Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies

Pamela Jakiela and Owen Ozier*

June 4, 2015

Abstract

This paper measures the economic impacts of social pressures to share income with kin and neighbors in rural Kenyan villages. We conduct a lab experiment in which we randomly vary the observability of investment returns to test whether subjects reduce their income in order to keep it hidden. We find that women adopt an investment strategy that conceals the size of their initial endowment in the experiment, though that strategy reduces their expected earnings. This effect is largest among women with relatives attending the experiment. Parameter estimates suggest that women anticipate that observable income will be "taxed" at a rate above four percent; this effective tax rate nearly doubles when kin can observe income directly. At the village level, we find an association between willingness to forgo expected return to keep income hidden in the laboratory experiment and worse economic outcomes outside the laboratory.

JEL codes: C91, C93, D81, O12

*Jakiela: University of Maryland, pjakiela@umd.edu; Ozier: Development Economics Research Group, The World Bank, oozier@worldbank.org. We are grateful to Felipe Dizon for excellent research assistance, to the staff at IPA-Kenya for their assistance and support in the field, and to the Weidenbaum Center and the Center for Research in Economics and Strategy at Washington University in St. Louis for funding. We also thank Kelly Bishop, Alexander Cappelen, Esther Duflo, Ray Fisman, Garance Genicot, Xavi Giné, Jess Goldberg, Shachar Kariv, Cynthia Kinnan, Ashley Langer, Karen Levy, Maggie McConnell, Justin McCrary, Ted Miguel, Sendhil Mullainathan, Alvin Murphy, Carol Nekesa, Silvia Prina, Gil Shapira, Erik Sørensen, Bertil Tungodden, Sergio Urzúa, Dean Yang, Andrew Zeitlin, numerous conference and seminar participants, the associate editor, and four anonymous referees for helpful comments. All errors are our own.

The findings, interpretations and conclusions expressed in this paper are entirely those of the authors, and do not necessarily represent the views of the World Bank, its Executive Directors, or the governments of the countries they represent.

1 Introduction

Risk is a pervasive aspect of the lives of individuals in many developing economies, and informal risk-pooling arrangements which help households cope with shocks can have significant welfare impacts when credit and insurance markets are incomplete. A substantial body of evidence documents the existence of mutual insurance arrangements throughout the world (cf. Townsend 1994, Coate and Ravallion 1993, Fafchamps and Lund 2003). Much of the literature focuses on mutual insurance arrangements which are efficient given constraints, characterizing the conditions under which self-interested households will enter risk-pooling schemes voluntarily *ex ante* and the participation constraints which keep households from defecting *ex post*.[1] Yet, as Kinnan (2014) demonstrates, models that allow agents to hide part of their income may better explain observed patterns of income and consumption in rural villages.[2] Moreover, the expectation of future transfers is only one of many reasons households offer assistance to those worse off: altruism, guilt, and social pressure to share income may also play a role (Geertz 1963, Scott 1976, Foster and Rosenzweig 2001, Alger and Weibull 2010). In fact, several recent studies suggest that individuals living in poor communities often feel obligated to make transfers to relatives and neighbors, and that successful families who do not make sufficient transfers to others can face harsh social sanctions (Platteau 2000, Barr and Stein 2008, Di Falco and Bulte 2011).[3] For example, Baland, Guirkinger, and Mali (2011) provide evidence of this type of behavior in Cameroon, where members of credit cooperatives take out loans to signal that they are liquidity constrained — even when they also hold substantial savings — in order to avoid sharing accumulated wealth with relatives. Collier and Garg (1999) argue that such pressures to share income with kin are most intense in Africa, and are common to many ethnic groups across the continent; they write that "African kin groups are distinctive both by their ubiquity and by the strength of their claims upon members." Yet, such social pressure to share is not confined to Sub-Saharan Africa, or even the developing world: Geertz (1963) describes a similar phenomenon in post-independence Indonesia, while Stack (1974) documents the existence of analogous kin obligations in poor, urban communities in the United States.

In this paper, we report the results of an experiment designed to measure social pressure to share income with relatives and neighbors. We use a controlled laboratory environment to explore behaviors which are difficult to document using survey data: the willingness to forgo profitable investment opportunities to keep income secret. Our experimental design yields simple theoretical predictions, and also provides the identification necessary for structural estimation of the underlying "kin tax" parameter in the presence of heterogeneous risk preferences.

[1] See Ligon (1998), Ligon, Thomas, and Worrall (2002), Albarran and Attanasio (2003) for examples. Foster and Rosenzweig (2001), which explores the impact of altruism on the set of self-enforcing insurance arrangements, is an important exception.

[2] A rich literature characterizes constrained optimal mutual insurance arrangements when agents can hide a portion of their income (cf. Cole and Kocherlakota 2001). Our work complements this literature, introducing more heterogeneity and noise in individual behavior, but at the cost of a reduced ability to characterize the constrained optimal arrangement.

[3] See Hoff and Sen (2006), Comola and Fafchamps (2011), and Dupas and Robinson (2013) for other examples.

We conduct economic experiments in 26 rural communities in western Kenya. Within the experiment, subjects receive an endowment which they divide between a risk-free savings account and a risky but profitable investment. The size of the endowment varies across subjects, and the distribution of endowment sizes is common knowledge. While the amount saved is always private information, we randomly vary whether the amount invested in the risky security can be observed by other subjects, creating an incentive for those receiving the large endowment to invest no more than the small endowment, thereby keeping their endowment size hidden. We also offer a subset of subjects the option of paying to keep their investment returns secret, allowing us to directly measure the willingness-to-pay to hide income.

In light of evidence that women and men have different risk preferences (cf. Croson and Gneezy 2009) and that women in poor communities may have more trouble accumulating savings (cf. Dupas and Robinson 2013), we stratify our experiment by gender, and report results for men and women separately.[4] Consistent with a simple model of decisions in the experiment when risk preferences are heterogeneous and subjects face pressure to share income, we find that women receiving the large endowment are more likely to make investment choices that obscure the amount of money they received when investment returns are observable. The effect we observe among women appears to be driven primarily by the behavior of women whose payoffs are most directly visible to their kin network: those with relatives attending the experiment. Interestingly, we find no similar tendency to hide income among men. Among women, impacts are unlikely to be driven by in-laws providing information to husbands: there is no direct impact of having one's husband present, and choice patterns are similar in the sub-sample of unmarried women. We also demonstrate that the gender difference we observe is unlikely to be driven by differences (across genders) in other observable characteristics. At the village level, women's tendency to hide income within the experiment is negatively associated with durable asset accumulation by households, skilled and formal sector employment, and the probability of using fertilizer on crops, suggesting that social pressure to share may hinder growth and development. Among subjects given the opportunity to pay a randomly-assigned price to keep income hidden, 30 percent of those able to afford the cost of hiding income choose to do so. These subjects pay an average of 15 percent of their gross payout from the experiment.

After presenting our reduced form results, we estimate the magnitude of the "kin tax" parameter via maximum simulated likelihood in a mixed logit framework. This is important because the size of the treatment effect of observability depends on both the level of social pressure and the distribution of individual risk preferences; decisions in our (control) private information treatments suggest substantial heterogeneity in risk aversion across subjects.[5] After controlling for

[4]Dupas and Robinson (2013) find evidence that female daily wage earners in western Kenya are more savings constrained than men in similar occupations. In a similar vein, De Mel, McKenzie, and Woodruff (2008), De Mel, McKenzie, and Woodruff (2009), and Fafchamps, McKenzie, Quinn, and Woodruff (2014) find lower returns to capital for microenterprises operated by women than for those operated by men.

[5]See Hey and Orme (1994); Choi, Fisman, Gale, and Kariv (2007); Andersen, Harrison, Lau, and Rutström (2008); Choi, Kariv, Müller, and Silverman (2014), and Von Gaudecker, van Soest, and Wengström (2011) for further experimental evidence on risk preference heterogeneity.

unobserved heterogeneity in risk preferences, we find a statistically significant tax for women, averaging four percent in general. For women whose kin attend the experimental session, consistent with the reduced form pattern, the tax appears to be twice as large. Simulations suggest that this level of social pressure may lead to a dramatic reduction in the likelihood of starting a business.

The main aim of this paper is to document the importance of social pressure in interhousehold transfer relationships within kin networks in poor communities. We make several contributions to the existing literature. First, we introduce a novel lab experiment designed to measure social pressure to share income in field settings.[6] The experiment is simple to understand, but provides subjects with a rich menu of investment options and multiple mechanisms for hiding income. Our design allows us to effectively rule out several alternative explanations: behavior is not consistent with models of attention aversion or a desire to avoid social sanctions against risk-taking, and the lack of a direct effect of having one's spouse present suggests that kin networks are not just passing information on to husbands. Second, treatment assignments within the experiment were randomized within villages, allowing us to explore the association between community outcomes and income hiding in the lab. Third, we link decisions in the experiment to a model of individual investment choices when risk preferences are heterogeneous, and estimate this model via maximum simulated likelihood. This allows us to recover an estimate of the kin tax parameter, and to simulate the magnitude of its impact on microentrepreneurship.

Our experimental approach is, however, not without drawbacks. One concern with all experimental work on distributional choices is that subjects might behave differently in non-laboratory settings when the money being shared or redistributed is someone's earned income. In fact, consistent with descriptive evidence suggesting that those living in rural villages in less developed countries often attribute success to luck and supernatural forces rather than individual effort (cf. Platteau 2000), previous experimental work suggests that this distinction between earned and unearned income is less salient in rural Kenya than in a developed country (Jakiela forthcoming). A second issue is that our one-shot, communication-free experiment intentionally minimizes the potential for risk-pooling; risk-pooling may facilitate risk-taking by cushioning downside risk. We view our work as a first step in the direction of characterizing and quantifying the extent of social pressure to share income; future work will be needed to explore the extent to which such pressure depends on the source of income and to better understand its impacts.

The rest of this paper is organized as follows: Section 2 describes our experimental design and procedures; Section 3 presents a simple theoretical framework for interpreting our results; Section 4 presents our main reduced-form empirical results; Section 5 presents our structural framework and estimates; and Section 6 concludes.

[6] Our design is most closely related to Barr and Genicot (2008) and Ligon and Schechter (2012), both of whom introduce original experiments capturing the desire to avoid social sanctions outside the lab. For recent examples of experimental work on risk-pooling with limited commitment, see Barr and Genicot (2008), Charness and Genicot (2009), Attanasio, Barr, Cardenas, Genicot, and Meghir (2012), and Barr, Dekker, and Fafchamps (2012).

2 Experimental Design and Procedures

2.1 Structure of the Experiment

The experiment was designed to introduce exogenous variation in the observability of investment returns. Within the experiment, each participant was given an initial endowment, either 80 or 180 Kenyan shillings.[7] Each subject divided her endowment between a zero-risk, zero-interest savings account and an investment that was risky but profitable in expectation. The subject received five times the amount that she chose to invest in the risky prospect with probability one half, and lost the amount invested otherwise. A coin was flipped to determine whether each risky investment was successful. Thus, the main decision subjects faced was how much of their endowment to invest in the risky security and how much to allocate to the secure, zero-profit alternative.[8]

Within the experiment, players were randomly assigned to one of six treatments. First, players were allocated either the smaller endowment of 80 shillings or the larger endowment of 180 shillings. Endowment sizes were always private information — experimenters did not identify those subjects who received the large endowment. However, the distribution of endowments was common knowledge, so all subjects were aware that half the participants received an extra 100 shillings.

Every player was also assigned to either the **private** treatment or one of two public information treatments, the **public** treatment or the **price** treatment. Participants assigned to the **private** treatment were able to keep their investment income secret: the decisions they made in the experiment were never disclosed to other participants. In contrast, those assigned to the **public** treatment were required to make an announcement revealing how much they had invested in the risky security and whether their investment was successful to all of the other participants at the end of the experiment.[9] The amounts that subjects invested in the zero-interest savings technology were never revealed, so those who received the larger amount could choose whether to invest 80 shillings or less so as to obscure their endowment size. Finally, those assigned to the **price** treatment were obliged to make the public announcement revealing their investment returns unless they preferred to pay a price, p, to avoid making the announcement. Prices ranged from 10 to 60 shillings, and were randomly assigned to subjects in the price treatment. Subjects were informed what price they faced *before* making their investment decisions, but decided whether to pay the price *after* investment returns were realized.[10]

[7] Experimental sessions were conducted between August 3 and October 1, 2009. We report dollar amounts using the exchange rate 75.9 shillings to the dollar, which is the average over that period. The two endowments were equivalent to 1.05 and 2.35 U.S. dollars, respectively. For comparison, the median monthly wage among subjects in full-time, unskilled employment is 2000 shillings (26.35 U.S. dollars).

[8] Variants of this portfolio choice design have been used for at least two decades in experimental economics, though without the variation in observability that is the focus of our study. Earlier examples include Loomes (1991), Alm, Jackson, and McKee (1992), Gneezy and Potters (1997), and Choi, Fisman, Gale, and Kariv (2007).

[9] Subjects were informed that they were allowed to delegate the task of making the public announcement to a member of the research team if they wished to avoid the public speaking aspect of the announcement process.

[10] Hence, subjects in the price treatment were not always able to afford to pay to avoid the public announcement:

Random assignment to treatment generated exogenous variation in the observability of investment returns and created costly opportunities to hide income. Assignment to the public information treatments meant that outcomes were verifiable, and might therefore facilitate risk-pooling and, consequently, risk-taking. On the other hand, if subjects face social pressure to share income with neighbors and kin, they might be willing to pay for obscurity when returns are visible.[11] In particular, the experiment creates two mechanisms through which subjects could incur a cost to hide income from others. First, those receiving the larger endowment could keep their endowment size secret by investing no more than 80 shillings. Second, subjects in the price treatment could pay the randomly-assigned price, p, to conceal their income entirely.

Random assignment to treatment was stratified by gender, allowing us to conduct our analysis separately for men and women. Such a disaggregated analysis is important for two reasons. First, as discussed in more detail in Section 3, behavior in our experiment depends on individual risk preferences. There is ample evidence that risk preferences are heterogeneous, but many studies find that, on average, women are more risk averse than men (see Croson and Gneezy (2009) for a detailed discussion of the literature). Second, as we discuss in detail below, women and men in our sample have quite different social realities: women are more likely to be living on their own because their spouse is working elsewhere, they are less likely to have other sources of cash income or a bank account, and those who are married typically live near their husband's blood relatives rather than their own.[12] Though studies to date have not directly measured the extent of social pressure to share income, gender differences in social pressure would be consistent with existing evidence on gender differences in savings constraints and returns to capital for microenterprises (Dupas and Robinson 2013, De Mel, McKenzie, and Woodruff 2009, Fafchamps, McKenzie, Quinn, and Woodruff 2014).

2.2 Experimental Procedures

Experiments were conducted in 26 rural, predominantly agricultural communities in western Kenya.[13] One day prior to each experimental session, the survey team conducted a door-to-door recruitment campaign, visiting as many compounds and households within each village as possible. Short baseline surveys were administered during these recruitment visits.[14] All adult

those who invested and lost a large fraction of their endowment did not always have enough experimental income left to pay the randomly-assigned price, p. Subjects were never allowed to use money from outside the experiment to pay to avoid the public announcement.

[11] Our design is related to those of Barr and Genicot (2008), Attanasio, Barr, Cardenas, Genicot, and Meghir (2012), and Ligon and Schechter (2012), who also use experiments which capture the desire to avoid social sanctions outside the lab.

[12] See Brabin (1984) and Luke and Munshi (2006) for discussion of the patrilocal traditions of the most common tribes in Western Kenya.

[13] Communities were selected to be at least five kilometers apart from one another, to prevent overlap in subject populations, and to avoid areas where IPA–Kenya had ongoing projects.

[14] As is typical in rural Kenya, survey teams moved from compound to compound with the assistance of local guides, typically village elders. When the recruitment team was unable to contact and survey anyone from a particular household prior to the experiment, village elders were asked to invite the adult members of that household to attend the subsequent experimental session. In those cases, baseline surveys were administered on the day of

residents of each village were invited to participate in the experimental session the following day. Approximately 80 percent of households contacted prior to the experimental sessions sent at least one adult to participate.[15]

Experimental sessions were conducted in empty classrooms at local primary schools. Sessions included an average of 83 subjects; no session included fewer than 65 or more than 100 subjects. Each session lasted approximately three hours.

Within each session, participants were stratified by gender and education level. There were six experimental treatments, corresponding to the three information conditions (private, public, price) interacted with the two endowment sizes. Within each stratum, players were randomly assigned to each of the six treatments with equal probability.[16] Players assigned to the price treatments were subsequently assigned a random price from the set of multiples of ten between 10 and 60.[17]

Experimental sessions were structured as follows. After a brief introduction, enumerators read the instructions and answered participant questions, illustrating the decisions that a subject might face with a series of wall posters.[18] Subjects were then called outside one at a time to make their investment decisions.[19] Enumerators began by asking a series of questions designed to make sure that subjects understood the experiment. Subjects were then informed whether they had received the large or small endowment and whether they were assigned to the private, public, or price treatment. Those assigned to the price treatment were also told what price they would need to pay if they wished to avoid the public announcement. Subjects then made their investment decisions: each subject was handed a number of 10 shilling coins equivalent to her endowment; the participant divided these coins between a "savings" cup and a "business" cup.[20] After recording a subject's investment decision, the enumerator would give the subject a one shilling coin to flip.[21] The outcome of the coin toss determined whether the money placed in the business cup was multiplied by five or removed from the subject's final payout. Subjects assigned

the experimental session (typically before but sometimes after the experiment). Our main results are all robust to the omission of subjects who were not surveyed prior to the experiment.

[15]In the Online Appendix, we examine the correlates of choosing to attend to the experimental session (in the sample of individuals who were surveyed in their homes prior to the experiment). Individuals from larger households and those more embedded in interhousehold transfer networks were more likely to attend, while ethnic minorities were somewhat less likely to attend. We find no evidence of differential selection by gender. Subjects in more isolated communities — as measured by distance from a paved road — were more likely to participate, though this cannot bias our results since assignment to treatment occurs within villages.

[16]The first experimental session did not include the two price treatments, so players were assigned to each of the other four treatments with equal probability.

[17]Note that each session included subjects assigned to all three information conditions, so it was impossible to infer which individuals chose to pay to avoid the public announcement, since these subjects could not be distinguished from those randomly assigned to the private treatments.

[18]Detailed instructions are included in the Online Appendix.

[19]Since some participants had limited literacy skills, decisions were recorded by members of the research team. To ensure that earnings not announced publicly remained private information, each enumerator sat at a desk in an otherwise empty section of the schoolyard.

[20]The two cups made of plastic and were identical in size and color.

[21]To limit the possibility of influencing the outcome of the coin flip, each subject placed the coin into a sealed, opaque container which she shook vigorously before opening it to reveal the outcome of the coin toss.

to the price treatment were then asked whether they wanted to pay the fee to avoid announcing their investment results. If they had enough money left to pay the fee, and they chose to do so, it was deducted from their payoff. After all decisions had been recorded, public announcements were made. Each subject received her payout in private at the end of the session, and was allowed to leave immediately after receiving her money. Figure 1 summarizes the progression of activities within the experiment.

2.3 Experimental Subjects

Sessions were conducted in Kenya's Western Province, in three adjoining districts: Bunyala, Samia, and Butula. All three districts are predominantly smallholder farming communities, though Samia and Bunyala also have ports on Lake Victoria. Summary statistics on experimental subjects are presented in Table 1.

61 percent of subjects are female, reflecting the fact that women account for more than half of the rural population. Men often reside in urban and peri-urban areas — where more jobs are available — and make occasional visits to their wives and children in the village.[22] 77 percent of subjects are married, while 12 percent are widowed, separated, or divorced. 97 percent of male subjects who are married live in the same household as their spouse, compared with only 76 percent of married female subjects. This pattern reflects the migratory patterns described above and, to a lesser extent, polygyny (9 male subjects report living with with more than one current spouse). The median household size is 6; less than 2 percent of subjects live alone, and less than 5 percent of subjects live in households that include more than 12 people.

Respondents ranged in age from 18 to 88. 9 percent of subjects had no formal schooling, while 12 percent had finished secondary school. Among men in our sample, 8 years of schooling (completing primary school) is both the median and the mode; among women, the median and the mode is 7 years of schooling. 23 percent of respondents live in households with at least one employed household member; most (65 percent of) employed subjects do agricultural work or other unskilled labor. The median monthly wage among participants with full-time employment was 2,050 Kenyan shillings, or 1.35 USD per day (assuming 20 work days per month). 35 percent of subjects operate their own business enterprise, but only 12 percent of these businesses have any employees. 17 percent of participants have bank accounts,[23] and 53 percent are members of rotating savings and credit associations (ROSCAs).[24] Men are more likely to have bank accounts, but less likely to participate in ROSCAs.

[22] As part of another project in other rural villages in Kenya's Western Province, we conducted a census of all the adults living in each of two communities. We found that females accounted for 58 percent of adults in each of these villages.

[23] Dupas and Robinson (2013) found that less than 3 percent of the daily wage earners sampled in Bumala, Kenya, had savings accounts. While Bumala is just a few kilometers from the region where the present study took place, their data were collected over two years before our household survey.

[24] Gugerty (2007) surveys ROSCA participants in Busia and Teso Districts in western Kenya; she argues that the social component of ROSCA participation helps individuals overcome savings constraints. Anderson and Baland (2002) show income-earning women living in Nairobi slums use ROSCAs to protect their savings from their husbands.

Most experimental subjects in our sample live amongst their kin, and are embedded in inter-household transfer networks. Participants have a median of seven relatives living outside their households but in the same village. Because most of Kenya's tribes are patrilocal, married women generally live near their in-laws rather than their own blood relations (Brabin 1984, Luke and Munshi 2006). This is borne out in our sample: 50 percent of male subjects live in the same village as their mother, compared with only 12 percent of female subjects. 90 percent of subjects reported making a transfer in the last three months, while 41 percent reported receiving a transfer from another household over the same period.[25] 40 percent of subjects reported making a transfer to relatives within the village in the last 3 months, and 9 percent of subjects reported receiving a transfer from relatives in the village over the same time period. In the three months prior to being surveyed, 42 percent of subjects' households had been asked for a gift or loan, and 90 percent of households had contributed money to a "harambee," a local fundraising drive.[26]

Table 1 reports tests of balance across our six experimental treatments (anticipating the nature of our analysis, we report balance checks separately by gender).[27] The randomization was unproblematic, generating typically small differences in observables across treatments. Of 64 tests (32 for each gender) reported, we observe only six significant differences in observables across treatments (two of which are only marginally significant). Though the number of distant family members living in one's village differs significantly across treatments in the sample of men, this variation is driven by outliers: the maximum number (within a treatment) of distant relatives reported to live in one's village ranges from 72 to 154. A quantile regression of the median number of distant family members on the set of treatment dummies does not find significant differences across experimental treatments (results not shown). Monthly wages conditional on employment are also not balanced (possibly because so few subjects report being employed). Among women attending the experiment, the indicator for having one's spouse present at the experiment is not balanced. However, as we demonstrate in the Online Appendix, this results from an unusually large number of women whose spouses were present in the public small endowment treatment; this variable is balanced across the three large endowment treatments that are the focus of our reduced form analysis. Among women, there is also a significant difference in the number of televisions owned across treatments, and a marginally significant difference in the number of

[25]There are several reasons for the asymmetry between the probability of making a transfer and the probability or receiving a transfer. First, households comprising only children and/or the ill or handicapped could not participate (because only those adults who could be physically present for the session and were able to hear the oral instructions could take part), but are likely to be net recipients of transfers. Another reason is that most households in any village will make a small transfer toward the funeral expenses of their neighbors; since funerals are relatively rare events, we would expect to see an asymmetry between the likelihood of giving and receiving in the short-term. Finally, transfer data are quite noisy, and households are more likely to report transfers made than transfers received (Comola and Fafchamps 2011).

[26]A *harambee* is a self-help effort in which community members contribute money or resources to assist a particular person in need. The concept existed within a number of different tribal groups in Kenya, but was made into a national rallying cry by Kenya's first president, Jomo Kenyatta (Ngau 1987).

[27]In the Online Appendix, we report additional balance checks, including both tests which pool the genders and tests which restrict attention to subjects receiving the large endowment treatment (which are the focus of our reduced form analysis). We also show that the randomization of the price of avoiding the public announcement was as balanced as one could expect.

cows, but no difference in the total value of durable household assets. Thus, the randomization appears to have been successful in generating treatment groups that were broadly similar in terms of observable characteristics.

3 Theoretical Framework

3.1 Individual Investment Decisions

Subjects in our experiment receive an endowment of $m \in \{m_{small}, m_{large}\}$ — either 8 or 18 Kenyan coins worth 10 shillings each; they then decide how many of these coins to allocate to the business cup, choosing one of a finite number of possible investment levels. We model individual choices by assuming that subject i's utility of investing $b_j \in \{0, 10, 20, \ldots, m\}$ takes an additive random utility form,

$$EU_{ij} = EV_{ij} + \varepsilon_{ij}, \tag{1}$$

where EV_{ij} is the expected utility of investing b_j given a constant relative risk aversion (CRRA) utility function parameterized by $\rho_i \geq 0$,

$$v_i(x) = \frac{x^{1-\rho_i}}{1-\rho_i}, \tag{2}$$

and ε_{ij} is an i.i.d. type 1 extreme value distributed preference shock.[28] Subjects in our experiment make a discrete choice from among a finite set of investment levels; the investment amount associated with the highest CRRA expected utility (given ρ_i) is the most likely to be chosen, but other outcomes occur with positive probability. The probability that subject i chooses to invest b_j takes the standard logit form:

$$P_{ij} = \frac{e^{EV_{ij}/\sigma_\varepsilon}}{\sum_{k=1,\ldots,J_t} e^{EV_{ik}/\sigma_\varepsilon}} \tag{3}$$

where σ_ε^2 is proportional to the variance of $\varepsilon_{ij} - \varepsilon_{ik}$.[29] Thus, the probability of choosing an investment amount b_j depends on the associated CRRA expected utility, EV_{ij}; on the CRRA expected utilities of the other possible investment options; and on the noise parameter, σ_ε^2, which measures the relative importance of the CRRA expected utilities vis-a-vis the stochastic preference shocks.[30]

[28] In our setting, EV_{ij} is what Train (2003) terms "representative utility." Loomes (2005) refers to the logit error terms as "Fechner errors." See Hey and Orme (1994) and Von Gaudecker, van Soest, and Wengström (2011) for examples of their use in modeling stochastic choices in individual decision-making experiments.

[29] When $EV_{ij} = X'\beta$, β and σ_ε are not separately identified; σ_ε is identified in our framework because EV_{ij} is a non-linear function of parameters.

[30] A higher σ_ε^2 indicates noisier (less deterministic) choices; the probability of choosing the investment amount associated with the highest CRRA expected utility goes to 1 as σ_ε^2 approaches 0. For example, the probability that a subject with a CRRA coefficient of 0.5 chooses the investment amount with the highest CRRA expected utility ($b = 60$) in the private small endowment treatment is 0.92 when $\sigma_\varepsilon = 0.001$, but drops to 0.14 when $\sigma_\varepsilon = 0.1$.

Subjects assigned to the public and price treatments may face pressure to share their payoffs after the experiment. Following Ashraf (2009) and Goldberg (2013), we model this social pressure as a proportional tax on observable income: subject i is obliged to transfer a proportion, τ, of her observable income to members of her social network — for example, her spouse or her relatives.[31] Income is observable when an individual is known to have it with probability one, whether it is announced or not; income is unobservable whenever a person can plausibly deny having received it.[32] Thus, an individual assigned to one of the private treatments can plausibly claim to have invested and lost all or most of her endowment, limiting the potential for social pressure. In contrast, a subject assigned to the public small endowment treatment cannot hide any of her income: every subject is known to have received at least 80 shillings (the amount of the smaller endowment, m_{small}), and she is obliged to announce her investment level and whether her investment succeeded. In both the public large endowment treatment and the price treatments, subjects can make choices that conceal all or part of their payoffs — but at a cost. A subject assigned to the public large endowment treatment can choose to invest $b_j \leq m_{small}$, thereby making 100 shillings of income ($m_{large} - m_{small}$) unobservable (because she can plausibly claim to have received the small endowment). A subject assigned to one of the price treatments has the option of paying p shillings to avoid making the public announcement, making her entire payout from the experiment unobservable.

In what follows, we characterize the decisions of individual subjects conditional on ρ_i, the CRRA risk aversion parameter. Given ρ_i, the CRRA expected utility of investing b_j in experimental treatment t depends on: the size of one's endowment, $m \in \{m_{small}, m_{large}\}$; τ, the level of social pressure to share income (in the public and price treatments); and p, the randomly-assigned price of paying to avoid the public announcement (in the price treatments). We expect risk preferences to be heterogeneous (Binswanger 1980, Choi, Fisman, Gale, and Kariv 2007), and our goal is to characterize expected differences in behavior across (randomly-assigned) treatments. An advantage of our approach is that our discrete choice model characterizes the entire data-generating process; this allows us to complement analytical predictions (derived algebraically) with numerical predictions found by calculating choice probabilities at a range of points in the parameter space. To derive predictions numerically, we calculate quantities of interest (e.g. choice probabilities) at each of 1,000,000 points spanning the space of reasonable parameter values — at each point in a $\rho \times \tau \times \sigma_\varepsilon$ grid where $\rho \in [0.001, 3]$, $\tau \in [0.001, 0.5]$ and $\sigma_\varepsilon \in [0.001, 0.1]$.[33] We then

[31] Though Ashraf (2009) and Goldberg (2013) do not characterize the optimal contract, we note that a fixed tax rate τ, and piecewise linear variations on it, are the forms taken by optimal contracts under conditions described in Lacker and Weinberg (1989).

[32] Our theoretical assumptions in this regard are again consistent with those made by Ashraf (2009) and Goldberg (2013). See De Mel, McKenzie, and Woodruff (2009) for a closely related modeling approach.

[33] The range of reasonable risk aversion parameters was chosen to span the spectrum from (almost) risk neutral to highly risk averse. Our chosen range includes, for example, all the estimated CRRA parameters reported in Cardenas and Carpenter (2008). As discussed above, investment decisions are almost deterministic at the lowest σ_ε values we consider and are close to random (uniform) at the higher end of our σ_ε range. We view values of τ above 0.5 as unlikely since subjects would have little incentive to participate in our experiment if they expected most of their payoff to be appropriated by members of their social network.

extend our individual-level predictions to a population of heterogeneous individuals by combining our analytical results with numerical results that hold at every grid point in our simulations.

3.1.1 Public and Private Small Endowment Treatments

To frame our discussion, we begin by examining behavior in the private and public small endowment treatments. Subject i's CRRA expected utility of investing b_j in the private small endowment treatment is given by:

$$EV_{ij}^{PRIVATE \times SMALL} = \frac{(m_{small} - b_j)^{1-\rho_i}}{2(1-\rho_i)} + \frac{(m_{small} + 4b_j)^{1-\rho_i}}{2(1-\rho_i)}. \tag{4}$$

In the public small endowment treatment, i's CRRA expected utility of investing b_j is given by:

$$\begin{aligned} EV_{ij}^{PUBLIC \times SMALL} &= \frac{\left[(1-\tau)(m_{small} - b_j)\right]^{1-\rho_i}}{2(1-\rho_i)} + \frac{\left[(1-\tau)(m_{small} + 4b_j)\right]^{1-\rho_i}}{2(1-\rho_i)} \\ &= (1-\tau)^{1-\rho_i} EV_{ij}^{PRIVATE \times SMALL}. \end{aligned} \tag{5}$$

Thus, CRRA expected utilities in the public treatment are scaled by the (positive) factor $(1-\tau)^{1-\rho_i}$. The implication is that, given a fixed level of risk aversion ρ_i, the ordering of the CRRA expected utilities associated with the possible investment amounts (and hence the ordering of the probabilities of being chosen) is the same in the private and public small endowment treatments. Intuitively, because every subject receives at least m_{small}, subjects in the public small endowment treatment cannot adopt any strategy that will hide a portion of their income; the proportional social pressure tax on observable income, τ, impacts the absolute attractiveness of the investment levels, but not their relative attractiveness.[34]

Subject i's expected investment level in the private small endowment treatment is given by:

$$\sum_{j=1,\ldots,J_t} b_j \cdot P_{ij} = \frac{\sum_{j=1,\ldots,J_t} b_j \cdot e^{EV_{ij}^{PRIVATE \times SMALL}/\sigma_\varepsilon}}{\sum_{k=1,\ldots,J_t} e^{EV_{ik}^{PRIVATE \times SMALL}/\sigma_\varepsilon}}, \tag{6}$$

while i's expected investment level in the public small endowment treatment is:

$$\frac{\sum_{j=1,\ldots,J_t} b_j \cdot e^{EV_{ij}^{PRIVATE \times SMALL}/\left(\sigma_\varepsilon/(1-\tau)^{1-\rho_i}\right)}}{\sum_{k=1,\ldots,J_t} e^{EV_{ik}^{PRIVATE \times SMALL}/\left(\sigma_\varepsilon/(1-\tau)^{1-\rho_i}\right)}}. \tag{7}$$

[34] In the specific case of log utility ($\rho_i = 1$), the probability of choosing any amount b_j is the same in private and public small endowment treatments: the $\ln(1-\tau)$ term cancels out of the expression for the probability of choosing b_j. For other values of ρ_i, investment probabilities in the small endowment private and public treatments differ, but the difference in expected investment level is typically small — especially for small to moderate values of τ. For $\rho < 1$, $(1-\tau)^{1-\rho_i} < 1$, so (CRRA) expected utility differences are scaled down in the public treatments relative to the private treatments. Such a scaling has effects that are equivalent to an increase in the noise parameter, σ_ε: probability is shifted towards the tails (away from the central tendency) of the distribution. For $\rho > 1$, $(1-\tau)^{1-\rho_i} > 1$, so the effect is the opposite: probability is shifted away from the tails of the distribution.

The only difference between Equations 6 and 7 is that the noise parameter, σ_ε, is scaled by a factor of $1/(1-\tau)^{1-\rho_i}$. As is clear from inspection, when $\rho_i = 1$ or $\tau = 0$, the expressions in Equations 6 and 7 are identical. As such, for any positive δ, we can find a neighborhood of parameter values $\{\rho_i, \tau\}$ for which the difference between the expressions is less than δ. Moreover, this change in the noise parameter chiefly alters the spread of the distribution of investment decisions, not the location of its peak. In consequence, even when the expected investment levels in the two treatments are not identical, the differences for reasonable values of τ and ρ_i are not large.[35]

3.1.2 Public and Private Large Endowment Treatments

The CRRA expected utility of investing b_j in the private large endowment treatment is given by:

$$EV_{ij}^{PRIVATE\times LARGE} = \frac{(m_{large} - b_j)^{1-\rho_i}}{2(1-\rho_i)} + \frac{(m_{large} + 4b_j)^{1-\rho_i}}{2(1-\rho_i)}. \tag{8}$$

In the public large endowment treatment, subjects have the option of investing 80 shillings (m_{small}) or less, thereby making 100 shillings ($m_{large} - m_{small}$) unobservable. We indicate this potentially hidden quantity as:

$$H_{ij} = (m_{large} - m_{small}) \cdot \mathbb{1}\{b_j \leq m_{small}\} \tag{9}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function. Thus, the CRRA expected utility of investing b_j is given by:

$$\begin{aligned} &EV_{ij}^{PUBLIC\times LARGE} = \\ &\quad \frac{\left[(1-\tau)(m_{large} - b_j) + \tau H_{ij}\right]^{1-\rho_i}}{2(1-\rho_i)} + \frac{\left[(1-\tau)(m_{large} + 4b_j) + \tau H_{ij}\right]^{1-\rho_i}}{2(1-\rho_i)}. \end{aligned} \tag{10}$$

If the large endowment treatment did not create the possibility of hiding income, we would observe the same relationship between investment probabilities in the large endowment treatments that we see in the small endowment treatments: scaling CRRA utilities by $(1-\tau)^{1-\rho_i}$ would lead to a small shift in the probability distribution toward or away from the central tendency, but no change in the ordering of expected utilities or probabilities. However, the possibility of hiding income makes all possible investment levels below 90 shillings relatively more attractive — because $\tau H_{ij} = 100\tau$ is, in essence, added to the payoff in each state.

The empirical consequence of this shift in the relative attractiveness of investment levels above the size of the small endowment is an increase in the probability of investing no more than 80 shillings, thereby (implicitly) hiding 100 shillings, the difference between the large and the

[35] Calculating the expected level of investment in the private and public small endowment treatments at every point in our $\rho\times\tau\times\sigma_\varepsilon$ grid, we find that the median (grid-point level) difference in investment levels across treatments is -0.62 shillings. For $\tau \leq 0.1$, the magnitude of the difference in investment levels is less than 2 shillings at all simulated grid points. Thus, if parameter estimates suggest a moderately-sized τ of 10 percent or less, our numerical results indicate that average levels of investment in the private and public small endowment treatments should be approximately equal (specifically, though the expected difference depends on the exact distribution of ρ_i, it is, at most, 2 shillings).

small endowments. In Proposition 1, we show analytically that the probability of investing 80 shillings or less is always higher in the public large endowment treatment than in the private large endowment treatment when $\tau > 0$.

Proposition 1. *For any $\tau > 0$ and any $\rho \geq 0$, the probability of investing 80 shillings or less is higher in the public large endowment treatment than in the private large endowment treatment.*

Proof: see Online Appendix.

In fact, for the vast majority of cases involving reasonable parameter values, the probability of investing every (discrete) amount at or below 80 shillings is weakly higher in the public large endowment treatment than in the analogous private treatment: this holds for 89.9 percent of the 1,000,000 $\rho \times \tau \times \sigma_\varepsilon$ grid points we consider.[36] In the entire parameter range we consider, the probability of investing exactly 80 shillings is always at least as high in the public large endowment treatment as in the private large endowment treatment.[37] Combining this empirical regularity with Proposition 1 generates our first empirical prediction.

Prediction 1. *For $\tau \in (0, 0.5)$ and $\rho_i > 0$,*

(i) the probability investing no more than 80 shillings (the amount of the small endowment) is strictly higher, and

(ii) the probability of investing exactly 80 shillings is weakly higher

in the public large endowment treatment than in the private large endowment treatment.

This prediction applies not just to individuals, but to populations of individuals who are heterogeneous in terms of ρ_i and are randomly assigned to experimental treatments (so that, in expectation, the distribution of risk aversion parameters is the same in all treatments).

Interestingly, while the expected investment level is lower in the public large endowment treatment than in the private treatment at the majority of grid points, the opposite is true in some cases (at 5.4 percent of grid points). Specifically, for high values of ρ_i and relatively low values of σ_ε, the probability of investing more than 80 shillings is quite low in the private treatment; in such cases, the expected investment level may be higher in the public treatment than in the private treatment. Intuitively, for those who are relatively risk averse and would almost never choose to invest more than 80 shillings, the possibility of hiding one's endowment size is essentially equivalent to a small increase in background consumption; it makes it slightly more attractive to take on profitable risk, shifting probability mass toward higher investment levels that are still no larger than the small endowment.[38] We therefore focus our empirical analysis on probabilities rather than investment amounts.

[36] To focus attention on the set of empirically relevant cases and avoid rounding issues that arise with extremely small probabilities, we treat probability differences less than 0.001 as equivalent to 0.

[37] This does not hold analytically across the entire range of feasible (if unreasonable) parameter values. Specifically, when τ is sufficiently close to 1 (e.g. above 0.95), we observe combinations of parameter values such that the probability of investing exactly 80 shillings is lower in the private large endowment treatment than in the public large endowment treatment.

[38] Please see our earlier working paper version for related discussion (Jakiela and Ozier 2012).

3.1.3 Investment Decisions in Price Treatments

In the price treatments, subjects may choose to pay a randomly chosen price, $p > 0$, to avoid making the public announcement. Forward-looking subjects make their investment decisions in this setting conditional on their beliefs about whether or not they will pay p to avoid the public announcement in each of the possible outcome (payoff) states (heads and tails). They choose their investment level expecting to follow one of three strategies: never paying to avoid the public announcement, paying to avoid the public announcement only if their investment is successful, or always paying to avoid the public announcement.[39] We refer to these strategies as *never*, *heads*, and *always*, respectively.[40] We assume that subject i in a price treatment chooses investment amount b_j to maximize:

$$EU_{ij} = \max\left\{EV_{ij}^{never}, EV_{ij}^{heads}, EV_{ij}^{always}\right\} + \varepsilon_{ij}. \tag{15}$$

In other words, subjects choose an investment level expecting to follow an optimal strategy.[41]

Subjects assigned to the large endowment price treatment have two mechanisms for hiding income: paying p obscures their entire payout, while investing no more than 80 shillings keeps 100 shillings of their endowment hidden. This makes paying to avoid the public announcement relatively less attractive at investment levels at or below 80 shillings, especially when one's investment is unsuccessful. In fact, at reasonable parameter values, it is only optimal to pay to avoid the public announcement in the unsuccessful investment state when τ is extremely high. When subjects do not anticipate paying to avoid the public announcement in both the high and low payoff states (heads and tails), the same forces are at work in the price large endowment treatments that drive investment decisions in the public large endowment treatments, albeit to a

[39] Abstracting from random utility shocks, subject i receiving gross payout $x_i > 0$ in the small endowment treatment prefers to pay p to avoid the public announcement whenever

$$\frac{(x_i - p)^{1-\rho_i}}{1-\rho_i} \geq \frac{(1-\tau)^{1-\rho_i} x^{1-\rho_i}}{1-\rho_i}, \tag{11}$$

or $\tau \geq p/x_i$. Similarly, a subject in the large endowment treatment prefers to pay p to avoid the public announcement whenever

$$\frac{(x_i - p)^{1-\rho_i}}{1-\rho_i} \geq \frac{[(1-\tau)x + \tau H_{ij}]^{1-\rho_i}}{1-\rho_i}, \tag{12}$$

or $\tau \geq p/(x_i - H_{ij})$. It is clear that, if equations 11 and 12 hold when $x_i = m_i - b_j$ (or $x_i = m_i - b_j + \tau H_{ij}$), they will also hold when $x_i = m_i + 4b_j$ (or $x_i = m_i + 4b_j + \tau H_{ij}$). Hence, the strategy of paying to avoid the public announcement when one's investment fails but not when it succeeds is never optimal.

[40] For each possible investment level, b_j, the CRRA expected utility of following the *never* strategy is the same as in the public treatments, as defined in Equations 5 and 10. The CRRA expected utilities of the *always* and *heads* strategies are:

$$EV_{ij}^{always} = \frac{(m_i - b_j - p)^{1-\rho_i}}{2(1-\rho_i)} + \frac{(m_i + 4b_j - p)^{1-\rho_i}}{2(1-\rho_i)} \tag{13}$$

$$EV_{ij}^{heads} = \begin{cases} \left\{[(1-\tau)(m_{small} - b_j)]^{1-\rho_i} + (m_{small} + 4b_j - p)^{1-\rho_i}\right\} / [2(1-\rho_i)] & \text{if } m_i = m_{small} \\ \left\{[(1-\tau)(m_{large} - b_j) + \tau H_{ij}]^{1-\rho_i} + (m_{large} + 4b_j - p)^{1-\rho_i}\right\} / [2(1-\rho_i)] & \text{if } m_i = m_{large} \end{cases} \tag{14}$$

[41] For some values of b_j and p, *always* may not be a viable strategy — for example, participant i cannot plan to pay p if she invests and loses her entire endowment. Whenever $m_i - b_j < p$, EV_{ij}^{always} is omitted from Equation 15.

lesser extent. This leads to clear numerical patterns: at all 1,000,000 simulated grid points, the probabilities of investing both no more than 80 shillings and exactly 80 shillings are weakly higher in the price large endowment treatment as in the private large endowment treatment, leading to our second empirical prediction.

Prediction 2. *For $\tau \in (0, 0.5)$ and $\rho_i > 0$, both*

(i) the probability investing no more than 80 shillings (the amount of the small endowment), and

(ii) the probability of investing exactly 80 shillings

are weakly higher in the price large endowment treatment than in the private large endowment treatment.

Taken together Predictions 1 and 2 indicate that, if subjects face a non-negligible amount of social pressure to share their income with others, we should expect to observe higher proportions of subjects investing no more than 80 shillings and exactly 80 shillings in the public and price large endowment treatments than in the private large endowment treatment.[42] We test this prediction in Section 4.1.

3.2 Paying to Avoid the Public Announcement

After the investment return is realized, if subject i assigned to one of the price treatments has a gross payout greater than p, she decides whether to pay p to avoid the public announcement. As in the case of investment decisions, we model the noise in the decision about whether to pay p in a random utility framework: given a realized gross payout of x_i, the utility of paying to avoid the announcement is:

$$\frac{1}{(1-\rho_i)}\left(x_i - p\right)^{1-\rho_i} + \zeta_{i1}. \tag{16}$$

This includes ζ_{i1}, participant i's idiosyncratic utility from paying to avoid the announcement. The analogous expression for the utility of making the announcement is:

$$\frac{1}{(1-\rho_i)}\left[(1-\tau)x_i + \tau H_{ij}\right]^{1-\rho_i} + \zeta_{i0} \tag{17}$$

where ζ_{i0} is idiosyncratic utility associated with announcing her payout and facing pressure to share a portion of her observable income. We assume that ζ_{i1} and ζ_{i0} are draws from a type 1 extreme value distribution, and that they are independent of each other and of the ε_{ij} terms.[43]

[42]Numerical simulations demonstrate that the probability of investing no more than 80 shillings is always weakly larger in the public large endowment treatment than in the price large endowment treatment; however, the probability difference is often quite small. More surprisingly, numerical simulations also show that the probability of investing exactly 80 shillings is not always weakly larger in the public large endowment treatment than in the private large endowment treatment.

[43]Our model also implicitly assumes that subjects do not anticipate this stochastic component of their decisions about whether to pay to avoid the public announcement. Though it would also be possible to model beliefs about one's future actions as a probability distribution, this simplification makes the model substantially more tractable.

The probability that subject i pays to avoid the public announcement is then given by:

$$P_i^{exit} = \frac{1}{1 + e^{\left([(1-\tau)x_i+\tau H_{ij}]^{1-\rho_i} - (x_i-p)^{1-\rho_i}\right)/[(1-\rho_i)\gamma]}}, \tag{18}$$

where γ^2 is proportional to the variance of the difference between ζ_{i1} and ζ_{i0}. It is apparent that P_i^{exit} is decreasing in p, the randomly-assigned price. P_i^{exit} is equal to one half when subject i is indifferent between paying p to avoid the public announcement and not paying p — i.e. whenever

$$\frac{[(1-\tau)\, x_i + \tau H_{ij}]^{1-\rho_i}}{1-\rho_i} = \frac{(x_i - p)^{1-\rho_i}}{1-\rho_i}. \tag{19}$$

The probability of paying to avoid the public announcement is less than one half when the left-hand side of Equation 19 exceeds the right-hand side; when the right-hand side is larger, the probability of paying to avoid the public announcement is greater than one half. Rearranging the terms in Equation 19 generates our final proposition, which links behavior in price treatments to the underlying value of τ.

Proposition 2. *Let z_i denote subject i's observable payout — her gross payout x_i minus H_{ij}, the 100 shillings hidden from view if a subject receives the large endowment and then chooses an investment level of no more than 80 shillings. For all i,*

$$P_i^{exit} \geq \frac{1}{2} \Leftrightarrow \tau \geq \frac{p}{z_i}.$$

The expected proportion of subjects choosing to pay to avoid the public announcement is greater than one half for values of p and z_i such that $\tau \geq p/z_i$ and less than one half otherwise.

This result follows directly from Equation 18; the proof is omitted to save space. In Section 4.2, we examine our data through the lens of this proposition; this allows us to calculate a back of the envelope estimate of the level of social taxation consistent with our willingness-to-pay results.

3.3 Summary of Predictions

Before proceeding to our empirical analysis, we briefly review the predictions of our theoretical model. Our key theoretical result is that, when subjects face pressure to share observable income, the probabilities of investing both no more than 80 shillings and exactly 80 shillings will be weakly higher in both the public and price large endowment treatments than in the private large endowment treatment. Our model does not generate clear predictions about the overall levels of investment in the large endowment treatments (though it does suggest that expected investment levels will be similar in the private and public small endowment treatments for moderately-sized τ). In the price treatments, the willingness-to-pay to hide income depends on the ratio of the randomly-assigned exit price to one's observable income: the probability of paying to avoid the public announcement will be above (resp. below) 0.5 when τ is greater than (resp. less than) this ratio.

4 Results

In this section, we estimate the impacts of observability on behavior in the experiment. We report results relating to investment decisions in Section 4.1, and describe decisions regarding whether to pay to avoid the public announcement in Section 4.2. Summary statistics on outcomes in the experiment are presented in Table 2 and in Online Appendix Table 3 (for outcomes specific to the price treatments). On average, subjects chose to invest just over half their endowments in the business cup; the fraction of the endowment invested is similar in the large and small endowment treatments.[44] Subjects earned an average of 3.17 US dollars (240 Kenyan shillings) in the experiment, which is equivalent to eight percent of mean monthly wages among subjects reporting paid employment. Thus, stakes were large, but not life-altering. There is substantial variation in payoffs across across treatments: subjects in the public, small endowment treatment received the lowest average payoffs (1.83 US dollars); those in the private, large endowment treatment earned the most (4.68 US dollars).

4.1 Individual Investment Decisions

Our analysis of individual investment decisions focuses on testing the predictions of our theoretical model. Subjects allotted the larger (180 shilling) endowment could avoid revealing the amount they received by investing no more than 80 shillings (the amount of the small endowment) in the business cup; as a consequence, our model predicts that both the probability of investing no more than 80 shillings and the probability of investing exactly 80 shillings will be higher in the public and price large endowment treatments than in the private large endowment treatment.[45]

To test these hypotheses, we estimate OLS regressions of the form

$$Y_i = \alpha + \beta PublicTreatments_i + X_i'\gamma + \epsilon_i \tag{20}$$

where Y_i is an outcome of interest, an indicator for investing no more than 80 shillings ($Y_i = LTE80_i$) or for investing exactly 80 shillings ($Y_i = EX80_i$); $PublicTreatments_i$ is an indicator equal to 1 if subject i was assigned to either the public or price treatment; X_i is a vector of individual characteristics; and ϵ_i is a conditionally mean-zero error term. We also report specifications including village fixed effects and, because our key outcomes of interest are binary, logit

[44] We note that the fraction of the budget invested is very similar in all three small endowment treatments, as predicted by the model when τ is not unduly large. Histograms of investment decisions are included in the Online Appendix. All but 5 subjects invest a positive amount, suggesting that subjects are not concerned with social sanctions against (at least moderate) risk-taking within the experiment. Across all three small endowment treatments, 2.9 percent of subjects chose to invest their entire endowment. This suggests another possible reason subjects receiving the large endowment might choose to invest amounts below 80 shillings rather than exactly 80 shillings: doing so would avoid the possibility that investing exactly 80 shillings might make others suspicious about one's endowment size.

[45] Our theoretical model makes predictions about investment decisions in the large endowment treatments, and our analysis throughout Section 4.1 focuses on the sample of subjects randomly assigned to the large endowment treatments. To keep our writing as concise as possible, throughout the rest of Section 4.1 we omit the descriptive "large endowment" when referring to our experimental treatments.

specifications of the form

$$\Pr\left[Y_i = 1\right] = \frac{e^{\left(\alpha + \beta PublicTreatments_i + X_i'\gamma\right)}}{1 + e^{\left(\alpha + \beta PublicTreatments_i + X_i'\gamma\right)}}. \tag{21}$$

Results are reported in Table 3. Given our stratified design, we present results that are disaggregated by gender: columns 1 through 4 report results for women, while columns 5 through 8 report results for men.[46]

In all specifications, women are significantly more likely to invest no more than 80 shillings in the public and price treatments than in the private treatment. For example, OLS coefficient estimates indicate that women are 9.6 to 10.9 percentage points more likely to invest 80 shillings or less when investment returns are observable (Table 3, Panel A, columns 1 and 2).[47] Logit results are similar (Table 3, Panel A, columns 3 and 4). These coefficient magnitudes mean that the probability that a woman invests no more than 80 shillings is approximately 25 percent higher in the public and price treatments than in the private treatment. Women are also significantly more likely to invest exactly 80 shillings in the public and price treatments, as our model predicts. Coefficient estimates suggest that the probability of investing exactly 80 shillings is 6.2 to 7.1 percentage points higher in the public and price treatments than in the private treatments (Table 3, Panel B, Columns 1 through 4).

Interestingly, we find that random assignment to the public and price treatments does not appear to have a significant impact on men's investment decisions: men are not significantly more likely to invest either no more than 80 shillings or exactly 80 shillings in the public and price treatments than in the private treatments (Table 3, columns 5 through 8). In fact, coefficient estimates suggest a small, statistically insignificant decline in the probability that a male subject invests no more than 80 shillings (Table 3, Panel A, columns 5 through 8), and a small and statistically insignificant increase in the probability of investing exactly 80 shillings (Table 3, Panel B, columns 5 through 8). There are many possible interpretations. One possibility is that men are less subject to unwanted social pressure to share their income with members of their community, making the possibility of paying to hide income less attractive. However, this need not be the case. It is also possible that the opportunities to hide income afforded by the experiment

[46] We report pooled specifications, some of which include an interaction between $PublicTreatments_i$ and the indicator for being female, in Online Appendix Table 4. In pooled specifications which do not include gender interactions, random assignment to the public or price treatments is positively associated with the probability of investing both no more than 80 shillings and exactly 80 shillings, but the effect is only marginally statistically significant. When $LTE80_i$ (the indicator for investing no more than 80 shillings) is used as the dependent variable, the interaction between the female dummy and the $PublicTreatments_i$ indicator is marginally significant in both specifications we consider, and the combined effect of $PublicTreatments_i + Female \times PublicTreatments_i$ is highly significant. When $EX80_i$ (the indicator for investing exactly 80 shillings) is used as the dependent variable, the interaction term is not statistically significant; it is only the sum of $PublicTreatments_i + Female \times PublicTreatments_i$ that is statistically significant.

[47] Coefficients and significant levels are similar if we restrict attention to either the private vs. public treatment comparison or the private vs. price treatment comparison. Specifically, when the dependent variable is the indicator for investing no more than 80 shillings, coefficient estimates and significance levels are nearly identical; when the dependent variable is the indicator for investing exactly 80 shillings, coefficient estimates are directionally similar but differ in both magnitude and significance.

may have been less valuable to men — if, for example, a larger fraction of their consumption occurs outside of the home in contexts where expenditures are observable. Alternatively, men may have failed to grasp the fact that investing no more than 80 shillings obscured the size of one's endowment. We discuss these possibilities in greater detail below, and place this pattern in the context of our model in Section 5.3.

Given our stratified design, it is appropriate to report results for men and women separately; this is particularly true given the extensive literature suggesting that men and women have different risk preferences and, at least in rural villages in Kenya, different social realities. Nonetheless, a concern is that gender may be proxying for some other household or individual characteristic that actually explains the differential gender impacts that we observe. Random assignment within gender strata means that this concern is related to the interpretation of our data, not to identification — but this does not minimize its potential importance. In the Online Appendix, we provide additional analysis exploring this issue. We show, first, that there is no evidence of differential selection into the experiment by gender; second, that other interactions between covariates and $PublicTreatments_i$ indicator do little to dampen observed gender differences; and finally, that the observable characteristics that differ systematically between men and women do not explain treatment effect heterogeneity within each gender separately, making it highly unlikely that gender is merely proxying for some other observable characteristic. This suggests that the effects we document are indeed gender-specific, and are not explained by other characteristics which differ across the sexes.

4.1.1 Treatment Effect Heterogeneity

If women are socially obligated to share their income, the extent of income hiding should be associated with factors predicting the level of social pressure an individual is likely to face after the experiment. We consider two sources of social pressure suggested by existing literature: relatives outside of one's household and, within the household, one's husband.[48] We first test whether the impact of the public information treatments on investment behavior (documented above) is concentrated among those women with close kin attending the experiment who would observe their payoffs. We test for such treatment effect heterogeneity by interacting the $PublicTreatments_i$ variable with indicators for whether or not a subject had any close kin attending the experiment, thereby disaggregating the impact of the public information treatments.[49] We estimate the OLS

[48] Hoff and Sen (2006) highlight the role played by kin networks in extracting surplus from successful relatives, while Baland, Guirkinger, and Mali (2011) provide evidence that individuals seek to hide income from family members. Robinson (2012) and Ashraf (2009) document limited commitment and observability effects *within* the household, while Anderson and Baland (2002) argue that Kenyan women use ROSCAs to hide savings from their husbands.

[49] We define close kin as a respondent's parents, in-laws, grandparents, siblings, grown children, and aunts and uncles. Note that we were unable to stratify treatment assignment by kin presence because the variable was constructed after experimental sessions took place using a name-matching algorithm that compared close relatives' names (recorded during pre-experiment surveys) with the roster of experimental subjects.

regression equation

$$\begin{aligned} Y_i = \alpha + \beta_1 KinPresent_i + \beta_2 PublicTreatments_i \times KinPresent_i \\ + \beta_3 PublicTreatments_i \times NoKinPresent_i + \nu_v + X_i'\zeta + \varepsilon_{it} \end{aligned} \tag{22}$$

where Y_i is one of our binary outcomes of interest — either the indicator for investing no more than 80 shillings ($LTE80_i$) or the indicator for investing exactly 80 shillings ($EX80_i$).[50] Our specification includes an indicator for having kin present, thereby controlling for any overall differences in investment behavior between subjects with relatives attending the experiment and those without. Results are reported in Table 4.

Women with close kin attending the experiment are more likely to invest both 80 shillings or less and exactly 80 shillings in the public and price treatments than in the private treatment, though only the first result is statistically significant. Women with relatives present are 41.8 percentage points more likely to invest 80 shillings or less (Table 4, Column 1, p-value < 0.001) and 15.2 percentage points more likely to invest exactly 80 shillings (Table 4, Column 4, p-value 0.101). In contrast, coefficient magnitudes are smaller — though still positive — and not statistically significant for women with no kin attending the experiment (p-values 0.125 and 0.126, respectively). Thus, our results suggest that relatives outside the household may be an important source of pressure to share income.

Women with relatives attending the experiment may, of course, differ from other female subjects in many ways; the heterogeneity that we observe could be driven by another variable that is correlated with kin presence. One factor that is strongly associated with kin presence at the experiment is, unsurprisingly, the number of close relatives residing in the village. To test whether observability effects are driven by simply having relatives in the village (whether or not they are present at the experiment), we estimate specifications which control for the presence of close kin in the village but not at the game, plus and interaction between this variable and $PublicTreatments_i$ (Table 4, Columns 2 and 5). In both specifications, we find that having kin in the village who are physically present at the experiment is the critical factor predicting differences in investment behavior across treatments.[51]

Women may also seek to hide income from their husbands. An alternative explanation of our result is that kin are relevant because they pass information about wives' incomes to their husbands. As discussed above, 76 percent of women in our sample are married, and the two main ethnic groups in the area — the Luhya and the Luo — are both patrilocal (Brabin 1984, Luke and Munshi 2006), so married women typically live near their in-laws rather than their blood relatives. To explore this possibility, we estimate Equation 22 replacing the kin variables with indicators for whether or not a subject's spouse attended the experiment (Table 4, Columns 3 and 6). If relatives were only important because they shared information with husbands, we would expect

[50] We report OLS results; logit results are nearly identical in terms of both coefficient magnitudes and statistical significance, and are omitted to save space. Village fixed effects and controls for age and education level, household size, marital status, and household durable asset holdings are included in all specifications.

[51] We explore this issue in greater detail in the Online Appendix.

the direct effect of a husband's presence to be at least as large as the impact of having kin at the experiment. However, in both specifications we consider, we cannot reject the hypothesis that $PublicTreatments_i \times SpousePresent_i$ is equal to $PublicTreatments_i \times NoSpousePresent_i$, i.e. than women are no more inclined to conceal the size of their endowment when their husband is present.[52] Thus, our evidence suggests that women in the experiment seek to shield information about their income from their close relatives rather than their husbands.

4.1.2 Income Hiding across Villages

Next, we explore the association between the level of income hiding, aggregated within a community, and village-level outcomes. If pressure to share income does, in fact, act like a tax on investment returns, we would expect to see a correlation between the level of social pressure within a community and a range of development indicators. We explore this by constructing two aggregate measures of income hiding motivated by the theoretical framework presented in Section 3: (i) the within-village difference in the proportion of women investing no more than 80 shillings in the public and price treatments and the proportion investing no more than 80 shillings in the private treatment, and (ii) an analogous measure of the difference in the proportion of women investing exactly 80 shillings across treatments. Larger positive differences suggest greater levels of income hiding. Across all 26 communities, the average difference in the proportion investing no more than 80 shillings is 11.4 percentage points, while the average difference in the proportion investing exactly 80 shillings is 7.2 percentage points.[53]

Having constructed our measures of women's propensity to hide income, we explore the associations between these measures and four outcomes: the mean of the log value of subjects' durable household assets, the fraction of subjects with any amount of formal or skilled employment, subjects' average level of wages from paid work, and the proportion of subjects that have used fertilizer on their plots in the past year. We report regression results in Table 5. Odd-numbered columns include only a session-level measure of income hiding as an independent variable, while even-numbered columns also include controls for for distance to the nearest paved road, the mean education level of subjects from that community, the mean number of close relatives in the village, and the mean number of community groups in which subjects participate (a measure of social capital). All estimated coefficients are negative, and the majority are statistically significant.

[52] To further explore the possibility that kin may simply act as a channel through which husbands exert control over their wives, we examine investment decisions among the small sample of women receiving the large endowment who were not married and had close kin attending the session. These women are 51.4 percentage points more likely to invest 80 shillings or less in the public treatments. These differences are not statistically significant in this twelve subject sub-sample, but it nonetheless appears that women with kin present are concerned about observability even in cases where the kin cannot possibly pass information on to husbands.

[53] The former measure is negative in 7 of 26 villages, the latter in 9 of 26 villages. This pattern is consistent with the types of behavior predicted by our stochastic choice model, as we discuss further in Section 5.3. Both measures of income hiding are positively but weakly correlated with the percentage of subjects who report (in the pre-experiment survey) making a transfer to another household in the last three months, and more strongly correlated with the percentage of households who report making a transfer but who do not report receiving a transfer.

Higher levels of income hiding within the experiment are associated with lower levels of household wealth, lower rates of skilled and formal employment, lower wages, and a reduction in the likelihood of using fertilizer.

These cross-community results do not merit a causal interpretation. Nonetheless, these findings are of interest for two reasons. First, results are consistent with the view that pressure to share income may act as a drag on savings and investment, and may therefore slow development. An alternative, though not mutually exclusive, explanation is that in the poorest areas, enforced sharing norms remain an essential tool for handling risk, and thus the pressure to share is the most acute. In either case, the evidence suggests that social pressure to share income is greatest in the worst-off areas, and should be viewed as an important social factor in models of poor, rural village economies. More generally, these results go some way toward addressing the concern that behavior in lab experiments may not be related to decisions and outcomes outside the lab (cf. Levitt and List 2007); in this case, there is a clear association between what happens within the experiment and economic outcomes outside the lab.

4.2 Paying to Avoid the Public Announcement

We now examine the willingness-to-pay to avoid the public announcement using data from the price treatments. Subjects assigned to the price treatments were offered the option of paying a randomly chosen price — 10, 20, 30, 40, 50, or 60 shillings — to avoid announcing their investment income to the other subjects.[54] Of 690 subjects assigned to the price treatments, 627 subjects could afford to pay p to avoid the public announcement, and 190 of them (30.3 percent of those able to pay) chose to do so. Subjects who chose to buy out paid an average price of 29.3 shillings — equivalent to an average of 15.3 percent of their gross earnings.

Table 6 reports OLS regressions of the probability of paying to avoid the public announcement on exogenous factors: the randomly assigned price and indicators for receiving the large endowment and having the coin land with heads facing up. As the model predicts, both men and women are less inclined to buy out at higher exit prices: the coefficient on the price variable is negative and significant in all specifications. Women with more observable income are more likely to pay to avoid the public announcement: the coefficient on *Heads* is positive and significant, suggesting that subjects with successful investments are approximately 22 percentage points more likely to pay to avoid announcing their investment return.

Our model predicts that the probability of paying to avoid the public announcement will be greater than one half when τ is greater than the randomly-assigned price divided by a subject's gross observable income from the experiment (her gross payoff or, if a subject was assigned to the large endowment treatments and invested no more than 80 shillings, her gross payoff minus 100 shillings). In Figure 2, we plot the relationship between the proportion of subjects paying to

[54]Though this randomization was not stratified, it was largely successful. Online Appendix Table 2 reports the results of a balance check exercise in which subject characteristics were regressed on dummies for prices 20 through 60. The price variables are jointly significant at the five percent level in only four of the 68 F-tests.

avoid the public announcement and the price as a fraction of the gross observable payout. We plot locally-weighted (lowess) regressions (separately for women and men). For both women and men, we observe a negative relationship between the price as a proportion of the observable payout and the probability of paying to avoid the public announcement. In fact, the two regression plots are quite similar. Both cross the horizontal line indicating that the probability of paying to avoid the public announcement is 0.5 at a price to observable payout ratio of approximately 0.039, suggesting that — for both men and women — we might expect to find a τ of around 4 percent. Unfortunately, the number of data points in the range of the crossover point is relatively small, so this exercise does not provide a particularly precise estimate of the level of social pressure to share income.

5 Estimating the Extent of Social Taxation

Our reduced form results indicate that social pressure to share income impacts individual investment decisions. However, the magnitude of the reduced form impact does not identify τ because the treatment effect depends on an individual's level of risk aversion. In this section, we estimate the magnitude of the social pressure parameter, τ, while controlling for unobserved heterogeneity in risk aversion. We begin with a discussion of the likelihood function that follows from the model detailed in Section 3, moving directly to the parameter estimates that result. The model permits simulation of data under estimated parameter values, and as a check on the appropriateness of the model, we compare its predictions under estimated parameter values to the empirical patterns we actually observe. As two extensions, we then discuss both what the model predicts in other potential lab experiments, and what the estimated parameters could imply for entrepreneurship and investment outside the laboratory.

5.1 Likelihood Function

Random assignment guarantees that, in expectation, the distribution of individual risk parameters in each of the six treatments is representative of the population distribution. We can therefore estimate the parameters of the risk preference distribution using data from the private treatments, and simultaneously estimate τ while controlling for unobserved heterogeneity in risk aversion in a structural framework.

Subjects are heterogeneous in terms of their risk preferences (Choi, Fisman, Gale, and Kariv 2007). Assumptions about the functional form of that heterogeneity allow us to estimate the magnitude of τ. As discussed in Section 3, we assume that subject i's expected utility of investing b_j is given by:

$$EU_{ij} = EV_{ij} + \varepsilon_{ij} \tag{23}$$

where EV_{ij} denotes the CRRA expected utility of investing b_j and ε_{ij} is an i.i.d. type 1 extreme value distributed preference shock. We further assume that individual ρ_i parameters are normally distributed with mean μ_ρ and variance σ^2_ρ.[55] Given these assumptions, the probability that subject i chooses to invest b_j takes a mixed logit form:

$$P_{ij} = \int \left(\frac{e^{EV_{ij}/\sigma_\varepsilon}}{\sum_{k=1,\ldots,J_t} e^{EV_{ik}/\sigma_\varepsilon}} \right) f(\rho)\, d\rho. \tag{24}$$

In other words, because individual risk preferences are unobserved, the probability that subject i chooses to invest b_j is the expectation of the probability integrated over the observed distribution of levels of risk aversion. Let y_{ij} be an indicator function equal to one if subject i chooses to invest amount b_j. The log-likelihood function for treatment t can be written as:

$$LL_t = \sum_{i \in I_t} \sum_{j \in J_t} y_{ij} \ln \left[\int \left(\frac{e^{EV_{ij}/\sigma_\varepsilon}}{\sum_{k=1,\ldots,J_t} e^{EV_{ik}/\sigma_\varepsilon}} \right) f(\rho)\, d\rho \right]. \tag{25}$$

The log-likelihoods can then be summed across treatments. We simulate the log-likelihood by taking one thousand random draws from the distribution of ρ_i for each individual, and maximize the simulated log-likelihood numerically.[56]

When estimating a mixed logit model which allows for heterogeneity in risk preferences, it is necessary to scale the CRRA utility function in order to ensure comparable utility magnitudes across the range of risk aversion levels (Goeree, Holt, and Palfrey 2003, Wilcox 2008, Von Gaudecker, van Soest, and Wengström 2011). We achieve this by re-writing the CRRA utility function as

$$v(x|\rho_i) = \frac{1}{900^{1-\rho_i} - 10^{1-\rho_i}} x^{1-\rho_i}, \tag{26}$$

where $900^{1-\rho_i} - 10^{1-\rho_i}$ represents the difference between the utility of the highest possible payout in the experiment and the utility of the lowest strictly positive payout.[57] Since VNM expected utility functions represent preference orderings over lotteries and are robust to positive, affine transformations, this utility function represents the same preferences as the more traditional CRRA formulation described in Section 3.[58]

Before proceeding to our estimates of the kin tax parameter, it is worth reviewing the sources of identification in our model. Because subjects choose from a rich menu of investment options, the distributions of investment decisions in the two private treatments allow us to identify μ_ρ,

[55] Results are nearly identical if we instead assume a triangular distribution for ρ_i, as shown in the Online Appendix.

[56] Our implementation follows the procedures outlined in Train (2003). We implement this using the MATLAB command `fminunc` using the default Broyden-Fletcher-Goldfarb-Shanno (BFGS) updating procedure. Standard errors are calculated using the inverse Hessian.

[57] See Goeree, Holt, and Palfrey (2003) and Wilcox (2008) for similar approaches. Using the lowest possible payout, zero, is not feasible because $u(x) \to -\infty$ as $x \to 0$ for subjects with $\rho \geq 1$. Results are almost identical when $10^{1-\rho}$ is replaced with $1^{1-\rho}$ or $0.01^{1-\rho}$.

[58] Our scaling generates estimates of μ_ρ and σ_ρ that are nearly identical to those generated by the contextual utility model of Wilcox (2008) and the certainty equivalent procedure used in Von Gaudecker, van Soest, and Wengström (2011). We refer the interested reader to the Online Appendix, where we discuss our scaling procedure in more detail and report the results of a sensitivity analysis.

σ_ρ, and σ_ε. The kin tax parameter, τ, is identified by the differences in the distribution of investment decisions between the private large endowment treatment and the public and price large endowment treatments; the relationship between the willingness to pay to avoid the public announcement and observable income provides a second source of identification. Finally, the variance of the logit error term in the buyout decisions, γ, is identified by the excess willingness to pay to avoid the public announcement at very low gross payouts and high exit prices.

5.2 Parameter Estimates

Summing the log likelihoods across all treatments, we can estimate all the parameters of the model using the full dataset. We do so separately for men and women, paralleling our reduced form analysis (since the experimental treatment randomization was stratified on gender to permit such analysis). Parameter estimates are reported in Table 7. In Column 1, we use only the data from the private treatments to estimate the parameters of the risk preference distribution (μ_ρ and σ_ρ) and the logit noise parameter associated with the investment decision (σ_ε). In Column 2, we report estimates from a simplified likelihood function which only uses data from investment decisions (in all 6 treatments) to estimate the risk preference parameters and τ. In Column 3, we report parameter estimates based on all the decisions within the experiment.

We estimate a mean ρ_i of 0.75 for women, and a slightly higher value, 0.77, for men (Table 7, Column 1). Estimated levels of risk aversion are higher than those typically reported for undergraduate subjects (cf. Holt and Laury 2002, Goeree, Holt, and Palfrey 2003), but are in line with existing estimates of risk aversion among non-student subjects (cf. Andersen, Harrison, Lau, and Rutström 2008, Cardenas and Carpenter 2008).[59] Somewhat surprisingly, we do not find evidence that women are more risk averse than men — this contrasts with the broad conclusion reported in Croson and Gneezy (2009). One possible explanation is that the village represents a somewhat selected sample: as discussed above, many Kenyan men migrate to urban areas in search of work; and women account for a disproportionately large fraction of the village population. Our results suggest that those men remaining behind may be a selected sample in terms of their risk preferences.

Next, we turn to our estimates of the level of social pressure to share income. For women, the estimated τ ranges from 0.0432 (in Column 2) to 0.0450 (in Column 3), and is significantly different from zero at the 99 percent confidence level. Among men, the point estimate is smaller, ranging from 0.0234 to 0.0267, but is also marginally significant. For both men and women, these estimated values of τ are broadly in line with the reduced form results suggested by Figure 2.

One alternative explanation for the relatively high willingness to pay to avoid the public

[59] Andersen, Harrison, Lau, and Rutström (2008) report mean CRRA parameters of 0.74 in a representative sample of the Danish population. Cardenas and Carpenter (2008) review experiments measuring risk preferences of in developing countries; a key conclusion is that different studies (using different designs in different populations) lead to divergent estimates of risk aversion. Harrison, Humphrey, and Verschoor (2010) estimate a mean CRRA parameter of 0.536 in a sample of subjects drawn from Ethiopia, India, and Uganda. Tanaka, Camerer, and Nguyen (2010) estimate the average CRRA parameter to be between 0.59 and 0.63 among their Vietnamese subjects. Schechter (2007), on the other hand, reports a CRRA coefficient above 2.

announcement is that subjects are attention averse — in other words, they wish to avoid public scrutiny, irrespective of whether their income from the experiment is revealed. A straightforward way to model participants' attention aversion is to incorporate an additive cost parameter, κ, into the utility function, subtracting it from the expected utility of any outcome that involves making a public announcement of one's investment decision. It is apparent that the inclusion of κ would not change the probability ordering of investment levels in the public treatments (since one is always obliged to make the public announcement). However, within the price treatment, $\kappa > 0$ would influence the willingness-to-pay to avoid making the public announcement.

We explore this by incorporating the κ additional parameter into the model, subtracting it from all EV_{ij} expressions any time a public announcement is being made. Intuitively, κ is identified by the frequency with which subjects pay to avoid the public announcement when they have little income to conceal or when a large fraction of their income is already obscured by their investment strategy. Recall that τ is identified by changes in investment behavior across public and private treatments, and by variation in willingness to pay in price treatments with respect to *observable* income. In contrast, κ is identified by excess willingness to pay in the price treatments, and its variation with respect to *total* income. As shown in Column 4 of Table 7, κ is precisely estimated and very close to zero in the sample of women, and we cannot reject the null that κ is zero for both women and men.

Our ability to explore potential heterogeneity in τ is limited by the relatively small samples in different demographic categories, particularly in the price treatment, where subjects faced different randomly-assigned prices and decisions depended on both prices and (not randomly-assigned) gross payouts. However, in light of our reduced form results, we take an initial step in this direction by allowing τ to depend on whether a subject's close relatives were present at the experiment. In Table 8, we report parameter estimates when we replace the single τ with two separate parameters, $\tau_{no\ kin\ present}$ and $\tau_{kin\ present}$. Among women, $\tau_{no\ kin\ present}$ is slightly smaller than in the full sample, at 0.043, but remains statistically significant. The estimated value of $\tau_{kin\ present}$ is 0.080 — substantially higher, and also statistically significant. Among men, as in the full sample, $\tau_{no\ kin\ present}$ remans about 2.7 percent, while the estimated $\tau_{kin\ present}$ is slightly below zero and imprecisely estimated. Thus, consistent with our reduced form results, we find evidence that women with kin present at the experiment are more exposed to social pressure to share their income.

5.3 Comparing Model Predictions to Reduced-form Patterns

Next, we ask whether the parameters we estimate in Section 5.2 can explain the reduced-form patterns we observed in Section 4.1. An intuitive approach is to simply compare the actual decisions made by subjects in our experiments to those predicted by our model given the estimated parameters. Online Appendix Figures 1, 2, and 3 allow for such comparisons. Online Appendix Figure 1 presents the empirical distributions of investment amounts in each of our six treatments (separately for men and for women); Online Appendix Figure 2 is the simulated analog (based on the parameter estimates reported in in Column 3 of Table 7. There is more sampling variation in the actual investment decisions, since the simulations use 10,000 subjects per treatment. The model also fails to capture the salience of particular investment amounts in the large endowment treatments: specifically, 80 and 100 at the expense of 70, 90, and 110. Otherwise, the model seems to fit the small endowment treatments very well, and also does a good job of modeling the choices of women in the large endowment treatment.

To examine behavior in the price treatments, in Online Appendix Figure 3, we compare the predictions of the model to the actual frequency with which subjects pay to avoid the public announcement. We disaggregate the data according to the implied tax rate facing subjects: the randomly-assigned exit price they faced divided by their gross payout. As the figure demonstrates, the stochastic component plays a large role in individual decisions within our experiment, particularly when the difference between the expected utility of paying p and the expected utility of sharing income with others is small — though the estimated τ parameter is less than 0.05, the model predicts the observed rates of paying to avoid the public announcement even when the price is more than five percent of the gross payout. For women, the estimated parameters fit the data well, though the model over-predicts the likelihood of paying p at implied tax rates between 30 and 50 percent. Thus, visual inspection suggests that our estimated model parameters do a relatively good job of explaining the patterns that observe in the data.

To explore this issue somewhat more formally and quantitatively, we carry out a numerical exercise, simulating two of our key reduced form regression results (from the OLS regressions of $LTE80_i$ and $EX80_i$ on the $PublicTreatments_i$ variable reported in Columns 1 and 5 of Table 3). Again, we use the estimated parameter values shown in Column 3 of Table 7: for women, for example, for example, $\mu_\rho = 0.7488$, $\sigma_\rho = 0.1992$, $\sigma_\varepsilon = 0.0125$, $\tau = 0.045$, and $\gamma = 0.0588$. For each of 1000 simulations, we create a sample that is exactly analogous to our actual sample in terms of size (within each treatment) and gender composition. The key dimension of unobserved heterogeneity in our model is risk aversion; for each subject in each of our 1000 simulations, we draw a ρ_i parameter from the (estimated) gender-specific distribution of risk aversion parameters. We then calculate exact decision probabilities for every decision option each subject would have faced in the treatment to which she was assigned. Given these probabilities, a draw from the uniform distribution is used to determine which choice each simulated subject actually makes. Having simulated the behavior of all individuals in each of our 1000 simulated samples, we then estimate the regression specifications reported in Columns 1 and 5 of Table 3 for each simulated sample and assemble a distribution of simulated regression coefficients. We are then able to locate the percentile that each empirical (i.e. actual) coefficient from Table 3 occupies in these distributions.

In Column 1 of Table 3, we find that women are 9.6 percentage points more likely to invest no more than 80 shillings in the public and price large endowment treatments than in the private large endowment treatments. This point estimate falls at the 65^{th} percentile of the simulated distribution of coefficient estimates, suggesting that the model parameters predict our key reduced form outcome quite well. In fact, for all four of the coefficients in columns 1 and 5 of Table 3, the coefficient in the actual data falls between the fifth and 95^{th} percentiles in simulations: the percentiles are 65^{th} and 89^{th} for women, and 51^{st} and 8^{th} for men. This suggests that under the parameters of the model, the estimates we obtain in actual data are well within the realm of sampling variation.

5.4 Simulating Alternative Laboratory Experiments

The reduced-form results depend on the specific parameters of the investment scenario, while the structural parameter estimates shown in Tables 7 and 8 do not. To illustrate this, we briefly show how the reduced form results would have differed with the same population, but using slightly different experimental parameters. To simulate the experiment for women, we begin by using the estimated parameters in Column 3 of Table 7: $\mu_\rho = 0.7488$, $\sigma_\rho = 0.1992$, $\sigma_\varepsilon = 0.0125$, $\tau = 0.045$, and $\gamma = 0.0588$. Given these parameters, we draw half a million individuals from the population, and simulate their investment decisions using the probabilities specified by the model.

We begin by exploring the robustness of our reduced form results to alternative rates of return on the risky investment. In our actual experiment, the investment amount was multiplied by five when successful, and zero when unsuccessful. We simulate the experiment for a range of values of the successful investment return on the interval [2,12], and show the predictions for one of our reduced form regressions — Table 3, Panel A, Column 1. For subjects receiving the large endowment treatment, we compare behavior in the private treatment to choices in the two public information treatments, and measure the change in the probability of investing no more than the small endowment (80 shillings).

For the value 5.0, highlighted in Figure 3, the predicted change in the fraction of participants investing 80 shillings or less is 8 percentage points. This is close to the point estimate of 9.6 percentage points shown in Column 1 of Table 3, Panel A. Figure 3 shows, however, that we might have seen even larger effects if successful investments had been multiplied by four rather than five, while we would have seen smaller effects if successful investments had been multiplied by three or ten, for example. The intuition behind the figure is straightforward: the largest reduced form effect is found when much of the population invests just over 80 shillings in the private treatment. Effects diminish as investment returns go up because the investment becomes so profitable that it is no longer worth sacrificing one's expected return to hide the large endowment. On the other side, effects diminish as investment returns decline because few participants invest over 80 shillings when the risky investment has a low expected return.

Taking our analysis a step further, we simulate a range of experiments under counterfactual scenarios in which populations are more or less risk averse. Populations with different risk aversions would change the picture: while the bold curve in Figure 3 shows simulations for the estimated parameter values ($\hat{\mu_\rho} \approx 0.75$), the figure also includes two other curves for populations with $\hat{\mu_\rho}$ one standard deviation higher or lower. For the more risk averse population, the largest change in investment behavior around 80 shillings would have been seen for an investment multiplier of about 5.5; for the less risk averse population, a lower investment multiplier of about 3.0 would have induced investments close to the small endowment amount.

Similar figures can be constructed for other reduced form results in the paper, by varying other features of the laboratory experiment, or by simulating different populations. There are two key insights from this exercise. The first is that, without knowing the precise risk aversion characteristics of the population in advance, our choice of investment multiplier was not a bad one: it would have elicited responses no smaller than five percentage points even with substantially different study populations. The second insight is that while the reduced form results would have varied across a range of closely related laboratory scenarios, the underlying structural parameters remain the same. It is these parameters that allow us to model behavior both in the lab, and beyond it.

5.5 Simulating the Impacts of Social Taxation on Entrepreneurship

Our results thus far indicate that women in the experiment behave as if they expect to be pressed to share four percent of their cash income with others, and substantially more if their close kin can observe their income directly. A 4–8 percent "kin tax" may have large disincentive effects if, for example, relatives observe wages or micro-enterprise revenues, but are not able to easily separate profits from total income by accounting for labor time and indirect costs, both of which may be unobserved.[60] We explore this possibility by simulating a simple model of individual entrepreneurship adapted from the theoretical framework in Banerjee, Duflo, Glennerster, and Kinnan (2015). In their model, individuals receive income y_i in each of two periods. y_i can be seen as an individual's income from subsistence agriculture. In the first period, each person decides whether to invest in a microenterprise which yields return

$$A\left(K_i - \underline{K}\right) \tag{27}$$

where K_i is the amount that i invests in her microenterprise. To keep the model as simple as possible, we assume that individuals can neither save nor borrow. Thus, i chooses K_i to maximize

$$\frac{1}{\eta_i}\left(y_i - K_i\right)^{1-\rho_i} + \delta \frac{1}{\eta_i}\left(y_i + (1-\tau)A\left(K_i - \underline{K}\right)\right)^{1-\rho_i}. \tag{28}$$

We assume that individuals are pressed to share a proportion of their business income, but can avoid sharing their subsistence income — for example, because it may never be converted into cash.

We simulate a village economy comprising equal numbers of poor and non-poor individuals. We set $y_i = 800$ shillings for poor individuals, and $y_i = 1500$ shillings for the non-poor. These are roughly equivalent to the 30^{th} and 70^{th} percentiles of the rural consumption distribution in

[60] This would be consistent with evidence that many micro-entrepreneurs are unable to calculate their own profits, and do not correctly deduct time and indirect costs (cf. Karlan, Knight, and Udry 2012).

Uganda, as reported in Uganda Bureau of Statistics (2006).[61] We discretize the decision problem by assuming that individuals can invest any multiple of 100 shillings in a microenterprise. We do not allow individual consumption to drop to zero in any period, and we set $\underline{K}$ equal to 100 shillings, so that the fixed costs of starting an enterprise are quite low. We assume that individuals have heterogeneous risk preferences, and that individual CRRA parameters are normally distributed according to the mean and variance reported in Column 3 of Table 7.

We explore values of A ranging from 1.5 to 2.5. For each A value, we take $10,000$ draws from the distribution of CRRA parameters and calculate the fraction of women who invest a positive amount in a microenterprise and the average amount invested. We report results for three values of τ: when τ equals zero (i.e. in the absence of social pressure to share income), when τ is equal to 4.5 percent (as in our pooled data), and when τ is equal to 8 percent (our estimated value of τ when relatives are able to observe income streams directly). Results from the simulations are presented in Figure 4.

In the absence of kin pressure, only 1 percent of women become entrepreneurs when A is equal to 1.5, while 98 percent of women become entrepreneurs when A is equal to 2.5. Averaging across all values of A between 1.5 and 2.5, our simulations suggest that 59.9 percent of women would become entrepreneurs in the absence of kin pressure; but that number drops to 44.2 percent when τ is equal to 8 percent. Similarly, our simulations suggest that women would invest an average of 341 shillings in their enterprises in the absence of kin pressure, but this drops to 247 shillings if τ is equal to 8 percent. Thus, an 8 percent kin tax leads to more than a 27 percent decline in overall investment.

Thus, the investment impacts of relatively small kin taxes can be quite large. However, as the figures suggest, the simulated impact of social pressure depends on the value of A, the return on business investment. At the highest value of A we consider, 2.5, moving from a τ of zero to a τ of 8 percent reduces the fraction of women starting businesses from 0.98 to 0.97, and reduces average business investment by only 7 percent (from 598 shillings to 555 shillings). Intuitively, when entrepreneurship is profitable enough, social pressure has little effect since the return to starting a business is large even after sharing a portion of one's revenue with kin. On the other hand, at the value of A which maximizes the impact of social pressure to share, 1.94, even the relatively lower τ of 4.5 percent reduces the percent of women starting businesses from 64 to 40, and reduces average business investment by more than a third (from 349 to 232 shillings). At this level of A, a τ of 8 percent would have even more dramatic impacts, reducing the share of women starting businesses to only 26 percent, and reducing the average investment level by 56.7 percent. Thus, a relatively moderate level of social pressure to share can have large disincentive effects and, potentially, consequences for growth when profitable opportunities (jobs, new businesses) are relatively observable. However, the overall impacts will clearly depend on the range of investments available to individuals, and their relative levels of observability. In general,

[61] We were unable to find analogous statistics for rural Kenya. Beegle, De Weerdt, and Dercon (2011) report similar figures for rural Tanzania.

we expect social pressure to share income to distort investment decisions toward activities that generate income streams that are relatively less observable; our simulations highlight the fact that, when more profitable activities are also more observable, even small amounts of social pressure can have large impacts.

6 Conclusions

We report the results of an economic experiment designed to measure social pressure to share income in Kenyan villages. Our approach is stratified to ensure balance, randomized within villages, and is conducted on a large sample. The design permits both reduced-form estimates to find results in line with theoretical predictions, and structural estimates to identify parameters of interest in the presence of heterogeneity.

Women who know that the outcome of their investments will be made public distort their choices to hide a portion of their income. Results are strongest for those who have relatives present at the experiment. When we offer some participants the opportunity to pay a fee to avoid making an announcement, they do so at substantial cost: 30 percent of those able to pay to avoid the public announcement choose to do so; these subjects sacrifice 15 percent of their gross payout, on average. Structural estimates of the average "kin tax" are significantly different from zero for women, estimated at roughly 4.3 percent for those whose relatives did not attend the experiment, but at 8.0 percent for those with relatives present. We also estimate a small and only modestly significant kin tax for men. Our model of stochastic choices in the experiment fits the data well, explaining both investment and exit decisions. We see no evidence that this behavior can be explained by household bargaining with a spouse, aversion to public announcements, or aversion to risk-taking in general.

We hypothesize that the behavior observed in this experiment is a sign that village sharing norms distort investment incentives towards less visible, but potentially less profitable, investments, and may consequently slow economic growth. The negative correlations we observe between the extent of income hiding at the village level and the level of prosperity in the village, measured several different ways, are in agreement with this interpretation. However, such results should be interpreted with caution, since the direction of causality is unclear. Moreover, the efficiency impacts of social pressure to share income will clearly depend on the range of income-hiding technologies available, and correlation between observability and profitability. There is still much to be learned about the range of observable and unobservable investment opportunities available to poor households, the set of technologies for hiding or protecting income from pressure to share, and the mechanics of that pressure. Nonetheless, simulating a simple model of entrepreneurship using our estimated structural parameters suggests that the levels of pressure to share documented in this paper may lead to large reductions in the probability of starting a business and the overall level of entrepreneurial investment. Whether the levels of social pressure we observe are large or small will depend on how social pressure is exerted outside of the lab: for example, whether individuals are pressed to share all cash income (such as business revenues) or only their profits after compensating themselves for indirect costs or labor time.[62]

Studies of mutual insurance typically assume that transfer arrangements are on the efficient frontier, though the analogous assumption has been called into question in intrahousehold bargaining contexts. Our work suggests that relationships with close kin outside the household may be similar to within-household interactions, and that social sanctions which encourage cooperation and sharing may also have important disincentive effects.

References

Albarran, P., and O. P. Attanasio (2003): "Limited Commitment and Crowding Out of Private Transfers: Evidence from a Randomised Experiment," *The Economic Journal*, 113(486), C77–C85.

Alger, I., and J. Weibull (2010): "Kinship, Incentives and Evolution," *American Economic Review*, 100(4), 1727–1760.

Alm, J., B. Jackson, and M. McKee (1992): "Institutional Uncertainty and Taxpayer Compliance," *American Economic Review*, 82(4), 1018–1026.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76(3), 583–618.

Anderson, S., and J.-M. Baland (2002): "The Economics of Roscas and Intrahousehold Resource Allocation," *Quarterly Journal of Economics*, 117(3), 963–995.

Ashraf, N. (2009): "Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines," *American Economic Review*, 99(4), 1245–1277.

Attanasio, O., A. Barr, J. C. Cardenas, G. Genicot, and C. Meghir (2012): "Risk Pooling, Risk Preferences, and Social Networks," *American Economic Journal: Applied Economics*, 4(2), 134–167.

[62] One such possibility is that individuals are not pressured to share earned income. However, Jakiela (forthcoming) reports experimental evidence that respect for "earned property rights" is less prevalent in rural villages than in university lab settings.

Baland, J.-M., C. Guirkinger, and C. Mali (2011): “Pretending to Be Poor: Borrowing to Escape Forced Solidarity in Cameroon,” *Economic Development and Cultural Change*, 60(1), 1–16.

Banerjee, A. V., E. Duflo, R. Glennerster, and C. Kinnan (2015): “The Miracle of Microfinance? Evidence from a Randomized Evaluation,” *American Economic Journal: Applied Economics*, 7(1), 22–53.

Barr, A., M. Dekker, and M. Fafchamps (2012): “Who Shares Risk with Whom under Different Enforcement Mechanisms?,” *Economic Development and Cultural Change*, 60(4), 677–706.

Barr, A., and G. Genicot (2008): “Risk Pooling, Commitment and Information: An Experimental Test,” *Journal of the European Economic Association*, 6(6), 1151–1185.

Barr, A., and M. Stein (2008): “Status and egalitarianism in traditional communities: an analysis of funeral attendance in six Zimbabwean villages,” CSAE Working Paper No. 26.

Beegle, K., J. De Weerdt, and S. Dercon (2011): “Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey,” *Review of Economics and Statistics*, 93(3), 1010–1033.

Binswanger, H. P. (1980): “Attitudes toward Risk: Experimental Measurement in Rural India,” *American Journal of Agricultural Economics*, 62(3), 395–407.

Brabin, L. (1984): “Polygyny: an indicator of nutritional stress in African agricultural societies?,” *Africa*, 54(1), 31–45.

Cardenas, J. C., and J. Carpenter (2008): “Behavioural Development Economics: Lessons from Field Labs in the Developing World,” *The Journal of Development Studies*, 44(3), 311–338.

Charness, G., and G. Genicot (2009): “Informal Risk Sharing in an Infinite-Horizon Experiment,” *The Economic Journal*, 119(537), 796–825.

Choi, S., R. Fisman, D. Gale, and S. Kariv (2007): “Consistency and Heterogeneity of Individual Behavior under Uncertainty,” *American Economic Review*, 97(5), 1921–1938.

Choi, S., S. Kariv, W. Müller, and D. Silverman (2014): “Who Is (More) Rational?,” *American Economic Review*, 104(6), 1518–1550.

Coate, S., and M. Ravallion (1993): “Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements,” *Journal of Development Economics*, 40(1), 1–24.

Cole, H. L., and N. R. Kocherlakota (2001): "Efficient Allocations with Hidden Income and Hidden Storage," *The Review of Economic Studies*, 68(3), 523–542.

Collier, P., and A. Garg (1999): "On kin groups and wages in the Ghanaian labour market," *Oxford Bulletin of Economics and Statistics*, 61(2), 133–151.

Comola, M., and M. Fafchamps (2011): "Are Gifts and Loans between Households Voluntary?," working paper.

Croson, R., and U. Gneezy (2009): "Gender Differences in Preferences," *Journal of Economic Literature*, 47(2), 1365–1384.

De Mel, S., D. McKenzie, and C. Woodruff (2008): "Returns to Capital in Microenterprises: Evidence from a Field Experiment," *Quarterly Journal of Economics*, 123(4), 1329–1372.

——— (2009): "Are Women More Credit Constrained? Experimental Evidence on Gender and Microenterprise Returns," *American Economic Journal: Applied Economics*, 1(3), 1–32.

Di Falco, S., and E. Bulte (2011): "A Dark Side of Social Capital? Kinship, Consumption, and Savings," *Journal of Development Studies*, 47(8), 1128–1151.

Dupas, P., and J. Robinson (2013): "Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya," *American Economic Journal: Applied Economics*, 5(1), 163–192.

Fafchamps, M., and S. Lund (2003): "Risk-sharing networks in rural Philippines," *Journal of Development Economics*, 71(2), 261–287.

Fafchamps, M., D. McKenzie, S. Quinn, and C. Woodruff (2014): "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana," *Journal of Development Economics*, 106, 211–226.

Foster, A. D., and M. R. Rosenzweig (2001): "Imperfect Commitment, Altruism and the Family: Evidence from Transfer Behavior in Low–Income Rural Areas," *Review of Economics and Statistics*, 83(3), 389–407.

Geertz, C. (1963): *Peddlers and Princes: Social Change and Economic Modernization in Two Indonesian Towns*. The University of Chicago Press, Chicago, IL.

Gneezy, U., and J. Potters (1997): "An experiment on risk taking and evaluation periods," *Quarterly Journal of Economics*, 112(2), 631–645.

Goeree, J. K., C. A. Holt, and T. R. Palfrey (2003): "Risk averse behavior in a generalized matching pennies game," *Games and Economic Behavior*, 45, 97–113.

Goldberg, J. (2013): "The Lesser of Two Evils: The Roles of Social Pressure and Impatience in Consumption Decisions," working paper.

Gugerty, M. K. (2007): "You Can't Save Alone: Commitment in Rotating Savings and Credit Associations in Kenya," *Economic Development and Cultural Change*, 55(2), 251–282.

Harrison, G. W., S. J. Humphrey, and A. Verschoor (2010): "Choice Under Uncertainty: Evidence from Ethiopia, India and Uganda," *Economic Journal*, 120(543), 80–104.

Hey, J. D., and C. Orme (1994): "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62(6), 1291–1326.

Hoff, K., and A. Sen (2006): "The Kin System as a Poverty Trap," in *Poverty Traps*, ed. by S. Bowles, S. Durlauf, and K. Hoff, Chapter 4, pp. 95–115. Princeton University Press, Princeton.

Holt, C. A., and S. K. Laury (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), 1644–1655.

Jakiela, P. (forthcoming): "How Fair Shares Compare: Experimental Evidence from Two Cultures," *Journal of Economic Behavior and Organization*.

Jakiela, P., and O. Ozier (2012): "Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies," Policy Research Working Paper 6085, The World Bank.

Karlan, D., R. Knight, and C. Udry (2012): "Hoping to Win, Expected to Lose: Theory and Lessons on Micro Enterprise Development," working paper.

Kinnan, C. (2014): "Distinguishing barriers to insurance in Thai villages," working paper.

Lacker, J. M., and J. A. Weinberg (1989): "Optimal Contracts under Costly State Falsification," *Journal of Political Economy*, 97(6), 1345–1363.

Levitt, S. D., and J. A. List (2007): "What Do Laboratory Experiments Measuring Social Preferences Tell Us About the Real World?," *Journal of Economic Perspectives*, 21(2), 153–174.

Ligon, E. (1998): "Risk Sharing and Information in Village Economies," *The Review of Economic Studies*, 65(4), 847–864.

Ligon, E., and L. Schechter (2012): "Motives for Sharing in Social Networks," *Journal of Development Economics*, 99(1), 13–26.

Ligon, E., J. P. Thomas, and T. Worrall (2002): "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *The Review of Economic Studies*, 69(1), 209–244.

Loomes, G. (1991): “Evidence of a New Violation of the Independence Axiom,” *Journal of Risk and Uncertainty*, 4(1), 91–108.

——— (2005): “Modelling the Stochastic Component of Behaviour in Experiments: Some Issues for the Interpretation of Data,” *Experimental Economics*, 8(4), 301–323.

Luke, N., and K. Munshi (2006): “Evidence of a New Violation of the Independence Axiom,” *Review of Economics and Statistics*, 88(2), 264–282.

Ngau, P. M. (1987): “Tensions in Empowerment: The Experience of the “Harambee” (Self–Help) Movement in Kenya,” *Economic Development and Cultural Change*, 35(3), 523–538.

Platteau, J.-P. (2000): *Institutions, Social Norms and Economic Development*. Harwood Academic Publishers, Amsterdam.

Robinson, J. (2012): “Limited Insurance Within the Household: Evidence from a Field Experiment in Western Kenya,” *American Economic Journal: Applied Economics*, 4(4), 140–164.

Schechter, L. (2007): “Risk Aversion and Expected-Utility Theory: A Calibration Exercise,” *Journal of Risk and Uncertainty*, 35(1), 67–76.

Scott, J. C. (1976): *The Moral Economy of the Peasant*. Yale University Press, New Haven.

Stack, C. (1974): *All Our Kin*. Basic Books, New York, NY.

Tanaka, T., C. F. Camerer, and Q. Nguyen (2010): “Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam,” *American Economic Review*, 100(1), 557–571.

Townsend, R. M. (1994): “Risk and Insurance in Village India,” *Econometrica*, 62(3), 539–591.

Train, K. (2003): *Discrete Choice Methods with Simulation*. Cambridge University Press, New York, NY.

Uganda Bureau of Statistics (2006): “Uganda National Household Surevy 2005/2006: Report on the Socio-Economic Module,” available online at `http://www.ubos.org/onlinefiles/uploads/ubos/pdf%20documents/UNHSReport20052006.pdf`.

Von Gaudecker, H.-M., A. van Soest, and E. Wengström (2011): “Heterogeneity in Risky Choice Behavior in a Broad Population,” *American Economic Review*, 101(2), 664–694.

Wilcox, N. T. (2008): “Stochastic Models for Binary Discrete Choice Under Risk: A Critical Primer and Econometric Comparison,” *Research in Experimental Economics*, 12, 197–292.

Table 1: Summary Statistics on Experimental Subjects

							— Balance —	
Variable	Mean	S.D.	Median	Min	Max	N	Women	Men
Female	0.61	0.49	1	0	1	2145	·	·
Years of schooling	6.74	3.36	7	0	16	2145	0.20	0.98
Age	36.82	14.28	34	18	88	2127	0.33	0.12
Currently married	0.77	0.42	1	0	1	2145	0.74	0.17
Spouse attended the experiment	0.08	0.27	0	0	1	2145	0.00^{***}	0.28
Ever married	0.88	0.32	1	0	1	2145	0.82	0.18
HH size	6.18	2.82	6	1	26	2145	0.18	0.40
Close relatives in village (outside of HH)	2.36	2.57	2	0	19	2145	0.19	0.05^{*}
Any close relatives attended the experiment	0.19	0.39	0	0	1	2145	0.64	0.73
Distant relatives in village	10.41	16.12	5	0	199	2145	0.05^{**}	0.01^{***}
No. chicken owned by HH	6.42	7.19	4	0	40	2145	0.66	0.78
No. cattle owned by HH	1.20	2.08	0	0	36	2144	0.08^{*}	0.41
No. bicycles owned by HH	0.83	0.76	1	0	6	2145	0.19	0.85
No. phones owned by HH	0.73	0.82	1	0	6	2145	0.82	0.92
No. televisions owned by HH	0.14	0.39	0	0	3	2145	0.03^{**}	0.92
Value of durable HH assets (in US dollars)	469.31	655.19	357.05	13.18	22695.65	2145	0.03^{**}	0.90
HH farms	0.99	0.12	1	0	1	2145	0.76	0.13
HH uses fertilizer on crops	0.46	0.50	0	0	1	2114	0.21	0.98
Has regular employment	0.08	0.28	0	0	1	2145	0.25	0.74
Monthly wages if employed (in US dollars)	39.28	61.59	19.76	1.32	434.78	178	0.19	0.02^{**}
Any HH member employed	0.23	0.42	0	0	1	2145	0.72	0.73
Self-employed	0.35	0.48	0	0	1	2145	0.27	0.81
Has bank account	0.17	0.37	0	0	1	2142	0.07^{*}	0.21
Member of ROSCA	0.53	0.50	1	0	1	2142	0.25	0.21
HH gave transfer in last 3 months	0.90	0.31	1	0	1	2145	0.45	0.91
Transfers to HHs in village (in US dollars)	6.79	21.96	1.98	0.00	480.90	2145	0.64	0.78
HH received transfer in last 3 months	0.41	0.49	0	0	1	2145	0.13	0.12
Transfers from HHs in village (in US dollars)	2.58	17.62	0.00	0.00	527.80	2145	0.39	0.57
Community groups	2.76	1.87	3	0	10	2145	0.30	0.11
Belongs to Luhya ethnic group	0.80	0.40	1	0	1	2145	0.33	0.99
Local minority ethnic group	0.20	0.40	0	0	1	2145	0.30	0.99
Christian	0.98	0.14	1	0	1	2145	0.28	0.57

Balance columns report p-values from F tests of the joint significance of treatment dummies in OLS regressions in which the variables listed in the first column are used as the dependent variable. Regressions are estimated separately by gender. *** indicates significance at the 99 percent level; ** indicates significance at the 95 percent level; and * indicates significance at the 90 percent level. Additional balance checks, including tests which pool male and female subjects and tests which restrict attention to the large endowment treatments, are included in the Online Appendix.

Table 2: Summary Statistics on Outcomes in Experiment, by Treatment

Panel A: All Subjects

Information Condition: *Budget Size:*	Private Small	Public Small	Price Small	Private Large	Public Large	Price Large	All All
Amount invested in business cup	41.44	42.59	42.00	93.35	91.98	90.26	66.68
	(0.82)	(0.86)	(0.81)	(1.90)	(1.88)	(1.84)	(0.80)
Fraction of budget invested in business cup	0.52	0.53	0.53	0.52	0.51	0.50	0.52
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.00)
Total payout (US dollar equivalent)	2.03	1.83	1.86	4.68	4.24	4.47	3.17
	(0.08)	(0.08)	(0.08)	(0.18)	(0.18)	(0.18)	(0.06)
Invested exactly m_{small} (80 Kenyan shillings)?	0.03	0.04	0.02	0.20	0.23	0.27	0.13
	(0.01)	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)
Invested m_{small} (80 Kenyan shillings) or less?	1.00	1.00	1.00	0.42	0.46	0.48	0.73
	(0.00)	(0.00)	(0.00)	(0.03)	(0.03)	(0.03)	(0.01)
N	369	370	345	358	358	345	2145

Standard errors in parentheses. One US dollar was equivalent to 75.9 Kenyan shillings at the time of the experiment.

Panel B: Women Only

Information Condition: *Budget Size:*	Private Small	Public Small	Price Small	Private Large	Public Large	Price Large	All All
Amount invested in business cup	40.49	41.31	41.82	96.30	91.28	90.82	66.82
	(1.00)	(1.15)	(1.10)	(2.30)	(2.33)	(2.29)	(1.02)
Fraction of budget invested in business cup	0.51	0.52	0.52	0.54	0.51	0.50	0.52
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Total payout (US dollar equivalent)	2.06	1.90	1.82	4.54	4.14	4.53	3.16
	(0.10)	(0.10)	(0.10)	(0.23)	(0.22)	(0.23)	(0.08)
Invested exactly m_{small} (80 Kenyan shillings)?	0.01	0.02	0.02	0.18	0.22	0.27	0.12
	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)	(0.01)
Invested m_{small} (80 Kenyan shillings) or less?	1.00	1.00	1.00	0.38	0.47	0.48	0.72
	(0.00)	(0.00)	(0.00)	(0.03)	(0.03)	(0.03)	(0.01)
N	224	221	209	219	218	207	1298

Standard errors in parentheses. One US dollar was equivalent to 75.9 Kenyan shillings at the time of the experiment.

Panel C: Men Only

Information Condition: *Budget Size:*	Private Small	Public Small	Price Small	Private Large	Public Large	Price Large	All All
Amount invested in business cup	42.90	44.50	42.28	88.71	93.07	89.42	66.47
	(1.41)	(1.29)	(1.18)	(3.26)	(3.17)	(3.08)	(1.27)
Fraction of budget invested in business cup	0.54	0.56	0.53	0.49	0.52	0.50	0.52
	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)	(0.01)
Total payout (US dollar equivalent)	1.98	1.74	1.91	4.89	4.39	4.38	3.19
	(0.13)	(0.13)	(0.12)	(0.28)	(0.29)	(0.28)	(0.10)
Invested exactly m_{small} (80 Kenyan shillings)?	0.06	0.06	0.01	0.23	0.24	0.26	0.14
	(0.02)	(0.02)	(0.01)	(0.04)	(0.04)	(0.04)	(0.01)
Invested m_{small} (80 Kenyan shillings) or less?	1.00	1.00	1.00	0.50	0.45	0.49	0.74
	(0.00)	(0.00)	(0.00)	(0.04)	(0.04)	(0.04)	(0.02)
N	145	149	136	139	140	138	847

Standard errors in parentheses. One US dollar was equivalent to 75.9 Kenyan shillings at the time of the experiment.

Table 3: Regressions of Investment Outcomes for Subjects in Large Endowment Treatments

Sample:	— Women Only —				— Men Only —			
Specification:	OLS	OLS	Logit	Logit	OLS	OLS	Logit	Logit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Dependent Variable = Indicator for Investing 80 Shillings or Less								
Public or price treatment	0.096^{**}	0.109^{***}	0.097^{**}	0.109^{***}	-0.025	-0.018	-0.025	-0.011
	(0.041)	(0.042)	(0.041)	(0.042)	(0.052)	(0.052)	(0.052)	(0.052)
Panel B: Dependent Variable = Indicator for Investing Exactly 80 Shillings								
Public or price treatment	0.062^{*}	0.069^{*}	0.064^{*}	0.071^{**}	0.018	0.019	0.018	0.026
	(0.033)	(0.035)	(0.036)	(0.036)	(0.044)	(0.046)	(0.045)	(0.046)
Village FEs	No	Yes	No	No	No	Yes	No	No
Additional Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	644	644	644	644	417	417	417	417

Robust standard errors in parentheses. *** indicates significance at the 99 percent level; ** indicates significance at the 95 percent level; and * indicates significance at the 90 percent level. Logit marginal effects reported in Columns 3, 4, 7, and 8 (probit results are nearly identical). Sample restricted to subjects receiving larger endowment (women in Columns 1 through 4, men in Columns 5 through 8). A constant is included in all specifications. Columns 2, 4, 6, and 8 include controls for all variables that are not balanced across genders (see Online Appendix Table 7): age, education level, the log value of household assets, the number of close relatives residing in the village, the number of distant relatives residing in the village, involvement in community groups, and indicators for having a bank savings account, participating in a ROSCA, having given a gift or loan to another household in the last 3 months, and belonging to a local ethnic minority; even-numbered columns also include controls for marital status and household size. Assignment to treatment was random within villages; similar results are obtained when standard errors are clustered at the village level. We report pooled specifications, some of which include an interaction between $PublicTreatments_i$ and the indicator for being female, in Online Appendix Table 4. In pooled specifications which do not include gender interactions, random assignment to the public or price treatments is positively associated with the probability of investing both no more than 80 shillings and exactly 80 shillings, but the effect is only marginally statistically significant. When $LTE80_i$ (the indicator for investing no more than 80 shillings) is used as the dependent variable, the interaction between the female dummy and the $PublicTreatments_i$ indicator is marginally significant in both specifications we consider (p-values without and with controls 0.065 and 0.066, respectively). When $EX80_i$ (the indicator for investing exactly 80 shillings) is used as the dependent variable, the interaction term is not statistically significant (p-values without and with controls 0.427 and 0.592, respectively); it is only the sum of $PublicTreatments_i + Female \times PublicTreatments_i$ that is statistically significant (p-values without and with controls 0.064 and 0.058, respectively).

Table 4: Impact Heterogeneity: Women's Investment Decisions and the Presence of Close Kin

Dependent Variable:	Invested 80 Shillings or Less			Invested Exactly 80 Shillings		
Specification:	OLS	OLS	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Close kin attended game	-0.245***	-0.299***	.	-0.093	-0.153*	.
	(0.09)	(0.11)		(0.073)	(0.088)	
Close kin at game × public	0.418***	0.42***	.	0.152	0.155*	.
	(0.109)	(0.109)		(0.093)	(0.093)	
No close kin at game × public	0.069	0.043	.	0.058	0.039	.
	(0.045)	(0.075)		(0.038)	(0.063)	
Close kin in village, but not at game	.	-0.066	.	.	-0.07	.
		(0.087)			(0.07)	
Close kin in village (not at game) × public	.	0.041	.	.	0.029	.
		(0.095)			(0.08)	
Spouse at game	.	.	-0.055	.	.	0.013
			(0.121)			(0.098)
Spouse at game × public	.	.	0.202	.	.	0.016
			(0.144)			(0.115)
No spouse at game × public	.	.	0.1**	.	.	0.073**
			(0.044)			(0.037)
Village FEs	Yes	Yes	Yes	Yes	Yes	Yes
Additional Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	642	642	642	642	642	642
R^2	0.117	0.118	0.107	0.089	0.091	0.088

Robust standard errors in parentheses. *** indicates significance at the 99 percent level; ** indicates significance at the 95 percent level; and * indicates significance at the 90 percent level. OLS specifications reported; logit and probit results are nearly identical. Sample restricted to women receiving the larger endowment. A constant is included in all specifications. All specifications include controls for all variables that are not balanced across genders (see Online Appendix Table 7): age, education level, the log value of household assets, the number of close relatives residing in the village, the number of distant relatives residing in the village, involvement in community groups, and indicators for having a bank savings account, participating in a ROSCA, having given a gift or loan to another household in the last 3 months, and belonging to a local ethnic minority; all specifications also include controls for marital status and household size.

Table 5: OLS Regressions of Village-Level Outcomes on Income Hiding by Women in Experiment

Dependent Variable:	Ln HH Assets		Has Regular Job		Wages from Work		Fertilizer Use	
Specification:	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Income hiding (investing no more than 80 shillings)	-0.009	-0.043	-0.048**	-0.05**	-2.276	-2.758*	-0.179	-0.027
	(0.114)	(0.116)	(0.022)	(0.023)	(1.858)	(1.669)	(0.229)	(0.213)
Observations	26	26	26	26	26	26	26	26
R^2	0.0003	0.222	0.167	0.339	0.059	0.423	0.025	0.36
Income hiding (investing exactly 80)	-0.254*	-0.312**	-0.09***	-0.096***	-5.197**	-6.024***	-0.7***	-0.504**
	(0.134)	(0.129)	(0.024)	(0.024)	(2.162)	(1.799)	(0.254)	(0.244)
Observations	26	26	26	26	26	26	26	26
R^2	0.13	0.394	0.368	0.546	0.194	0.58	0.24	0.472
Additional Controls	No	Yes	No	Yes	No	Yes	No	Yes
Observations	26	26	26	26	26	26	26	26

Standard errors in parentheses. *** indicates significance at the 99 percent level; ** indicates significance at the 95 percent level; and * indicates significance at the 90 percent level. Sample includes one observation per village. A constant is included in all specifications. Ln HH Assets is the average of the log value of durable assets owned by households. Has Regular Job is the fraction of participants with formal, skilled, and/or professional employment. Wages from Work is the average of wages received from paid work over the last month in US dollars; wages are set to zero for subjects with no paid employment. Fertilizer Use denotes the fraction of households engaged in agricultural that used fertilizer over the previous twelve month period. Even-numbered specifications include controls for the distance to the nearest paved road and the mean education level, mean number of close relatives, and mean number of community groups across all experimental subjects from a given village.

Table 6: Paying to Avoid the Public Announcment

Dependent Variable:	Paid to Avoid Public Announcement			
Sample:	— Women Only —		— Men Only —	
Specification:	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)
Price of exit	-0.006***	-0.006***	-0.005***	-0.009***
	(0.001)	(0.002)	(0.002)	(0.002)
Large budget	0.165***	0.158	0.12**	-0.242
	(0.044)	(0.115)	(0.055)	(0.149)
Price × large budget	.	-0.00008	.	0.007*
		(0.003)		(0.003)
Coin flip lands heads	.	0.224***	.	0.065
		(0.058)		(0.069)
Heads × large budget	.	-0.015	.	0.267**
		(0.083)		(0.105)
Village FEs	Yes	Yes	Yes	Yes
Additional Controls	Yes	Yes	Yes	Yes
Observations	415	415	273	272
R^2	0.218	0.268	0.247	0.323

Robust standard errors in parentheses. *** indicates significance at the 99 percent level; ** indicates significance at the 95 percent level; and * indicates significance at the 90 percent level. OLS specifications reported; logit and probit results are similar. Sample restricted to subjects assigned to the price treatments. Results are unchanged if the sample is restricted to subjects who can afford to pay to avoid the public announcement. A constant is included in all specifications. All specifications include controls for all variables that are not balanced across genders (see Online Appendix Table 7): age, education level, the log value of household assets, the number of close relatives residing in the village, the number of distant relatives residing in the village, involvement in community groups, and indicators for having a bank savings account, participating in a ROSCA, having given a gift or loan to another household in the last 3 months, and belonging to a local ethnic minority; all specifications also include controls for marital status and household size. We report specifications that pool men and women assigned to the price treatments in Online Appendix Table 5.

Table 7: Parameter Estimates

	(1)	(2)	(3)	(4)
Panel A: Women in All Treatments				
μ_ρ	0.7562	0.7498	0.7488	0.7504
	(0.0163)	(0.0108)	(0.0107)	(0.0114)
σ_ρ	0.1994	0.2000	0.1992	0.2006
	(0.0170)	(0.0116)	(0.0115)	(0.0118)
σ_ϵ	0.0107	0.0125	0.0125	0.0124
	(0.0017)	(0.0011)	(0.0011)	(0.0011)
τ		0.0432	0.0450	0.0419
		(0.0124)	(0.0113)	(0.0097)
γ			0.0588	0.0577
			(0.0088)	(0.0114)
κ				0.0022
				(0.0118)
Obs.	1298	1298	1298	1298
Panel B: Men in All Treatments				
μ_ρ	0.7747	0.7555	0.7557	0.7557
	(0.0233)	(0.0131)	(0.0132)	(0.0134)
σ_ρ	0.2657	0.2385	0.2391	0.2395
	(0.0225)	(0.0125)	(0.0125)	(0.0125)
σ_ϵ	0.0107	0.0101	0.0102	0.0102
	(0.0022)	(0.0012)	(0.0012)	(0.0012)
τ		0.0267	0.0234	0.0242
		(0.0139)	(0.0134)	(0.0146)
γ			0.0623	0.0917
			(0.0122)	(0.0334)
κ				-0.0325
				(0.0292)
Obs.	847	847	847	847

Model parameters estimated via mixed logit maximum likelihood as described in Section 5. Standard errors (calculated using the inverse Hessian) in parentheses. In Column 1, we use only the data from the private treatments to estimate the parameters of the risk preference distribution (μ_ρ and σ_ρ) and the logit noise parameter associated with the investment decision (σ_ε). In Column 2, we use data from investment decisions in all 6 treatments to estimate the risk preference parameters and τ. In Column 3, we use data from both investment decisions and decisions about whether to pay to avoid the public announcement. In Column 4, we estimate a model which includes an attention-aversion parameter, κ.

Table 8: Parameter Estimates Allowing for Heterogeneity in τ

Sample:	WOMEN (1)	MEN (2)
μ_ρ	0.750^{***}	0.760^{***}
	(0.011)	(0.013)
σ_ρ	0.199^{***}	0.241^{***}
	(0.011)	(0.012)
σ_ϵ	0.013^{***}	0.010^{***}
	(0.001)	(0.001)
$\tau_{no\ kin\ present}$	0.043^{***}	0.027^{*}
	(0.012)	(0.015)
$\tau_{kin\ present}$	0.080^{**}	-0.011
	(0.032)	(0.022)
γ	0.058^{***}	0.062^{***}
	(0.009)	(0.012)
Obs.	1298	847

Model parameters estimated via mixed logit maximum likelihood as described in Section 5. Standard errors (calculated using the inverse Hessian) in parentheses.

Figure 1: Structure of Experiment

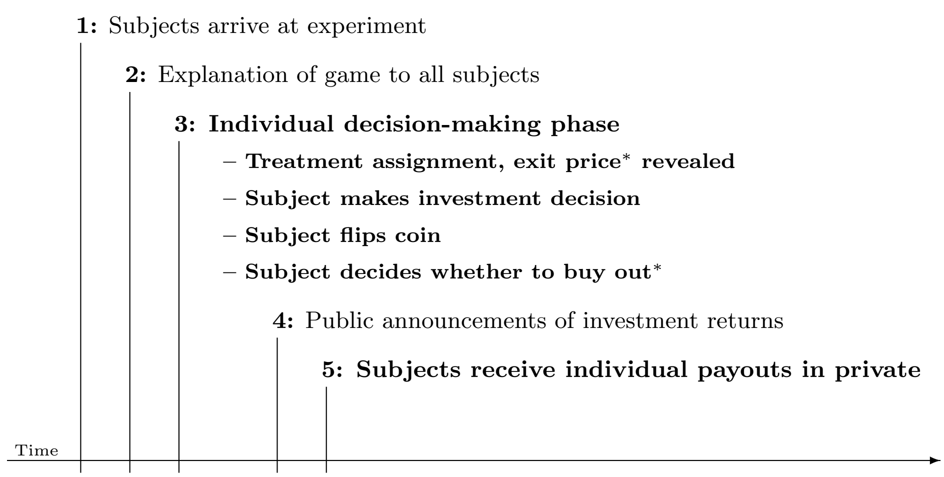

Activities in plain text took place in primary school classrooms, with all subjects seated together. Activities in bold text took place in one-on-one interactions between individual subjects and enumerators; during these interactions, subjects and members of the research team were seated at desks in private locations in the schoolyard.

Figure 2: Paying to Avoid the Public Announcement

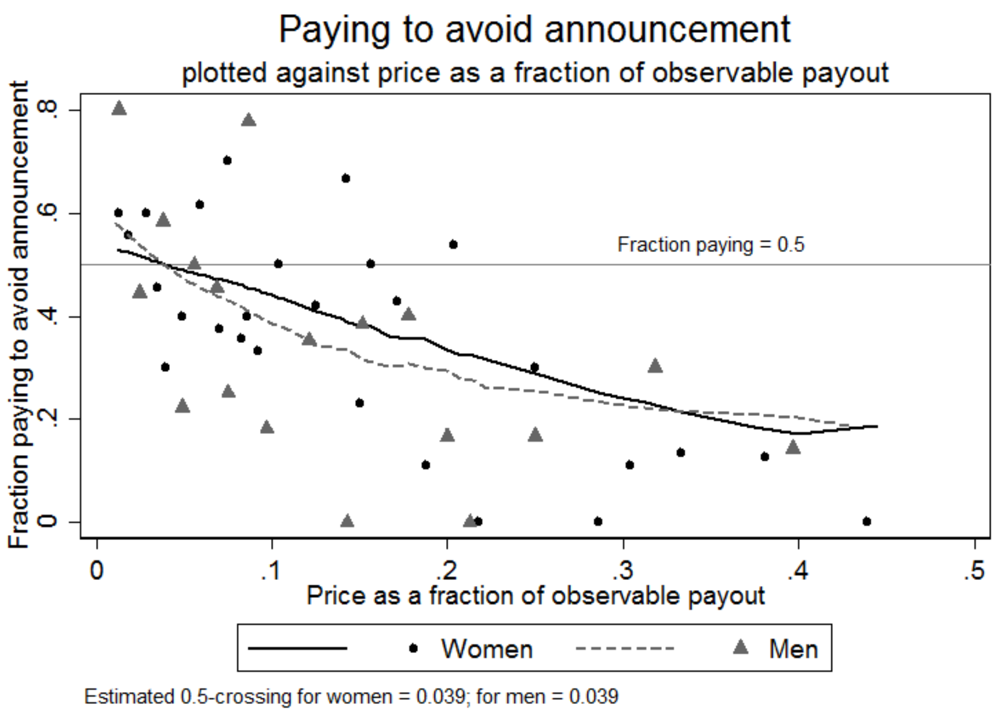

Figure 3: Simulated Alternative Laboratory Experiments

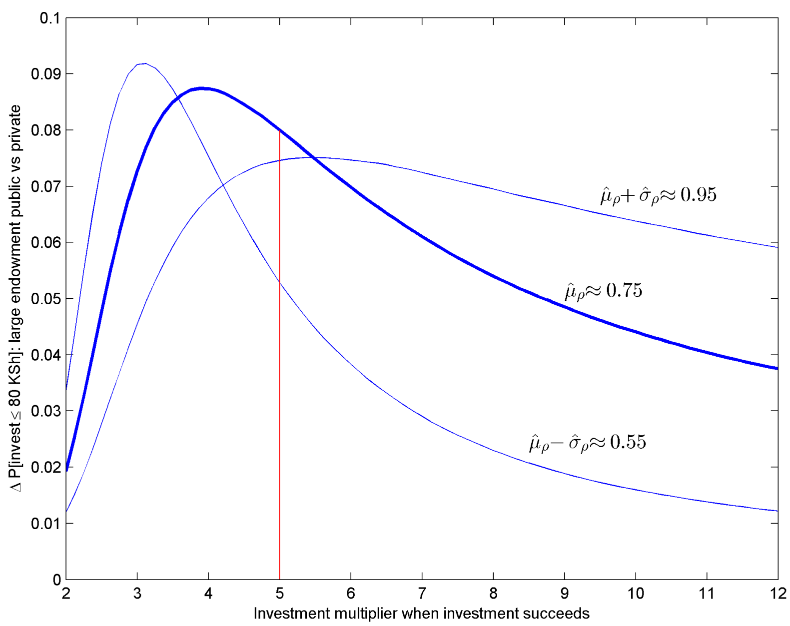

Results from 10 million simulations of players' behavior in alternative laboratory experiments in which a successful investment sees its value multiplied by a value between 2.0 and 12.0. The value 5.0, used in the actual experiment we perform, is highlighted. Parameters used in these simulations are taken from Column 3 of Table 7 for women.

Figure 4: Simulated Entrepreneurship Decisions