

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

## Journal of Economic Behavior and Organization

journal homepage: [www.elsevier.com/locate/jebo](http://www.elsevier.com/locate/jebo)

# One size does not fit all: Plurality of social norms and saving behavior in Kenya



Hanna Fromell<sup>a,b</sup>, Daniele Nosenzo<sup>b,\*</sup>, Trudy Owens<sup>c</sup>, Fabio Tufano<sup>c</sup>

<sup>a</sup> Department of Economics, Econometrics, and Finance, University of Groningen, the Netherlands

<sup>b</sup> Department of Economics and Business Economics, Aarhus University, Denmark

<sup>c</sup> School of Economics, University of Nottingham, UK

## ARTICLE INFO

### Article history:

Received 9 March 2021

Revised 17 September 2021

Accepted 18 September 2021

### Keywords:

Sharing norms

Forced solidarity

Social pressure

Savings

Social norms

Krupka–Weber method

Lab-in-the-field experiment

## ABSTRACT

Using the Krupka–Weber norm-elicitation technique in a lab-in-the-field experiment in rural Kenya, we measure the social norms that regulate the trade-off between wealth accumulation through saving and sharing income with kin and neighbors. We find a plurality of norms: from a strict sharing norm prohibiting any form of wealth accumulation to a norm that allows moderate wealth accumulation. We show that several individual and social network characteristics predict the norms perceived and that the pro-saving norm becomes majoritarian when an individual can conceal their income from kin and neighbors. In further exploratory analysis, we find some evidence that the type of norm individuals perceive mediates the effect of income secrecy on actual saving behavior. Taken together, our results highlight the importance of measuring social norms when devising pro-saving policy interventions.

© 2021 The Authors. Published by Elsevier B.V.  
This is an open access article under the CC BY license  
(<http://creativecommons.org/licenses/by/4.0/>)

## 1. Introduction

Informal redistribution of resources is commonplace in many developing countries, particularly in sub-Saharan Africa. Reciprocal gift-giving and informal exchanges of money and goods play an important role in these communities by insuring individuals against natural risk, as well as by fostering and cementing social ties with relatives, friends and neighbors (Platteau, 2014). It has been argued that these practices are so ingrained in these communities that sharing one's wealth with others is often perceived as a social obligation – a *social norm* that is actively enforced within the community by means of social and economic sanctions (e.g., Platteau, 2000; Barr and Stein, 2008).<sup>1</sup> Importantly, the literature has pointed out that there may also be a dark side to this social norm of sharing: the obligation to share the yields of one's work with others may discourage effort, entrepreneurship, investments in human capital, and the accumulation of assets and wealth in these communities (e.g., Platteau, 2000, 2014; Hoff and Sen, 2006).

Much of the evidence on the existence, and negative effects, of sharing norms in developing countries comes from field studies showing that individuals are willing to incur material costs in order to keep their income private (e.g., Baland et al.,

\* Corresponding author.

E-mail address: [daniele.nosenzo@econ.au.dk](mailto:daniele.nosenzo@econ.au.dk) (D. Nosenzo).

<sup>1</sup> The focus here is on *injunctive* norms that define what a person ought to do in a specific situation, as opposed to *descriptive* norms that describe what people normally do. See Bicchieri (2006) for a discussion of the differences between injunctive and descriptive norms.

2011; Beekman et al., 2015; Jakiela and Ozier, 2016; Squires, 2018; Boltz et al., 2019), and that mechanisms that allow individuals to shield wealth from public scrutiny enable them to accumulate more savings (Dupas and Robinson, 2013; Jakiela and Ozier, 2016) and to reduce the rate at which they spend resources (Goldberg, 2017).<sup>2</sup> While these behavioral patterns are consistent with the existence of sharing norms, little is known about the actual pervasiveness of these norms, or the shape that they take. When individuals trade off wealth accumulation and sharing with others, are there levels of wealth accumulation that are considered socially acceptable, or alternatively, does the social norm proscribe any form of wealth accumulation? How strong is the social disapproval for deviating from the level of wealth accumulation prescribed by the norm? Does the possibility to shield one's wealth from public scrutiny affect the perception of what is appropriate, in addition to lowering the cost of deviation? Finally, are these normative prescriptions widely shared within a community, or is there instead substantial heterogeneity in normative views?

In this paper, we open the black-box of sharing norms and provide direct empirical evidence of the existence, shape and pervasiveness of these norms among rural communities in Western Kenya. We run a lab-in-the-field experiment where we exploit a recent advance in the experimental literature on social norm compliance (Krupka and Weber, 2013) to measure the prevailing perception of the social appropriateness of wealth accumulation through saving when an individual is confronted with demands to share income with other members of the community. In the experiment, participants are described a vignette in which a hypothetical person receives unanticipated windfall income that they plan to save to grow their own business. The person, however, is asked to share the money with a friend in need. Participants are incentivized to report their perception of the social (in)appropriateness of the individual keeping various amounts of the windfall income rather than sharing it with the friend.

We find that a *plurality of norms* exists in our context. Only a minority of participants perceive a “strict” norm of sharing that proscribes any form of wealth accumulation through saving. Most participants recognize moderate accumulation of wealth as socially acceptable, although they do view full sharing as most appropriate. About a quarter of participants even perceive a “pro-saving” norm, whereby keeping most of one's wealth for oneself is the most appropriate course of action.

We show that a number of individual characteristics, as well as features of the social network in which the individual is embedded, are predictive of the type of norm that a person perceives. Specifically, the strict sharing norm that prohibits wealth accumulation is positively associated with age, lower levels of income and liquid assets, as well as traits such as unwillingness to contribute to public goods and negative reciprocity. In contrast, network characteristics, such as being connected to a larger set of people and being embedded in a more cliquish social network, work against the strict norm of sharing in favor of the pro-saving norm.

We then ask how mechanisms that allow individuals to keep their income private – a prominent feature in the development literature on sharing norms – affect the perceived obligation to share. We describe to a subset of experimental participants a modified version of the vignette in which the hypothetical person can keep the income secret from others, and can thus plausibly refuse to share it with the friend without appearing unwilling to do so. We find that in this setting the share of participants who endorse the pro-saving norm doubles to 50%. This is an important finding because it suggests that mechanisms that promote income privacy can be useful to counteract the potentially negative effects of sharing norms not only by lowering the cost of norm deviations (as argued by the previous literature), but also by changing the perception of the social norm in such situations.

In the last part of the paper, we conduct an exploratory analysis to investigate the effectiveness of privacy in counteracting pressures to share. We study how giving participants the opportunity to keep their income secret affects actual accumulation of wealth. In collaboration with a local commercial bank, we open bank accounts for each individual who participates in our experiment and pay their entire experimental earnings directly into these new accounts. We then offer a generous interest rate to create a strong incentive to save the experimental earnings. We promise participants that we will pay an additional 50% on the share of earnings that they keep in the account for an entire month from the day of the experiment (which we can monitor via the bank).

We study levels of savings across two treatments that differ in the procedure used to announce earnings at the end of the experimental sessions. In one treatment (PUBLIC), each participant's earnings were publicly announced to the whole session. In the other treatment (PRIVATE), earnings were communicated to each individual in private. The latter group of participants therefore had the opportunity to keep their earnings secret – very much like the person in the vignette – and may therefore be better able to escape the social pressure to share income with other village members.

We find that on aggregate privacy has only a small and insignificant average treatment effect on savings. However, we also find some indication that there is heterogeneity in treatment effects across participants who perceive different norms of sharing.

Our paper contributes to the growing literature on the impact of sharing norms on the economic decision-making of individuals in developing countries (e.g., Di Falco and Bulte, 2011, 2015; Dupas and Robinson, 2013; Jakiela and Ozier, 2016; Goldberg, 2017; Squires, 2018; Boltz et al., 2019). One message that emerges from this literature is that easing the pressure of a sharing norm can have beneficial effects for the economic and financial decisions of individuals who live in poor com-

<sup>2</sup> However, see also Brune et al. (2016), who find less conclusive evidence for this effect. Another strand of the literature shows that there is a negative relation between savings and investment and the size of one's social network (e.g., Di Falco and Bulte, 2011, 2015; Grimm et al., 2017). For income hiding within the household, see Ashraf (2009) and Bulte et al. (2018). Relatedly, Baseler (2021) finds that urban migrants strategically underreport their income to their parents and home contacts to reduce their remittance obligations.

munities. Our results confirm the existence of the pressure to share in these rural communities. However, we also show that a plurality of norms may exist, and that these different normative ideals may carry different implications for the effectiveness of interventions that aim to mitigate social pressure by other community members. This highlights the importance of furthering our understanding of the interplay between policy interventions and norms.

Our results also show that strong sharing norms are not very pervasive – at least not in the context of our lab-in-the-field experiment. A common understanding of the appropriateness of sharing does not appear to exist at the “global” community level in our villages. Rather, different understandings seem to emerge within subgroups of individuals that share similar characteristics or are part of the same intra-village networks. In the aggregate, only about a quarter of our participants perceive a very strict sharing norm that prohibits any form of wealth accumulation. Most people find it acceptable to resist the pressure to share and instead accumulate some wealth. Moreover, this stricter normative ideal is easily displaced: introducing income privacy changes the majority view of what is most acceptable in favor of wealth accumulation. More research is needed to assess whether our findings are specific to the villages we study, or whether the existence of a single common, social norm that favors sharing over individual savings is largely absent in today’s poorer Kenyan communities.

Finally, from a methodological perspective, our findings highlight the importance of studying the effects of social norms *empirically* – by measuring directly the norms that prevail in a particular setting rather than inferring them from observed behavior (on this point see also [Bicchieri, 2017](#)). We show that combining behavioral data with data on the normative beliefs that individuals hold about their behavior can lead to important and novel insights about the mechanisms that underlie human decision-making. We show that measuring norms empirically is of particular importance in predicting the impact of policies aimed at changing norm-driven behaviors: a detailed understanding of the precise shape of the norms that prevail in a community is essential in order to design effective interventions and mitigate the risk of the interventions not only failing but potentially backfiring.<sup>3</sup>

The rest of the paper is organized as follows. [Section 2](#) illustrates the design of the study and the lab-in-the-field experiment. [Section 3](#) reports summary statistics and balance checks. [Section 4](#) presents our results. [Section 5](#) concludes.

## 2. Design of the study

The study took place in six villages outside Eldoret town in Uasin Gishu County, Western Kenya. The villages were randomly selected conditional on containing approximately 100 households, being located at a short distance from Eldoret town (reachable within 30 min by car), and not lying in direct proximity to one another. The study ran between April and June in 2016 with the help of a team of 10 local enumerators.

In each village, we initially conducted a household survey aimed at gathering background characteristics, including social network information, for all village households. We then ran a lab-in-the-field experiment with a randomly selected subsample of 228 members of the surveyed households.<sup>4</sup> The experiment measured subjects’ perceptions of prevailing social norms of sharing in their villages through the telling of a vignette in which subjects rated the appropriateness of sharing different amounts of windfall income with another villager, instead of saving the money to grow one’s own business. Across two treatments (either PUBLIC or PRIVATE), we varied whether the windfall income described in the vignette was made public or kept private. Moreover, we used the treatment variation to match the way we announced experimental earnings to the subjects at the end of the session, namely, in public or in private, and studied the effect this had on their use of the experimental money after the experiment.

In line with our norm-elicitation vignette, we are particularly interested in the trade-off between sharing the experimental money with others in the community and using it for savings/investment purposes. In an ideal setting, we would have measured, in the weeks or months after the experiment, the fraction of experimental money that subjects decided to share with others, instead of saving it. However, this measurement would have relied on self-reports, which may have been imprecise, noisy and potentially biased by socially desirable responding. Therefore, we decided to rely instead on a behavioral measure of subjects’ ability to save the money earned in the experimental session.

One concern with this approach is that, in principle, subjects may decide not to save the experimental earnings for a multitude of reasons, not just because they face pressures to share the money with others. However, given that subjects are randomly assigned to treatment, and the treatments only vary in the extent to which subjects could credibly conceal their experimental earnings from others (low in PUBLIC, high in PRIVATE), we interpret treatment differences in the amounts saved as a *ceteris paribus* measure of how income privacy affected the trade-off between saving the experimental money vs. sharing it with others.

To measure subjects’ saving behavior, we collaborated with a local commercial bank (Kenya Commercial Bank, KCB), whereby, at the end of the experiment, each participant was helped to open a bank account in which we paid their experi-

<sup>3</sup> On this point see for instance [Kölle et al. \(2020\)](#), who in their Study I show that an intervention designed to promote voter registration in the UK (rewarding registrations with monetary lottery prizes) is ineffective. In their Study II, by eliciting social norms, Kölle et al. show that this is because the intervention weakened the norm of registering to vote. This exemplifies the importance of measuring the impact of interventions on norms in order to accurately predict their effectiveness.

<sup>4</sup> For each household that we selected to invite to the experiment, we randomly selected whether to invite the female head of household or the male head of household. Sometimes the selected household-head was not available, in which case the household-head of the opposite gender was interviewed and invited to the experiment instead. In a few cases, a member of the household who was not its head participated.

mental earnings. We tracked withdrawals from the account for a month after the experiment to measure the impact of the two treatments on their saving behavior.<sup>5</sup> Below we describe in detail each component of our study.

### 2.1. Household survey

Our aim was to run a household survey to collect general information about all 755 resident households in the six villages (and not just those who were then selected to participate in the experiment). We achieved this for 74% of households across all villages. However, in 89 cases, the households had time constraints and we had to run a reduced version of the survey, collecting only basic information about the household; and in 111 cases we could not gather information for a variety of reasons, predominately because household members were not present.

For most households we were able to collect information on a range of socio-demographic characteristics (including, age, gender, marital status of subjects being interviewed),<sup>6</sup> preferences and attitudes (risk, patience, trust, altruism, and positive and negative reciprocity – based on the general questions outlined by Falk et al., 2018), perceptions of social pressures from family and friends,<sup>7</sup> as well as their social networks.

Social networks were measured using the approach introduced by Banerjee et al. (2013). The module, described in more detail in appendix A of the Supplementary Materials (SM), consists of a series of questions asking respondents to name people with whom they have various types of relationships within the village. We use this rich dataset to reconstruct the network characteristics of the individuals who participated in the experiment.

### 2.2. Lab-in-the-field experiment: social norms of sharing

In each village, on the day after the household survey, a random sample of individuals participated in a lab-in-the-field experiment designed to elicit social norms of sharing. Following the economics literature on norms, we define *social norms* as shared understandings, among a reference group of people, of the appropriateness or inappropriateness of certain actions in a given context (Akerlof and Kranton, 2000; Ostrom, 2000; Bicchieri, 2006; López-Pérez, 2008; Krupka and Weber, 2013; Bicchieri, 2017; Fehr and Schurtenberger, 2018; d'Adda et al., 2020). Krupka and Weber (2013) introduced an incentivized experimental module to empirically measure an individual's perceptions of the norms that pertain to a given situation.<sup>8</sup> Social norms elicited using this methodology have been shown to predict actual behavior in a variety of contexts, from reciprocity to bribery and discrimination (e.g.; Gächter et al., 2013; Krupka and Weber, 2013; Nikiforakis et al., 2014; Banerjee, 2016; Gächter et al., 2017; Krupka et al., 2021; Barr et al., 2018; Chang et al., 2019; Bicchieri et al., 2020; Dimant et al., 2020).

Adapting this norm-elicitation module, we use an *incentivized vignette experiment* to elicit the perceptions, among our sub-sample of participants, of the *norms of sharing* that prevail in their village. Compared to other modules to elicit normative opinions (e.g., short survey items or questionnaires), the use of incentivized vignette experiments presents several advantages. First, the use of incentives reduces the issue of socially desirable responding, which may be particularly serious in regard to the evaluation of behaviors that are subject to social approval or disapproval. Furthermore, short survey items often lack the level of detail that is needed to gain a deeper understanding of norms. Norms are, in fact, typically context-dependent (Bicchieri, 2006; Krupka and Weber, 2013) and it is therefore important to precisely define the situation being evaluated in order to obtain accurate measurements of the norms that apply to it. Moreover, it is important to measure opinions about a rich range of behaviors that pertain to a given situation, including deviations from the action prescribed by the norm, so as to reconstruct the precise structure of social incentives that sustain the norm (Görges and Nosenzo, 2020b). It is established in the literature that carefully constructed vignettes are particularly well suited for achieving the level of detail and specificity that are necessary to measure norms.

In our vignette experiment, participants were read one of two vignettes describing a hypothetical situation faced by a person in their own village. Each vignette was identical except in whether the knowledge of extra income was publicly known. In one vignette (PUBLIC), the person has recently earned extra income to the value of 900 Ksh, which everyone in the village is aware of (equivalent to approximately two weeks per capita income among rural Kenyans at the time of the experiment). The vignette does not specify the source of the extra income, but it makes clear that the person intends to save the money to grow their own business with the purpose of increasing their future household wealth (we did not specify what this business investment was, although in the context of our villages it may have been interpreted as investment in

<sup>5</sup> Subjects were told at the first contact with the enumerator that they had to agree to open an account free of charge with our partner bank in order to be eligible to take part in the study. Moreover, during the experiment we obtained written consent from all participants to give us access to their bank account statements for the relevant study period. Participants were free to withdraw from the study or deny access to their bank account while still receiving their full experimental earnings (no one did so). The study received ethical approval from the University of Nottingham's Research Ethics Committee and research permission from the Kenyan National Commission for Science, Technology and Innovation.

<sup>6</sup> Moreover, we also measured income and savings only for the individuals who were randomly selected to participate in the experiment (due to the length of this part of the questionnaire).

<sup>7</sup> These included an estimate of the claims that would be made by others in the event that the person would have cash savings at home (John, 2020), an estimate of the gifts and loans that have been transferred to others in the past month, and a measure of the individual's feeling of obligation to give money to other villagers asking for it (Dupas and Robinson, 2013).

<sup>8</sup> See Bicchieri and Xiao (2009) for an alternative method and Görges and Nosenzo (2020a, 2020b) for a discussion of the pros and cons of the different methods.

agricultural or farming activities, which most of our participants would have had some experience with). However, a friend approaches this person and asks for the money as they need it to pay their children's school fees. The person has to decide how much money to give to the friend, choosing among six actions: from giving the whole 900 Ksh to the friend and saving nothing, to giving nothing to the friend and saving the whole 900 Ksh for themselves. The six possible actions described in the vignette are: share 0%, 20%, 40%, 60%, 80% or 100% of the 900 Ksh (these actions were described to participants as amounts in Kenyan Shillings and not percentages, although on their answer sheets participants could also see pie charts indicating the corresponding percentage amounts).

Participants were asked to rate the appropriateness of each of these six actions using a four-point scale, ranging from "very socially appropriate" to "very socially inappropriate". Participants were told that, by "socially appropriate", we meant behavior that most people would agree is the "correct thing to do" (see SM B for details about the instructions). These ratings inform us of subjects' perception of the social appropriateness of saving money for oneself instead of sharing the money with another villager who is in need. As such, the vignette measures subjects' perception of the social norm of sharing in the presence of a trade-off between sharing and wealth accumulation, and not about the social norm of sharing in general (or the social norm that regulates wealth accumulation in general).

Importantly, this vignette focuses on situations where an individual's income is *observable* by other members of the village. This is common in rural, close-knit communities, such as the one where we ran our lab-in-the-field experiment, where individuals can, and continuously do, closely monitor each other's yield from agricultural production (Platteau, 2006).

To examine the extent to which mechanisms or institutions helping individuals to keep information about their income private influences the social norm of sharing, and hence saving behavior, we designed a second version of the vignette in which the person's income was *not observable* by others in the village. This second vignette (which we call PRIVATE) was identical to the vignette described earlier, except participants were told that no other person in the village (including the friend asking for money) knows the amount of extra income that the person has earned. In this version, no one can infer from the person's decision how much they have saved for themselves. All other aspects of the vignette were kept exactly the same, including the description of the windfall income and its amount, the reason why the person in the vignette wanted to save the income, the reason why the friend needed the money, and the actions available to the decision-maker. Thus, our vignette keeps constant – insofar as possible – the material aspects of the decision situation, and only varies the social information available to others in the village about the person's income.

This subtle variation in the context of the decision situation across the two vignettes allows us to study how social norms of sharing are influenced by the information that others have on an individual's income.<sup>9</sup> The reason why unobservability of income may influence the perception of what constitutes appropriate behavior is that it may change the way the friend in need would rationalize a rejection of their request for help. When income is observable, a rejection may signal that the person in the vignette is unwilling to share, but when income is unobservable, it may be attributed to lack of financial means. Psychologically, a rejection may be less painful for the friend in need in the latter case. If norms respond to a "lesser harm principle", whereby the inappropriateness of an action is proportional to the harm caused, then refusing to share may be viewed as less inappropriate when income is unobservable rather than observable.<sup>10</sup>

Importantly, the task was incentivized both for the PUBLIC and PRIVATE conditions. Participants were told that, at the end of the experiment, one of the six actions they had rated would be randomly selected and their rating for that action would be compared to the rating of another randomly selected villager, who was also participating in the same experimental session and rating the same vignette. If the two ratings were the same, then both subjects were paid 900 Ksh each, otherwise they were paid nothing.<sup>11</sup>

Note that, although participants are therefore incentivized to coordinate their ratings with other subjects who are in the same session as themselves, they are not allowed to communicate with each other, and so they need to rely on focal points present in the task to solve the coordination problem. As discussed in Krupka and Weber (2013), if a social norm exists for the behavior described in the vignette, this constitutes a very salient focal point in the task, which participants may refer to when rating the actions in the vignette. Participants' ratings therefore allow us to (indirectly) measure the social norms that are perceived to apply in situations that involve a tension between wealth accumulation and sharing money with other members of one's village.

In conjunction, we were also interested in studying how decreasing observability of one's own income affects *actual* saving behavior. As we discuss in the next section, we can observe the saving decisions of our participants for one month

<sup>9</sup> Our vignettes focus on extreme cases where either no one or everyone in the village knows about the extra income earned by the person. While these extreme cases may not be the most realistic in practice, they help create a stark contrast between the two information conditions, thus providing the ideal testing ground to detect any existing differences between conditions.

<sup>10</sup> The experimental literature on norm compliance suggests that this principle does indeed apply to normative judgments. Krupka and Weber (2013), for instance, show that giving nothing in the dictator game is viewed as relatively more acceptable if the recipient is kept in the dark about the possibility of receiving positive amounts. In a labor context, Nikiforakis et al. (2014) show that it is relatively more acceptable for senior workers to exploit junior workers if the latter are not fully aware of the extent of the exploitation.

<sup>11</sup> Alternatively, we could have asked participants to match the most common (modal) response by other subjects. Both elicitation methods have been used in the experimental literature. Note that our method also incentivizes subjects to match the modal response, since individuals are paired at random. The advantage of using this method is that it is easier to implement in the context of a pen-and-paper field experiment as it does not require data aggregation and analysis (to compute the mode) in order to pay participants.

after the lab-in-the-field experiment. To study how income observability affects their saving decisions, we extend the two treatments by varying the way participants were paid at the end of the experiment.

In the PUBLIC treatment, payments were publicly announced to everyone in the session. We did this by writing the income of every session participant on a board placed in the front of the room where the experiment was conducted. We then called participants one-by-one to the front of the room to collect a paper slip on which their payment amount was written. In contrast, in the PRIVATE treatment, payments were made in a separate room, and each participant was privately informed of their own earnings.<sup>12</sup> No one else in the session was informed of other participants' earnings and so this information could be kept private as long as a participant chose not to disclose.

Importantly, in order for participants in the PRIVATE treatment to be truly able to maintain private information over their earnings, it was essential that the amount earned by subjects in the experiment varied sufficiently across individuals. This ensured that – just as the person in our norm-elicitation task – participants in PRIVATE could credibly deny having received high earnings and thus possibly escape the demands that others in the village may make over their income. To achieve this variation in earnings we embedded the norm-elicitation task into a larger study consisting of four tasks in total (a public goods game, a cognitive ability task, a financial literacy task, and the norm-elicitation task – see below and SM C for details). Participants were not informed about the results of any of these tasks, including the norm-elicitation task, until the very end of the experiment. Even then, participants were only informed about their earnings from one randomly-selected task that they were paid for. Participants did not learn what others chose in this (or any other) task, although in some cases they could infer it (e.g., in the two-person public goods game). In the norm-elicitation task, earnings also depended on another random component, namely, which of the six actions was selected and whether it matched that of another participant. Overall, we achieved sufficient variation in subjects' earnings, which ranged from 300 Ksh to 1500 Ksh (including a 300 Ksh participation fee) in both treatments, averaging 940 Ksh in PRIVATE and 863 Ksh in PUBLIC (see SM C for the distributions of earnings in the two treatments). We believe this guaranteed sufficient privacy of earnings for participants in the PRIVATE treatment.

Using community buildings within each village, the two treatments were run on the same day, in two separate sessions, one after the other to avoid contamination, with each lasting for approximately 4.5 h with a break with refreshments in between.<sup>13</sup> Sessions contained between 16 and 23 participants each, with participants having on average 1.86 other people belonging to their social network present in the same experimental session (the median was 2). Participants were randomly assigned to either the PUBLIC or the PRIVATE treatment. The treatment assignment was determined at the individual level: each person was randomly assigned to one of the two treatments and invited to attend the corresponding treatment session. Individuals were turned away if they showed up to the wrong session. This precludes any possible selection into session and ensures that the treatment randomization occurred at the individual level. Participants assigned to the PUBLIC treatment received the PUBLIC vignette in the norm-elicitation task and were paid in public at the end of the experiment, while those assigned to the PRIVATE treatment received the PRIVATE vignette and were paid in private. Technically, this means that the effect of the PUBLIC/PRIVATE treatment measures the compound effect of being shown a norm-elicitation vignette in a public/private setting and of having been paid in public/private. However, note that subjects did not receive any information about the ratings of other subjects throughout the experiment. Thus, being exposed to the vignette cannot convey any information about the prevailing norm. We have 106 participants in PRIVATE and 122 in PUBLIC, the small imbalance being due to variation in the number of people who showed up to the sessions.

As mentioned above, the experiment also included three other incentivized tasks in addition to the norm-elicitation task: a two-person, one-shot linear public goods game; a cognitive ability task where participants solved three Raven's matrices; and a financial literacy task based on Cole et al. (2011). Table 1 shows the timeline of the experiment, while SM C contains a more detailed description of the three additional tasks. These tasks were identical across treatments. Choices in these tasks are included as controls in the analysis below.

All instructions were read to participants aloud by an experimenter. Instructions for each game were complemented by publicly-displayed visual aids that two other experimenters used to illustrate the content of the instructions. In all games, before collecting decisions, participants were asked a few control questions about the game to test their understanding. The experimenters checked participants' answers on an individual basis so that any clarification needed could be made in private. The experiment was run with pen and paper. During the experiments all choices and answers were taken individually. Anonymity was ensured by providing each participant with a squared box with tall edges. Participants had to make all choices inside their box without ever removing any material from the box. (See SM B for further details on the experimental instructions, visual aids and the script used to run the experiments). These procedures ensure that no information transmission or communication were possible among participants within a session.

<sup>12</sup> Participants in PRIVATE also received a paper slip with their payment amount written on it, but we then made them destroy it so that no one could know how much they had earned.

<sup>13</sup> The morning session lasted between 9:30 am and 2:00 pm and the afternoon session between 2:00 pm and 6:30 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.

**Table 1**  
Timeline of a session.

Sequence of activities	Activity description
1	Participants welcomed and signed off the session register; <b>session starts</b> .
2	<b>General Instructions</b> read aloud.
3	Instruction for Game 1 ( <b>Public Good Game</b> ) read aloud using visual aids. Public Good Game decisions collected.
4	Instructions for Game 2 ( <b>Raven's matrices task</b> ) read aloud using visual aids. Raven's matrices responses collected.
5	<b>Break</b> with refreshments.
6	Instructions for Game 3 ( <b>Financial literacy questions</b> ) read aloud using visual aids. Financial literacy questions responses collected.
7	Instructions for Game 4 ( <b>Krupka–Weber task</b> ) read aloud using visual aids. Krupka–Weber task responses collected.
8	<b>Payment</b> stage (in private or public according to treatment).
9	<b>Bank account setup</b> explained by experimenter; subjects escorted one by one to private meeting with bank officer where they could open bank account.
10	Participants dismissed; <b>session ends</b> .

### 2.3. Saving behavior

After all tasks had been played, participants were informed of their experimental earnings, with payment procedures varying between treatments as described above. Participants then left the experimental room and privately met an officer of our partner bank who helped them open a standard bank account in their own name (at no cost).<sup>14</sup> On the following day, their bank account was activated, and their experimental earnings were paid directly, and in full, into their bank account.

Subjects were promised a very competitive interest rate for the experimental money deposited in their new account. They were told that throughout the month following the experiment we would track any experimental money they withdrew from the account, and that at the end of the period, we would deposit an additional sum to the value of 50% of the money remaining in the account.<sup>15</sup> This generous interest rate was chosen so that it would be unlikely that subjects could find a more lucrative way to save the money. We use the share of experimental money *not* withdrawn by a subject at the end of the first month as a measure of their saving behavior. In Section 4, we check whether our treatment intervention (PUBLIC or PRIVATE payments) had an impact on the amount saved in the month after the experiment.

We took great care to explain the mechanism for calculating the interest to subjects at the end of the experiment (see SM B). We also sent each participant two text messages (one on the day when the money was paid into the account, the other 15 days later) to remind them how the interest would be determined and when it would be paid out. Other practical information about account usage, including the fees associated with account activity such as depositing, withdrawing and transferring money via the mobile-money-transfer service M-Pesa, was explained to participants directly by bank officers when they opened the account.<sup>16</sup>

One possible limitation of our measure of saving behavior is that it does not capture potential pressures to share the money *after* the interest has been paid into subjects' accounts. That is, subjects may have been pressured to keep their money in their account for the first month so that they could then share a larger amount with other community members afterwards. To the extent that such a strategy may have been used, our measure of saving behavior captures a lower bound of the effect of the treatment intervention on savings. In Section 4.2, we report some robustness checks to test the extent of this issue.

### 3. Attrition, balance check and subject characteristics

Table 2 presents balance checks and subject characteristics that will be used in the analysis in the next section. We use data on socio-demographic characteristics: age, gender, marital status, education level, income, and liquid assets reported

<sup>14</sup> Bank usage is generally low in our sample. Only 50 of our 228 subjects (22%) report either loans or savings with a bank at the time of the experiment.

<sup>15</sup> Very few subjects deposited extra money in the account during the period of the study. The instructions made it clear ex-ante that we would not pay any interest on extra money deposited in the account (including withdrawing and re-depositing the experimental earnings).

<sup>16</sup> For example, the most common way of withdrawing money, i.e., through M-Pesa, was charged a fee of 55 Ksh. Since subjects were informed of these usage fees, we include them in the calculation of the share of experimental money withdrawn. Our results are robust to not counting fees as withdrawals.

**Table 2**  
Balance check and sample characteristics.

Variable Column	Variable Mean [SD] (1)	Balance check: Participate		Balance check: PRIVATE	
		Coefficient (2)	(SE) (3)	Coefficient (4)	(SE) (5)
1 if Male	0.31 [0.46]	−0.11	(0.07)	−0.10	(0.06)
1 if Married	0.79 [0.41]	−0.03	(0.05)	0.01	(0.05)
Age	41.27 [12.75]	2.70	(1.87)	−0.61	(1.71)
Years in Education	8.63 [3.54]	−1.04**	(0.45)	−0.08	(0.47)
Liquid Assets (in Ksh1000)	22.94 [129.63]	2.50	(11.60)	1.56	(17.28)
Monthly Income (in Ksh1000)	23.63 [31.21]	−6.24	(4.38)	4.11	(4.28)
Patience	6.93 [1.94]	0.27	(0.23)	0.18	(0.26)
Risk Aversion	3.23 [2.11]	−0.03	(0.26)	−0.13	(0.28)
Altruism	7.05 [2.27]	−0.23	(0.25)	0.17	(0.30)
Positive Reciprocity	7.51 [1.92]	0.20	(0.28)	0.12	(0.26)
Negative Reciprocity	4.35 [3.20]	0.21	(0.43)	−0.56	(0.43)
Trust	6.47 [1.93]	0.17	(0.27)	−0.04	(0.26)
PGG Contribution (% endow.)	0.32 [0.241]			−0.27	(0.64)
Cognitive Ability	1.86 [0.82]			0.06	(0.11)
Financial Literacy	2.72 [1.06]			−0.09	(0.14)
Feel Obligated to Give	7.62 [1.56]	0.06	(0.20)	0.13	(0.21)
Claim Estimate (% of 900 Ksh)	0.57 [0.28]	0.02	(0.04)	0.01	(0.04)
Gifts & Loans (in Ksh1000)	0.25 [1.75]	−0.30	(0.38)	−0.06	(0.23)
Clustering	0.18 [0.14]	−0.03	(0.02)	−0.00	(0.02)
Degree	12.25 [4.94]	0.57	(0.62)	0.18	(0.65)
<i>P-value – Joint Test</i>		0.347		0.950	

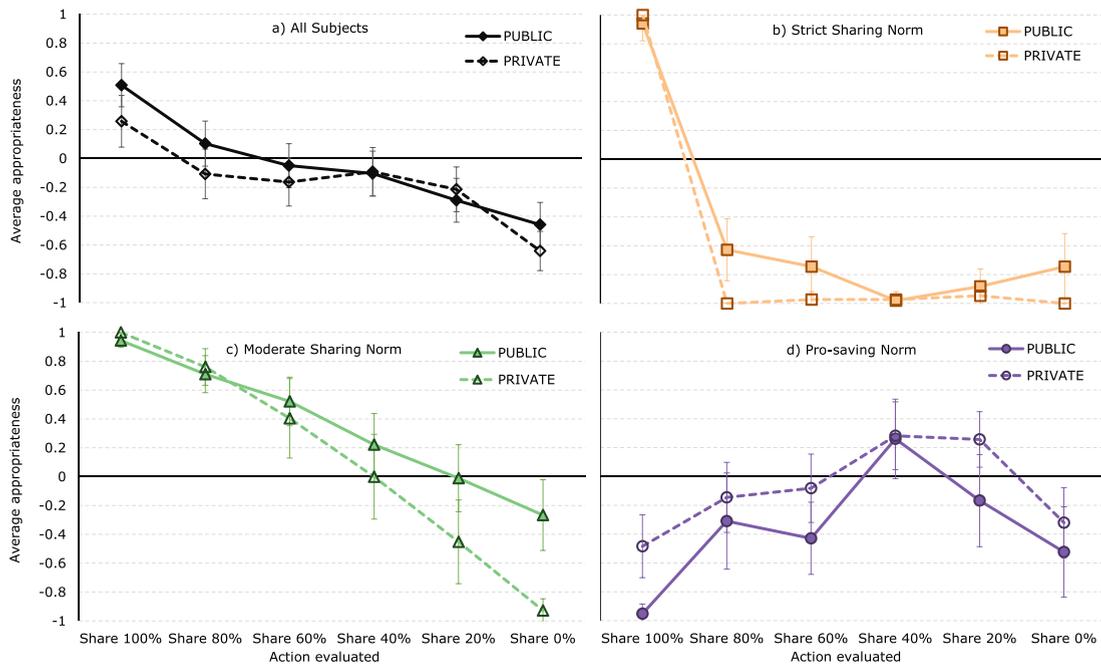
Note: Column 1 reports averages of the characteristics of the 228 participants in the experiment. Standard deviations in brackets. Columns 2–5 report the results of the balancing checks. Robust standard errors in parentheses, derived from OLS regressions of the variable listed in the corresponding row of the table on whether an individual participated in the experiment (Columns 2 and 3) or was assigned to the PRIVATE treatment (Columns 4 and 5). \*\*Significant at the 5% level. \*Significant at the 10% level. Due to missing data for specific variables (negative reciprocity, trust, feel obliged to give), some variable means/regressions are based on slightly smaller samples.

in the household survey. Variables measuring behavioral preferences collected in the household surveys include: patience, risk taking, altruism, positive and negative reciprocity, trust – all unincentivized and measured using Likert scales from 0 to 10 where higher values indicate stronger traits. They are described in more detail in SM D. Variables measuring behavioral preferences collected in the experimental tasks include: fraction of endowment contributed in the public goods game, number of correct answers to the three Raven matrices and four financial literacy questions – all incentivized. We also include variables measuring subjects' perceptions of the financial demands on their income by other village members: a measure of their feeling of obligation to give in to others' demands, rated between 0 and 10; an estimate of the percentage of money they think others would ask for if they had an extra 900 Ksh cash; and an estimate of the gifts and loans transferred to relatives and friends over the past one month, net of any gifts and loans received. Finally, we also include data on network characteristics: clustering and degree – explained in Section 4 in more detail.

Table 2 (Column 1) shows that the 228 individuals who participated in the experiment were on average 41 years of age and predominantly female (69%). They had on average 8.6 years of formal education. Their aggregate monthly income was around 24,000 Ksh (US\$240), which compares with an average per-capita income of 84,000 Ksh in Kenya at the time of the study.

Regarding attrition, although 228 individuals participated in our experiment, 302 were initially invited to take part. The 74 subjects who did not participate dropped out of the study without knowing the treatment they were assigned to. Columns 2 and 3 test whether those 228 who participated were significantly different from the 74 that did not. We report coefficients and robust standard errors from regressions where the independent variable takes value 1 if an individual participated in the experiment and the dependent variable is the corresponding characteristic listed in each table row. We find that only one variable is significantly different between participants and non-participants. Individuals with more years of education were less likely to participate in the experiment. We also find that males were marginally less likely to participate, but the effect is only significant at the 10% level in a logit regression and not significant in a linear probability model. Despite these differences, the joint test of the hypothesis that all coefficients are equal to zero cannot be rejected ( $p = 0.347$ ). We conclude that surveyed individuals who participated in our experiment are not systematically different from those who did not participate.

In Columns 4 and 5, we test whether the experimental participants in the PUBLIC treatment are different from those in the PRIVATE treatment. Similarly, we report coefficients and robust standard errors from regressions of the corresponding characteristic listed in each table row on whether a subject was assigned to the PRIVATE treatment. As one would expect, given the random assignment to treatment, none of the individual characteristics are significantly different across treatments. The joint test confirms this ( $p = 0.950$ ).



**Fig. 1.** Social norms of sharing in PUBLIC and PRIVATE. Note: Each panel plots the average social appropriateness of the six actions evaluated by subjects in the vignette (ranging from 1, “very socially appropriate”, to –1, “very socially inappropriate”). The solid lines indicate the PUBLIC treatment, the dashed lines the PRIVATE treatment. Whiskers represent 95% confidence intervals. The top-left panel contains data for the whole sample, while the remaining three panels show averages constructed on subgroups of subjects defined using hierarchical cluster analysis (see SM E for details).

## 4. Results

### 4.1. Social norms of sharing

We start with the analysis of subjects’ responses in the norm-elicitation task.<sup>17</sup> As is common practice in the social norms literature, we convert subjects’ ratings in the task to numerical values, by assigning evenly-spaced values of +1 to the rating “Very socially appropriate”, +0.33 to “Somewhat socially appropriate”, –0.33 to “Somewhat socially inappropriate”, and –1 to “Very socially inappropriate”. This allows us to construct a *norm function*, which assigns a value of social appropriateness to each action. The function takes positive values for actions that, on average, are evaluated as socially appropriate and negative values for inappropriate actions.

Fig. 1 plots the norm functions elicited in the experiment (the full distribution of norm ratings are reported in Table 3). It consists of four panels. Panel a (top-left, black lines) shows the norm functions constructed by averaging the appropriateness ratings across the whole sample, separately for subjects in the PUBLIC (solid lines) and PRIVATE (dashed lines). The panel shows that, in both treatments, the most appropriate action is for the person in the vignette to give the whole 900 Ksh to the friend in need and thus save nothing of the extra money earned. Small (20%) amounts of savings are perceived as somewhat appropriate in PUBLIC, but any larger amount is viewed as inappropriate in both treatments. The action that is least appropriate, in either treatment, is to save the whole 900 Ksh and give nothing to the friend. Overall, this suggests the existence of a norm of sharing in the villages of our study. Moreover, there seems to be only small differences in the norm across information conditions given under the public and private treatments. Using non-parametric tests, we do not detect any significant differences in the ratings of any action between PRIVATE and PUBLIC (adjusted *p*-values > 0.192, Mann-Whitney rank-sum test).<sup>18</sup> Using regression analysis (see SM Table F1), we find a significant difference for the action

<sup>17</sup> Although we did not register a pre-analysis plan for this study, our ex-ante hypotheses concerned the effects of income privacy on the elicited norms and subjects’ saving behavior in the month following the experiment. We did not have ex-ante hypotheses about the heterogeneity in norms and treatment effects. Additional analyses that were conducted and do not appear in the final version of the paper include: an analysis of treatment effects on normative ratings based on the “slope” of the social norm functions (results are similar to those discussed in connection to Fig. 1); an analysis of treatment effects of savings disaggregated by experimental earnings and gender, and using quantile and probit regression analysis (results are similar to those discussed in Section 4.2); an analysis of homophily within norm types using network measures (this analysis was excluded from the final version of the paper following comments by an anonymous referee).

<sup>18</sup> Tests are conducted using the individual as the unit of observation and so they are based on 228 observations (106 in PRIVATE and 122 in PUBLIC). Reported *p*-values are two-sided. When we test for differences in norm functions between treatments, we perform multiple comparisons across treatments

**Table 3**  
Distribution of social appropriateness ratings.

		PRIVATE					PUBLIC					
		mean	-1	-0.33	+0.33	+1	mean	-1	-0.33	+0.33	+1	
All	Share 100%	<b>0.26</b>	33%	4%	5%	58%	<b>0.51</b>	22%	2%	4%	72%	
	Share 80%	-0.11	47%	3%	19%	31%	0.10	34%	5%	21%	39%	
	Share 60%	-0.16	46%	9%	17%	27%	-0.05	39%	8%	23%	30%	
	Share 40%	-0.09	42%	11%	15%	31%	-0.10	44%	6%	21%	29%	
	Share 20%	-0.21	45%	11%	24%	20%	-0.29	53%	11%	12%	24%	
	Share 0%	-0.64	76%	7%	4%	13%	-0.46	70%	2%	4%	24%	
Strict sharing norm	N = 25							N = 34				
	Share 100%	<b>1.00</b>	0%	0%	0%	100%	<b>0.94</b>	3%	0%	0%	97%	
	Share 80%	-1.00	100%	0%	0%	0%	-0.63	71%	6%	21%	3%	
	Share 60%	-0.97	96%	4%	0%	0%	-0.74	79%	12%	0%	9%	
	Share 40%	-0.97	96%	4%	0%	0%	-0.98	97%	3%	0%	0%	
	Share 20%	-0.95	92%	8%	0%	0%	-0.88	88%	6%	6%	0%	
Moderate sharing norm	N = 28							N = 60				
	Share 100%	<b>1.00</b>	0%	0%	0%	100%	<b>0.94</b>	0%	0%	8%	92%	
	Share 80%	0.76	0%	0%	36%	64%	0.71	5%	0%	28%	67%	
	Share 60%	0.40	14%	7%	32%	46%	0.52	12%	2%	33%	53%	
	Share 40%	-0.00	29%	14%	36%	21%	0.22	27%	7%	23%	43%	
	Share 20%	-0.45	57%	18%	11%	14%	-0.01	38%	13%	10%	38%	
Pro-saving norm	N = 53							N = 28				
	Share 100%	-0.48	66%	8%	9%	17%	-0.95	93%	7%	0%	0%	
	Share 80%	-0.15	47%	6%	19%	28%	-0.31	54%	14%	7%	25%	
	Share 60%	-0.08	40%	13%	17%	30%	-0.43	50%	18%	29%	4%	
	Share 40%	<b>0.28</b>	23%	13%	12%	52%	<b>0.26</b>	18%	7%	43%	32%	
	Share 20%	0.26	17%	9%	42%	32%	-0.17	43%	11%	25%	21%	
Share 0%	-0.32	58%	8%	8%	26%	-0.52	71%	4%	7%	18%		

Note: “-1” means “very socially inappropriate”, “-0.33” means “somewhat socially inappropriate”, “+0.33” means “somewhat socially appropriate” and “+1” means “very socially appropriate”. The action deemed to be on average most appropriate is shown in bold in the table. The modal appropriateness rating for each action is shown in a gray-shaded cell.

“Share 100%” (adjusted  $p$ -value = 0.060), indicating that full sharing is viewed as relatively less appropriate in PRIVATE than PUBLIC.

However, this aggregate result masks substantial heterogeneity in the way that different subgroups of participants rated the actions in the vignette. This is shown in the remaining three panels of Fig. 1. To divide subjects in subgroups, we use hierarchical cluster analysis (e.g., Fallucchi et al., 2019), a technique that does not rely on arbitrary ex-ante classifications of types, but is instead data-driven and exploits patterns of similarity in the data to classify subjects into homogeneous groups. Specifically, the analysis constructs subgroups of individuals (“clusters”) based on the similarity between their norm functions: clusters are selected in order to minimize the differences between norm functions *within* a cluster and to maximize the differences *between* clusters. Further details on the clustering procedure are provided in SM E.

The analysis reveals three key insights. First, three distinctly different social norms emerge, showing that different groups of subjects hold different beliefs about what is appropriate in the vignette. These three norms are identified in each of the two treatments, with remarkable similarities. This is despite the fact that we perform the cluster analysis for each treatment separately in order not to impose ex-ante any similarities of norms across treatments. Second, within each group and treatment, normative perceptions are quite homogeneous, as shown by the rather small confidence intervals displayed in Fig. 1b–d. Third, the proportion of subjects that perceive each of these three social norms differs across treatment conditions (see Table 4).

The first norm (top-right, orange lines with squares) is represented by a subgroup of subjects who perceive the existence of what we call a “strict sharing norm”. According to these subjects, saving nothing and giving everything to the friend is the *only* appropriate action, every other action is rated as highly inappropriate. This norm maximally penalizes any type of wealth accumulation. The norm is similar in the two treatments: we only detect one significant difference in the norm function across treatments, for action “Share 80%”, which is rated as inappropriate in both treatments, but slightly less so in PUBLIC (adjusted  $p$ -value = 0.030, Mann-Whitney rank-sum test; in all other cases adjusted  $p$ -values > 0.141).

The second norm (bottom-left, green lines with triangles) is represented by subjects who also view full sharing as the most appropriate action, but who are more permissive regarding small deviations from full sharing. Moderate deviations from full sharing (i.e., sharing 80–60% of the total) are not stigmatized, but are viewed as acceptable. Because the structure of

at various points of each function (for actions “Save 100%”, “Save 80%”, etc.). To correct for the increased likelihood of type I errors, we use the false discovery rate procedure and report “adjusted  $p$ -values” in the paper (Simes, 1986; Benjamini and Hochberg, 1995).

**Table 4**  
Sharing norm types in PUBLIC and PRIVATE.

Norm type	PUBLIC	PRIVATE	Bilateral tests for differences in norm types
Strict Sharing Norm	27.9%	23.6%	0.545
Moderate Sharing Norm	49.2%	26.4%	0.001
Pro-Saving Norm	22.9%	50.0%	0.000
Test for overall difference in type distributions	0.000		

*Note:* The table shows the percentage of participants classified as one of the three norm types in each treatment. Adjusted  $p$ -values (Simes, 1986; Benjamini and Hochberg, 1995) from Fisher's exact tests reported in the last column/row of the table (column: bilateral tests for treatment differences in the frequency of each norm type; row: overall tests for treatment differences in the distribution of types).

social incentives (stigma/honor that accrue to norm violators/compliers) that sustain compliance with this norm is different from that of the strict sharing norm, we call it a “*moderate sharing norm*”. In this cluster we detect a difference between norms in PUBLIC and PRIVATE for the action “*Share 0%*”, which is penalized more harshly in the PRIVATE case (adjusted  $p$ -value = 0.030, Mann-Whitney rank-sum test). No other action is evaluated differently across treatments (adjusted  $p$ -values > 0.114).

Finally, the third norm (bottom-right, blue lines with dots) is represented by subjects who do *not* perceive full sharing as the most appropriate action. Here, the norm functions are hump-shaped: in both treatments, the most appropriate action is to save 60% of the money and give the remaining 40% to the friend. Full sharing is viewed as the *least* appropriate action. Sharing nothing is also viewed as inappropriate. Given these patterns in the norm function, we call this a “*pro-saving norm*” since saving most of the money is encouraged, and not proscribed, by the norm. We only find a difference between treatments for the action “*Share 100%*” (adjusted  $p$ -value = 0.031, Mann-Whitney rank-sum test; in all other cases adjusted  $p$ -values > 0.114).

We next examine the percentage of subjects who fall into each subgroup in the two treatments (Table 4). In PUBLIC most subjects (49%) recognize the existence of a moderate sharing norm, 28% perceive the strict sharing norm, and 23% the pro-saving norm. In PRIVATE these shares differ markedly: the most popular norm is now the pro-saving norm (50% of subjects), while 26% perceive a moderate sharing norm and 24% a strict sharing norm. These differences in proportions across treatments are statistically significant ( $p < 0.001$ , Fisher's exact test). In bilateral tests, the norm types that differ significantly between treatments are “moderate sharing norm” and “pro-saving norm”, whereas there are no significant differences in the share of “strict sharing norm” types between PRIVATE and PUBLIC (see Table 4).

These results provide evidence that the information condition has an impact on social norms that was only partially visible in the aggregate data.<sup>19</sup> When privacy over income is guaranteed, and other village members cannot distinguish whether deflecting requests for money are due to lack of financial means or an unwillingness to share, most subjects view saving (and refusing to share) as socially appropriate. In contrast, in the absence of privacy, most subjects believe a stricter social norm applies. This suggests that interventions and mechanisms that allow individuals to shield their income from the eyes (and claims) of others have the potential of increasing saving behavior by tempering the norm of sharing that would otherwise prevail in the community. We will return on this point later, in Section 4.2.

The results also show that it would be incorrect to talk about the existence of a norm of sharing – at least not at a “global” level. A plurality of norms clearly exists in our setting, both in the PUBLIC and PRIVATE information conditions. Although in the PUBLIC condition a majority of people perceive that full sharing is the most appropriate thing to do, the share of people who disagree with this is not negligible (about a quarter of subjects). Moreover, the full sharing norm is no longer majoritarian in PRIVATE. It is also remarkable that this heterogeneity emerges in all villages, i.e., in all cases we observe subjects from each norm type present in each of our villages (see SM E). Thus, the observed heterogeneity does not reflect differences between villages as much as disagreement about what is appropriate behavior within a village. In this sense, our study reveals that, at least in our context, the strict norm of sharing hypothesized in the literature is not dominant.

An interesting question that follows from this result is whether individual and network characteristics can systematically explain the differences in norms across different subgroups of subjects. We address this question by performing multinomial logit regression analysis. We have two models. In both regression models, the dependent variable is a categorical variable

<sup>19</sup> Note that the shift in the proportion of norm types across treatments is consistent with the aggregate treatment effect on average appropriateness of sharing we found in Fig. 1a. The main shift in norm types between PRIVATE and PUBLIC occurs for the share of individuals classified as moderate and pro-saving norm types (the former halves, the latter doubles as we move from PUBLIC to PRIVATE). The norm functions of these two groups are similar for half of the actions (i.e. from share 40% to share 0%), and only differ sharply in the appropriateness of sharing 100%, 80% or 60% of the total. Indeed, this is the region where we see the largest differences in the aggregate norm functions (Fig. 1, panel a).

**Table 5**  
Determinants of sharing norm types.

	Pro-saving norm vs. Moderate sharing norm		Pro-saving norm vs. Strict sharing norm		Moderate sharing norm vs. Strict sharing norm	
	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2
1 if PRIVATE	574.0*** (0.000)	1314.3*** (0.000)	443.7*** (0.001)	1261.9*** (0.001)	-19.3 (0.587)	-3.7 (0.948)
Age		-10.3 (0.671)		-58.4*** (0.008)		-53.6** (0.011)
1 if Married		189.3** (0.033)		-47.9 (0.410)		-82.0** (0.012)
Liquid Assets (in Ksh1000)		25.4 (0.169)		99.6*** (0.009)		59.2* (0.073)
Monthly Income (in Ksh1000)		-9.7 (0.606)		86.2* (0.095)		106.3** (0.043)
Positive Reciprocity		107.1* (0.073)		59.1 (0.290)		-23.2 (0.377)
Negative Reciprocity		35.7 (0.448)		-82.6*** (0.001)		-87.2*** (0.000)
PGG Contribution		8.7 (0.704)		76.9* (0.052)		62.8* (0.062)
Financial Literacy		35.8 (0.342)		-55.1** (0.033)		-66.9*** (0.000)
Gifts & Loans (in Ksh1000)		-19.6 (0.302)		20.3 (0.487)		49.7* (0.095)
Degree		27.9 (0.342)		150.2** (0.027)		95.7* (0.062)
Clustering		56.4** (0.049)		167.1*** (0.007)		70.88 (0.112)
Share of Males in Session		28.8 (0.459)		228.2*** (0.005)		154.8** (0.012)

Note: Multinomial logit regressions with robust standard errors. The dependent variable is a categorical variable measuring the type of norm that a subject perceives (as defined by the hierarchical cluster analysis). Each column shows the percentage changes in the odds ratios of being classified as type *m* (listed first in the column heading) rather than type *n* (listed second in the column heading). For dummy variables (PRIVATE and Married in the table), we compute the change in the odds for a unit increase in the variable. For continuous variables, we compute changes for a standard deviation increase in the explanatory variable (standard deviations are: 12.8 for Age, 31.4 for Income, 130.8 for Liquid Assets, 1.9 for Positive Reciprocity, 3.2 for Negative Reciprocity, 4.8 for PGG Contributions, 1.1 for Financial Literacy, 1.8 for Gifts & Loans, 4.9 for Degree, 0.1 for Clustering and 11.9 for Share of Males in Session). Other control variables included in Model 2 but not reported in the table are: gender, education, savings, patience, risk aversion, altruism, trust, number of correct answers in the Raven matrix task, feeling of obligation to share, perceived claims by others. A constant is also included in both models, but omitted from the Table. Both models also include village and enumerator fixed effects. Similar results are obtained when standard errors are clustered at the village level. P-values are reported in parentheses. \*\*\* Significant at the 1% level. \*\* Significant at the 5% level. \* Significant at the 10% level. The number of observations is 226 in Model 1 and 224 in Model 2 because we have missing values in the explanatory variables of some subjects. The pseudo R2 value is 0.258 in Model 1 and 0.412 in Model 2.

describing whether a subject perceives a strict sharing norm, a moderate sharing norm, or a pro-saving norm. In both models, we include village and enumerator fixed-effects (similar results hold if we additionally cluster standard errors at the village level). As explanatory variables, in the first model we only include a treatment dummy (1 if PRIVATE, 0 otherwise). In the second model we add the socio-demographic variables, behavioral preferences measurements, and social pressure indicators listed in Table 2, plus a variable measuring the fraction of male participants present in a session (this varies between 10% and 52%, averaging 31%).

In the second model, we also include the social network data collected in the household survey. For each subject, we have data on 12 types of interactions with other members of their village: social activities with relatives or friends; borrowing/lending money or food; visiting or being visited by other village members; receiving or giving advice to others; going to church or medical facilities with others. We combine these data to construct, for each individual, a measure of the *degree* and *clustering* of their aggregate social network.<sup>20</sup> Degree captures the number of direct connections that each subject has in the village, that is, their immediate network. Clustering measures the fraction of these connections that are themselves connected with one another, that is, a measure of how interconnected an individual's network is. Intuitively, degree may matter for the type of norm one perceives because the higher the degree, the more people (and normative ideals) a subject may be exposed to. We suspect that this may work against the more traditional and strict norms, since individuals may be tempted to abandon these norms once they become aware of the existence of the pro-saving norm. Clustering may also matter for the formation of normative ideals, especially in a situation with plurality of social norms, since it may be easier to support one's norm within self-contained cliques of like-minded people.

Table 5 reports the regression results. Although the second model has been estimated with the full list of regressors mentioned above, for ease of readability, the table only reports variables with statistically significant coefficients in at least

<sup>20</sup> See SM A for details on the formal derivations.

one of the three possible bilateral comparisons (the full regression is reproduced in SM, Table F2). Results are reported as percentage changes in the odds of being classified within a certain norm type rather than an alternative type: a “pro-saving norm” type rather than a “moderate sharing norm” type in the first column; a “pro-saving norm” type rather than a “strict sharing norm” type in the second column; and a “moderate sharing norm” type rather than a “strict sharing norm” type in the third column.

We see three clear patterns in the data. First, the regression confirms that the private information condition increases dramatically the odds of being classified as a pro-saving norm type rather than either of the two stricter norm types. These effects are significant at the 1% level and confirm the results of our earlier univariate analysis.

Second, focusing on variables measuring socio-demographic characteristics, preferences and social pressure, we find that subjects who are classified as perceiving the strict sharing norm are on average older, poorer, more negatively reciprocal, less willing to contribute to public goods, and more financially literate. The regressions show that these factors increase the odds of being classified as a strict sharing norm type as opposed to both the moderate sharing norm and the pro-saving norm types. We find it intuitively plausible that the strict sharing norm, which harshly penalizes deviations from full sharing, is positively associated with age and traits like the willingness to retaliate against others at a cost to oneself, and negatively associated with income and willingness to contribute to public goods (implying that these subjects are probably on the demand side of the sharing market, namely, they are in need of financial help from others).<sup>21</sup>

We also find some factors that explain the differences between the moderate sharing norm type and the other types. Perceptions that one is exposed to greater financial demands by other village members (captured by the variable *Gifts & Loans*) increases (perhaps self-servingly) the odds of being a moderate sharing norm type rather than a strict sharing norm type. Moreover, we find some negative associations between the moderate sharing norm type and positive reciprocity and marital status. We also find that the probability of being classified as a strict sharing norm type decreases in the fraction of males present in a subject’s session, although the subject’s own gender is not significantly associated with norm type.

Finally, the third set of results that emerges from the regressions concerns the network data. First, we find that individuals who are connected to a higher number of people (i.e., higher degree) in the village are less likely to be classified as a strict sharing norm type – possibly reflecting the fact that, as discussed above, more connected people are more likely to come into contact with the plurality of norms and opinions that exist in their villages, and switch away from the more traditional norms. In support of this interpretation, the regression models estimated separately for each type of network show that this aggregate effect is partly driven by the number of people a subject turns to in seeking advice. Moreover, we find that individuals who belong to networks that are more clustered are more likely to be classified as the pro-saving norm type. This could reflect the fact that it is easier to sustain such a norm – especially one that goes against tradition – in self-contained cliques.<sup>22</sup>

Overall, these results suggest that, although a unique norm of sharing does not exist at a global community level, a collective understanding of what is appropriate and inappropriate behavior exists at a more “local” level, within sub-networks of individuals within the village. Moreover, the results show that our classification of subjects into norm types reflects some intuitive characteristics of the individual as well as the network in which they are embedded. We find this reassuring since it reinforces our interpretation that the classification captures some genuine underlying differences between individuals assigned to different social norm types, rather than mere noise. This interpretation is further reinforced by the fact that the shape of the norm functions in the three subgroups seem to reflect intuitive and sensible normative principles (share all the wealth with others; divide wealth roughly equally), which we would not necessarily expect to observe if the source of heterogeneity was mere noise.

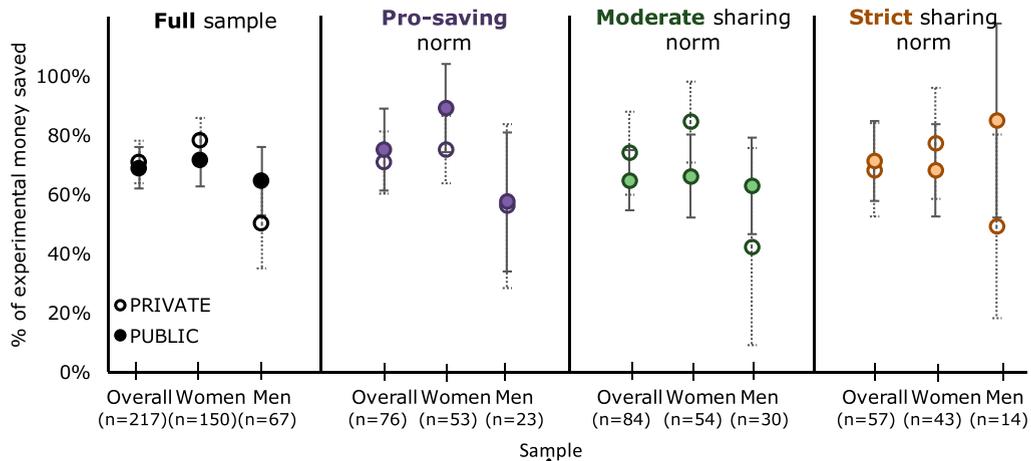
#### 4.2. Saving behavior and the impact of observability of earnings

The previous section has shown that observability of income affects the perception of the obligations to share with others. When income can be kept private, so that one can plausibly refuse to share with others without appearing unwilling to do so, nearly half of the subjects think that a person ought to resist pressures to share and instead should save one’s income. When income is observable, however, a majority of people think that full sharing is the most appropriate behavior. In this section, we ask how observability of income affects actual saving behavior.

To do so, we use the data that our partner bank shared with us regarding bank account usage of the subjects who successfully opened an account with the bank at the end of the experiment. Note that, in this analysis of savings behavior, we include 217 of the 228 participants. Eleven of the participants experienced problems with their bank accounts: 1 subject experienced a delay in the activation of their bank account of more than one month resulting in no available data; and 10

<sup>21</sup> We were instead surprised by the positive association between strict sharing norm type and financial literacy. Upon further investigation, we found that the effect is driven by answers to two of the four questions that underlie the variable, about borrowing from money lenders and risk associated with planting crops.

<sup>22</sup> We also ran a version of Model 2 interacting degree and clustering with the PRIVATE dummy to explore whether there are differences in the influence of the network variables across treatments. We find that the positive effect of clustering on the likelihood of being classified as a pro-saving norm or moderate sharing norm instead of strict sharing norm type is significantly larger in PUBLIC than PRIVATE. This is in line with the idea that self-contained cliques may be particularly beneficial to those who go against the traditional norm. We find little differences in the effect of degree between treatments.



**Fig. 2.** Saving behavior, by norm type and gender *Note:* The figure shows the share of experimental money saved by subject in the month following the experiment, disaggregated by treatment (full dots = PUBLIC, hollow dots = PRIVATE). Whiskers represent 95% confidence intervals. The first panel shows aggregate data from the full sample. The second panel focuses on subjects who perceive the pro-saving norm, the third panel on subjects who perceive the moderate sharing norm and the fourth on subjects who perceive the strict sharing norm. In each panel we report average saving behavior across all subjects in the first column, for women in the second, and men in the third. The number of observations underlying each comparison are reported below the figure.

subjects had loan arrears from previous financial activities that were automatically deducted from their balance after the opening of their account.<sup>23</sup>

We paid the money that a subject had earned in the experiment directly into their new bank accounts, and promised that, after one month, we would increase these earnings by 50% after first subtracting any money they had withdrawn during the period. Thus, we use the share of experimental money that a subject left in the bank account as a measure of their saving behavior in the month following the experiment. We examine whether subjects who were assigned to the PUBLIC treatment, in which experimental earnings were publicly observable by other participants in their session, are more or less likely to save their earnings compared to subjects in the PRIVATE treatment, whose earnings were kept completely private during the experiment.

Two effects are potentially at play when income is observable by other village members. On the one hand, pressures to share income may be stronger when others in the village know the amount of money that a person has earned in the experiment. This is because there may be more demands to share and/or it may be harder to resist these demands without being able to plausibly deny having earned much money in the experiment. This mechanism would lead to lower savings in PUBLIC than in PRIVATE. On the other hand, observability may also have a positive effect on savings, as argued in a recent literature showing that making saving amounts observable to others can have beneficial effects on savings by creating reputational pressures on the individual to stick to one's saving goals (e.g., Kast et al., 2018; Breza and Chandrasekhar, 2019). Although in our setting individuals do not set any specific saving goal for themselves, it is likely that the very generous interest rate that we offered on their savings could trigger analogous reputational pressures, since one may be viewed as irresponsible and short-sighted if they withdrew money from the account.

Importantly, the relative strength of these two opposing effects may depend on the type of social norm that an individual perceives being upheld in their social network. In particular, pressures to share income should have less bite for subjects who are classified as being a pro-saving norm type. As we have seen, resisting the pressures to share is considered socially acceptable by these subjects, even when others have full observability of income. This suggests that, among pro-saving norm subjects, the negative effect of observability may be small and the PUBLIC treatment may have a positive net effect. In contrast, the negative effect of observability may be stronger for the remaining subjects, who perceive either a moderate or strict saving norm; such a negative effect could compensate, or even outweigh, any positive effect of observability, and could be amplified by their social network characteristics.

We start by examining these specific effects in Fig. 2, which plots the share of experimental money that a subject had in the account at the end of the month following the experiment (i.e., the amount that we increased by 50%) across the PRIVATE and PUBLIC treatments. Generally, subjects managed to save a fairly large share of their experimental money (around 70% on average) – which is a likely reflection of the substantial interest rate we paid on their savings.

<sup>23</sup> Comparison of variables in our balance test between these 11 and the 217 participants were insignificant, except for two variables: liquid assets and obligation to share. Here the 11 were significantly more likely to report fewer liquid assets ( $p = 0.053$ ) and to respond with a lower rating on how obliged they felt to share money with someone else ( $p = 0.011$ ). Excluding these 11 from the previous analysis does not alter the results in any material way.

**Table 6**  
Effect of observability of earnings on savings.

	All subjects		Women only		Men only	
	(1)	(2)	(3)	(4)	(5)	(6)
1 if PRIVATE	0.02 (0.084)	-0.13 (0.148)	0.09 (0.107)	-0.38** (0.182)	-0.14 (0.129)	-0.05 (0.222)
Moderate sharing norm		-0.17 (0.159)		-0.59*** (0.203)		0.12 (0.205)
Moderate sharing norm * PRIVATE		0.24 (0.208)		0.69*** (0.248)		-0.06 (0.333)
Strict sharing norm		-0.01 (0.186)		-0.38 (0.237)		0.45 (0.352)
Strict sharing norm * PRIVATE		0.15 (0.234)		0.47 (0.299)		-0.41 (0.372)
N. Observations	215	215	148	148	67	67
Pseudo R <sup>2</sup>	0.179	0.186	0.246	0.276	0.449	0.473

Note: Tobit regressions with robust standard errors in parentheses. The dependent variable is the share of experimental money left in the bank account at the end of the month following the experiment. The first two columns contain data from all subjects; the subsequent two columns contain data from women; and the last two columns contain data from men. \*\*\* Significant at the 1% level. \*\* Significant at the 5% level. \* Significant at the 10% level. Control variables: gender (only in columns 1 and 2); married; age; education; income; savings; patience; risk aversion; altruism; positive and negative reciprocity; trust; contributions to public goods; number of correct answers in the Raven matrix and financial literacy tasks; feeling of obligation to share; perceived claims by others; amount spent on gifts and loans in the past month; degree; clustering; total earnings from the experiment. All models also include village fixed effects (we do not include enumerator fixed effects to avoid overfitting, especially in the male-only regressions where we have only 67 observations). A constant is also included in all models, but not displayed in the table. We have 215 observations (instead of 217) due to missing values in some of the control variables.

The first column in the left-most panel of the figure shows the aggregate treatment effect among our 217 subjects. Income privacy has a slightly positive effect on savings but is very small in magnitude. In acknowledgement of the existing empirical literature which shows the effect of social pressure to share is particularly marked among women (e.g., Jakiela and Ozier, 2016; Boltz et al., 2019), we also report data disaggregated by gender. The first finding is that overall women save more than men do. The second finding is that the effect of privacy differs by gender: women save more in the PRIVATE treatment, whereas men save more in the PUBLIC treatment.

The remaining three panels show the effect of income privacy disaggregated by the type of social norm that a subject perceives. For subjects perceiving the pro-saving norm, we find that income privacy has a *negative* effect on savings, which is driven by the large effect among women. Privacy has instead a *positive* effect among women who perceive the moderate or strict sharing norm. We observe the reverse pattern among men, although it should be noted that these averages rely on a very small number of observations.

To explore the statistical significance of these effects, we conduct regression analysis of saving behavior. To account for the fact that savings are censored (nearly 50% of observations are censored from above because subjects saved the whole amount), we use Tobit regression models where the dependent variable is the share of experimental money saved by the subject in the month after the experiment. As regressors, we use a treatment dummy and the controls shown in Table 2 as well as a variable measuring the total earnings a subject made in the experiment. To account for the fact that – as seen in Fig. 2 – income privacy has different effects depending on subjects' norm type and gender, in a second specification we also include dummy variables for the moderate and strict sharing norm types, with the pro-saving norm type serving as base category, and interact these dummies with the treatment dummy PRIVATE. We report models that include all subjects, as well as separate models for women and men. In the SM, Table F3, we report the same regression models but without controls.

Table 6 presents the results. In the table, we only report the coefficients of the treatment dummy, the norm type dummies and the interaction terms, but all models have been run with the full set of covariates (the full regression is reproduced in SM, Table F4). In all regressions without interaction terms (columns 1, 3 and 5), the treatment dummy is not significantly different from zero. Thus, without taking into account the heterogeneity in norms, the aggregate evidence is that removing observability of earnings has no effect on saving behavior.

In columns 2, 4 and 6, we report regressions that account for the plurality of norms that we observe in the experiment. Starting with column 2, where we include all subjects, we confirm that privacy has a negative effect for the subjects classified as pro-saving norm type. The coefficient, however, does not reach statistical significance ( $p = 0.394$ ). The coefficients of the interaction terms indicate that privacy has the opposite effect for subjects who perceive the moderate sharing norm, while it has an effect close to zero for subjects who perceive the strict sharing norm. The interaction terms, however, are not significant.

The picture becomes sharper when we consider the effect of privacy among women (column 4), which the literature has repeatedly found to be a subgroup particularly susceptible to social pressure to share. Here, for pro-saving norm subjects, we estimate a negative effect of privacy on savings which is statistically significant ( $p = 0.040$ ). In contrast, the interaction terms for subjects who perceive the moderate and strict sharing norms take the opposite sign. Two F-tests reveal a small

positive effect for moderate sharing norm subjects ( $p = 0.068$ ), whereas the effect is not significantly different from zero for strict sharing norm subjects ( $p = 0.700$ ).

Finally, column 6 reports estimates of the treatment effect among men. The sample is very small (only 67 subjects) and we find no significant effect of privacy among pro-saving norm subjects, nor any significant interaction term for the other norm types.

Overall, these results suggest that – in line with our conjecture – privacy can have different effects among subjects who perceive different types of norms. For subjects, particularly women, who represent around 70% of our sample, who perceive a moderate norm of sharing, making one's earnings unobservable to others is beneficial for savings. For subjects who however perceive a pro-saving norm the effect is detrimental: privacy substantially reduces their ability to save their income.<sup>24,25</sup>

While this analysis suggests some intriguing interactions between privacy and social norms, a number of caveats should be borne in mind when interpreting the results. First, the way we classified individuals into norm types is based on their responses to the experimental vignette in *one* of the two possible treatment conditions (private/public). However, it is reasonable to assume that a person's "true" norm type is actually defined by how that person evaluates the appropriateness of saving *both* in private and in public. For example, someone may perceive a strict sharing norm both in private and public situations. Another person may perceive a strict sharing norm in public, but a pro-saving norm in private. It is reasonable to think of these two individuals as different "types", who may have intrinsically different dispositions to save. The analyses of Fig. 2 and Table 6 pool together subjects who have been classified as a specific norm type according to responses to *different* vignettes. Therefore, it potentially pools together different types of individuals (e.g., both individuals described above would be pooled in the strict sharing norm category if the former was assigned to the PRIVATE treatment and the latter to PUBLIC). This implies that part of the (average) treatment effects we observed may reflect an underlying heterogeneity in the propensities to save of the pooled groups.

A second concern is that in the analysis of Table 6 we could only partially account for interdependences that may have arisen at the village level, by including village fixed-effects. We did not cluster standard errors at the village level because we only have observations from six villages. If we nevertheless use village-level clustered standard errors (using wild cluster bootstrapping), all effects turn insignificant.

A third concern is that the small treatment differences that emerge in Table 6 may reflect the relatively weak manipulation that differentiated the PRIVATE and PUBLIC conditions. The differences in payment procedures implied that, in the PUBLIC treatment, experimental earnings were observable by other participants in a subject's session. The extent to which this translated in different pressures to share depended on whether the information about one's earnings spread more widely in the village, beyond the people present in the session. Given that subjects on average only knew approximately 2 other participants in their session, these informational spillovers may have been modest. A stronger manipulation (e.g., publicly announcing the earnings at the village level by posting the experimental results on a public message board) may have produced stronger effects.

Finally, another possible concern with the analysis is that it may fail to fully capture the extent to which subjects were exposed to pressures to share money with others in the PUBLIC treatment. In fact, subjects may have been pressured to share money *after* having received the interest on the amount saved at the end of the first month, so that a larger amount of money would be available for sharing. If so, our estimates above represent a lower bound of the effects of observability on saving behavior.

To check whether this is the case, we collected data on bank account usage for one more month from the day that subjects received their interest payment (i.e. the second month after the experiment). Note that, while bank account usage in the first month after the experiment can be interpreted as a measure of subjects' willingness to save, the interpretation of bank account usage in the second month is less straightforward. The decision to withdraw money may in fact reflect the availability of better saving opportunities, which we do not observe (this is not the case for the first month after the experiment since we pay a very high interest rate on savings).

Despite this limitation, looking at bank account usage in the second month can be informative. In particular, if there were pressures to share after having collected the interest on the first month savings, we would expect to observe differences between the PUBLIC and PRIVATE treatments in the second month after the experiment. We therefore run the same regressions as those presented in Table 6, but using as dependent variable the share of experimental money including the

<sup>24</sup> Among the control variables included in the regressions of Table 6 (see SM F), it is worth noting that there is a negative relation between experimental earnings and savings: those who earn more in the experiment save a smaller fraction of their experimental earnings. This is consistent with the idea that these individuals may have faced stronger demands on their saving. In further analysis, we interact the experimental earnings variable with the treatment dummy and find that the relation between experimental earnings and savings is more negative in PUBLIC than PRIVATE, which further corroborates the idea that subjects who earned more in the experiment faced stronger pressures to share.

<sup>25</sup> An interesting question is whether the effect of treatment varies with the strength of the pressure to share that a subject may have experienced after the experiment. As proxy for the strength of the pressure to share we use a variable measuring the number of other individuals who are in the same session as the subject and are in the subject's social network (the average number of network connections in the same session is 1.86; the median is 2). The idea is that the larger the number of network connections who are in the same session as the subject, the higher the likelihood that information about their participation in the experiment and, in PUBLIC, earnings, may spread in the subject's social network, hence leading to stronger pressures to share. However, when we add this variable and its interaction with treatment to the regressions of Table 6, we find that both are statistically insignificant in all regression models.

interest that a subject had left in the account at the end of the second month following the experiment (we neglect any deposits by the subjects to keep comparability with our previous dependent variable).

The estimates, reported in SM Table F5, reveal essentially no treatment effects for any norm type in the regressions using all subjects or only women. For men, however, we find significant treatment effects that vary across norm types. Specifically, we find that income privacy has a significant negative effect for pro-saving norm types and a significant positive effect for subjects who perceive the moderate and strict sharing norms. Thus, in the second month after the experiment we find the analogous pattern for men that we found for women in the first month. Bearing in mind the small sample size of only 67 observations, this suggests that our main analysis may have underestimated the effect of income privacy for men and that they may have indeed adopted a strategy of postponing sharing with others until after they had collected the generous interest we paid on their experimental earnings. We can only speculate as to why women did not adopt such a strategy.

## 5. Conclusions

Our paper contributes to the emerging empirical and experimental literature on the social norms that govern the tension between wealth accumulation and sharing with kin, friends and neighbors among communities in developing countries (e.g., Baland et al., 2011; Di Falco and Bulte, 2011, 2015; Dupas and Robinson, 2013; Jakiela and Ozier, 2016; Goldberg, 2017; Squires, 2018; Boltz et al., 2019). One of the key findings of this literature is that individuals, especially those in poor communities in sub-Saharan Africa, are prepared to incur significant monetary costs in order to keep income and wealth accumulation unobservable by others in their community – a behavioral pattern that is consistent with the existence of strong “sharing norms” in these communities. The main contribution of our paper is to complement this evidence by eliciting these social norms directly in several small rural communities in Kenya.

Unexpectedly, we find evidence of a plurality of sharing norms in our setting: while some individuals perceive any form of wealth accumulation at the expense of kin, friends and neighbors as completely inappropriate, others find moderate accumulation of wealth acceptable, and a third group even view keeping most of one’s earnings for oneself as the most appropriate behavior. This pro-saving norm becomes even majoritarian in settings where income remains unobservable by others.<sup>26</sup>

It is interesting to compare our findings with those of Dupas and Robinson (2013), Jakiela and Ozier (2016), and Squires (2018) who also ran their experiments in Kenya and found strong social pressure to share income with kin and neighbors, which one could interpret as suggesting that there exist pervasive strict sharing norms in their settings. Dupas and Robinson (2013) show that informal saving technologies that allow individuals to turn down requests to share with others substantially increase investment in preventative health. Jakiela and Ozier (2016) find that women in their lab-in-the-field experiment are prepared to follow financially suboptimal investment strategies if these allow them to conceal their income – particularly when relatives are present in the same session. Squires (2018) find strong distortions in decision making for a third of their sample.

We can only speculate as to why our evidence points towards that strict sharing norms are not so prevalent, especially given that neither Dupas and Robinson (2013) nor Jakiela and Ozier (2016) nor Squires (2018) measured the norms that prevailed in their settings. We cannot therefore examine how different these norms are relative to those elicited in our context. One interesting difference between our study and those by Dupas and Robinson (2013) and Jakiela and Ozier (2016) is that their studies were run with subjects with lower income and education levels. The average weekly income in Dupas and Robinson (2013) and Jakiela and Ozier (2016) was approximately US\$8 and individuals had on average 6, 7 years of schooling. The sample of Squires (2018) was richer (US\$18) but less educated (0.3 years). Our sample is substantially richer and more educated. The average weekly income reported in the survey is around US\$80, earned mostly from farming activities. On average subjects had 8.5 years of schooling, with 70% of the sample reporting at least 8 years of schooling. These differences could potentially explain why the stricter and more traditional sharing norms may have a weaker hold in the communities we visited, compared to those in the previous studies. Indeed, in our analysis we do find evidence that having a lower level of income correlates with the likelihood of perceiving a stricter norm.

Another observation to be made is that the effects of sharing norms are often found to be stronger in some subsamples than others. For example, Jakiela and Ozier (2016) report effects for women but not for men; Squires (2018) report substantial heterogeneity in the impact of kinship tax in his sample. It is possible that this heterogeneity may be associated with a pluralism of normative ideals, which went undetected in those experiments as the researchers did not directly measure social norms. An interesting avenue for further research would be to directly measure sharing norms across a variety of locations within the same country in order to exploit natural variations in socio-demographic and economic characteristics.

When interpreting the results of our norm-elicitation experiment, one should keep in mind that our methodology relies on vignettes that are very specific about the decision situation under evaluation. The person described in our vignettes faces a particular trade-off between investing money to grow their own business and helping a friend in need of money to pay for school fees. The advantage of providing specific details about the decision setting is that we can control subjects’

<sup>26</sup> This result relates to a larger literature on the pluralism of normative views and personal attitudes showing that in many situations individuals may hold heterogeneous and even conflicting perceptions of what constitutes appropriate or fair behavior. See, for instance, Cappelen et al. (2007) for fairness ideals, Burks and Krupka (2012) for norms of ethical conduct of financial advisers, Winter et al. (2012) for bargaining norms, and Reuben and Riedl (2013) for contribution norms.

beliefs about these aspects of decision-making that may matter for normative considerations (e.g., refusing to share money with someone in need may be less acceptable than with someone who is not in need). The enhanced control allows us to measure with more precision the norm that applies to the decision situation described in the vignette, thus improving the internal validity of the experiment. However, the high degree of specificity may make it harder to generalize from our results to other type of decision situations where some of the key aspects of decision-making differ (e.g., the person asking for money is not in need; he/she is a family member instead of a friend; etc.). An interesting avenue for further research would be to systematically map out the norms that apply to different decision situations in order to obtain a more complete understanding of sharing norms among the poor.

Finally, our results suggesting the presence of heterogeneity in treatment effects may point towards possible threats to the efficacy and scalability of interventions aimed at encouraging investment and savings (e.g., Al-Ubaydli et al., 2017). In our exploratory analysis, we show that income secrecy, a feature that the empirical literature has suggested as a potentially effective remedy against sharing norms to encourage investment and savings, can have very different effects on saving behavior depending on the type of norm that an individual perceives. For the sizeable fraction of individuals (particularly women) who perceive a pro-saving norm, reducing income observability has a *perverse* effect on savings. While mindful of the discussed shortcomings, this analysis tentatively indicates that introducing saving technologies that reduce observability of wealth accumulation decisions may actually backfire in communities where sharing norms may have a weaker hold. It is therefore important to be able to observe – and not just assume – the type of norms that prevail in a community before designing and implementing interventions that are engineered to withstand those norms.

### Declaration of Competing Interest

None.

### Acknowledgments

We thank the Editor, Eugen Dimant, and three anonymous referees for their helpful comments. We also thank Abigail Barr, Erwin Bulte, Kristina Czura, Francesco Fallucchi, Michael Grimm, Anett John, Robert Lensink and Xueheng Li for their inputs and comments on the project. We also received helpful comments from seminar audiences at the University of Aarhus and the University of Groningen, and participants at the EUDN's Scientific Conference in Toulouse, the Nordic Conference on Behavioral and Experimental Economics in Gothenburg, the CSAE Conference in Oxford, the ESA European meeting in Bergen, and the Nordic Conference on Behavioral and Experimental Economics in Oslo. This work was supported by the British Academy & Leverhulme Trust (grant number SG140436) and by the University of Nottingham.

### Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:[10.1016/j.jebo.2021.09.028](https://doi.org/10.1016/j.jebo.2021.09.028).

### References

- d'Adda, G., Dufwenberg, M., Passarelli, F., Tabellini, G., 2020. Social norms with private values: theory and experiments. *Games Econ. Behav.* 124, 288–304.
- Akerlof, G.A., Kranton, R.E., 2000. Economics and identity. *Q. J. Econ.* 115 (3), 715–753.
- Al-Ubaydli, O., List, J.A., Suskind, D.L., 2017. What can we learn from experiments? Understanding the threats to the scalability of experimental results. *Am. Econ. Rev.* 107 (5), 282–286.
- Ashraf, N., 2009. Spousal control and intra-household decision making: an experimental study in the Philippines. *Am. Econ. Rev.* 99 (4), 1245–1277.
- Baland, J.M., Guirkinger, C., Mali, C., 2011. Pretending to be poor: borrowing to escape forced solidarity in cameroon. *Econ. Dev. Cult. Chang.* 60 (1), 1–16.
- Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O., 2013. The diffusion of microfinance. *Science* 341 (6144), 1236–1248.
- Banerjee, R., 2016. On the interpretation of bribery in a laboratory corruption game: moral frames and social norms. *Exp. Econ.* 19 (1), 240–267.
- Barr, A., Lane, T., Nosenzo, D., 2018. On the social appropriateness of discrimination. *J. Public Econ.* 164, 153–164.
- Barr A., and Stein M. 2008. Status and egalitarianism in traditional communities: an analysis of funeral attendance in six Zimbabwean villages. (mimeograph).
- Baseler T. 2021. Hidden income and the perceived returns to migration. SSRN Working Paper 10.2139/ssrn.3534715.
- Beekman, G., Gatto, M., Nillesen, E., 2015. Family networks and income hiding: evidence from lab-in-the-field experiments in rural Liberia. *J. Afr. Econ.* 24 (3), 453–469.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B* 57 (1), 289–300 (Methodological).
- Bicchieri, C., 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press, Cambridge.
- Bicchieri, C., 2017. *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford University Press, Oxford, New York.
- Bicchieri C., Dimant E., Gächter S., and Nosenzo D. 2020. Social proximity and the erosion of norm compliance. IZA Discussion Paper No. 13864.
- Bicchieri, C., Xiao, E., 2009. Do the right thing: but only if others do so. *J. Behav. Decis. Mak.* 22 (2), 191–208.
- Boltz, M., Marazyan, K., Villar, P., 2019. Income hiding and informal redistribution: a lab-in-the-field experiment in Senegal. *J. Dev. Econ.* 137, 78–92.
- Breza, E., Chandrasekhar, A.G., 2019. Social networks, reputation, and commitment: evidence from a savings monitors experiment. *Econometrica* 87 (1), 175–216.
- Brune, L., Giné, X., Goldberg, J., Yang, D., 2016. Facilitating savings for agriculture: field experimental evidence from Malawi. *Econ. Dev. Cult. Chang.* 64 (2), 187–220.
- Bulte, E.H., Lensink, R., Winkel, A.B., 2018. The impact of a gender and business training on income hiding: an experimental study in Vietnam. *J. Econ. Behav. Organ.* 148, 241–259.
- Burks, S.V., Krupka, E., 2012. A multimethod approach to identifying norms and normative expectations within a corporate hierarchy: evidence from the financial services industry. *Manag. Sci.* 58 (1), 203–217.

- Cappelen, A.W., Hole, A.D., Sørensen, E.Ø., Tungodden, B., 2007. The pluralism of fairness ideals: an experimental approach. *Am. Econ. Rev.* 97 (3), 818–827.
- Chang, D., Chen, R., Krupka, E., 2019. Rhetoric matters: a social norms explanation for the anomaly of framing. *Games Econ. Behav.* 116, 158–178.
- Cole, S., Sampson, T., Zia, B., 2011. Prices or knowledge? What drives demand for financial services in emerging markets? *J. Financ.* 66 (6), 1933–1967.
- Di Falco, S., Bulte, E., 2011. A dark side of social capital? Kinship, consumption, and savings. *J. Dev. Stud.* 47 (8), 1128–1151.
- Di Falco, S., Bulte, E., 2015. Does social capital affect investment in human capital? Family ties and schooling decisions. *Appl. Econ.* 47 (2), 195–205.
- Dimant, E., van Kleef, G.A., Shalvi, S., 2020. Requiem for a Nudge: framing effects in nudging honesty. *J. Econ. Behav. Organ.* 172, 247–266.
- Dupas, P., Robinson, J., 2013. Why don't the poor save more? Evidence from health savings experiments. *Am. Econ. Rev.* 103 (4), 1138–1171.
- Falk A., Becker A., Dohmen T., Enke B., Huffman D.B., and Sunde U. 2018. Global evidence on economic preferences quarterly journal of economics. Q. J. Econ. 133(4), 1645–1692.
- Fallucchi, F., Luccasen, R.A., Turocy, T.L., 2018. Identifying discrete behavioural types: a re-analysis of public goods game contributions by hierarchical clustering. *J. Econ. Sci. Assoc.* 5, 238–254.
- Fehr, E., Schurtenberger, I., 2018. Normative foundations of human cooperation. *Nat. Hum. Behav.* 2 (7), 458–468.
- Gächter, S., Gerhards, L., Nosenzo, D., 2017. The importance of peers for compliance with norms of fair sharing. *Eur. Econ. Rev.* 97 (C), 72–86.
- Gächter, S., Nosenzo, D., Sefton, M., 2013. Peer effects in pro-social behavior: social norms or social preferences? *J. Eur. Econ. Assoc.* 11 (3), 548–573.
- Goldberg, J., 2017. The effect of social pressure on expenditures in Malawi. *J. Econ. Behav. Organ.* 143, 173–185.
- Görges, L., Nosenzo, D., 2020a. Social norms and the labor market. In: Zimmermann, K.F. (Ed.), *Handbook of Labor, Human Resources and Population Economics*. Springer International Publishing, Cham, pp. 1–26.
- Görges, L., Nosenzo, D., 2020b. Measuring social norms in economics: why it is important and how it is done. *Anal. Krit.* 42 (2), 285–311.
- Grimm, M., Hartwig, R., Lay, J., 2017. Does forced solidarity hamper investment in small and micro enterprises? *J. Comp. Econ.* 45 (4), 827–846.
- Hoff, K., Sen, A., Bowles, S., Durlauf, S.N., Hoff, K., 2006. The Kin System as a Poverty Trap? In *Poverty Traps*. Princeton University Press.
- Jakiela, P., Ozier, O., 2016. Does africa need a rotten kin theorem? Experimental evidence from village economies. *Rev. Econ. Stud.* 83 (1), 231–268.
- John, A., 2018. When commitment fails - evidence from a field experiment. *Manag. Sci.* 62 (2), 503–529.
- Kast, F., Meier, S., Pomeranz, D., 2018. Saving more in groups: field experimental evidence from Chile. *J. Dev. Econ.* 133, 275–294.
- Kölle, F., Lane, T., Nosenzo, D., Starmer, C., 2020. Promoting voter registration: the effects of low-cost interventions on behaviour and norms. *Behav. Public Policy* 4 (1), 26–49.
- Krupka, E., Leider, S., Jiang, M., 2021. A meeting of the minds: informal agreements and social norms. *Manag. Sci.* 63 (6), 1657–2048.
- Krupka, E., Weber, R.A., 2013. Identifying social norms using coordination games: why does dictator game sharing vary? *J. Eur. Econ. Assoc.* 11 (3), 495–524.
- López-Pérez, R., 2008. Aversion to norm-breaking: a model. *Games Econ. Behav.* 64 (1), 237–267.
- Nikiforakis, N., Oechssler, J., Shah, A., 2014. Hierarchy, coercion, and exploitation: an experimental analysis. *J. Econ. Behav. Organ.* 97, 155–168.
- Ostrom, E., 2000. Collective action and the evolution of social norms. *J. Econ. Perspect.* 14 (3), 137–158.
- Platteau, J.P., 2000. *Institutions, Social Norms, and Economic Development*. Psychology Press.
- Platteau, J.P., Serge-Christophe, K., Jean Mercier, Y., 2006. Solidarity norms and institutions in village societies: static and dynamic considerations. In: *Handbook of the Economics of Giving, Altruism and Reciprocity*, 1. Foundations Elsevier, pp. 819–886.
- Platteau, J.P., 2014. Redistributive pressures in Sub-Saharan Africa: causes, consequences, and coping strategies. In: Akyeampong, E., Bates, R.H., Nunn, N., Robinson, J. (Eds.), *Africa's Development in Historical Perspective*. Cambridge University Press.
- Reuben, E., Riedl, A., 2013. Enforcement of contribution norms in public good games with heterogeneous populations. *Games Econ. Behav.* 77 (1), 122–137.
- Simes, R.J., 1986. An improved Bonferroni procedure for multiple tests of significance. *Biometrika* 73 (3), 751–754.
- Squires M. 2018. Kinship taxation as a constraint to microenterprise growth: experimental evidence from Kenya. Working paper available at [https://www.dropbox.com/s/4arw5frgjt9pzk/Squires\\_2018\\_KinshipTaxation.pdf?dl=0](https://www.dropbox.com/s/4arw5frgjt9pzk/Squires_2018_KinshipTaxation.pdf?dl=0).
- Winter, F., Rauhut, H., Helbing, D., 2012. How norms can generate conflict: an experiment on the failure of cooperative micro-motives on the macro-level. *Soc. Forces* 90 (3), 919–946.