instead analysed in the proceeding section.

²⁰ The non-significance among males could partly be due to a loss of power due to the relatively small sample size (70 men compared to 157 women). However, the interaction term for males is in fact weakly negative, which leads us to think this is unlikely to be the case. In particular, using a post regression F-test, we can exclude the possibility that the interaction term equals the one estimated in our female sample, which equals .34 (Prob. > F = 0.06).

Table 2 –OLS regressions of norm ratings on treatment and savings

	All subjects (1)	Women (2)	Men (3)
PRIVATE	-.20** (.10)	-.21* (.12)	-.17 (.18)
Savings	-.87*** (.11)	-1.03*** (.12)	-.56*** (.19)
PRIVATE × Savings	.20 (.16)	.34* (.18)	-.03 (.30)
Constant	.23 (.09)	.28 (.10)	.13 (.20)
Village FE	Yes	Yes	Yes
Observations	1,361	941	420
R ²	.10	.13	.09

Notes: Robust standard errors clustered at the individual level in parentheses. ***Significant at the 1 percent level. **Significant at the 5 percent level. *Significant at the 10 percent level.

4.2 Saving Behaviour

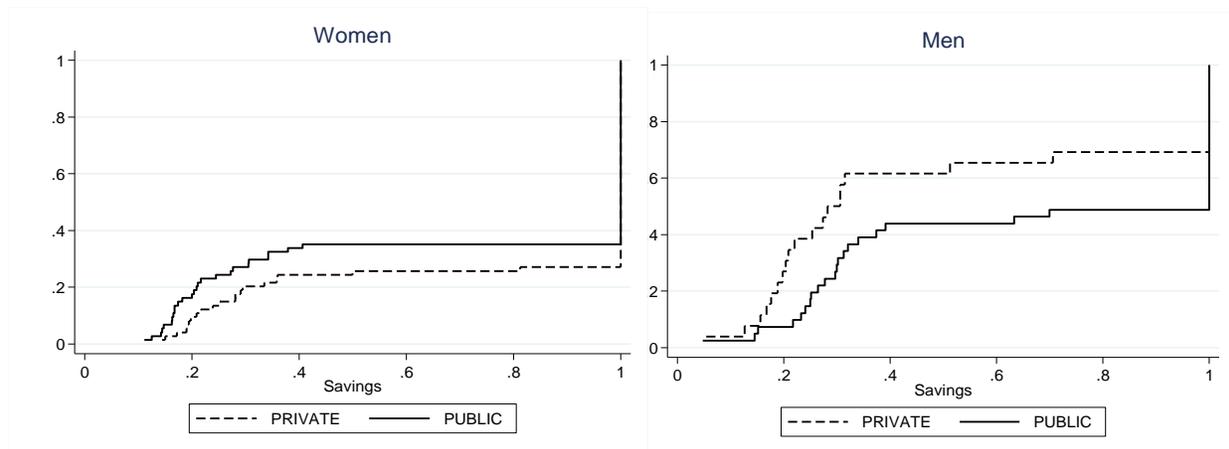
So far we have focused on analysing social norms among our studied subjects and have demonstrated the existence of a sharing norm that prescribes sharing at the expense of personal savings. We found evidence that, for women, the sharing norm is less pronounced if information about earnings and the individual's decision of how much to share and to save is private information. We next test whether we find a similar effect in actual saving behaviour by varying information about individual earnings and studying the effect on personal savings.

Hypothesis 3: *Saving behaviour follows norms. Individuals save more of their extra income if this amount is private information compared to if it is publicly known to others in their social network.*

Savings are measured as the percent of initial experimental earnings left in the account one month after the experiment. Average savings among our subjects is 71%. Although similar studies on savings in poor communities have found considerably lower levels of savings, our relatively higher level of savings is not surprising given that we pay a very high interest on savings for the period of study (equal to 50% on experimental earnings left in the account after one month). We find a clear difference in savings across gender where women save on average 77% while men save on average 60% (Wilcoxon Rank-sum test, p-value < .01).

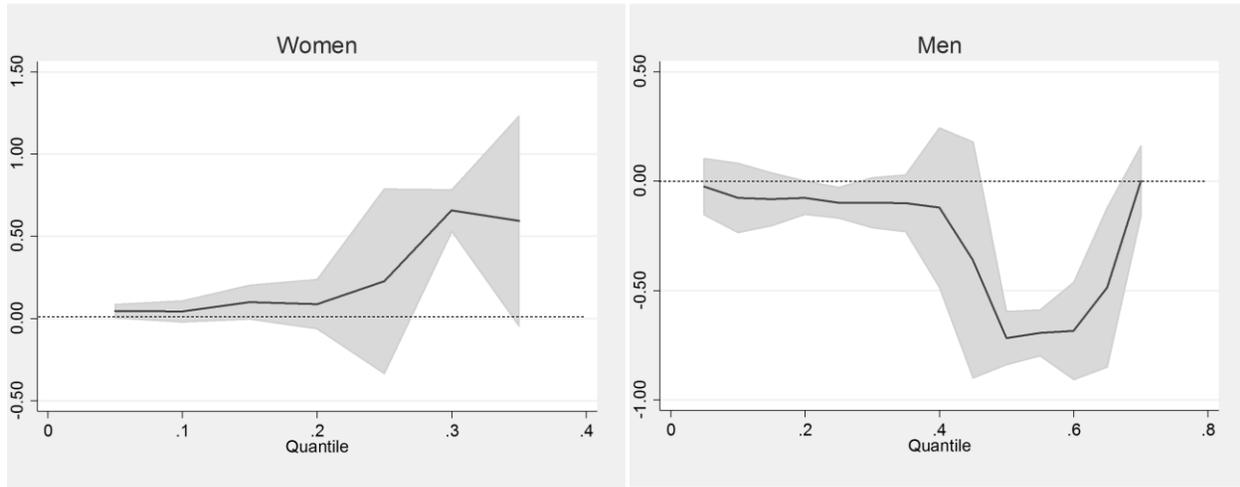
The distribution of savings is highly skewed: the share that saved 100% was 69% for women and 44% for men. Figure 2 illustrates the cumulative distribution functions of savings in the two treatments, PRIVATE and PUBLIC, for women and men separately. These graphs illustrate a clear difference in savings in the two treatments. However, comparing these differences between women and men highlight a contrast in this treatment difference. Due to the skewness in the savings variable for both men and women, OLS regressions may not capture possible treatment effects that take place at different parts of the distribution.

Figure 2 – Cumulative Distribution Functions of Savings by treatment



To explore treatment effects across the distribution of savings we follow the method used by Schaner (2016) and study the quantile treatment effects on our dependent variable. Figure 3 reports coefficients from quantile regressions of savings on the PRIVATE treatment for women and men separately. The graphs report estimates from the 5th quantile up to the quantile at which the savings variable equals 100% in both treatments, i.e. at the 36th percentile for women and the 70th percentile for men. For 14 of the 30 quantiles reported for women, the treatment coefficient is positive and significant. This suggests that the PRIVATE treatment did have some meaningful positive impact among a large fraction of women. For 21 of the 66 quantiles reported for men, 21 give a negative and significant coefficient. This, contrary to what we find for women, suggests that men saved less in the PRIVATE treatment compared to men in the PUBLIC treatment.

Figure 3 – Distributional impact of the PRIVATE treatment on savings



Notes: In each graph, the solid line represents coefficients from quantile regressions of savings on the PRIVATE treatment. The shaded area indicates 95 percent confidence intervals.

Our results thus suggest that women with relatively low savings manage to save more if they have private information about earnings compared to if this is known by others in the community. As we find that the sharing norm among women is weaker when information is kept private, our results indicate that women may benefit from a weaker sharing norm in the sense that this may enable them to save more. Our results are consistent with a series of studies that small changes in the context may induce important changes in the social norm, which may in turn affect behaviour among individuals who have a preference for norm compliance (Bicchieri, 2006; Krupka and Weber, 2013). Further literature in line with our results for women argues that a change in the social norm may have particularly strong impact on behaviour if it allows the individual to act more in their own self-interest (Dana et al., 2007; Bicchieri and Xiao, 2009).

A similar explanation seems however not to be consistent with our findings among men who instead save less when information about earnings is private. There is no difference in the sharing norm for men between the PRIVATE and the PUBLIC treatment. Therefore, we cannot explain the difference in saving behaviour for men across treatments with a change in the social norm. However, it is possible that the treatment difference identified among male savers results from intra-household differences in preferences over income allocation. It has commonly been argued that, in the case for poorer countries, women tend to have the main responsibility of the

household and allocate a greater share of their income for family needs whilst men allocate more income towards personal expenditures.²¹ It is thus possible that men in our sample save less when earnings are private because this allows them to avoid bargaining over this money with their spouse who instead prefers that the money is saved.²²

5. CONCLUSIONS

In this study we present results from a field experiment run in rural Kenya. We provide empirical evidence of the existence of sharing norms that prescribe extensive sharing of income at the expense of own savings. We identify these norms using a modified version of the norm-elicitation method developed by Krupka and Weber (2013).

A series of studies in developing countries find that individuals are willing to make sub-optimal economic choices for themselves in order to keep information about earnings private (e.g. Baland et al., 2011; Jakiela and Ozier, 2015). We contribute to this literature by showing that, for women, sharing norms are weaker when information about earnings is private compared to when this information is known by everyone in the community. Our results suggest that women may feel obliged to share their income with others when earnings are public information as this is strongly prescribed by the social norm, but that this sharing norm is less strong when earnings are private and thus enables individuals to make more self-interested choices.

Studying subjects' bank account usage we show that these findings translate into effects on actual saving behaviour. We find that women with relatively low savings save more if they have private information about experimental earnings compared to if their earnings have been publicly announced to others in the community.

As formal means of savings often is expensive and difficult to access for individuals living in poorer communities, they are often left with informal alternatives such as saving in

²¹ Bruce (1989) presents a series of case studies suggestive of this view. Duflo (2012) argues that there is evidence for gender differences in preferences and by extension relevant outcomes for the welfare of the household. For example women in control of income have a positive effect on child nutrition.

²² The result that men save more when information about earnings is public is also in line with findings presented by Castilla and Walker (2013). They run a field experiment in Ghana and find that men allocate more money into assets and investments from having received income publicly observed by others in the community whilst no such effect is found for privately observed income.

groups, through a friend or at home. These informal alternatives tend to imply that others know about the amount of money that is being saved (with the possible exception for savings kept at home for individuals who live on their own). Making formal means of savings that can ensure privacy more accessible to the poor may thus help a substantial group of individuals to save more.

REFERENCES

- Angelucci, M., and G. De Giorgi. 2009. Indirect effects of an aid program: how do cash transfers affect ineligibles' consumption? *The American Economic Review* 99(1), 486–508.
- Baland, J.-M., C. Guirkinger, and C. Mali. 2011. Pretending to be poor: Borrowing to escape forced solidarity in Cameroon. *Economic Development and Cultural Change* 60(1), 1–16.
- Barr, A., and M. Stein. 2008. Status and Egalitarianism in Traditional Communities: An Analysis of Funeral Attendance in Six Zimbabwean Villages.(mimeograph).
- Beekman, G., M. Gatto, and E. Nillesen. 2015. Family networks and income hiding: evidence from lab-in-the-field experiments in rural Liberia. *Journal of African Economies*, ejv008.
- Bicchieri, C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.
- Bicchieri, C., and E. Xiao. 2009. Do the right thing: but only if others do so. *Journal of Behavioral Decision Making* 22(2), 191–208.
- Boltz, M., K. Marazyan, and P. Villar. 2015. *Income hiding and informal redistribution: A lab-in-the-field experiment in Senegal*. Citeseer.
- Bruce, J. 1989. Homes divided. *World Development* 17(7), 979–991.
- Brune, L., X. Giné, J. Goldberg, and D. Yang. 2015. *Facilitating Savings for Agriculture: Field Experimental Evidence from Malawi*. National Bureau of Economic Research.
- Castilla, C., and T. Walker. 2013. Is ignorance bliss? The effect of asymmetric information between spouses on intra-household allocations. *The American Economic Review* 103(3), 263–268.
- Coate, S., and M. Ravallion. 1993. Reciprocity without commitment: Characterization and performance of informal insurance arrangements. *Journal of development Economics* 40(1), 1–24.
- Cole, S., T. Sampson, and B. Zia. 2011. Prices or knowledge? What drives demand for financial services in emerging markets? *The Journal of Finance* 66(6), 1933–1967.
- Dana, J., D.M. Cain, and R.M. Dawes. 2006. What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes* 100(2), 193–201.
- Dana, J., R. Weber, and J. Kuang. 2007. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33(1), 67–80.
- Di Falco, S., F. Feri, P. Pin, and X. Vollenweider. 2016. Ties that Bind: Redistributive Pressure and Economic Decisions in Village Economies.
- Duflo, E. 2012. Women empowerment and economic development. *Journal of Economic Literature* 50(4), 1051–1079.
- Dupas, P., and J. Robinson. 2013. Why Don't the Poor Save More? Evidence from Health Savings Experiments. *American Economic Review* 103 (June), 1138–1171.
- Fafchamps, M., and S. Lund. 2003. Risk-sharing networks in rural Philippines. *Journal of development Economics* 71(2), 261–287.
- Fafchamps, M., and B. Minten. 1999. Relationships and traders in Madagascar. *The Journal of Development Studies* 35(6), 1–35.

- Fafchamps, M., and B. Minten. 2002. Returns to social network capital among traders. *Oxford economic papers* 54(2), 173–206.
- Falk, A., A. Becker, T.J. Dohmen, D. Huffman, and U. Sunde. 2016. The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences.
- Gächter, S., B. Herrmann, and C. Thöni. 2004. Trust, voluntary cooperation, and socio-economic background: survey and experimental evidence. *Journal of Economic Behavior & Organization* 55(4), 505–531.
- Hoff, K., and A. Sen. 2011. THE KIN SYSTEM AS A POVERTY TRAP? *Poverty Traps*.
- Jakiela, P., and O. Ozier. 2015. Does africa need a rotten kin theorem? experimental evidence from village economies. *The Review of Economic Studies*, rdv033.
- John, A. 2016. When Commitment Fails-Evidence from a Regular Saver Product in the Philippines. *EOPP Discussion Papers* 55.
- Kenya, F. 2014. Kenya Financial Diaries Shilingi Kwa Shilingi—The Financial Lives of the Poor. *Nairobi: FSD Kenya, August*. http://www.fsdkenya.org/pdf_documents/14-08-08_Financial_Diaries_report.pdf.
- Krupka, E., and R.A. Weber. 2013. Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association* 11(3), 495–524.
- Lazear, E.P., U. Malmendier, and R.A. Weber. 2012. Sorting in Experiments with Application to Social Preferences. *American Economic Journal: Applied Economics* 4(1), 136–163.
- Mani, A., S. Mullainathan, E. Shafir, and J. Zhao. 2013. Poverty impedes cognitive function. *science* 341(6149), 976–980.
- Olken, B.A., and M. Singhal. 2011. Informal taxation. *American Economic Journal: Applied Economics* 3(4), 1–28.
- Platteau, J.-P. 2000. *Institutions, social norms, and economic development*. Vol. 1. Psychology Press.
- Raven, J. 2000. The Raven’s progressive matrices: change and stability over culture and time. *Cognitive psychology* 41(1), 1–48.
- Schaner, S. 2016. The Persistent Power of Behavioral Change: Long-Run Impacts of Temporary Savings Subsidies for the Poor. *NBER Working paper series*, Working Paper 22534.
- Townsend, R.M. 1994. Risk and insurance in village India. *Econometrica: Journal of the Econometric Society*, 539–591.