

Kinship Taxation as an Impediment to Growth: Experimental Evidence from Kenyan Microenterprises

Munir Squires*

November 8, 2017

Abstract

Many developing country entrepreneurs face pressure from relatives to share income. This pressure, a kinship tax, can distort productive decisions including capital allocation. I combine evidence from a lab experiment—which allows me to estimate an individual-level sufficient statistic for the distortion—with data I collected on a sample of Kenyan entrepreneurs, to quantify the importance of the tax. My data reveal high distortions for a third of entrepreneurs in my sample. A simple structural model of input allocation fitted to my data implies that removing distortions from kinship taxation would increase total factor productivity by a quarter, and increase the share of inputs used in the largest firms.

*munir.squires@ubc.ca. VSE, University of British Columbia. I thank my PhD advisers Greg Fischer, Maitreesh Ghatak and Gharad Bryan, for their guidance and encouragement. This paper has also benefited from the thoughts of Tim Besley, as well as Siwan Anderson, Orazio Attanasio, Oriana Bandiera, Robin Burgess, Alex Clymo, Sergio De Ferra, Laura Derksen, Patrick Francois, Ethan Ilzetzki, Pamela Jakiela, Anders Jensen, Henrik Kleven, Margaret Leighton, Sam Marden, Luis Martinez, Jamie McCasland, Gerard Padró-i-Miquel, and Erik Snowberg. Financial support from STICERD and PEDL is gratefully acknowledged. The protocol for this study was approved by the IRB of the London School of Economics (REC328).

1 Introduction

Communities of people living near subsistence often engage in a remarkable amount of resource sharing. Early development economists such as Arthur Lewis worried that this could act as an informal tax, diluting productive incentives.¹ Anthropologists and sociologists have similarly identified this as a key characteristic of poverty, and a barrier to accumulation and development.² Since the 1980s, however, economists have focused on the insurance benefits of such communal arrangements, largely abstracting from potential implications for incentives.³ Indeed, such distortions may be minimal: Transfers are made between households with familial or social ties, and altruism towards recipients in need may mean the value of those transfers is internalized. The key question to understand distortions then is not the fraction of marginal income individuals transfers, but extent to which these transfers are tax payments or gifts.

I show that these transfers do distort productive activity, and that their aggregate cost in my setting is greater than that of credit constraints. A 40 year old woman from a slum near Nairobi provides an illustrative anecdote:

“I sell second-hand clothes without anyone knowing, far from home. I hide from my friends because I believe not all friends will be happy with my success, and from family to create a picture that I have no money, for them to work hard for their own money. My previous business, a street-side restaurant, failed due to my in-laws using me for money, yet I wanted to expand it.”

Measuring the cost of these distortions involves two important challenges this paper addresses. The first is difficulty measuring marginal tax rates. I show that the relevant parameter for understanding productive distortions cannot be drawn from data on transfers alone, which cannot distinguish gifts from tax payments. Instead the correct measure, which I refer to as the marginal kinship tax rate, relies on the insight that an optimizing agent will evade taxes so long as the cost of doing so is lower than the marginal tax rate. I then elicit this measure for 1805 rural Kenyans using a lab experiment inspired by Jakiela and Ozier (2016). The second challenge is reverse causation: poverty may lead some groups to rely

¹See Lewis (1955), as well as Wolf (1955), Yamey and Bauer (1957), Moore (1963), and Kindleberger (1965).

²See for example Foster (1965); Stack (1974); Farber (1975); Hydén (1980); Kennedy (1988); Harvey (1993). The idea goes back even further to Max Weber, who felt that strong economic links within kin groups (“fettters of the sib”) discouraged entrepreneurial activity (Weber, 1915).

³Rosenzweig and Stark (1989), Morduch (1991), Townsend (1994), Kurosaki and Fafchamps (2002) and Fafchamps and Lund (2003) study consumption smoothing achieved through informal transfers. Recent evidence suggests not all these transfers are voluntary (Comola and Fafchamps (2014); Kinnan (2017)).

on each other more, increasing measured rates of kinship taxation. I directly test the effect of kinship taxation on investment decisions using a cash transfer field experiment with the same sample, and estimate the aggregate efficiency cost of this distortion using a model of input misallocation.

I find that kinship taxation is highly distortionary for a minority of my sample. For the quarter of female participants and third of males who face positive marginal kinship tax rates, the mean marginal rate is above 50%. Since those who face these high tax rates are better educated, have higher cognitive ability, and are more likely to own a firm, aggregate costs may be disproportionately high. Focusing on firm owners, I use a simple structural model that suggests aggregate firm productivity would be 26% higher in a counterfactual input allocation with no distortions from kinship taxation. This is striking given that only a minority of entrepreneurs face positive kinship tax rates, and it is driven by the most able entrepreneurs expanding their firms.

I then directly test whether kinship taxation affects investment decisions exploiting a cash transfer randomized controlled trial with the same sample. I find that low returns to female-owned microenterprises, consistent with a number of papers in similar settings, are not explained by kinship taxation. However kinship tax rates are a strong predictor of whether male entrepreneurs invest monthly cash transfers into their firm. This suggests that kinship taxation is an important drag on investment, without relying on the structural assumptions of my model. Combining the structural model with the cash transfer experiment, I also show that kinship taxation is a more costly distortion in this setting than credit constraints.

Sociologists and anthropologists have written widely about distortions from reliance on kin, most notably in ethnographies. While Platteau (2000) documents this across Africa, it has also been described in the US, both among poor urban blacks and rural whites, as well as in Indonesia, Vietnam and the Philippines.⁴ This idea has entered the more recent economics literature mostly through the work of Platteau (2000), introducing the idea that kinship taxation could explain otherwise puzzling behaviour. The emerging literature on this topic illustrates the wide range of productive distortions that may emerge from this institution. Member of these networks may exert lower effort (Hadnes et al., 2011), spend income quickly instead of saving it (Boltz, 2013; Goldberg, 2017), invest less in their firms (Di Falco and Bulte, 2011; Grimm et al., 2016; Fiala, 2017), take costly and unnecessary loans (Baland et al., 2011), and choose lower levels of human capital (Wantchekon et al.,

⁴See Stack (1974); Lewis (1978); Harvey (1993); Geertz (1963); Malarney (1998); Szanton (1998). Unfortunately analysis of this phenomenon has been done across disjointed theoretical frameworks and using inconsistent terms. One useful concept is Farber's (1975) 'centrifugal kinship,' where strong ties to kin tend to spread wealth among a broad group rather than concentrate it in smaller, more tightly-knit families. See also Hydén's (1980) 'economy of affection' and Foster's (1965) idea of the 'limited good.'

2014; Baland et al., 2016).⁵

The use of a lab experiment to measure of kinship taxation as I do in this paper was pioneered by Jakiela and Ozier (2016), where participants make incentivized investment decisions, which entail different levels of observability of winnings. They back out an implicit kinship tax rate from this, and show that at the village level, average tax rates correlate with lower incomes. Following on this approach, Boltz et al. (2015) explicitly elicit willingness-to-pay to hide income among lab participants in Senegal. The focus of this paper and Jakiela and Ozier (2016) is on whether people make costly hiding decisions, and not on the economic distortions these may lead to. Fiala (2017) uses a similar experiment to study spousal capture among Ugandan microentrepreneurs, which is related to but conceptually distinct from kinship taxation.

My contribution to this literature is twofold. First, I derive a simple and general expression for the distortionary pressure a given individual faces from kinship taxation. This motivates the use of willingness-to-pay to hide income as a measure of marginal kinship tax rates. Second, I use the results of the lab-in-the-field experiment to estimate aggregate economic costs from kinship taxation, whereas the literature has focused on illustrating behaviors distorted by kinship taxation. I find that kinship taxation reduces investment of the most able entrepreneurs and thereby lowers aggregate productivity by a quarter.

Why, if it is so costly, would the norms that underpin kinship taxation emerge and persist? First, these have obvious redistributionary and consumption-smoothing benefits, in settings where neither the state nor financial markets provide insurance or social security. Indeed, I show in a stylized model that under a reasonable range of parameters, the insurance benefits of kinship taxation in my setting seem to outweigh the efficiency costs. That is, the welfare gain from income pooling is greater than the cost to aggregate production.⁶

The following section describes precisely how I define kinship taxation, and presents a simple model which demonstrates that willingness-to-pay to hide income is equivalent to the marginal kinship tax rate. This result motivates the lab experiment described in section 3. The results of the experiment, in section 4, provides individual-level measures of kinship tax rates. I use these in sections 5 and 6 for my estimation of the losses from this dilution of incentives, and finally in section 7 consider the welfare implications of having insurance at the cost of lower output.

⁵Prior theoretical work helps rationalize why pressure from kin to redistribute may be distortionary, notably Hoff and Sen (2006).

⁶Platteau (2000) cites additional reasons: Groups inhibit private accumulation to discourage exit by its most successful members; in pre-scientific settings economic success is typically associated with luck, not effort; and absent the rule of law, any attempt to upset the established social order and hierarchy is sanctioned.

2 Conceptual framework

Taxes on producers are distortionary if they create a wedge between private and social returns from a productive activity. In formal tax systems, this gap is simply the marginal tax rate. In the case of informal taxation, the obvious analogue would be to measure the change in net transfers (outgoing minus incoming) with a change in income. A person who increases net transfers by fifty cents with every dollar of income faces a 50% marginal transfer rate.⁷ However, this marginal transfer rate may not be a correct measure of the wedge between private and social returns. Unlike payments made to tax collectors, transfers to your social network likely are valued in their own right. An entrepreneur who works to help support her parents values the transfers he makes to them, and may internalize the full social value of her productive decisions. The gap between the private and social returns from productive activity depends not just on the transfers made, but also on the extent to which these transfers are valued by the person making them. For kinship taxation to be distortionary, it must therefore be the case that, on the margin, transfers are valued less than private consumption. This is what I refer to as the marginal kinship tax rate: the part of the marginal transfer rate which is not valued by the person making the transfers.

I propose a simple model of utility maximization with productive decisions, and allow for a kinship tax rate that depends flexibly on income. In contrast with a formal tax, which would appear directly in the budget constraint, I model kinship taxation as an additional constraint which specifies a minimum level of net transfers, given an agent's income. This minimum transfer may be negative, and indeed will be for much of the population to have total net transfers sum to zero.⁸

2.1 Model

In this static model, a single entrepreneur makes productive decisions and divides her income between individual consumption c , and net transfers d to her social network. Both consumption and transfers enter directly into the utility function $u(c, d)$. Utility from transfers may come from warm-glow altruism, increased social status, and anticipated reciprocity. I do not distinguish between these.

Production decisions consist of how much capital k to hire at rental cost r , which is the only input in a value-added production function $f(k)$, with heterogeneous entrepreneurial ability A .

⁷The literature on crowding-out of formal transfers to the poor because of changes in informal transfers refers to this as the 'transfer derivative,' as in Cox and Fafchamps (2008).

⁸Marginal tax rates are just as likely to be positive for net transfer recipients, as in the analogous case of having welfare benefits withdrawn as income increases.

Kinship transfer obligation

The key part of this model is an economic obligation to an agent's social network. Given some income y , agents are required to make a minimum net transfer $T_i(y)$, such that $d \geq T_i(y)$. Entrepreneurs may choose to transfer more than they are required to, but not less. The minimum transfer function T_i is individual specific, and may be a function of personal characteristics such as gender or age, as well as ability A_i .

While stark, I believe this way of modelling transfer obligations is useful and appropriate. Take, as a benchmark, a world of anonymous relationships, where no demands are made on anyone's income. In this case, $T(y) = 0$ such that, for any y , agents are free to set $d = 0$. If one of these agents is altruistic, she may choose to set $d > 0$ despite the lack of social obligation—this would depend on the shape of her utility function with respect to d . But she would have no recourse to make demands of others, and hence cannot choose $d < 0$.

Consider if now, by begging, an agent can receive transfers from others. In this case, $T_i(y) = -B$, where B is the amount that can be collected by someone who chooses to beg full time. If agents can costlessly pretend to be poor, B is not a function of income, though it may still be individual-specific because of differences in begging skill, or because of characteristics such as age, sex, visible disability, etc.

Consider alternatively if the T function were to represent transfer obligations to a state's fiscal authority rather than to one's social network. $T(y)$ would now be increasing in y . With no direct utility from paying taxes and no disutility from receiving government transfers, $\partial u / \partial d = 0$ and hence the constraint $d \geq T(y)$ would bind for all agents.

Entrepreneur's choice of inputs

The agent's problem consists of:

$$\max_{c,d,k} u(c,d),$$

s.t.

$$\begin{aligned} c + d &= A_i f(k) - rk, \\ d &\geq T_i(y), \end{aligned} \tag{1}$$

where $y \equiv A_i f(k)$. The optimal choice of each input depends on whether the kinship tax constraint (equation 1) binds at the optimum. If $d > T_i(y)$, then the choice of capital satisfies $(1 - \tau_i^y) A_i f_k = r$. If it does bind, $d = T_i(y)$, and the choice of capital instead satisfies:

$$\left[1 - \frac{\partial T_i}{\partial y} \left(1 - \frac{u_d}{u_c}\right)\right] A_i f_k = r,$$

where u_c and u_d are marginal utilities of consumption and transfers, respectively.

The insight here is that distortion from the first best takes the form $\partial T_i / \partial y \times (1 - u_d / u_c) \equiv t_i$. That is, the marginal kinship tax rate t_i is equal to the marginal transfer rate $\partial T_i / \partial y$, attenuated by the marginal utility of transfers u_d . If an entrepreneur receives no direct benefits from transfers at the margin, then $u_d = 0$ and the kinship tax rate is simply the transfer rate, $\partial T_i / \partial y$. If on the other hand the kinship tax constraint is slack, such that $u_d = u_c$, then the marginal kinship tax rate is zero. This is a crucial point, and highlights the problem with measuring transfers to analyze distortions from kinship taxation. Doing so ignores the distinction between transfers increasing with income without the tax constraint binding (insofar as transfers are a ‘normal good’), and transfers that increase with income because of a binding kinship tax constraint. Measuring an individual’s tax rate then is not straightforward—it requires knowledge both of an implicit transfer obligation ($\partial T_i / \partial y$) that may be individual-specific, as well as marginal utilities u_d / u_c .

Conveniently, an individual’s marginal kinship tax rate is equal to their willingness-to-pay to hide income. Since this can be elicited in a lab setting, it allows me to measure marginal kinship tax rates directly, without separately estimating the marginal transfer rate and the marginal utility from transfers.

2.2 Willingness-to-pay to hide income

To measure marginal kinship tax rates I make use of a common insight: a costly tax evasion technology will be used up to the point where its marginal cost equals the marginal tax rate. This point is made explicitly in Slemrod (2001), and underlies the Allingham and Sandmo (1972) model of tax evasion. Here tax evasion is modelled as a choice to hide income h at cost $g(h) \geq 0$. In practice, hiding income in this setting can take many forms: choosing an occupation or workplace where income flows are harder to observe, hiding cash at home despite a risk of theft. Paying to hide income reduces the agent’s ‘taxable income’ by h , such that the agent’s problem becomes:

$$\max_{c,d,k} u(c, d),$$

s. t.

$$c + d + g(h) = A_i f(k) - rk,$$

$$d \geq T_i(y - h).$$

Optimal input choices remain the same, and an entrepreneur chooses to hide income up to the point where

$$\frac{\partial g}{\partial h} = \frac{\partial T_i}{\partial y} \left(1 - \frac{u_d}{u_c} \right) = t_i.$$

This is again equal to an individual's marginal transfer rate discounted by their marginal utility of transfers. Offered a choice to hide income from those who might make demands on it, this person would choose to hide so long as the fraction lost by doing so were less than the marginal kinship tax rate. Willingness-to-pay to hide income, then, is equal to the distortion to productive incentives caused by kinship taxation. Crucially, this method is not meant to measure the amount of hiding outside the lab, or its costliness. Participants may never choose to hide income, perhaps because they have no access to a technology to do so. What is important is that they are willing to hide at a marginal cost equal to their marginal kinship tax rate.

This measure of kinship taxation is robust to allowing for distortions that come not from having to make transfers, but from avoiding envy or jealousy.⁹ It captures transfers in the broadest sense, including not just cash payments but also in-kind transfers, preferential loans, and services of any kind. It also accounts for the fact that productive distortions come from expected future transfers rather than realized past transfers.¹⁰

However, one important concern is that earned and unearned income may face different kinship tax rates. This could be either because pressure to share ($\partial T_i / \partial y$) or benefits from sharing (u_d / u_c) depend on the source of income. Intuitively, these could act in opposite directions. On the one hand, it may be that the disutility cost of making transfers is lower for windfall gains. On the other, groups may tax earned income less heavily than a windfall, to attenuate moral hazard. This is particularly concerning given the finding that in some

⁹It also correctly accounts for a desire to be seen to be rich, as in Charles et al. (2009). Conspicuous consumption would reduce marginal kinship tax rates, rather than increase them. Communities which face large distortions from kinship taxation may find it optimal to encourage their members to show off their income, which is consistent with the observation that conspicuous consumption is more common among the poor.

¹⁰Given risk aversion this would imply measured transfers underestimate the distortions from anticipated transfers, which my measure captures correctly. For example, you may choose to consume your income rather than invest it, if you know that whenever your brother is sick he comes and asks for help. If your brother remains healthy, this will not be observable by the econometrician despite having distortionary effects. Realized transfers do not accurately measure the pressure to share resources when agents are risk averse.

settings, framing income as having been earned has been shown to affect giving in dictator games. However, a recent literature finds that in Africa, and especially among the poor, whether income is earned or not does not affect allocation decisions (Cappelen et al., 2013; Barr et al., 2015; Jakiela, 2015).

While the distortions of primary economic importance would come from taxes on earned income, lab experiments typically rely on unearned income. I address this concern in part by having participants complete a real effort task to earn the income they can choose to hide. I also directly test in section 4.4 whether, in this setting, people have different willingness-to-pay to hide earned versus windfall income, and find that they do not.

The following intuition underpins the proposed method to identify kinship taxation: paying to hide income from one’s social network is a sign that the value of potential future income is not fully internalized when making productive decisions.

3 Eliciting kinship tax rates

This section describes a lab-in-the-field experiment which elicits willingness-to-pay to hide income. This measure provides an individual-specific marginal kinship tax rate, which I use to identify and describe those who face distortionary pressure to share income. In subsequent sections, I use these kinship tax rates to estimate the economic cost of kinship taxation.

3.1 Experimental setup

In each location, eligible participants gather in a designated area outside the local primary school, and are registered by the team supervisor. Each participant is assigned to a separate enumerator. Enumerators are distributed in teams of two to four, each in a separate classroom. These sessions are synchronized, so that once all the participants assigned to one team of enumerators complete the experiment, they leave at the same time and a new set of participants are ushered in. From the beginning to the end of a session, participants work one-on-one with their enumerator, first signing a consent form, then answering survey questions, receiving experimental instructions, making decisions in the experiment, and finally receiving any experimental winnings. These steps are described in turn below.

Consent forms are read aloud to participants in Somali, emphasizing the confidentiality of their answers. Participants then go through a short survey, answering questions about basic demographics and family characteristics. They are asked for their relationship to all other participants in their session (whether family, friend, or stranger). They are administered a

short version of the Raven’s Progressive Matrices test, to measure cognitive ability.¹¹ Next they are presented with hypothetical situations and asked whether, in these situations, they would choose to pay to hide income. If they own a business, they are administered an additional microenterprise module, discussed in section 5.2 of this paper. After completing the survey, they move on to the experimental section of the interview. All responses are recorded using data entry software on smartphones, which allows randomization of the order in which some questions are asked to participants.¹²

Before starting the experiment, participants are told they will be given the chance to earn money if they agree to do some work. This is done to frame the money they might receive as having been earned. The task is to sort a plate of dry beans, which usually took less than five minutes.¹³

The experimental instructions explain the game in two steps.¹⁴ Each participant is first told they will soon have a chance to receive a prize, but that before determining whether they do, they will have some choice over what the prize will be. They are told they will make a series of binary decisions, and that one of their decisions will be chosen as the prize they might win. They are told that their best strategy is to make each decision based on their preference over the options offered, without thinking of the other choices they might make. Unlike in a typical bargaining situation they may be more familiar with, it is important they understand that in this context there is no strategic advantage to falsely reporting preferences.

The following question is asked eight times to each participant: “If the prize were either \$5 announced or \$ x in secret, which would you choose?”, where x is a different value for each question. Before being asked these questions, they are told that a prize ‘announced’ will be given to them in full view of the other participants in their session. Receipt of a prize ‘in secret’ will be known only to them.

The value of the ‘announced’ prize is held fixed at \$5 for each of the eight questions. The value ‘in secret’ goes from \$1.50 to \$5 in \$0.50 increments. Before making these choices, each participant is shown eight small paper tickets. They are told that each represents one of the decisions they will have to make. After a participant makes the first of their eight binary decisions (as in Table 1), the associated ticket for that decision is folded and placed into a small plastic basket, and so on for each of the questions.¹⁵ This is done to signal the

¹¹See <https://www.raventest.net/> for an example of this type of test.

¹²The software used is SurveyCTO (<http://www.surveycto.com>), a variant of Open Data Kit (ODK) for the Android platform.

¹³Tasks similar to this are commonly used in lab-in-the-field experiments to induce effort. See for example Jakiela (2015); Barr et al. (2015).

¹⁴See Appendix D for detailed experimental instructions.

¹⁵The order in which they make these eight decisions is randomized, as described in section 4.4.

Table 1: Choice list for Willingness-to-pay to hide income

If the prize were either...	Person A	Person B	Person C	Person D
1) \$5 announced or \$5.00 secret	Secret	Secret	Secret	Secret
2) \$5 announced or \$4.50 secret	Announced	Secret	Secret	Secret
3) \$5 announced or \$4.00 secret	Announced	Secret	Secret	Secret
4) \$5 announced or \$3.50 secret	Announced	Secret	Secret	Announced
5) \$5 announced or \$3.00 secret	Announced	Secret	Secret	Secret
6) \$5 announced or \$2.50 secret	Announced	Announced	Secret	Announced
7) \$5 announced or \$2.00 secret	Announced	Announced	Secret	Announced
8) \$5 announced or \$1.50 secret	Announced	Announced	Secret	Secret
Imputed WTP to hide income	0%	40%	70%	20%

Notes: Given four hypothetical sets of choices in the lab experiment, the bottom row specifies the corresponding imputed willingness-to-pay to hide income. In each case, the price paid for the last contiguous “secret” choice, starting from the first row, is the highest price we know with certainty the person is willing to pay. Person D is treated as if they chose “announced” for rows 4 to 8.

independence of each decision.

After making these choices, they are asked to confirm their decisions and told that once the next stage begins they will no longer be able to amend them. One of their eight decisions is then chosen at random by drawing paper slips from a small basket, a simple and transparent way to implement a Multiple Price List elicitation procedure.¹⁶ The participant is shown the ticket they drew, and their choice for that ticket (whether \$5 in public or the specified value in secret) is emphasized.

Participants are allowed to make dominated choices. That is, they can say they prefer \$3 in secret to \$5 in public, but also that they prefer \$5 in public to \$4 in secret. (Hence, by transitivity, that they prefer \$3 to \$4, both in secret.) After all eight decisions are made, if any decisions are inconsistent the enumerator explains that one of their choices implies choosing less cash instead of more cash. They are asked to make each of the eight choices again. If they again make inconsistent choices, these are kept and the lottery progresses using their inconsistent choices.

The next step determines whether or not the participant receives the prize they chose, and is added to provide deniability to anyone who chooses to take money in secret. This is done by putting small differently-colored tokens in a second opaque plastic container. Before putting them in the container, the participant is asked to choose one color, which becomes

¹⁶This method is used instead of a BDM elicitation mechanism (Becker et al., 1964) to emphasize the non-strategic nature of the decisions. In this context, asking about the maximum price a participant is willing to pay might prime them to act as they would when bargaining in a marketplace, a setting which does not reward truth-telling.

the winning color. The container is then lifted above the eye-level of the participant, and they are asked to reach in and choose one of the tokens.

If a participant wins the lottery and their chosen prize is to receive some amount of money in secret, it is given to them discreetly by their interviewer. Prize money is given in the form of cash, to more easily hide it.¹⁷ If they win and their chosen prize is to have \$5 announced, this is announced to the other participants in their session and the cash prize given to them in full view of these participants. It is not announced more generally, to members of other sessions. Since the probability of winning the lottery is low, participants can plausibly deny having won anything. It is important that someone who does not publicly win a prize is thought to have lost the lottery, rather than have won but chosen to hide.

To help ensure comprehension, immediately before the lab experiment each participant plays an almost identical “airtime” game. In this game, participants choose between a \$0.50 mobile airtime card and an equal or smaller amount of cash. Participants who choose the smaller cash prizes are effectively paying for liquidity, since airtime cards can in principle be resold at face value. This is not framed as a practice game, but rather a real game with a real prize. All details of the practice game, from the randomizing device to the text of the instructions, is as similar as possible to the subsequent income hiding game. If they win airtime or coins in the practice game, the prize is given to them on the spot. Once completed, participants are encouraged to ask any questions they have about the experiment before proceeding to the income hiding game.

Calculating tax rates

Before interpreting the results of the lab experiment, the eight binary decisions of whether to hide income at a given price need to be converted into a single willingness-to-pay. Table 1 demonstrates how I do this with four hypothetical sets of responses. Person A never chooses to pay to keep their income hidden, and hence has a willingness-to-pay of 0%. Person B is willing to pay \$2 to hide \$5. That is, they are willing to accept a 40% reduction of the potential prize to keep it hidden. Person C prefers as little as \$1.50 in secret to \$5 publicly, which means they have at least a 70% willingness-to-pay to hide income.

In each of these cases the imputed willingness-to-pay is conservatively chosen as the lower bound on their true willingness-to-pay.¹⁸ Person A, for example, may be indifferent between

¹⁷When piloting this experiment in slums of Nairobi, some participants preferred receiving their prize through a mobile payment (MPesa) for fear of theft. Participants in Garissa told us they preferred the liquidity of cash.

¹⁸The choice to use the lower-bound imputed tax rate for a given set of choices, as is done throughout the rest of this paper, is of course somewhat arbitrary. Appendix A shows results are insensitive to using alternative methods of imputing willingness-to-pay to hide income.

\$4.80 in secret and \$5 in public, implying a willingness-to-pay of 4%. This is especially true for person C, whose revealed willingness-to-pay is in the range of 70-100%.

Of the 1805 participants, 23 made dominated choices, as illustrated by Person D. In such cases, they are coded conservatively as having permanently switched to ‘Announced’ after their first such decision. In the case of Person D, this means all decisions below the first Announced decision (at \$5 announced versus \$3.50 secret) are treated as if they were also Announced.¹⁹

3.2 Sample description

Data collection for this study took place from mid-November 2014 to mid-January 2015, in 17 villages across Garissa County in Eastern Kenya. The population of this area is mostly Somali, Islamic and agro-pastoralist. Subjects were recruited from the pool of participants in a separate cash transfer field experiment, which aimed to increase primary school enrollment and attendance. The cash transfer experiment is not the subject of this study, though it provides an exogenous shock to income that I use in my analysis of the distortions from kinship taxation.

Since my study drew from participants in the cash transfer experiment (“CT participants”) and their spouses, it is useful to describe these in more detail. For purposes of the cash transfer program, CT participants were chosen by village committees under the guidance that each had to be responsible for at least one school-age child, and be from a relatively poor household. The condition related to having a child was chosen because the cash transfers were aimed at increasing school enrollment. Targeting relatively poorer households was done to target households where the cash transfers would be most likely to have an effect on schooling decisions.

Each household could have at most one CT participant, and village committees were told to prioritize women over men. This was done because of the belief that transfers to women may benefit children more than transfers to men. However, in order to examine to what extent kinship tax distortions come from intra-household conflict over resources, spouses of CT participants were also recruited for the lab-in-the-field experiment. Of the 1805 participants in the experiment which is the focus of this paper, 1,365 are themselves CT participants, while 440 are spouses of a CT participant.²⁰

¹⁹The box around the first three choices (the set of contiguous “secret” choices, starting from the top), is meant to illustrate how inconsistent choices are interpreted.

²⁰See Appendix B for an analysis of who chooses to participate in the lab experiment.

Descriptive statistics

Mean household daily income of participants is 251 Kenyan Shillings, or approximately \$2.51.²¹ Subjects earned an average of \$1.60 through their participation, including a \$1.10 show-up fee (paid in the form of 1kg of sugar) and experimental winnings ranging from \$0 to \$5.50.

Advertisement for participation in the lab-in-the-field experiment was done in each village by a local leader shortly before I arrived in that village with my team of enumerators. This person, usually contacted by mobile phone, was told that we would be at the village's elementary school on a given day.²² They were asked to notify potential participants that we would like to interview them, and that there would be some compensation. No details were provided at this stage about the content of the interviews. 1805 participants were recruited across the 17 villages, of whom 70% were female.²³ Descriptive statistics of all participants are presented in panel A of Table 2.

4 Anatomy of kinship taxation

Willingness-to-pay to hide income, elicited from the lab experiment, provides a measure of the marginal kinship tax rate each entrepreneur faces. Before using these tax rates to estimate their economic costs, this section provides an anatomy of kinship taxation: how do participants who face kinship taxation differ from those who do not? And how do marginal kinship tax rates change with income? I then present various pieces of evidence which support the notion that participants understood the game, and that their choices reflect their preferences.

4.1 Distribution of kinship tax rates

Table 3 presents the distribution of willingness-to-pay to hide for all participants, of whom 23% are willing to pay to hide income. For these participants, the average willingness-to-pay is \$2.59/\$5.00, or 52%. As described in section 2, each participants kinship tax rate is determined by their willingness-to-pay to hide income in the lab-in-the-field experiment. I take the fact that some participants are willing to pay to hide income as prima facie evidence that, within these social networks, transfer arrangements have an efficiency cost. People are

²¹All monetary values are converted from KES to USD at the rate of 100 to 1. The exchange rate in January 2015 was 90.2, which I approximate to 100 for simplicity.

²²This experiment was done between terms, when schools were not in use.

²³The sex-imbalance of the sample is related to the fact that 78% of the cash transfer beneficiaries were women, though we also interviewed spouses of beneficiaries.

Table 2: Descriptive Statistics

Variable	Mean	Std. Dev.	N
<i>(A) Individual characteristics</i>			
Female	0.70	0.46	1805
Age	40.83	13.45	1805
Married	0.88	0.33	1805
Separated	0.07	0.26	1805
Single (never married)	0.01	0.1	1805
Divorced	0.03	0.18	1805
Widowed	0.08	0.27	1805
Raven's (SD)	1.45	1	1726
Education (years)	0.31	1.39	1805
Islamic educ (years)	0.99	1.34	1805
Number of siblings	7.4	4.13	1805
<i>(B) Session co-participants</i>			
Co-participants: Spouse	0.01	0.1	1805
Co-participants: Family	0.74	0.78	1805
Co-participants: Friends	0.13	0.45	1805
Co-participants: Acquaintances	0.87	0.89	1805
Co-participants: Strangers	0.07	0.3	1805

Notes: All data are from the lab experiment and its accompanying survey. Panel A: 'Married,' 'Separated,' 'Single,' 'Divorced' and 'Widowed' are dummy variables equal to 1 for the appropriate marital status and 0 otherwise. (Married and Separated are not exclusive categories, but the others are.) 'Raven's (SD)' is the score from a cognitive ability test, normalized to a standard deviation of 1. 'Number of siblings' refers to siblings who are still alive. Panel B: Identity of co-participants in experimental session, as reported by the person being interviewed.

Table 3: Frequency of kinship tax rates

Kinship tax rate	Number	Percent
0%	1,382	77
10%	37	2
20%	60	3
30%	25	1
40%	36	2
50%	22	1
60%	18	1
70%	225	12
Total	1,805	100

willing to reduce the total income of their network to reduce the visibility of their own income.²⁴

Conditional on facing strictly positive tax rates, most people are willing to forgo at least 70% of a transfer to keep it hidden. This distribution of responses, with most respondents either not willing to pay to hide, or willing to pay a very high amount, is striking. This suggests the population is divided into two types, the more common type with a mode kinship tax rate at zero and a second type with a modal tax rate above 70%. To understand what determines a person’s type, the analysis which follows focuses on predicting an indicator variable for having a tax rate strictly greater than zero. That is, what predicts being a high-tax versus a low-tax type?

4.2 Kinship tax and income

By analogy to formal progressive tax schedules, the primary source of heterogeneity in marginal tax rates may be income. Given the high tax rates measured in the experiment, however, an important concern is reverse causation: high tax rates could lead to lower effort and investment, and hence lower income. An income shock from the cash transfer field experiment allows us to address this concern, isolating the causal impact of income on marginal tax rates.

Surprisingly, kinship tax rates appear to be flat with respect to this income shock. Column 1 of Table 4 shows that tax rates of participants who received cash transfers are no different from those who did not. Adding a range of individual-level characteristics as controls, in column 2, does not alter this result. The absence of a difference between the control

²⁴This is analogous to the literature on non-cooperative intra-household bargaining, which finds non-Pareto efficient outcomes even within households (Udry, 1996; Duflo and Udry, 2004).

Table 4: Kinship tax rates and Income

<i>Sample:</i>	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	Control	Control	Treatment	Treatment
Cash transfer	0.00633 (0.0239)	0.00131 (0.0237)				
Household income (log)			0.00178 (0.0269)	-0.00682 (0.0266)	0.00271 (0.0148)	0.00636 (0.0146)
Observations	1805	1805	343	343	1216	1216
Includes controls		X		X		X

Notes: Robust standard errors in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All columns are OLS regressions, with a dependent variable equal to 1 if kinship tax rate is strictly positive, and equal to 0 otherwise. Controls are: Age, sex, years of secular and Islamic education, number of siblings, and type of co-participants. ‘Cash transfer’ is equal to 1 if the participant is in the cash transfer RCT treatment group, and equal to 0 if in the control group. Household income is measured as total monthly household income in USD, prior to randomization into the cash transfer treatment or control groups. Household income is winsorized at 1% to attenuate the effect of outliers. Columns (3) and (4) include only participants in the RCT control group, and columns (5) and (6) include only participants in the RCT treatment group.

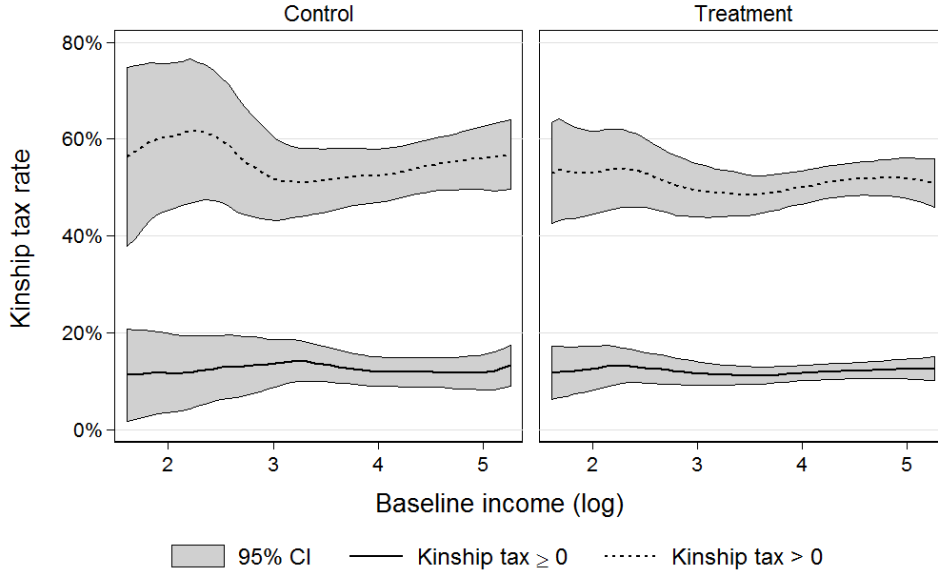
and treatment group is not simply a result of an insufficiently large treatment: The sum of cash transfers received by the treatment group is about 9 months of median household income.

An alternative to using the cash transfer treatment as a proxy for income is to use household income measured before participants were randomly assigned to the control or treatment groups.²⁵ Separating the sample into control and treatment in columns 3 and 5 and regressing the tax rate on the log of household income suggests that the slope is close to zero. Adding controls, in columns 4 and 6, does not alter this result.

Similarly, Figure 1 shows, non-parametrically, how kinship tax rates change in the cross-section, for different values of household income. In both panels, the solid line plots mean tax rates for all participants in a given subgroup, while the dashed line only includes participants with a kinship tax rate strictly greater than zero. The left panel includes members of control households in the cash transfer RCT, while the right panel includes only members of the treatment group. The solid lines in both panels are remarkably flat, and do not differ in height for the control and treatment groups. This suggests that kinship tax rates do not vary with income, either in the observed cross section or due to an exogenous increase in income. Restricting the sample to just those with non-zero kinship tax rates (dashed lines) shows a relationship which is less flat, but does not suggest kinship tax rates increase or

²⁵All income measures in this section are household income, winsorized at 1% to attenuate the effect of outliers.

Figure 1: Kinship tax rates and Income



Notes: Baseline income is measured as total monthly household income in USD, prior to randomization into the cash transfer treatment or control groups, winsorized at 1%. Kinship tax rates are taken from willingness-to-pay to hide income as measured in the lab experiment. The solid lines show the relationship between household income and kinship tax rates for all participants. The dashed lines include only participants with kinship tax rates greater than zero ($t_i > 0$). Participants in the cash transfer control group are in the left panel, while the treatment group are in the right panel.

decrease systematically with income.

There are a few caveats to this result. Respondents with high kinship tax rates may have understated their household income, which would be reasonable if they felt their answers may not be fully private. Household income is also hard to measure well in this context, given that much of household production is not marketed, and hence difficult to value accurately. Also there may be a disconnection between individual kinship tax rates and household (rather than individual) income. Each of these issues might attenuate the slope of the relationship. Despite these concerns, the randomized cash transfer study provides compelling evidence that marginal kinship tax rates do not vary substantially with income.

4.3 Who faces distortionary kinship taxation?

The large sample of participants allows for description of the individual-level predictors of kinship taxation, shown in Table 5. Column 1 starts with the most basic of these, age and gender. Column 2 includes characteristics that may proxy for earning ability: education, cognitive ability, and being a firm owner. Column 3 instead includes various family and

household characteristics. Column 4 combines these, and column 5 adds variables that capture the relationship to co-participants in the experimental session.²⁶ These are meant to control for the fact that participants face the decision to hide from people of varying degrees of closeness, which could influence their decision to pay to hide. The effect of relationships to co-participants is discussed below, in Section 4.4.2.

Female participants face substantially lower marginal kinship tax rates, a finding which is robust across specifications. The relationship with age, in contrast, is weak and not robust. Lower kinship tax rates for women is perhaps surprising if we expect that an important part of redistributive pressure is intra-household, and that women are less able to resist demands for their income. It is consistent however with the idea that, at least in some contexts, women have higher altruism than men, and hence derive greater utility from transfers.²⁷ It may also be that women in this setting have fewer opportunities than men to invest or save, and therefore are more willing to make transfers now in the expectation of future reciprocity. The finding that women have a lower kinship tax rate is consistent with the findings in Boltz et al. (2015) and Fiala (2017). In contrast to this, Jakiela and Ozier (2016) find that women are more likely to pay to hide income than men.²⁸

Each of the four ability measures is positively associated with kinship taxation. An added year of either formal (secular) or Islamic schooling is associated with a 2 percent increase in kinship taxation.²⁹ In addition to education, I also collect a measure of cognitive ability, using Raven’s matrices. These are standard nonverbal puzzles that measure what is termed ‘general intelligence’ in psychometric studies. Kinship taxation seems to be increasing with performance in this test, though the link is not robust to controlling for co-participants. Owning a firm is also associated with higher kinship taxation, and the choice of being an entrepreneur may be a proxy for ability.³⁰ Together, these results suggest that kinship tax rates are increasing in ability.

²⁶Each person was asked about their relationship to every other participant in their session, classified as one of the following categories: spouse, family, friend, acquaintance and stranger. The most common relationships are family and acquaintance, as shown in Table 2.

²⁷Andreoni and Vesterlund (2001) find that women display more altruism than men in situations where it is sufficiently costly.

²⁸This may be related to cultural differences between our settings. Gender inequality seems more significant in Garissa County than in Western Kenya, where their study takes place. For example, average years of education for adults in the region of Kenya they study is almost equal for men and women, whereas in mine, men have almost three times more years of education than women. (Using the 2008/2009 DHS.) This is consistent with the idea that the reason women hide less than men in this setting is that they have fewer opportunities outside the home.

²⁹Islamic education (mainly memorizing the Quran) is more common in this setting than formal secular education. These two types of education are not mutually exclusive: they are often undertaken simultaneously, as Islamic education is done mostly in evenings and weekends, and during school holidays.

³⁰Indeed as shown in Section 5.2, entrepreneurs in this sample have higher education and cognitive ability.

Table 5: Correlates of kinship taxation

<i>Dependent variable:</i>	(1)	(2)	(3)	(4)	(5)
		Marginal kinship tax rate > 0			
Female	-0.111*** (0.0237)	-0.0628** (0.0251)	-0.107*** (0.0240)	-0.0721*** (0.0272)	-0.0675** (0.0269)
Age	-0.00134* (0.000731)			-0.000314 (0.000985)	-0.000249 (0.000970)
Education (yrs)		0.0193** (0.00876)		0.0196** (0.00886)	0.0199** (0.00872)
Islamic educ (yrs)		0.0203** (0.00856)		0.0191** (0.00873)	0.0205** (0.00877)
Raven's score (SD)		0.0234** (0.0102)		0.0231** (0.0104)	0.0140 (0.0105)
Microentreprise owner		0.0661** (0.0266)		0.0588** (0.0268)	0.0664** (0.0264)
Married			0.0123 (0.0249)	0.0101 (0.0267)	0.00226 (0.0266)
Num children in household			0.00853 (0.00530)	0.0122** (0.00530)	0.0135*** (0.00511)
Num own children			-0.00335 (0.00456)	-0.00515 (0.00494)	-0.00617 (0.00470)
Num siblings			0.00605** (0.00258)	0.00501* (0.00261)	0.00570** (0.00262)
Num parents (living)			0.0217 (0.0136)	0.0138 (0.0149)	0.0150 (0.0147)
Minority clan			0.0353 (0.0506)	0.0124 (0.0518)	-0.000370 (0.0548)
Observations	1805	1726	1805	1726	1726
Control for co-participants					X

Notes: Robust standard errors in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All columns are linear probability models (OLS), with a dependent variable equal to 1 if kinship tax rate is greater than zero, and equal to 0 otherwise. Type of co-participants included as control in column 5. 'Education (yrs)' is years of formal secular education. 'Raven's score (SD)' is the score from a cognitive ability test, normalized to a standard deviation of 1. 'Married' is equal to 1 if married and not separated. 'Num children in household' is the number of children living with the respondent, whether they are related to them or not. 'Num own children' is the number of biological children the respondent has who are living. 'Num parents (living)' takes the value of 2 if both the mother and father are alive, and 1 or 0 if either or both are not. Minority clan is an indicator for being member of a clan which is in the minority in the respondent's village.

The family and network characteristics are more mixed. Being married is not a significant predictor of kinship taxation.³¹ To capture subtleties that are important in this setting, number of children is measured in two distinct ways. The first is the number of children living with the respondent, presumably under their care. This often includes children who are not their immediate offspring, such as nieces/nephews, grandchildren, or even fostered children of non-kin. The second is the number of biological children who are alive, but are not necessarily living with the respondent, including adult children. Kinship taxation increases with the first, but is uncorrelated with the latter. While not very robust, this suggests there is some channel through which intra-household issues do affect kinship taxation.³²

Having more siblings is associated with higher kinship tax rates, which provides some support for the use of siblings as a proxy for kinship taxation as in ? and Baland et al. (2016). However, whether both, only one, or neither of a respondent’s parents are alive does not predict kinship taxation. This mirrors the result on one’s own children, and suggests that, to the extent that parents and children rely on each other for support, these transfers are not distortionary. Finally, larger group identity does not seem of primary importance. Each of the 17 villages in this sample has a dominant clan, which includes more than half that village’s population. The ‘minority clan’ variable indicates whether the respondent is outside this dominant clan, and is not predictive of kinship taxation.

The most consistent conclusion from this set of results is that marginal kinship tax rates increase with earning ability. This can be reconciled with the finding that marginal rates are flat with respect to income in the following three ways. First, it is reasonable to assume that the relationship between income and kinship taxation varies within and across groups. Within groups, those with the highest income should face the highest rates (consistent with the results on ability). However the groups with the highest mean incomes may in turn face lower average tax rates, as their members need less assistance, and when assistance is needed there are more members who can help. When looking at individuals across many groups, these two forces go in opposite directions, and attenuate the within-group relationship between income and tax rates. The second way to reconcile these findings is that households vary in the amount of consumption needed to maintain comparable levels of welfare, and this is correlated with the household income. If tax rates are extracted on income in excess of the required level of consumption, then this correlation with the size or the needs of the household would also attenuate the relationship between tax rates and

³¹While not included in this table, the interaction of indicators for being female and being married is also not significant.

³²An alternative explanation however is that people partly fulfill their kinship tax obligations by fostering children of their relatives (and hence causality runs in the opposite direction). This is a common practice in much of Africa, and has been documented more broadly in Serra (2009).

income. Finally, there may simply be reverse causation: high kinship tax rates discourage production and reduce income, which again attenuates the relationship between these two. Ability is a measure of potential income, and hence that relationship is not distorted by this reverse causal channel. Indeed the following sections of this paper discuss precisely the strength of the causal relationship going from marginal kinship tax rates to income. First however, this section ends with a discussion of the reliability of the lab experiment used to elicit kinship tax rates.

4.4 Reliability of lab measure of kinship taxation

Results from any lab experiment rely on participants understanding the questions they are asked and choosing answers carefully and truthfully. This is of particular concern in a context where participants are unaccustomed to such exercises and have little formal education. In this section I provide evidence that participants understood the lab experiment, that we can interpret their choices as reflective of their true willingness-to-pay to hide income, and that this measure does reflect their kinship tax rate.

First, to ensure they understood the lab experiment, participants played a practice ‘air-time’ game, as described in section 3.1. While its purpose was simply to teach participants how to play the hiding game, there is evidence that they understood the airtime game: willingness-to-pay to convert the airtime card into cash is small on average, but almost doubles in locations without reliable mobile network coverage. This is what we would expect given lower liquidity of airtime cards in locations without a reliable network, and suggests an understanding of the mechanics of the experiment.

4.4.1 Stated reason for hiding

After completing the lab experiment, participants were asked why they chose to hide or not. In particular, we might be worried about the 225 people who reported that they prefer receiving \$1.50 in secret to \$5 in public. Responses were open-ended, and combined into categories ex-post. Of these 225 participants, 89 said they were hiding because they did not want to share with others. The following quotes are illustrative: “I don’t want to share with others that is why I prefer \$1.50 not announced”; “Nobody will ask me to give something”; “The \$5 announced might get finished due to demands from friends and family”; “There are so many people who are poor, they will ask you to give them something.”

A further 69 said they chose to hide because they did not want others to see what they received. While they did not explicitly say they did not want to share the income, it seems reasonable to believe that the reason they did not want others to know they received money

is that they would be asked to share some of it. “I don’t want people to know that I got money, I just want it to be secret”; “I don’t want people to see what I am receiving”; “I never want people to see what I am receiving.”

37 people said they were hiding to avoid an ‘evil eye’: “I don’t want people to notice, they may curse me with evil eye.” This use of magical curses to enforce redistribution of income is consistent with LeMay-Boucher et al. (2012) and Gershman (2015).

Of the remaining 30 participants with high willingness-to-pay to hide income, 12 simply said they preferred having money in secret, with no further explanation, and 7 gave confused answers. The latter include “I don’t know”; “If God wishes that I get that money, it is ok.” The remaining responses are: wanting to choose who to give the money to (4 people), being worried about safety (4), and wanting to avoid hatred, suspicion or humiliation (3). Taken together, it is hard to reconcile these responses with simple confusion or misunderstanding.

4.4.2 No hiding from strangers

The experiment can also help us understand who participants are hiding their income from, which may suggest who they are being taxed by. We can do this by focusing on the relationship between people in a given experimental session. Participants were told that if they chose \$5 ‘in public’ as a prize and they won, this would be announced to the other participants in their session. These sessions had significant variation in the relationship between participants. Panel B of Table 2 summarizes the number of co-participants of a given type of relationship (spouse, family, friend, acquaintance, or stranger). The categories were left somewhat ambiguous, to let each person decide how they thought of their relationship with a given person. They were each asked about every other participant in their experimental session, so in a given pair one person could say they were friends, while the other said they were acquaintances. Friends included all close relationships between people unrelated by blood or marriage. Acquaintances were people known to respondents, but with no close ties. Strangers were people who did not know each other.

Table 6 presents results from a regression of kinship tax rates on the number of co-participants who are members of the participant’s social network, though this should not be interpreted causally.³³ I find that an additional close friend or relative in the session is correlated with an increase in the probability of hiding in the experiment. An additional stranger in the experimental session does not seem to increase the probability of paying to hide money. While these relationships are not causal, they suggest the choice to hide in the

³³The experimental design calls for the supervisor to record the relationship between each participant in a session, before they are randomly allocated an enumerator. This would have allowed me to estimate the causal impact of having an additional relative or friend as a co-participant. Because of logistical issues, this relationship data was recorded for only 242 of 1805 participants, and is not used in this analysis.

Table 6: Kinship Tax correlations: Co-participants

	(1)	(2)	(3)
<i>Sample:</i>	All participants	Female only	Male only
<i>Dependent variable:</i>	Kin tax > 0	Kin tax > 0	Kin tax > 0
Co-participants: Spouse	0.00694 (0.0931)	-0.0789 (0.0975)	0.0832 (0.155)
Co-participants: Relatives	0.0339* (0.0193)	0.0234 (0.0220)	0.0705* (0.0384)
Co-participants: Friends	0.156*** (0.0297)	0.136*** (0.0363)	0.192*** (0.0512)
Co-participants: Acquaintances	0.0126 (0.0168)	0.00735 (0.0191)	0.0353 (0.0337)
Co-participants: Strangers	-0.0230 (0.0320)	-0.0137 (0.0357)	-0.0380 (0.0703)
Observations	1805	1272	533

Notes: All data are from the lab experiment and the accompanying survey. Robust standard errors in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All columns are linear probability models (OLS), with a dependent variable equal to 1 if kinship tax rate is strictly positive, and equal to 0 otherwise. Each ‘Co-participants’ variable is equal to the number of people in the same experimental session who were identified by the participant as being a member of the given category.

experiment is driven by a desire to prevent demands from one’s social network, rather than a fear of theft.

Since information can spread, hiding from a particular person is not evidence that this person will tax the recipient. It may be for example that when anyone wins, their co-participants will immediately tell all relevant people about this, and hence the relationship between participants should have no effect on tax rates. To the degree that information does spread, these results represent the sum of taxation by co-participants and anyone they convey this information to.

It may seem surprising that people hide as much from friends as from relatives, given that it seems clear from the literature that the majority of informal transfers are made between relatives (hence ‘kinship’ taxation). It may be that while transfers to relatives are larger than those to friends, it is the latter that are more distortionary because the value of these transfers is not internalized, as it is for close family members. Qualitative evidence supports this idea. For example, a 53 year-old man who chose to hide his earnings from ‘acquaintances’ reported that he did so because: “I want to use [the money] alone with my

family.”

4.4.3 Simplified hiding decision

As part of the survey and before doing the experiment, each participant answered a simple unincentivized question about willingness-to-pay to hide income. Specifically they were asked:

Imagine that I offer you \$5 today. Imagine also I did not offer money to the other participants in this room. Now what if I gave you the choice of not telling anyone that I gave you money. Then the others would not know that you received any money from me. If you could choose either, (1) I give you \$5 and I announce to the other participants in this room, or (2) I give you \$4 and do not tell the others, which would you prefer?

Responses to this simpler and more straightforward question are closely correlated to those in the more complex lab experiment. Table 7 shows in 82% of cases, participants made the same decision of whether to hide or not in this hypothetical question as in the incentivized lab experiment.³⁴ This suggests choices in the lab experiment are not driven by its complexity.

4.4.4 Random order of experimental questions

The order of the eight binary choices in the experiment, shown in Table 1, was randomized for each participant. For half of them, the first choice was \$5 announced or \$5 in secret and moved down, top to bottom. For the other half, the first choice was between \$5 announced or \$1.50 in secret and moved up, bottom to top. If people were making decisions based on their true preferences, the order of questions ought not to matter.

Panel B of Figure 2 shows that revealed willingness-to-pay to hide in the incentivized lab experiment is insensitive to the order of questions. Participants did not, for example, make the first binary decision carefully, and then stick to their decision for the remaining seven questions. This supports the conclusion that participants were not systematically making mistakes when choosing to pay to hide income.

³⁴For those whose decisions in the lab did not match their hypothetical answer, most switched to hiding in the lab experiment after having said they would not hide in the hypothetical question. This is consistent with the idea that some participants are embarrassed to admit they want to hide income from their group, and will only choose to do so when the stakes are real.

Table 7: Decision to pay to hide, incentivized versus hypothetical

		Choose to pay to hide (Hypothetical question)		
		Hide	Don't hide	Total
Choose to pay to hide (Incentivized question)	Hide	181 12%	164 11%	345 23%
	Don't hide	95 6%	1036 70%	1131 77%
	Total	276 19%	1200 81%	1476 100%

Notes: The rows in this table (Incentivized) correspond to a choice in the lab experiment to pay a positive amount to hide \$5. The columns (Hypothetical) correspond to an unincentivized stated preference for receiving either \$5 announced to co-participants, or \$4 in private. The shaded cells highlight the 82% of participants whose decision to hide income is consistent across the incentivized and hypothetical questions.

4.4.5 Few inconsistent choices, and educated hide more

As described in the experimental design, participants were allowed to make dominated choices in the lab experiment. Only 23 of 1805 participants made inconsistent decisions in the experiment. That these are so infrequent suggests a high degree of understanding from participants.

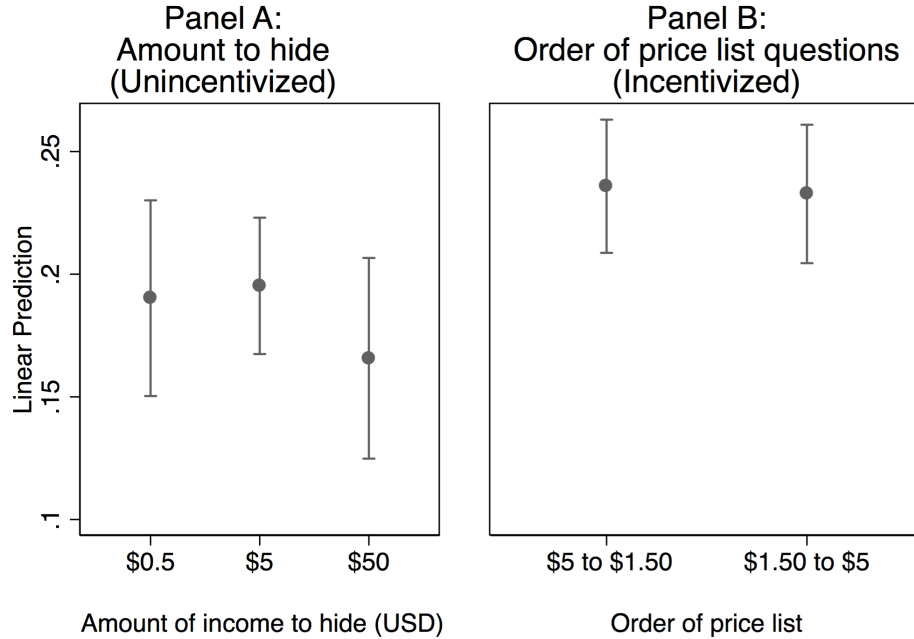
Similarly, if choosing to pay to hide were a product of confusion, we would expect more educated participants to be less willing to pay to hide. Instead, as noted above, I find that education and cognitive ability are both higher among those who pay to hide.

4.4.6 Avoiding embarrassment

It may be that participants pay to hide income not to reduce transfers to others, but because of a desire to avoid being the subject of an announcement. The preference for money received in secret comes for these people from this social embarrassment, rather than any pressure to share the announced income. If this were the case, the desire to avoid this announcement would be roughly independent of the amount received. That is, willingness-to-pay to hide (in percentage terms) should be decreasing in the amount to be hidden.

To test whether hiding of income is fixed rather than proportional, participants were asked one of three questions, chosen at random. One of the three is the hypothetical income-hiding question discussed above, where the choice is between receiving \$4 in secret or \$5 in public. The other two are identical, except for the fact that the amounts participants chose between

Figure 2: Willingness to pay to hide income, incentivized and unincentivized



Notes: Each bar represents a mean and 95% confidence interval for whether participants choose to pay to hide income in a given condition. Panel A shows the difference in responses across amounts to hide (in each case the cost of hiding was 20% of the total amount) in an unincentivized, hypothetical question. Panel B shows differences across order of price list questions in the incentivized lab experiment.

were either (1) \$0.40 in secret or \$0.50 announced, or (2) \$40 and \$50. That is, they were asked whether they would be willing to give up 20% of potential income to keep it hidden, across three orders of magnitude.

Panel A of Figure 2 shows that the decision to pay to hide seems relatively insensitive to the size of the sum to be hidden. Given the consistency of decisions across the three very different amounts, it seems that the motivation to hide is roughly proportional in the amount.

4.4.7 Earned versus windfall income

Kinship tax rates may vary for different types of income. Specifically, participants may feel differently about paying to hide earned versus windfall income. This is important, as all estimates in this paper of the cost of kinship taxation relate to income which would be earned, either through judicious investment or additional effort.

As mentioned in section 3.1, the money to be hidden in the lab experiment was framed as payment for an effort task, to induce participants to see the income they might receive as

having been earned. However, the need to ensure deniability for participants who chose to hide meant there had to be some element of chance to winning the prize.

To see whether this matters, a subset of participants were asked whether they would pay to hide in a situation involved earned income. After giving an answer to the hypothetical question described in section 4.4.3, they were also given the following scenario:

Now, imagine that I asked you to do a difficult job, and offered \$5 as payment for work. Imagine also I do not offer this job to the other participants in this room.

They were then again given the same choice between receiving \$4 in secret or \$5 in public. The results, for people who were asked both questions, are in table 8. First, the fraction of respondents choosing to hide in both scenarios is almost equal (18% versus 17%). Second, the shaded cells in the table show that these choices are highly consistent: for 88% of participants, the hiding decision is unaffected by whether the income is a windfall or is earned.

This result is surprising in light of evidence from dictator games that, when income is framed as having been earned, people generally choose to give less of it away (Cherry et al., 2002; Cappelen et al., 2007). However, recent evidence suggests that in poor African settings, behavior may not be affected by whether the endowment to be transferred has been earned (Cappelen et al., 2013; Barr et al., 2015; Jakiela, 2015). This may be because of egalitarian norms which focus on equality of outcomes, rather than what might be thought of as concern for fairness.

Together, the evidence from this section suggests that participants understood the game and revealed their true willingness-to-pay to hide income. Using these as measures of kinship tax rates, we can estimate the distortions in productive activity from kinship taxation.

5 Aggregate cost of misallocation

Using the measures of kinship taxation elicited in the lab, this section estimates the aggregate economic consequences of these taxes. I proceed in three stages. First, I modify the model from section 2 to include idiosyncratic distortionary wedges. Next, I describe firm-level data, and back out productivity and wedges for each. I discuss the counterfactual used to estimate the effect of distortions, and provide evidence which suggests the productivity and wedges that I measure are reasonable both in terms of my setting and compared to other studies on microenterprises. Finally, I present and discuss the effect of removing distortions from kinship taxation on efficiency and the firm size distribution.

Table 8: Hiding earned versus windfall income

		Choose to pay to hide EARNED income		
		Hide	Don't hide	Total
Choose to pay to hide WINDFALL income	Hide	79 11%	38 6%	117 17%
	Don't hide	46 7%	533 77%	579 83%
	Total	125 18%	571 82%	696 100%

Notes: The rows in this table (Windfall) correspond to an unincentivized stated preference for receiving either \$5 announced to co-participants, or \$4 in private. The columns (Earned) correspond to a similar choice, but where the income which can be hidden comes as payment for “a difficult job”. The shaded cells highlight the 88% of participants for whom the hiding decision is unaffected by whether the income is a windfall or is earned.

5.1 Model with misallocation

The structural model used to evaluate the costs of kinship taxation relies on the assumption that the tax rates measured in the lab affect productive decisions as modelled in section 2. That is, while the decision to pay to hide income faces the same degree of distortion, on the margin, as does the decision to increase the use of firm inputs. To make better use of the data available, I start by adding labor as a second productive input alongside capital. Next, to study the effect of kinship taxation in the presence of multiple other distortions, I add a constraint on access to credit, which places an upper-bound on the amount of capital that can be used in production. I allow this limit, \bar{k}_i , to be individual specific, as it may differ according to the assets an agent owns, their collateral, their ability, or their repayment history.

Finally, I add distortionary wedges τ_i^k (on capital) and τ_i^y (on output), as is common in the literature on input misallocation across firms.³⁵ These wedges are intended to capture, in summary form, all distortions facing entrepreneurs. While credit constraints are explicitly modelled, capital wedges here correspond to all other distortions that affect the relative use of capital versus labor, such as labor market distortions or differential input costs. Distortions which affect the use of both inputs, are output wedges. These include, for example, advantageous political connections for some entrepreneurs, or limits on the types

³⁵Hsieh and Klenow (2009) is an early and influential paper in this literature.

of businesses women or particular ethnic groups can operate. Either wedge can be positive or negative, depending on whether they represent a net tax or subsidy for that entrepreneur.

The entrepreneur's problem is now:

$$\max_{c,d,k,l} u(c, d),$$

s.t.

$$c + d = (1 - \tau_i^y)A_i f(k, l) - wl - (1 + \tau_i^k)rk, \quad (2)$$

$$d \geq T_i(y),$$

$$k \leq \bar{k}_i, \quad (3)$$

where $y \equiv A_i f(k, l)$. The optimal choice of each input depends on whether the kinship tax constraint (equation 2) and the credit constraint (equation 3) bind at the optimum. If neither is binding, $d > T_i(y)$ and $k < \bar{k}_i$, the choice of inputs satisfies the following familiar conditions: $(1 - \tau_i^y)A_i f_l = w$, and $(1 - \tau_i^y)A_i f_k = (1 + \tau_i^k)r$.

If the kinship tax and credit constraints do bind, $d = T_i(y)$ and $k = \bar{k}_i$, the choice inputs instead satisfies:

$$\begin{aligned} [1 - \tau_i^y - t_i] A_i f_l &= w, \\ [1 - \tau_i^y - t_i] A_i f_k &= (1 + \tau_i^k)r + \frac{\mu_i}{u_c}, \end{aligned} \quad (4)$$

where $\partial T_i / \partial y \times (1 - u_d / u_c) \equiv t_i$ is the kinship tax rate as described in section 2. μ_i is the Lagrange multiplier on (3) and is a measure of the degree to which credit constraints bind.

The counterfactual analysis of the consequences of kinship taxation relies on the observation that, for both capital and labour, the distortion imposed by kinship taxation is analogous to an output wedge τ^y , and enters additively with other output wedges. Conversely, the distortion from credit constraints enters additively with the capital wedge τ^k . For ease of notation, I define $\tilde{\tau}_i^k \equiv \tau_i^k + \frac{\mu_i}{u_c r}$ and use $\tilde{\tau}_i^k$ throughout this paper as my measure of capital distortions. I do this because the data I have do not allow me to distinguish credit constraints from other capital-labor wedges.

Tax on value added

The quantity y on which the kinship tax is levied is not gross output but value added—that is, it is net of the cost of intermediates as implied by their absence in the production function.

I argue that the use of value added, rather than gross output or pure economic profits, is appropriate in this setting.³⁶

If an entrepreneur owns their capital stock, as is standard for owners of small businesses, payments to capital are paid to themselves and are part of their income, which is the quantity to be taxed. Payments to labor, similarly, should not be deducted from value added in the kinship taxation function. This is true in a context where labor is provided by the entrepreneur and their family. Payments for the entrepreneur’s labor are, like payments to capital, part of their income and hence taxed. Payments to family labor, likewise, are a function of the firm’s value added, since it is conceptually difficult to separate payments to a family member for their labor and transfers made to them because of kinship obligations. Indeed hiring unproductive relatives is a common form of transfers in such contexts.

More fundamentally, a goal of this paper is to test the hypothesis that kinship taxation is distortionary. This model allows for the possibility of a particular type of distortion, and the results indeed support the idea that entrepreneurs who face kinship taxation keep their firms inefficiently small.

5.2 Firm data

In this section I describe entrepreneurs and their firms. Respondents are considered entrepreneurs if they answered yes to the following question: “Do you personally own a business of any kind, for example a kiosk, a shop, selling charcoal, vegetables, a restaurant?”³⁷ These entrepreneurs, as shown in Table 9, differ systematically from non-entrepreneurs in the sample. They are more likely to be female, have more years of education, are less likely to be married, and have more siblings. They also face higher kinship tax rates. The following analysis on the cost of kinship taxation is restricted to entrepreneurs, so these differences signal the need for caution in generalizing these findings to non-entrepreneurs.

Entrepreneurs were asked about firm-level outcomes, using questions adapted from the Sri Lanka Microenterprise Survey of De Mel et al. (2008). Their responses are summarized in Table 10. 20% of the sample own a microenterprise, which provides a sample of 361 entrepreneurs. Firms are on average just under 4 years old, and 15% of them are registered with the local government.³⁸

³⁶If taxes were proportional and calculated on economic profits, they would be non-distortionary. The choice of inputs which maximizes profits π also maximizes $(1 - t)\pi$.

³⁷Respondents were told to include any business owned jointly with others. Since the effect of kinship taxation on firm-level decisions is unclear in the case of a firm with multiple owners, I show in Appendix A that results are robust to dropping microenterprises with more than one owner.

³⁸Despite registration, none of these firms are sufficiently formal to report their earnings to the government, and do not face direct formal taxation.

Table 9: Entrepreneurs

<i>Dependent variable:</i>	(1) Entrepreneur	(2) Entrepreneur	(3) Entrepreneur	(4) Entrepreneur	(5) Entrepreneur	(6) Entrepreneur
Female	0.0679*** (0.0194)	0.0752*** (0.0195)	0.0578*** (0.0207)	0.100*** (0.0212)	0.0579*** (0.0197)	0.0923*** (0.0238)
Cash transfer treatment	-0.0140 (0.0229)					-0.0203 (0.0235)
Kin Tax > 0		0.0716*** (0.0235)				0.0569** (0.0244)
Age			-0.00116* (0.000650)			0.000150 (0.000787)
Education (yrs)				0.0287*** (0.00824)		0.0277*** (0.00816)
Islamic educ (yrs)				0.00651 (0.00729)		0.00339 (0.00724)
Raven's score (SD)				0.0196** (0.00955)		0.0198** (0.00958)
Living with spouse					-0.0442* (0.0253)	-0.0571** (0.0278)
Number of siblings					0.00829*** (0.00234)	0.00782*** (0.00244)
Mean of dep var	0.200	0.200	0.200	0.207	0.200	0.207
Observations	1805	1805	1805	1726	1805	1726

Notes: Robust standard errors in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. Data are from the lab experiment and the accompanying survey. All columns are OLS regressions, with a dependent variable equal to 1 if the participant answered Yes to whether they “personally own a business of any kind” and 0 otherwise. ‘Cash transfer treatment’ is equal to 1 if the participant is in the cash transfer treatment group, and equal to 0 if they are in the control group. ‘Kin Tax > 0 is equal to 1 if kinship tax rate is strictly positive, and equal to 0 otherwise. ‘Education (yrs)’ is years of formal secular education. ‘Raven’s score (SD)’ is the score from a cognitive ability test, normalized to a standard deviation of 1. ‘Living with spouse’ is equal to 1 if married and not separated.

Using data on capital, labor, and output, I can estimate entrepreneur-specific parameters for productivity and distortionary wedges. I measure capital as the sum of the value of any inventory, equipment and structure used for the business. Labor use is the sum of the number of wage workers, unpaid workers (family members, for example) and any co-owners who provide labor to the business, including the owner.

To match the model presented in section 2, the remaining necessary variable is value added.³⁹ The measure I use for this is total income earned from the business, including any

³⁹As is well known, including intermediates in production functions can introduce bias because of the effect of productivity shocks on the use of intermediates, a problem referred to as transmission bias. A common approach to this problem is to omit intermediates from the set of inputs, and instead estimate a value added production function as I do. Gandhi et al. (2016) however show that doing this results in a different, and perhaps worse, bias in estimating production functions. Their proposed solution to this problem, and to the transmission bias that comes with including intermediates, relies on the use of panel data. With only a cross section, I cannot use this technique, or the alternatives used in this literature (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Akerberg et al., 2015). Note however that the goal of this exercise is not to estimate production functions (which I do not attempt with these data, given the absence of a panel and widespread distortionary wedges) but instead to see how much of the distortions in this setting come from

Table 10: Descriptive Statistics (Firms)

Variable	Mean	Std. Dev.	N
Microentreprise owner	0.2	0.4	1805
Business age (yrs)	3.91	4.13	317
Hours worked daily	8.01	3.32	361
Business registered	0.15	0.36	361
Capital: Inventory	251.56	706.13	358
Capital: Equipment	113.67	563.31	357
Capital: Structures	176.91	395.24	354
Monthly sales	301.1	1541.12	348
Monthly profits	85.98	191.68	346
Labor use	1.66	1.4	361

Notes: All data are from the lab experiment and the accompanying survey. Hours worked daily only includes work related to their business. Capital, sales and profits measures are in USD. Labor use is the sum of owners/partners who work in the business, wage workers, and unpaid workers.

wage entrepreneurs pay themselves. The precise question is: “What was the total income the business earned last month after paying all expenses including wages of employees, but not including any income you or another owner paid yourselves. That is, what were the profits of your business last month?” For similar firms in Sri Lanka, De Mel et al. (2009b) find that this question is a more accurate measure of business income than detailed questions on revenues and expenses. Unlike in more formal settings, these entrepreneurs do not keep formal records, and hence finding a measure of output which is reasonably well-measured is crucial.

Using the “total income” question for value added is reasonable, since in this context we cannot distinguish between payments entrepreneurs earn from the use of their labor, from entrepreneurial rents, and from returns to capital. Instead of trying to do this, I interpret their answer to the question above to be equal to value added, which includes all of these, but excludes payment for intermediate inputs.

Of the 361 firms in my sample, 326 have non-zero, non-missing data on capital, labor, and output. I restrict analysis to these firms, and winsorize these values at 1% (top and bottom) to attenuate the effect of outliers.⁴⁰ I discuss the sensitivity of results to the exclusion of outliers in Appendix A.

kinship taxation.

⁴⁰That is, I set all values above the 99th percentile equal to the 99th percentile, and set values below the 1st percentile equal to the 1st percentile.

5.3 Firm-specific parameters

In this section, I use data on capital, labor and output to back out firm-level measures of productivity A_i and distortionary wedges $\tau_i^y + t_i$ and $\tilde{\tau}_i^k$. Since the kinship tax rate t_i enters as an output wedge in the entrepreneur's first order condition, we cannot separately identify τ_i^y and t_i using firm level data. This is why a separate measure of t_i from the lab experiment is necessary to perform the desired counterfactual analysis.

5.3.1 Productivity

One of the key assumptions underlying this analysis is that distortionary wedges, whether from kinship taxation or any other source, do not affect the technical productivity of entrepreneurs. That is, the technical ability to convert inputs into output is not itself a function of t_i , τ_i^y , or $\tilde{\tau}_i^k$. In the case of kinship taxation, this may not be self-evident. For example, entrepreneurs faced with kinship taxation may be induced to hire relatively unproductive relatives, which would reduce the output produced with a given amount of labor. This is something I cannot test directly, though I provide evidence below that, compared to entrepreneurs in general, those with positive kinship tax rate do not seem to distort their relative use of capital and labor.⁴¹ I take this as evidence that kinship taxation does not directly affect the productivity of one input relative to the other. Further, if kinship taxation does negatively affect the ability of entrepreneurs to convert inputs into output, this implies that the resulting estimate of the gain from removing kinship taxation is a lower-bound on the true gain. This is because I would be underestimating the true productivity of entrepreneurs who face kinship taxation, and hence the gains from removing this distortion.

Given this assumption, we only need to impose a production function f to back out firm-specific productivity A_i . To pin down firm size, I assume entrepreneurs face a common Cobb-Douglas production function with a Lucas span-of-control parameter $\sigma < 1$,

$$y_i = A_i(k_i^\alpha l_i^{1-\alpha})^\sigma. \quad (5)$$

Given the production function and values of α and σ , data on output, capital and labor allow us to solve exactly for each entrepreneur's productivity, $A_i = y_i(k_i^\alpha l_i^{1-\alpha})^{-\sigma}$. Values for the capital share α and returns to scale σ are taken from the literature, since estimation using my sample is biased by precisely the distortions of interest. I use capital share α of 0.3 and span of control $\sigma = 0.7$. See Appendix A.2 for a discussion on the choice of these parameters, and sensitivity to alternative values.

⁴¹Also, the share of workers in a firm who are unpaid, which is a proxy for the use of family labor, is the same for entrepreneurs who face kinship taxation and those who do not.

Table 11: Distribution of productivity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Microenterprises		Agriculture		Manufacturing Plants		
	Squires (2015)	DMW (2008)	Shenoy (2015)	R&S (2015)	Hsieh & Klenow (2009)		
Setting	Kenya	Sri Lanka	Thailand	Malawi	India '94	China '05	US '97
SD	1.09	1.05		1.19	0.67	0.63	0.49
75/25	1.55	1.37	1.81	1.15	0.81	0.82	0.53
90/10	2.85	2.72	3.09	2.38	1.60	1.59	1.19

Notes: The first column is my own data. DMW (2008) is my own calculation, using data from De Mel et al. (2008). R&S is Restuccia and Santaaulalia-Llopis (2015). Data from Hsieh and Klenow (2009) are reported for 1994 for India, 2005 for China, and 1997 for the US. SD is the standard deviation of log productivity; 75-25 is the log difference between the 75 and 25 percentile and 90-10 the 90 to 10 percentile difference in productivity.

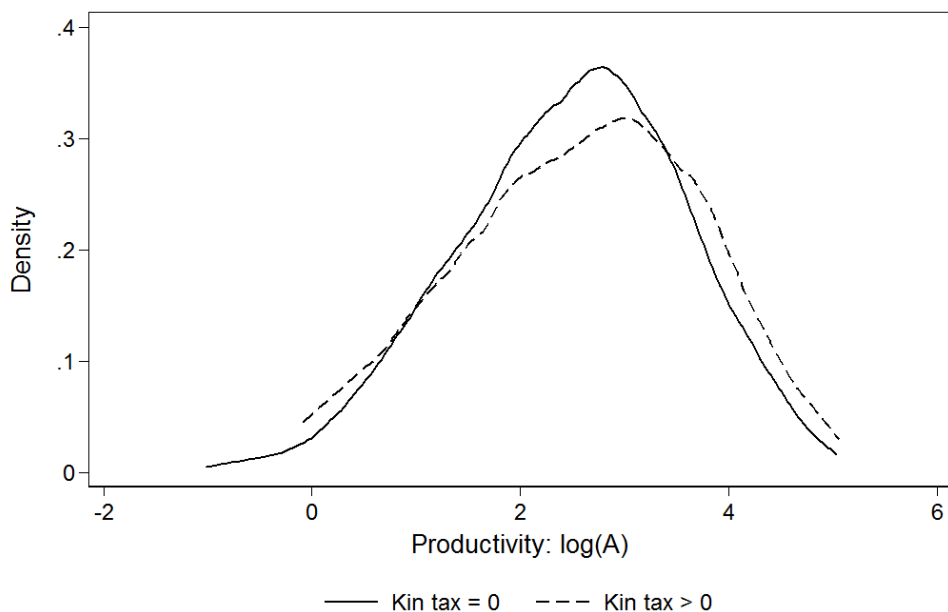
Using these values of α and σ , I estimate A_i for each entrepreneur from equation 5. The distribution of firm-level productivity that I derive is presented in Table 11. I report the standard deviation of the log of measured TFP (A_i), as well as the ratio of the 75th to 25th percentiles, and the ratio of the 90th to 10th percentiles. Dispersion of firm-level TFP is important in this analysis because it amplifies the potential costs of misallocation. To this end, I compare my results to corresponding values from similar papers. The distribution of productivity in my sample is comparable to what I calculate using De Mel et al. (2008) microenterprise data from Sri Lanka. Column 3 reports dispersion of TFP for farmers in Thailand from Shenoy (2015), where dispersion of productivity is greater than in my setting, while in column 4, dispersion amongst Malawian farmers (Restuccia and Santaaulalia-Llopis, 2015) is comparable to what I find. Hsieh and Klenow (2009) report substantially lower dispersion of productivity in manufacturing firms in India and China, and again even lower in the US. These results broadly suggest that dispersion of productivity is lower in more modern sectors which are better integrated into markets. This is reasonable if, in these sectors, unproductive firms are more likely to shut down, such that surviving firms all have relatively high TFP and dispersion is reduced.

The distribution of productivity across entrepreneurs in my sample is presented graphically in Figure 3. Entrepreneurs here are divided into two groups: those who face a strictly positive kinship tax rate and those who face a zero rate. The distribution of log productivity for entrepreneurs facing kinship taxation is mostly to the right of other entrepreneurs.

5.3.2 Distortionary wedges

Given a set of distortions $\tau_i^y + t_i$ and $\tilde{\tau}_i^k$, entrepreneurs choose inputs k and l such that they solve the entrepreneur's problem described in section 2. This means that using their choice of inputs, we can back out a measure of each entrepreneur's wedges from their first

Figure 3: Dispersion of productivity



Notes: Kernel density plot of the A_i productivity parameter estimated for each entrepreneur. As in equation 5, this is backed out from firm data on output, capital and labor using $A_i = y_i(k_i^\alpha l_i^{1-\alpha})^{-\sigma}$. The solid line represents the distribution of entrepreneurs who do not face binding kinship taxation ($t_i = 0$). The dashed line represents entrepreneurs who face a binding kinship tax constraint ($t_i > 0$).

order conditions. Incorporating the production function described above, this implies the following two conditions must hold:

$$1 + \tilde{\tau}_i^k = \frac{\alpha}{1 - \alpha} \frac{wl_i}{rk_i}, \quad (6)$$

$$1 - \tau_i^y - t_i = \frac{wl_i}{(1 - \alpha)\sigma y_i}. \quad (7)$$

These allow us to solve for $\tilde{\tau}_i^k$ and $\tau_i^y + t_i$. Recall that $\tilde{\tau}_i^k$ affects the entrepreneur's relative use of capital and labor, which includes credit constraints as well as other capital or labor market frictions. Firms with a low capital-to-labor ratio have correspondingly higher values of $\tilde{\tau}_i^k$. Conversely $\tau_i^y + t_i$ wedges affect the scale of the firm by reducing the use of both inputs symmetrically. An increase in τ_i^y or in t_i reduces the use of both capital and labor, but leaves the ratio unchanged.

To back out wedges $\tilde{\tau}_i^k$ and $\tau_i^y + t_i$, we also need values for w and r . I use \$30 as the monthly wage w . This is the median wage paid by firms in my sample that hire wage labor, where each firm's wage is calculated by dividing their wage bill by the number of wage

workers. For comparison, median household income in my sample is \$60 per month, and median firm profits are \$41 per month. I use $r = 2\%$, which implies a 27% compound annual rate. In this context, r is equal to the sum of the monthly interest rate and depreciation rate.

Crucially, the procedure used below to calculate wedges and the gains from misallocation means that the results are completely insensitive to the interest rate used. The only effect of a change in r is to scale the estimates of capital wedges. Recall that the actual price of capital faced by each firm is $(1 + \tilde{\tau}_i^k)r$, and so the two parameters cannot be separately identified. A change in r causes a shift to the right or left of the distribution of $\tilde{\tau}_i^k$. A change in the wage w does affect the estimated gain from removing kinship taxation, but results are relatively insensitive to the choice of w , for the same reason a change in r has no effect.⁴²

Using equations 6 and 7, we can now back out both wedges separately for each firm. Figures 4a and 4b show the distribution of $\tau_i^y + t_i$ and $\tilde{\tau}_i^k$ in my sample, again divided by whether the participants face a positive kinship tax or not.

5.4 Reallocation

5.4.1 Reallocation procedure

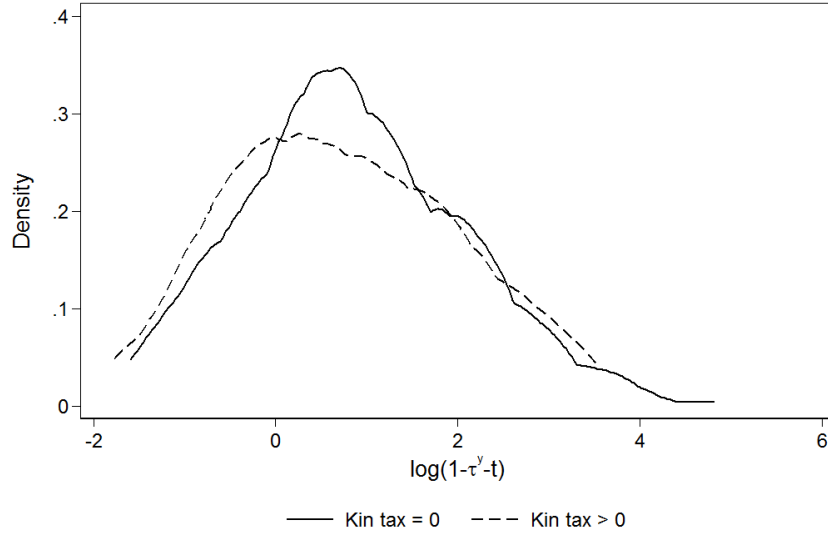
The vector of productivity and distortionary wedges for each entrepreneur allows us to consider counterfactual allocations of inputs with different sets of wedges. The goal of this exercise is to remove the distortion from kinship taxation, t_i , while keeping other distortions in place, and see what effect this has on output. To do so, I use a reallocation procedure that keeps the total amount of capital and labor constant, as well as the number of firms. Importantly, I keep fixed the technical productivity A_i of all entrepreneurs. Any gains or losses from reallocation come simply from reducing or increasing gaps in marginal products of capital and labor across firms.

Given some vector of productivity and wedges, we can use the entrepreneur's first order conditions to solve for their optimal choice of capital and labor:

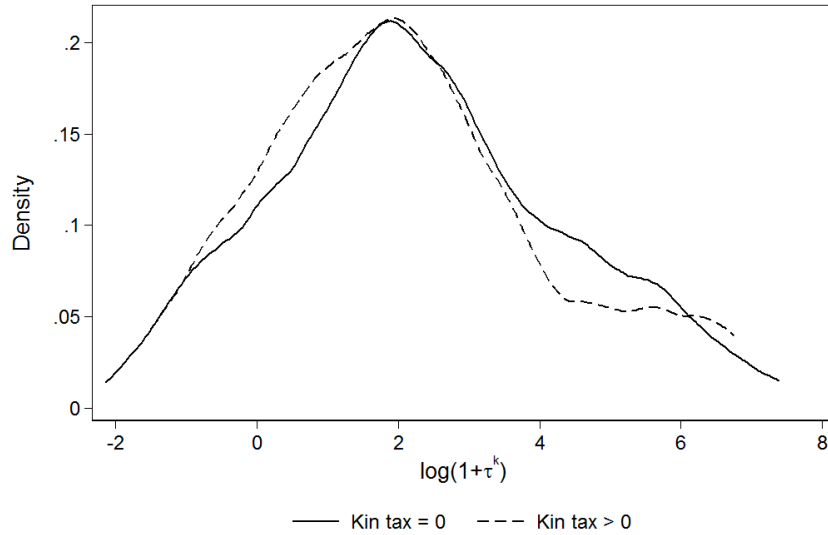
⁴²The reason a change in w does have an effect is it appears in equation 7, which determines the value of $\tau_i^y + t_i$. This might matter since most of the labor used in these firms is unpaid family labor, and it is reasonable to think that the opportunity cost of much of it is lower than \$30 per month. However the main result of this paper, the gain in aggregate TFP from removing kinship taxation, goes down by only 4% (1 percentage point) when the wage is reduced to only \$15 per month.

Figure 4: Dispersion of distortionary wedges

(a) Dispersion of output wedges, $\ln(1 - \tau_i^y - t_i)$



(b) Dispersion of capital-labor wedges, $\ln(1 + \tilde{\tau}_i^k)$



Notes: The solid lines represent the distribution of entrepreneurs who do not face binding kinship taxation ($t_i = 0$). The dashed lines represent entrepreneurs who face a binding kinship tax constraint ($t_i > 0$). Facing a costly distortion here moves a firm towards the left of the distribution, while facing a subsidy moves it to the right. (a) Kernel density plot of the $\tau_i^y + t_i$ output wedge. This is calculated for each firm using data on output, capital and labor: $1 - \tau_i^y - t_i = \frac{wl}{(1-\alpha)\sigma y_i}$. (b) Kernel density plot of the $\tilde{\tau}_i^k$ capital-labor wedge. This is calculated for each firm using data on output, capital and labor $1 + \tau_i^k = \frac{\alpha}{1-\alpha} \frac{wl}{rk}$.

$$k(A, \tau^y, \tau^k) = \left[\sigma(1 - \tau^y)A \left(\frac{\alpha}{(1 + \tau^k)r} \right)^{1-(1-\alpha)\sigma} \left(\frac{1 - \alpha}{w} \right)^{(1-\alpha)\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (8)$$

$$l(A, \tau^y, \tau^k) = \left[\sigma(1 - \tau^y)A \left(\frac{\alpha}{(1 + \tau^k)r} \right)^{\alpha\sigma} \left(\frac{1 - \alpha}{w} \right)^{1-\alpha\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (9)$$

If we use the vectors $\{A_i, \tau_i^y + t_i, \tilde{\tau}_i^k\}$ backed out from the data in section 5.3, as well as the same prices w and r , then by construction equations 8 and 9 return the observed values of k_i and l_i . Once we modify wedges, prices w and r need to float to keep the sum of capital and labor across entrepreneurs constant. For simplicity, note that many of the terms are common across all entrepreneurs, and hence:

$$k(A, \tau^y, \tau^k) \propto \left(\frac{(1 - \tau^y)A}{(1 + \tau^k)^{1-(1-\alpha)\sigma}} \right)^{\frac{1}{1-\sigma}} \equiv z^k, \quad (10)$$

$$l(A, \tau^y, \tau^k) \propto \left(\frac{(1 - \tau^y)A}{(1 + \tau^k)^{\alpha\sigma}} \right)^{\frac{1}{1-\sigma}} \equiv z^l, \quad (11)$$

where z^k and z^l are proportional to the input allocation of an entrepreneur given some productivity A and set of wedges τ^y and τ^k .⁴³ Given a set of z_i^k 's and z_i^l 's, the actual allocation of inputs for entrepreneurs is equal to the following equations, where $K \equiv \sum_i k_i$ and $L \equiv \sum_i l_i$ are the total amount of capital and labor being reallocated across entrepreneurs:

$$k_i = \frac{z_i^k}{\sum_i z_i^k} K, \quad l_i = \frac{z_i^l}{\sum_i z_i^l} L. \quad (12)$$

This procedure implicitly allows prices w and r to float to satisfy the resource constraint, while ensuring each entrepreneur chooses their optimal k and l given their productivity, wedges, and input prices. The first-best allocation would be to set all wedges to zero, such that $z_i^k = z_i^l = A_i^{\frac{1}{1-\sigma}}$ for all i . With decreasing returns to scale, $\sigma < 1$, this implies a non-degenerate distribution of firm size even in the absence of distortions.

⁴³This notation follows Restuccia and Santaaulalia-Llopis (2015).

5.4.2 Confirm wedges are reasonable

By using the reallocation algorithm described above to remove either all capital-labor wedges or all output wedges, I provide evidence that both the measured wedges and the reallocation procedure produce reasonable results. To do this, I first exploit variation in access to finance between the control and treatment groups of the cash transfer RCT. Recall that $\tilde{\tau}^k$ wedges include credit constraints, among other potential distortions that affect relative use of capital and labor. If credit constraints are important, removing all capital-labor wedges should disproportionately benefit entrepreneurs in the control group. This is because entrepreneurs in the treatment group who faced credit constraints were presumably able to use their cash transfers to increase their capital stock.

Column 1 of Table 12 shows the change in aggregate output from setting all τ^k wedges to zero. What this procedure does is reallocate capital and labor such that the capital-labor ratio becomes equal across all firms, and calculates the change in aggregate output from this reallocation. Row 1 reports an 8% gain in output from this procedure.⁴⁴ Panel C splits the set of 326 entrepreneurs into control and treatment groups. As expected, removing capital-labor wedges disproportionately benefits members of the control group. Their aggregate output goes up by 20%, while the output of the treatment group (who received the cash transfer RCT) only goes up by 5%.

Panel B of Table 12 splits the sample by whether or not the entrepreneur faces a positive kinship tax rate, as measured in the lab experiment. Here, the difference between the two groups is much smaller (7% and 10%). This supports the idea that kinship taxation does not have a large effect on the relative use of capital and labor, and hence does not primarily enter the entrepreneur's problem as a τ^k wedge.

Column 2 of Table 12 reports results from removing all output wedges τ^y , while keeping capital-labor wedges in place. First, we see from the first row that removing these distortions increases aggregate TFP by 69%, which is substantially larger than the gains from removing capital-labor wedges. This suggests that distortions which affect the scale of firms are, in this setting, more important than those which distort their mix of inputs. Second, Panel B shows that, in contrast with the results from removing all capital-labor distortions, removing output wedges disproportionately increases the output of entrepreneurs facing kinship taxation. This suggests that kinship taxation does contribute to output wedges, which supports the decision to model kinship taxation as an output wedge.⁴⁵

⁴⁴Or rather an increase in TFP, since total capital and labor are held constant.

⁴⁵See Appendix C for a discussion of the plausible sources of residual output wedges.

5.4.3 Counterfactual: Kin tax rates set to zero

In this section I estimate the effect on output of setting kinship tax rates to zero. Column 3 of Table 12 reports the change in aggregate TFP from removing the distortionary effect of kinship tax rates, by setting $t_i = 0$ for all entrepreneurs. I find that aggregate output increases by 26%. This is achieved by shifting capital and labor away from entrepreneurs who do not face kinship taxation, and giving it to those who do. Entrepreneurs who face positive kinship tax rates more than double their output, while other entrepreneurs shrink by 65%.

Unlike the procedure of setting all output wedges to zero, the effect of removing only the wedges from kinship taxation (such that $t_i = 0$ for all entrepreneurs) is a fundamentally distortionary exercise. This is because misallocation is caused not by the presence of wedges but their dispersion. So if every entrepreneur has the same output wedge, no matter its value, then it has no distortionary effect on choice of inputs and hence on output. Increasing or decreasing τ wedges for only part of the population will, in general, increase misallocation. However, given some set of pre-existing wedges, it is possible that total output increases after adding further distortions if the new wedges counteract existing ones.

Consider for example an environment where output wedges are largely driven by political connections. People with favorable connections face negative τ^y wedges, which represents their implicit subsidies, and those without connections are effectively taxed and face positive τ^y wedges. The kinship tax distortions could be a mechanism for the community to moderate or counteract these distortions, by taxing entrepreneurs with valuable political connections. This may imply that removing kinship taxation reduces aggregate output rather than increasing it—whether it does or not is an empirical question.

In addition to increasing aggregate output, this counterfactual allocation of inputs from setting kinship tax rates to zero increases the share of capital and labor used by larger firms. Because the entrepreneurs who gain are somewhat more productive, removing their kinship tax distortions increases the concentration of inputs in the hands of the best entrepreneurs. Figure 5 shows the decreasing share of workers in small firms (fewer than five workers) after removing distortions from kinship taxation. Before reallocation the share of workers who are in firms of size five or more is 9.3%. After setting kinship tax rates to zero, this share goes up to 56%.

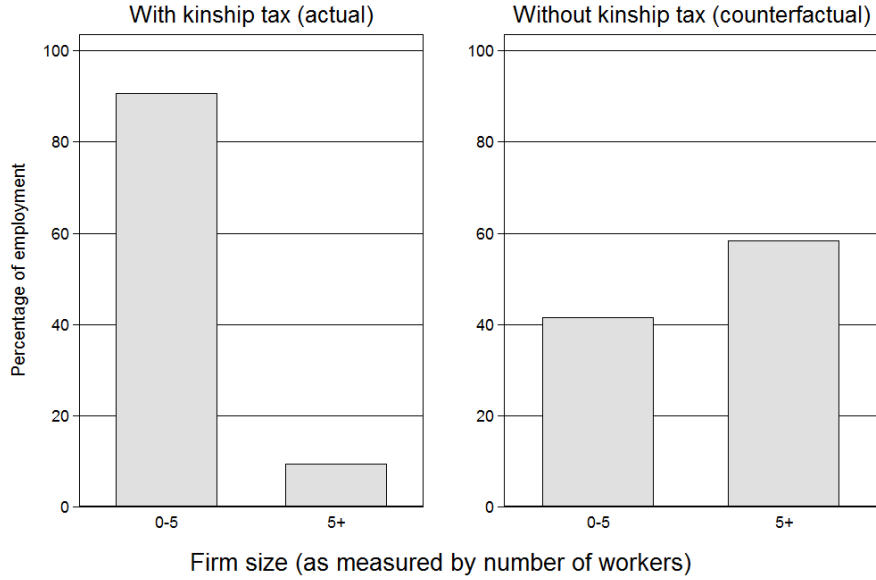
Figure 6 shows a similar increasing concentration of capital stock. The left-hand panel shows that, prior to reallocation, there is little correlation between an entrepreneur's productivity and their capital stock. Given decreasing returns to scale, an efficient allocation of capital is upward-sloping, where more productive firms have more capital. Other studies have found similar patterns across firms in developing countries. The right-hand panel

Table 12: Change in Output from Reallocation: Removing Capital-Labor and Output Wedges, and Kinship Tax Distortions

<i>Name of counterfactual:</i>	(1) No Capital-labor ($k-l$) wedges	(2) No Output wedges	(3) No Kinship tax	(4) No $k-l$ wedges or Kinship tax
<i>Counterfactual output wedge:</i>	$\tau_i^y + t_i$	0	τ_i^y	τ_i^y
<i>Counterfactual capital-labor wedge:</i>	0	$\tilde{\tau}_i^k$	$\tilde{\tau}_i^k$	0
<i>Panel A: Change in agg TFP ($\Delta Y/Y$)</i>				
Entire sample	0.078	0.694	0.265	0.747
<i>Panel B: Split sample by kinship tax</i>				
Kinship constraint does not bind ($t_i = 0$)	0.065	0.478	-0.650	-0.324
Kinship constraint binds ($t_i > 0$)	0.102	1.112	2.038	2.823
<i>Panel C: Split sample by RCT assignment</i>				
Control	0.202	-0.088	-0.044	0.841
Treatment	0.045	0.903	0.348	0.722
N	326	326	326	326

Notes: Each column reports the change in output from a particular counterfactual reallocation procedure. In each case, capital and labor for a given entrepreneur is given by equations 10 and 11, after replacing their observed output wedge ($\tau_i^y + t_i$) and their capital-labor wedge ($\tilde{\tau}_i^k$) with the corresponding counterfactual values. Column (1) reports outcome from removing all capital-labor wedges. That is, capital wedges are set to zero (those distortions are removed), while output wedges are kept as derived in equation 7. Column (2) reports the complementary test: all output wedges are set to zero, while capital-labor wedges are kept as in the equation 6. Column (3) removes distortion from the kinship tax rate t_i , as measured in the lab experiment. Here the procedure is to set the output wedge to be equal to τ_i^y and keep the capital-labor wedge unchanged. Column (4) reports the outcome from removing the kinship tax rate and all capital-labor wedges simultaneously. Panel A reports the percent change in output from each counterfactual reallocation for the entire set of entrepreneurs in the sample. Since total capital and labor are held constant, this is equal to a change in aggregate TFP. Panel B reports the change in output separately for entrepreneurs who have a zero kinship tax rate, and for those who have a strictly positive tax rate. Panel C splits the sample into entrepreneurs who are in the cash transfer RCT control group and those in the treatment group.

Figure 5: Distribution of Employment by Firm Size



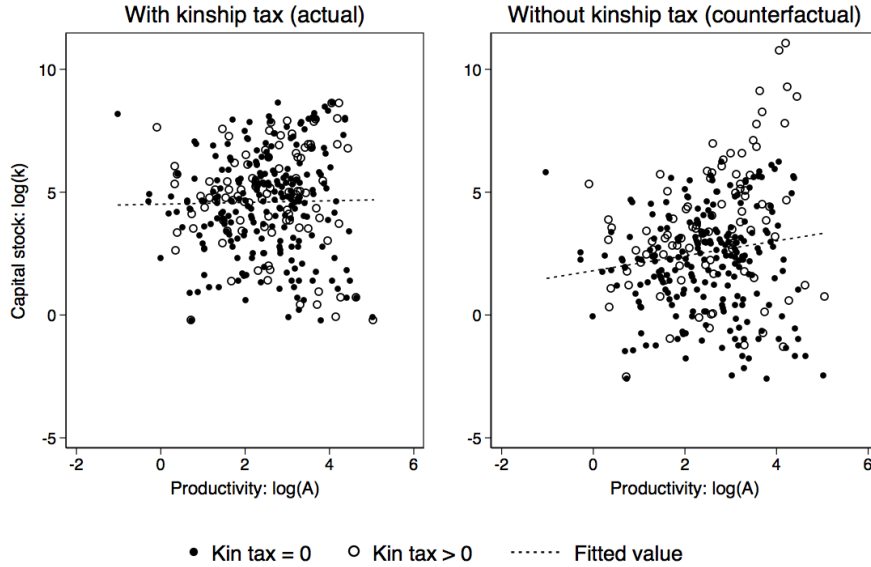
Notes: Number of workers is defined as total labor used by the firm, including the owner’s labor. The ‘0-5’ bars include firms with labor use $l \in [0, 5)$. Firms with five or more workers are in the second category. The left-hand panel uses observed (actual) labor allocations in my sample of microenterprises. The right-hand panel uses the counterfactual labor allocations after setting kinship tax rates to zero for all entrepreneurs. $N = 326$ and $\Sigma l = 535$ in both panels.

shows the allocation of capital after removing the distortion from kinship taxation. This has two effects. The first is that entrepreneurs with positive kinship tax rates (hollow dots) move up, as they now use more capital. Since total capital stock across all entrepreneurs is held constant, this means a reduction of capital used by other entrepreneurs (solid dots). Crucially, the entrepreneurs with high ability, on the right of the graph, show the most dramatic increase in capital stock after setting kinship tax rates to zero. These entrepreneurs now use a significant fraction of the total capital stock, in addition to using a greater share of the labor force as described above.

Recall that this counterfactual does not change the productivity of any firm, which is held fixed as A_i . The gain in output from reallocation towards larger firms comes from the fact that entrepreneurs who face kinship taxation make better use of the reallocated capital and labor than the firms who would otherwise be using them.

There are two ways to interpret the size of the results from this reallocation exercise. The first is that the key outcome of interest is the aggregate TFP of this economy, and that holding capital and labor inputs fixed allows us to measure this directly. This speaks directly to the literature which tries to identify sources of misallocation, in order to explain TFP differences observed across economies. The second interpretation is that the 26% increase in output is

Figure 6: Capital use and Productivity



Notes: Productivity is measured as the log of the A_i , entrepreneur-level TFP. Capital stock values are winsorized at 1% to attenuate the effect of outliers. Solid dots are individual entrepreneurs who do not face a binding kinship tax constraint ($t_i = 0$). Hollow dots are entrepreneurs who do face a binding kinship tax. Dotted lines are OLS linear predictions, using all entrepreneurs. The left-hand panel uses observed (actual) capital data, and the right-hand panel uses counterfactual capital allocations after setting kinship tax rates to zero for all entrepreneurs. $N = 326$ in both panels.

a lower bound on the gains from removing distortions from kinship taxation, since we are not allowing any inputs to flow into this sector of the economy. Increasing efficiency in this sector implies higher wages and returns to capital, which would draw labor and capital from other sectors. If increases in aggregate TFP increase the pool of capital and labor in the microenterprise sector, then the total output in this sector would increase by more than the estimates presented here.

6 Credit constraints

The ability to construct counterfactual outcomes also allows us to explore distortions from credit constraints in relation to those from kinship taxation. In my sample, relaxing credit constraints has a smaller effect on efficiency than removing distortions from kinship taxation. Further, theory would suggest that these two constraints are complementary, since the distortions affect each other multiplicatively in the entrepreneur's problem. That is, removing an output wedge, and thereby increasing the scale of a firm, has a larger effect on output when the capital-labor ratio is undistorted. Indeed, I find evidence that this holds

empirically in my setting. That is, removing both constraints simultaneously has a greater effect than removing both separately. I show this both by using the structural framework from section 5 and by directly using the cash transfer field experiment.

6.1 Evidence from structural model

As with kinship taxation, a direct measure of credit constraints for each entrepreneur would allow us to calculate counterfactual allocations after removing only the wedges from credit constraints while keeping any other wedges unchanged. In the absence of such a measure, the best we can do is to remove all capital-labor wedges, since these include, amongst other distortions, the effect of credit constraints. Doing so, as in Table 12, increases output by only 8% for the sample as a whole. This measure of the gains from removing credit constraints includes both entrepreneurs who received cash transfers as part of the RCT, and those who did not. I interpret the cash transfers as increasing \bar{k}_i , the upper bound on an entrepreneur's capital stock, and hence removing part of the capital-labor wedge.⁴⁶ The control group from the cash transfer RCT may then be a better counterfactual, and their gain from removing capital-labor wedges is a more substantial 20% increase in output.

In column 4 of Table 12, I show the gain in output from simultaneously removing all capital-output distortions (as in column 1) and setting all kinship taxes to zero (as in column 3). That is, the vector of productivity and distortionary wedges that I use for each entrepreneur goes from $\{A_i, \tau_i^y + t_i, \tilde{\tau}_i^k\}$ to $\{A_i, \tau_i^y, 0\}$. Doing both simultaneously increases aggregate TFP by 75%, which I interpret as evidence of complementarity.⁴⁷

6.2 Reduced form

Analysis of the differential impact of cash transfers for entrepreneurs who face kinship taxation or not corroborates the structural results presented above. In this section, I make direct use of the cash transfer RCT to test the interaction of credit constraints and kinship taxation as barriers to firm growth. I focus on a subset of firms which are more likely to be credit constrained, and look at the effect of the cash transfer on their stock of capital. We should expect to find no effect of the cash transfer RCT if entrepreneurs are not credit constrained, since the only effect of the income shock from the cash transfer on the choice

⁴⁶In the model presented in section 2, this is equal to μ_i/u_c from equation 4.

⁴⁷If there were no interaction effect between the 8% gain from removing capital-labor wedges (including credit constraints) and the 26% gain from removing kinship tax wedges, removing both would mean a 36% gain rather than 75%.

of capital stock is through μ_i , the shadow price on the credit constraint in equation 4.⁴⁸

Since I do not have a measure of credit constraints prior to the cash transfer treatment, I rely on a proxy to identify which firms are likely to be constrained. There is substantial evidence that while male-owned microenterprises see sustained improvements in business outcomes, female-owned ones do not. See Kevane and Wydick (2001); De Mel et al. (2008), and Fafchamps et al. (????) for evidence of this phenomenon in Guatemala, Sri Lanka, and Ghana. De Mel et al. (2009a) discuss this finding at greater length.

The finding that male-owned microenterprises invest cash transfers in capital stock, while female owned ones do not, holds in my data. First, Figure 7a shows that while treatment seems to have increased the capital stock of entrepreneurs, the effect is modest at best. This is a cumulative density plot, where the dotted line is mostly to the right of the solid line, indicating entrepreneurs in the treatment group have somewhat higher amounts of capital. Second, figures 7b and 7c show the effect separately for female and male entrepreneurs. I find that cash transfers do not increase the capital stock of female-owned microenterprises, while I do find a substantial effect for male-owned firms. Female entrepreneurs in this setting face constraints to firm growth that are not alleviated by cash transfers alone, or are already operating at their optimal size.

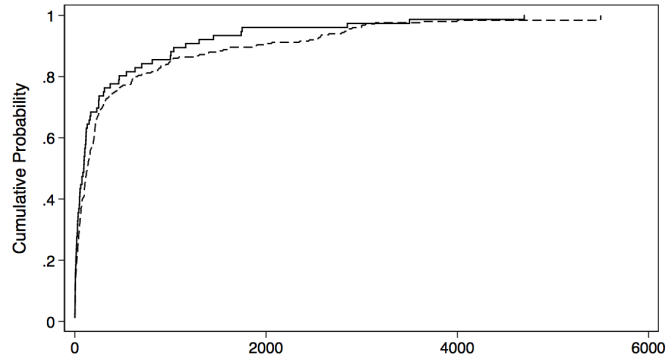
Exploiting this distinction between male and female-owned microenterprises, Table 13 reports the effect of the cash transfer treatment separately for both groups. Column 1 shows that while treatment may have increased capital stock on average, the gain is modest and not statistically significant. For women, the results are small and again insignificant. Further, it does not seem to be the case that the reason women do not invest cash transfers in their microenterprises is that they face kinship taxation, since the coefficient on the interaction between kinship taxation and treatment is also small and not significant. For men (column 3) the results are strikingly large. Male-owned microenterprises more than double their capital stock when they receive cash transfers. However, summing coefficients from row 1 and row 3 shows that, conditional on facing a positive kinship tax rate, the effect of the cash transfer is small and insignificant.

I interpret this as evidence that kinship taxation discourages productive entrepreneurs from investing in their firms. This result supports the finding from the structural analysis that there are substantial interactions between kinship taxation and credit constraints. This suggests that kinship taxation discourages productive entrepreneurs from investing in their firms, and hence providing them with access to finance is likely to have small returns.

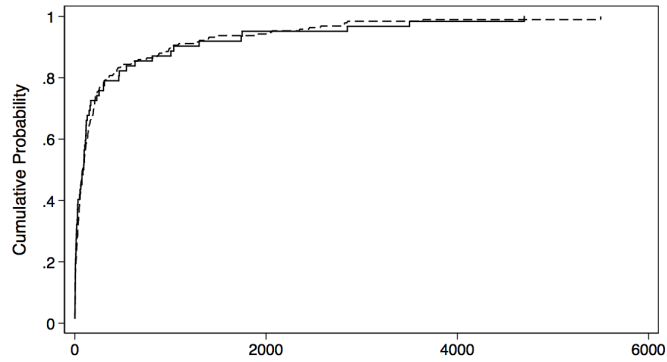
⁴⁸The intuition behind this exercise is closely related to Banerjee and Duflo (2014), who test for credit constraints by measuring whether easier access to credit for a subset of Indian firms causes them to increase their capital stock.

Figure 7: Effect of cash transfer on capital stock of firms

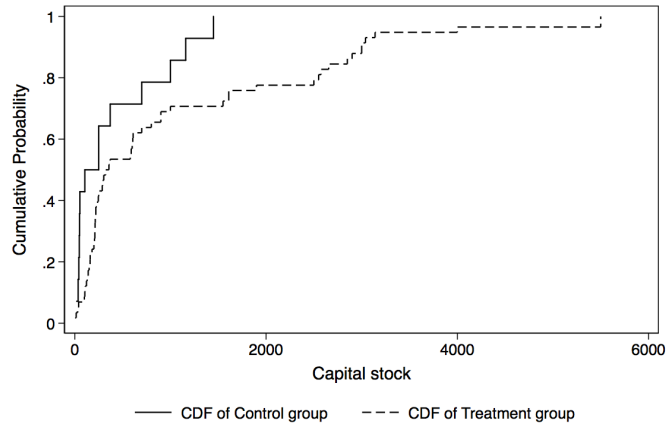
(a) All firms



(b) Female-owned firms



(c) Male-owned firms



Notes: Cumulative density function of capital stock. Capital is measured as the sum of the values of equipment, structures and inventory, in USD, as described in section 5.2. Total values are then winsorized at 1% to attenuate the effect of outliers. The solid lines represent the distribution of capital for entrepreneurs from control group households, who did not receive the cash transfers as part of the RCT. The dashed lines represent those who did receive the cash transfer. A distribution first-order stochastically dominates another if it is everywhere below it. (a) $N = 326$, (b) $N = 254$, (c) $N = 72$.

Table 13: Effect of Cash Transfer Treatment on Capital Stock

	(1)	(2)	(3)
<i>Sample:</i>	All entrepreneurs	Female entrepreneurs only	Male entrepreneurs only
<i>Dependent variable:</i>	Capital (log)	Capital (log)	Capital (log)
Treatment	0.144 (0.283)	-0.122 (0.306)	1.822** (0.822)
Kin Tax > 0	0.0867 (0.485)	0.0987 (0.576)	1.852** (0.889)
Treatment \times (Kin Tax > 0)	0.0827 (0.558)	0.235 (0.660)	-2.580** (1.158)
<i>N</i>	326	254	72

Notes: Robust standard errors in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All columns are OLS regressions, with a dependent variable equal to the (log) capital stock of the entrepreneur.

7 Welfare analysis

While the aim of this paper is to estimate the efficiency cost of kinship taxation in the microenterprise sector, the data also allow speculation on the welfare consequences of kinship taxation. To do so, I compare the observed allocation decisions to a counterfactual where entrepreneurs do not face kinship taxation, but correspondingly do not benefit from the insurance the institution provides.

The strongest assumption I make is that kinship taxation allows for full income pooling, removing all idiosyncratic income risk. Conversely, in the absence of kinship taxation, individuals have no access to alternative forms of consumption smoothing. I consider welfare comparisons between these two states across values of two parameters: risk aversion and income variance. The first is a conventional risk-aversion coefficient in a Constant Relative Risk Aversion (CRRA) utility function, which I take to represent agents' preferences. The second parameter determines the variance of income in autarky. My sample is a cross-section, rather than a panel, which can only provide a measure of the between-person variance of income. In order to calculate the gains from insurance, I need within-person (across time) variance, and hence I use a parameter to scale the variance of the distribution I observe to match the unobserved variance individuals would face without insurance.

As stated above, the welfare calculation I perform compares two stark alternatives. In the first, "income pooling," production values are the observed ones, and I assume that output is shared such that each entrepreneur consumes the mean output. In the second alternative, "autarky," production is as in the counterfactual allocation of inputs without kinship taxation (setting $t_i = 0$ for all i), such that aggregate output is 26% higher than in the data. Here however each person consumes exactly their output. This implies a standard

mean-variance tradeoff, where the “income-pooling” outcome has lower mean and lower variance.

I compare the unweighted sum of CRRA utilities in both scenarios to judge the welfare consequences of kinship taxation,

$$W = \sum_i \frac{c_i^{1-\gamma}}{1-\gamma}.$$

In the “income pooling” scenario, c_i is equal to \bar{y} , that is, the mean observed microenterprise income as measured in the firm survey. In the “autarky” scenario, $c_i = y_i^a$, where y_i^a is equal to the output of a given firm under the counterfactual (reallocated) input decision after setting kinship taxes to zero.

Since my data do not include a measure of risk aversion, I consider values of γ from the literature. Specifically, I use results from Chetty (2006), who considers data from a number of empirical estimations, and finds that the plausible range for γ is 0.15 to 1.78, and the most likely value is 0.71. Since the value of insurance is increasing in risk-aversion, lower measures of γ will make autarky relatively more attractive.

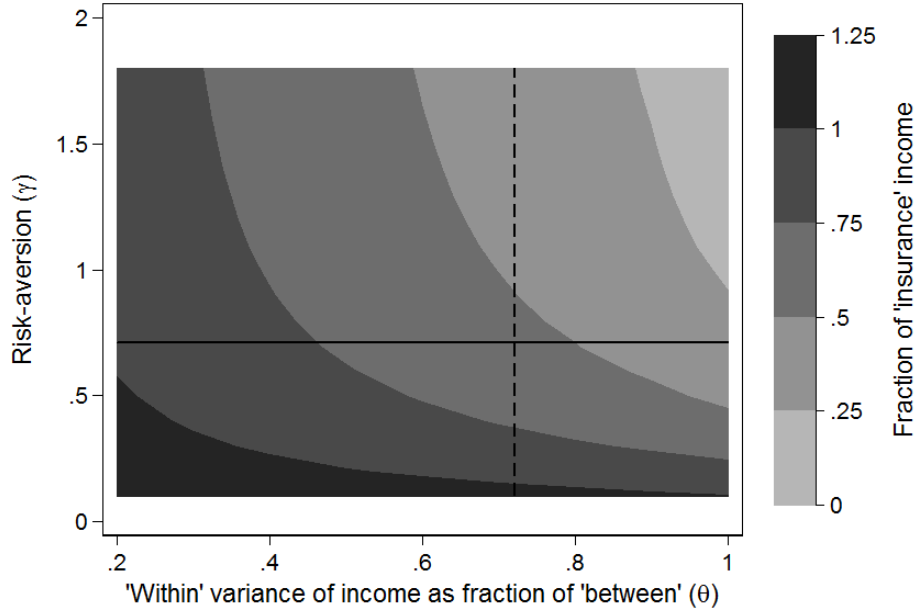
To determine the variance of income in autarky (which is also the variance of consumption), I can use the observed variance of income in my sample. However, this between-person variance of income is almost certainly an overestimate of the within-person variance of income an individual would face if they were to enter autarky. To get a sense of how far we should expect these two variances to be from each other, I use comparable data from De Mel et al. (2008). This sample resembles mine, and the questions I use to elicit firm-level data were taken from theirs. Using these data, which come in eleven waves, I perform a simple decomposition of variance.⁴⁹ I find that in this data, the ratio of between variance to within variance is 0.72. This suggests that the measure of variance in my sample (which is purely ‘between’) is an overestimate of the income variance an entrepreneur would face in autarky (‘within’ variance).

Since it is not clear to what extent this relationship between ‘within’ and ‘between’ variance holds in my data, I present results for welfare under a range of such values. For convenience, I define θ as the ratio of within-person variance to between-person variance, and apply a mean-preserving contraction of the income distribution in my data to match the specified income variance.

Figure 8 reports welfare results under a range of values for γ and θ . The outcome in this graph is the fraction of income the insurance group needs to equalize the welfare of

⁴⁹I winsorize the data at 1%, to make it more comparable to my data, and drop entrepreneurs who have missing profits for one or more waves.

Figure 8: Welfare under income pooling relative to autarky



Notes: Welfare comparison of the sample of entrepreneurs under the ‘insurance’ and ‘autarky’ scenarios. The outcome is the share of observed income which would be given to the insurance group to equalize welfare across the two groups. An outcome of 1 means the two scenarios give equal income, above 1 means welfare is higher under autarky, and below 1 means welfare is higher with insurance. Utility is calculated for each individual using a CRRA utility function: $u_i = c_i^{1-\gamma}/(1-\gamma)$, where the value of γ is on the vertical axis. The horizontal solid line represents the focal value of $\gamma = 0.71$. The horizontal axis represents the scalar used to adjust the variance of income under autarky to adjust for the fact that within-person variance of income is lower than between-person variance, and the data which are used in this analysis only have between-person variance. A value of $\theta = 0$ therefore represents the case of equal between- and within-person variance. For θ below 1, ‘within’ variance is lower than ‘between’. The vertical dashed line represents the focal value of $\theta = 0.72$.

the two scenarios. That is, an outcome of 0.5 means we would need to halve the income of the insurance group to equalize welfare across the two groups. As the graph shows, there is a very limited range of parameters where the autarky group has higher utility than the insurance group (where the outcome is greater than 1). This is when both γ and θ are low, that is, when both the level of risk aversion and the amount of risk in autarky are minimal. At the focal values of $\gamma = 0.71$ and $\theta = 0.72$, the outcome is 56%, implying that under reasonable parameter values, welfare is substantially higher with insurance than in autarky.

This result supports the idea that groups choose to engage in risk-sharing, and it would be hard to sustain such cooperation if the institution were not welfare-improving. These results, however, should be qualified. First, the stark assumption of full versus zero consumption smoothing is unrealistic. Without access to informal insurance, individuals could still save, and might choose income streams that are more diversified. Similarly, even with transfers

among kin, perfect consumption smoothing is unlikely. The scale of the difference in utilities across the two scenarios nonetheless suggests this institution is welfare-improving.

8 Conclusion

This paper documents the cost of kinship taxation among business owners in a poor rural setting. Pressure to share income reduces the productive incentives of some of these entrepreneurs, leading them to invest less in their business than they otherwise would. Specifically, entrepreneurs facing kinship taxation choose inputs as if they have lower productivity than they do. Since this disproportionately affects high-ability entrepreneurs, the losses from misallocation due to this distortion are substantial. Further, gains in efficiency from removing this distortion come largely from increasing the share of inputs going to the most able entrepreneurs. This implies a substantial shift towards an economy where much of the labor force works in larger firms.

The focus on input misallocation across firms is important given that differences in income across countries seem to be in large part driven by differences in productivity rather than factor accumulation.⁵⁰ An emerging literature has shown a significant share of aggregate TFP differences can be attributed to misallocation of inputs across productive units, as in Hsieh and Klenow (2009). While early work in this literature focused on large national surveys of manufacturing plants, a more recent set of papers have studied misallocation across the smaller, less formal productive units that occupy the majority of the labor force in most poor countries. See for example Shenoy (2015) and Restuccia and Santaella-Llopi (2015).

In keeping with this broad literature, I find relatively modest losses in aggregate productivity from credit constraints.⁵¹ However, the losses I find from kinship taxation are comparatively large. This may be because the distribution of kinship tax rates matches almost exactly the conditions laid out in Restuccia and Rogerson (2008) as conducive to factor misallocation: marginal tax rates are both very heterogeneous, and increasing in ability.

The degree to which these findings can be generalized to the broader economy is not yet clear. This depends in part on whether the phenomenon of kinship taxation is more or less important in urban settings than in rural ones. The greater mobility of urban workers and the erosion of traditional identities may lead to a reduction of kinship tax rates, and formal

⁵⁰See Parente and Prescott (1994), Klenow and Rodriguez-Clare (1997), Prescott (1998), Hall and Jones (1999).

⁵¹Most studies find relatively small reductions in aggregate TFP from specific taxes or policy distortions, as described in Restuccia and Rogerson (2013) and Hopenhayn (2014). For example Hopenhayn and Rogerson (1993) find small effects from firing costs, while Midrigan and Xu (2014) find that credit constraints seem to account for very little of the total misallocation among existing manufacturing firms.

substitutes for insurance provided by informal transfers may also be more readily available. On the other hand, whatever distortions there are from kinship taxation among the urban poor may be more costly if they have more opportunities for growth available. Piloting this study in Nairobi, I found that participants living in urban slums faced similar kinship tax rates as the rural sample. This may be because in an urban environment, the perceived benefits from making transfers are relatively small, so what transfers do happen are more distortionary. Arthur Lewis describes this potential problem well:

“[The extended family] is at any time a deterrent to making superior effort, and it is especially so at times when the family concept is narrowing, and the community is passing from wider to narrower recognition, since it is then that men are least likely to accept claims which they would previously have taken for granted.” (Lewis, 1955, p 114)

Measuring marginal kinship tax rates across societies by eliciting willingness-to-pay to hide income would provide a more reliable way of measuring where we would expect this distortion to matter most. It may also shed light on the societal transitions that weaken the grip of the extended family on economic activity.

References

- ACKERBERG, D. A., K. CAVES, AND G. FRAZER (2015): “Identification Properties of Recent Production Function Estimators,” *Econometrica*, 83, 2411–2451.
- ALLINGHAM, M. G. AND A. SANDMO (1972): “Income Tax Evasion: A Theoretical Analysis,” *Journal of Public Economics*, 1, 323–338.
- ANDREONI, J. AND L. VESTERLUND (2001): “Which is the Fair Sex? Gender Differences in Altruism,” *The Quarterly Journal of Economics*, 116, 293–312.
- ATKESON, A. AND P. J. KEHOE (2007): “Modeling the Transition to a New Economy: Lessons from two Technological Revolutions,” *The American Economic Review*, 97, 64–88.
- BALAND, J.-M., I. BONJEAN, C. GUIRKINGER, AND R. ZIPARO (2016): “The Economic Consequences of Mutual Help in Extended Families,” *Journal of Development Economics*, 123, 38–56.

- 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, 1–16.
- BANERJEE, A. V. AND E. DUFLO (2014): “Do Firms Want to Borrow More? Testing Credit Constraints Using a Directed Lending Program,” *The Review of Economic Studies*, 81, 572–607.
- BARR, A., J. BURNS, L. MILLER, AND I. SHAW (2015): “Economic Status and Acknowledgement of Earned Entitlement,” *Journal of Economic Behavior & Organization*, 118, 40–54.
- BASU, S. AND J. G. FERNALD (1997): “Returns to Scale in US Production: Estimates and Implications,” *Journal of Political Economy*, 105, 249–283.
- BECKER, G. M., M. H. DEGROOT, AND J. MARSCHAK (1964): “Measuring Utility by a Single-Response Sequential Method,” *Behavioral Science*, 9, 226–232.
- BOLTZ, M. (2013): “How can I say no? Savings and Forced Solidarity in the Extended Family: Evidence from Senegal,” *Working Paper*.
- BOLTZ, M., K. MARAZYAN, AND P. VILLAR (2015): “Preference for Hidden Income and Redistribution to Kin and Neighbors: A Lab-in-the-field Experiment in Senegal,” *Working Paper*.
- BUERA, F. J., J. P. KABOSKI, AND Y. SHIN (2011): “Finance and Development: A Tale of Two Sectors,” *The American Economic Review*, 101, 1964.
- CAPPELEN, A. W., A. D. HOLE, E. Ø. SØRENSEN, AND B. TUNGODDEN (2007): “The Pluralism of Fairness Ideals: An Experimental Approach,” *The American Economic Review*, 97, 818–827.
- CAPPELEN, A. W., K. O. MOENE, E. Ø. SØRENSEN, AND B. TUNGODDEN (2013): “Needs versus Entitlements: An International Fairness Experiment,” *Journal of the European Economic Association*, 11, 574–598.
- CHARLES, K. K., E. HURST, AND N. ROUSSANOV (2009): “Conspicuous Consumption and Race,” *The Quarterly Journal of Economics*, 124, 425–467.
- CHERRY, T. L., P. FRYKBLOM, AND J. F. SHOGREN (2002): “Hardnose the Dictator,” *The American Economic Review*, 92, 1218–1221.

- CHETTY, R. (2006): “A New Method of Estimating Risk Aversion,” *The American Economic Review*, 96, 1821–1834.
- COMOLA, M. AND M. FAFCHAMPS (2014): “Testing Unilateral and Bilateral Link Formation,” *The Economic Journal*, 124, 954–976.
- COX, D. AND M. FAFCHAMPS (2008): “Extended Family and Kinship Networks: Economic Insights and Evolutionary Directions,” Elsevier, vol. 4 of *Handbook of Development Economics*, chap. 58, 3711–3784.
- DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2008): “Returns to Capital in Microenterprises: Evidence from a Field Experiment,” *The Quarterly Journal of Economics*, 123, 1329–1372.
- (2009a): “Are Women More Credit Constrained? Experimental Evidence on Gender and Microenterprise Returns,” *American Economic Journal: Applied Economics*, 1, 1–32.
- DE MEL, S., D. J. MCKENZIE, AND C. WOODRUFF (2009b): “Measuring Microenterprise Profits: Must we ask how the Sausage is Made?” *Journal of Development Economics*, 88, 19–31.
- DI FALCO, S. AND E. BULTE (2011): “A Dark Side of Social Capital? Kinship, Consumption, and Savings,” *Journal of Development Studies*, 47, 1128–1151.
- DUFLO, E. AND C. UDRY (2004): “Intrahousehold Resource Allocation in Cote d’Ivoire: Social Norms, Separate Accounts and Consumption Choices,” *National Bureau of Economic Research Working Paper*.
- FAFCHAMPS, M. AND S. LUND (2003): “Risk-Sharing Networks in Rural Philippines,” *Journal of Development Economics*, 71, 261–287.
- FAFCHAMPS, M., D. MCKENZIE, S. QUINN, AND C. WOODRUFF (????): “Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana,” *Journal of Development Economics*, 106.
- FARBER, B. (1975): “Bilateral Kinship: Centripetal and Centrifugal Types of Organization,” *Journal of Marriage and the Family*, 37, 871–888.
- FIALA, N. (2017): “Business is Tough, but Family is Worse: Household Bargaining and Investment in Microenterprises in Uganda,” *Working Paper*.

- FOSTER, G. M. (1965): "Peasant Society and the Image of Limited Good," *American Anthropologist*, 67, 293–315.
- GANDHI, A., S. NAVARRO, AND D. RIVERS (2016): "On the Identification of Production Functions: How Heterogeneous is Productivity?" *University of Western Ontario Working Paper*.
- GEERTZ, C. (1963): *Peddlers and Princes: Social Development and Economic Change in two Indonesian Towns*, vol. 318, University of Chicago Press.
- GERSHMAN, B. (2015): "The Economic Origins of the Evil Eye Belief," *Journal of Economic Behavior & Organization*, 110, 119–144.
- GOLDBERG, J. (2017): "The Effect of Social Pressure on Expenditures in Malawi," *Working Paper*.
- GOLLIN, D. (2002): "Getting Income Shares Right," *Journal of Political Economy*, 110, 458–474.
- (2008): "Nobody's Business but my Own: Self-Employment and Small Enterprise in Economic Development," *Journal of Monetary Economics*, 55, 219–233.
- GRIMM, M., R. HARTWIG, AND J. LAY (2016): "Does Forced Solidarity Hamper Investment in Small and Micro Enterprises?" *Journal of Comparative Economics*.
- HADNES, M., B. VOLLAN, AND M. KOSFELD (2011): "The Dark Side of Solidarity: An Experimental Approach to the Impact of Solidarity on Work Output," *University of Frankfurt Working Paper*.
- HARVEY, D. (1993): *Potter Addition: Poverty, Family, and Kinship in a Heartland Community*, Social Institutions and Social Change, Transaction Publishers.
- HOFF, K. AND A. SEN (2006): "The Kin System as a Poverty Trap?" in *Poverty Traps*, ed. by S. Bowles, S. N. Durlauf, and K. Hoff, Princeton University Press.
- HOPENHAYN, H. AND R. ROGERSON (1993): "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, 101, 915–938.
- HOPENHAYN, H. A. (2014): "Firms, Misallocation, and Aggregate Productivity: A Review," *Annual Review of Economics*, 6, 735–770.
- HSIEH, C.-T. AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 124, 1403–1448.

- HYDÉN, G. (1980): *Beyond Ujamaa in Tanzania: Underdevelopment and an Uncaptured Peasantry*, University of California Press.
- JAKIELA, P. (2015): “How Fair Shares Compare: Experimental Evidence from two Cultures,” *Journal of Economic Behavior & Organization*, 118.
- JAKIELA, P. AND O. OZIER (2016): “Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies,” *The Review of Economic Studies*, 83, 231–268.
- KENNEDY, P. T. (1988): *African Capitalism: The Struggle for Ascendancy*, African Society Today, Cambridge University Press.
- KEVANE, M. AND B. WYDICK (2001): “Microenterprise Lending to Female Entrepreneurs: Sacrificing Economic Growth for Poverty Alleviation?” *World Development*, 29, 1225–1236.
- KINDLEBERGER, C. (1965): *Economic Development*, Economics handbook series, McGraw Hill Kogakusha Limited.
- KINNAN, C. (2017): “Distinguishing Barriers to Insurance in Thai Villages,” *Northwestern University, Department of Economics Working Paper*.
- KREMER, M., J. LEE, J. ROBINSON, AND O. ROSTAPSHOVA (2013): “Behavioral Biases and Firm Behavior: Evidence from Kenyan Retail Shops,” *The American Economic Review (Papers and Proceedings)*, 103, 362–368.
- KUROSAKI, T. AND M. FAFCHAMPS (2002): “Insurance Market Efficiency and Crop Choices in Pakistan,” *Journal of Development Economics*, 67, 419–453.
- LEMAY-BOUCHER, P., J. NORET, AND V. SOMVILLE (2012): “Facing Misfortune: Expenditures on Magico-Religious Powers for Cure and Protection in Benin,” *Journal of African Economies*, 22.
- LEVINSOHN, J. AND A. PETRIN (2003): “Estimating Production Functions Using Inputs to Control for Unobservables,” *The Review of Economic Studies*, 70, 317–341.
- LEWIS, M. (1978): *The Culture of Inequality*, A Meridan book, University of Massachusetts Press.
- LEWIS, W. A. (1955): *The Theory of Economic Growth*, Homewood, Illinois: RD Irwin,.

- MALARNEY, S. K. (1998): “State Stigma, Family Prestige, and the Development of Commerce in the Red River Delta of Vietnam,” in *Market Cultures: Society and Morality in the New Asian Capitalisms*, ed. by R. W. Hefner, Westview Press, 268–289.
- MIDRIGAN, V. AND D. Y. XU (2014): “Finance and Misallocation: Evidence from Plant-Level Data,” *American Economic Review*, 104, 422–58.
- MOORE, W. E. (1963): *Social Change*, Prentice Hall.
- MORDUCH, J. J. (1991): *Risk and Welfare in Developing Countries*, Harvard University.
- OLLEY, S. AND A. PAKES (1996): “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 64, 1263–97.
- PLATTEAU, J.-P. (2000): *Institutions, Social Norms, and Economic Development*, vol. 1, Psychology Press.
- RESTUCCIA, D. AND R. ROGERSON (2008): “Policy Distortions and Aggregate Productivity with Heterogeneous Establishments,” *Review of Economic Dynamics*, 11, 707–720.
- (2013): “Misallocation and Productivity,” *Review of Economic Dynamics*, 16, 1–10.
- RESTUCCIA, D. AND R. SANTAELALIA-LLOPIS (2015): “Land Misallocation and Productivity,” *University of Toronto, Department of Economics Working Paper*.
- ROSENZWEIG, M. R. AND O. STARK (1989): “Consumption Smoothing, Migration, and Marriage: Evidence from Rural India,” *The Journal of Political Economy*, 97, 905–926.
- SERRA, R. (2009): “Child Fostering in Africa: When Labor and Schooling Motives may Coexist,” *Journal of Development Economics*, 88, 157–170.
- SHENOY, A. (2015): “Market Failures and Misallocation,” *Working Paper*.
- SLEMROD, J. (2001): “A General Model of the Behavioral Response to Taxation,” *International Tax and Public Finance*, 8, 119–128.
- STACK, C. B. (1974): *All our Kin: Strategies for Survival in a Black Community*, Harper colophon books, Harper & Row.
- SZANTON, D. L. (1998): “Contingent Moralities: Social and Economic Investment in a Philippine Fishing Town,” in *Market Cultures: Society and Morality in the New Asian Capitalisms*, ed. by R. W. Hefner, Westview Press.

- TOWNSEND, R. M. (1994): "Risk and Insurance in Village India," *Econometrica*, 62, 539–591.
- UDRY, C. (1996): "Gender, Agricultural Production, and the Theory of the Household," *Journal of Political Economy*, 104, 1010–1046.
- VALENTINYI, A. AND B. HERRENDORF (2008): "Measuring Factor Income Shares at the Sectoral Level," *Review of Economic Dynamics*, 11, 820–835.
- WANTCHEKON, L., M. KLAŠNJA, AND N. NOVTA (2014): "Education and Human Capital Externalities: Evidence from Colonial Benin," *The Quarterly Journal of Economics*, 130, 703–757.
- WEBER, M. (1915): *The Religion of China, Confucianism and Taoism*, Free Press, english translation 1951, by H H Gerth.
- WOLF, C. (1955): "Institutions and Economic Development," *The American Economic Review*, 45, 867–883.
- YAMEY, B. S. AND P. T. BAUER (1957): *The Economics of Under-Developed Countries*, Cambridge University Press.

Appendix A Robustness

This appendix confirms the robustness of the counterfactual result which removes the distortion from kinship taxation. In the baseline specification, as described in section 5.4.3, the gain from removing distortions from kinship taxation is a 26.5% increase in aggregate efficiency.

A.1 Estimate of gains from reallocation

A.1.1 Outliers

The calculation of gains from counterfactual allocation of inputs is potentially sensitive to outliers. For example with returns to scale sufficiently close to 1, only the most productive firm would have substantial output. To deal with this, all results above are found after winsorization of the top and bottom 1% values of capital, labor and profits. Indeed, without any winsorization, gains from removing kinship taxation explode to more than 100%.

To test whether results are driven by extreme values despite 1% winsorization, I present results with more aggressive attenuation of potential outliers. With 2% winsorization, the gains from removing kinship taxation increase slightly to 27%. With 5% winsorization they drop to 13%, much smaller but still substantial. An alternative to winsorization, which reduces sample size, is to drop the top and bottom 1% values of capital, labor and profits. Doing so means dropping 4% of observations (lower than the maximum of 6% because of some overlap). After dropping these observations, I find gains of 24% from removing kinship taxation with the remaining 314 entrepreneurs.

A.1.2 Efficiency units of labor

The measure of labor used in this analysis does not account for human capital. This is fundamentally hard to measure in an environment with little formal education or training. However, I do have years of education of each entrepreneur, which I can use to ‘augment’ the measure of labor use. To test this, I let each year of education increase efficiency units of labor by 10%, a focal benchmark in the literature on returns to education. By doing so I let labor supplied by entrepreneurs to their own business be $l' = 1.1^h$, where h is years of education. The results following this transformation are almost unchanged. The gains from removing kinship taxation decrease from 26% to 25%. Going even further and assuming education increases effective units of labor by 20% moves estimates of gain back up to 27%.

A.1.3 Firms with more than one owner

Of the 326 firms used in the analysis, 50 have more than one active owner (that is, who participates in the working of the firm). Recall however that profits are measured as the total firm profits, and not those accruing to the person being interviewed. Since kinship tax rates are heterogeneous at the individual level, it is not clear how they affect input choices for firms with multiple owners. To test whether this influences my results, I measure the gains from reallocation after dropping all 50 firms with more than one owner. I find that the gain goes up from 27% to 29%.

A.1.4 High tax rates

The counterfactual exercise which estimates the loss from kinship taxation uses the tax rates as measured in the lab experiment. A substantial fraction of entrepreneurs are willing to accept \$1.50 in secret rather than \$5 in public, which implies a 70% tax rate. The estimated gains from removing kinship taxation do not rely on these high tax rates. To show this, I consider assigning a maximum tax rate of \bar{t} . That is, for this exercise, $t'_i = \min\{t_i, \bar{t}\}$. I find that gains from reallocation are robust to setting $\bar{t} = 0.5$, which reduces gains from 27% to 26%.

Conversely, recall that the willingness-to-pay as measured in the lab experiment is censored above 70%. The bunching of density mass at 70% in the data suggests a large fraction of these participants would have been willing to pay more than 70% to hide income. Given this, we can construct an alternative counterfactual where participant with a measured kinship tax rate of 70% all have a true tax rate of 90%. Doing so changes the gains from reallocation to 25%.

A.1.5 Drop inconsistent choices

Recall that some participants made inconsistent (dominated) choices in the lab experiment. As discussed in section 4, the counterfactual analysis uses the lower-bound of possible imputed tax rates for these participants. Instead, however, we can exclude them from this analysis. Of the 326 firms in the sample, 4 are owned by participants who made inconsistent choices in the lab experiment, and are hence dropped in this robustness check. Dropping these 4 firms reduces gains from reallocation very little, from 27% to 26%.

A.1.6 Alternative imputation of tax rates

The tax rate assigned to each participant is the lower-bound of the range of rates consistent with their eight binary decisions in the lab experiment, as described in section 4. That is, if

someone chooses to pay to hide when the cost was 20% but is not willing to pay at a cost of 30% or more, they are assigned a tax rate of 20%. However, any tax rate within the range $[0.2, 0.3]$ is consistent with their decision. I consider the gains from reallocation using either the midpoint (in this case, a 25% tax rate) or the upper bound (30%).

When imputing tax rates to be equal to the midpoint of the range, I assign a tax of 80% to those who are willing to pay at least 70% to hide. This increases the gains from reallocation to 27%. Instead, imputing tax rates to be the upper-bound of the consistent range, and using a tax of 90% for those with the maximum willingness-to-pay to hide, also results in a 27% gain from reallocation.

A.2 Sensitivity to span parameter (σ)

As mentioned in section 5.3, the benchmark structural results are done using a span-of-control parameter σ equal to 0.7, and a capital share α equal to 0.3. These are taken not from my data but instead I see them as broadly representative of the literature.

A.2.1 Using out-of-sample parameters

One of the key challenges in choosing parameters for the production function is that the wedges we are interested in will bias estimation of these parameters. For example, firms with positive or negative capital-labor wedges τ^k will have factor shares which are not representative of the underlying technology they use. The usual response to this challenge is to use a relatively undistorted counterfactual setting where estimation will be less biased. Notably, Hsieh and Klenow (2009) use parameters taken from US data to estimate misallocation across Indian and Chinese firms. One significant advantage in this case is that they can match manufacturing industries across countries, and use industry-specific US factor shares to estimate productivity and wedges.

In my context there are no obvious US firms which are equivalent to the ones in my sample in terms of production technology. For my benchmark results, I use factors shares which are representative of estimates from the literature on national accounts more broadly, with a capital share of 0.3 and hence a labor share of 0.7. While factor shares derived from national accounts often find higher capital shares for developing countries, (Gollin, 2002) argues this is due to systematic mismeasurement. As a sanity check, I find that naive OLS estimation using my sample of firms give values for the capital share which range from 0.25 to 0.35, depending on the span parameter chosen.

Finally, in section A.3 I consider an alternate version of my model where firms use only capital in production, and I abstract from labor use. This eliminates the need to choose a

factor share parameter, and provides results which are consistent with my benchmark choice of parameters.

A.2.2 Cross-sector variation

Throughout the analysis, I use a single common capital share for all firms, rather than using industry-specific shares. In contrast, when calculating the amount of misallocation in Indian and Chinese manufacturing plants, Hsieh and Klenow (2009) use industry-specific factor shares taken from US data.

In my sample it is almost impossible to match firms to their sectoral equivalents in the US. Unlike in the case of more formal manufacturing firms, the small owner-operated microenterprises of interest do not have obvious counterparts in the US in terms of production functions. Reassuringly, Valentinyi and Herrendorf (2008) find that differences across (non-agricultural) sectors in the US are small, and capital shares vary from 0.2 to 0.4. For my purposes, varying factor shares are hard to reconcile with heterogeneity in $\tilde{\tau}_i^k$. Individuals with high values of $\tilde{\tau}_i^k$ would choose sectors with low capital intensity, so I would underestimate the degree of distortions.

A.2.3 Span parameter

Values of the span parameter σ used in the literature vary widely, from 0.5 to 0.9. Midrigan and Xu (2014) use 0.85, as do Basu and Fernald (1997) and Atkeson and Kehoe (2007). Buera et al. (2011) use 0.79, while Gollin (2008) uses 0.9. I therefore think of my choice of $\sigma = 0.7$ as conservative, given that larger values of σ increase gains from reallocation, while smaller values reduce it. In the limit, if $\sigma = 1$, output is maximized by giving all inputs to the single most productive firm, while if $\sigma = 0$ then the allocation of inputs across firms has no effect on aggregate output.

Using Hsieh and Klenow (2009) as a benchmark is somewhat complicated by the fact that they use a model of monopolistic competition with constant returns to scale, instead of perfect competition and decreasing returns to scale. Importantly, these two models are isomorphic given the procedure used to calculate wedges and measure the costs of misallocation. Their elasticity of substitution parameter of 3 is equivalent in my model to a span of control parameter of $\sigma = 2/3$. They consider in their robustness check an alternative elasticity parameter of 5, which is equivalent to a span parameter of $\sigma = 4/5$.

A.2.4 Alternative parameters

To provide an alternative set of parameters, I conduct a simple empirical exercise using a panel of micro enterprises in another developing country. Publicly available data from the Sri Lanka Microenterprise Survey (De Mel et al., 2008) allow for a simple empirical exercise to estimate the span and factor share parameters in a setting similar to mine. Usefully, the questionnaire used in my study is a shortened version of theirs, so the questions for capital, labor and value-added are comparable. Also, they conduct eight survey rounds, which allow for using fixed effects to estimate firm-level productivity. The major caveat to this exercise is that parameter estimates from microenterprises in this Sri Lanka sample likely face substantial distortions, like the ones in my sample.

In this sample, a random subset of firms are given substantial transfers to allow them to increase their capital stock.⁵² The purpose of these transfers is to relax credit constraints, and since these distortions are an obvious source of bias in estimating the production function parameters (especially the factor shares), I restrict my analysis here to the firms that are in the treatment group and exclude the control group.

I run an OLS regression to estimate the same Cobb-Douglas production function as in my setup, and find that the coefficients on log of capital and of labor are 0.42 and 0.39, respectively. This is equivalent to capital share $\alpha = 0.52$ and span $\sigma = 0.81$. When I do the reallocation procedure from section 5.4 using these parameters, I find that the gain from reallocation is 22.4%, rather than the 26.5% in the benchmark case.

A.3 Single input production function

As an alternative to the reallocation procedure I perform in section 5, I consider here a similar procedure using only capital as an input, thereby abstracting from labor use.

One benefit of this approach is that it negates the need to choose a value for the capital share in this economy (α). Estimating the productivity term A_i now requires only one parametric assumption (given our decreasing-returns-to-scale Cobb-Douglas production function).

$$y_i = A_i k_i^\sigma. \tag{13}$$

Only one wedge can be uniquely identified given a single input. Following Hsieh and

⁵²Some of these transfers are in cash, and others I kind. Transfers also differ in size: some are equivalent to \$100 USD, and others \$200. I do not distinguish between these various treatments.

Klenow (2009), I model the one-input maximization problem as:

$$\max_{c,d,k} u(c,d), \quad (14)$$

s.t.

$$c + d = (1 - \tau_i^y)A_i f(k) - rk, \quad (15)$$

$$d \geq T_i(y). \quad (16)$$

The $\tau_i^y + t_i$ wedge can now be calculated using only k_i and A_i (which itself comes from 13).

$$1 - \tau_i^y - t_i = \frac{r}{A_i \sigma} k_i^{1-\sigma}. \quad (17)$$

Given this new set of $\tau_i^y + t_i$ wedges, the counterfactual without kinship taxation can now be calculated.

$$k(A, \tau^y) \propto ((1 - \tau^y)A)^{\frac{1}{1-\sigma}} \equiv z^k, \quad (18)$$

This procedure, using the same fraction of returns going to capital (which means that in this case $\sigma = 0.21$, equivalent to $\alpha\sigma = 0.3 \times 0.7$) leads to substantially lower gains from reallocation. The increase in TFP is now 5%, rather than the 26% found in the benchmark case with both capital and labor. However, this choice of parameter implies that the returns to scale are now 0.21, or equivalently that the share of output going to profits is 79%, which seems high. An alternative value of σ comes from simply keeping the returns to scale constant at 0.7, as in the two-input case. In that case the gains from reallocation go up to 321%, much larger than in the benchmark model. In between these two parameter estimates, to match the gains from reallocation of the two-input model, the corresponding return to scale in the single-input model is $\sigma = 0.36$.

Appendix B Balance of CT participants

CT participants were not chosen to be representative of their villages, as discussed above. Without data on non-participants, I cannot know exactly how important these differences are. However, within each village the share of households with a CT participant is between

x and y%.

What my data can speak to is the selection of CT participants into the lab-in-the-field experiment. All data in Table 14 are household characteristics from the baseline survey of the cash transfer study. Table 14 shows that people who participated in the lab-in-the-field experiment do indeed differ systematically from the CT participants who did not.

CT participants who chose to take part of the lab experiment come from households which are less likely to be female-headed, and are larger. Their primary source of household income is more likely to be a household-owned business, wages from casual labor, or farming. It is less likely to be from livestock holdings. This is consistent with what people told us when we asked why certain CT participants could not be found. The most common response was that they were grazing their livestock away from the village. People whose income comes from a business, from casual labor or from farming are presumably less likely to be in this situation.

Household income seems to be comparable, though participants have slightly more assets. This is also consistent with the livestock explanation: three of the six possible assets are only useful for people whose income does not come from livestock. They also live closer to their local school, which is reasonable given that the lab experiments were carried out in said schools, and a not-insignificant number of households live a fair distance from their school.

Appendix C Residual output wedge distortions

The large gains in column 2 from setting all output wedges to zero merits further discussion. While kinship taxes account for 38% of this gain,⁵³ more than half comes from other sources. This section discusses what could be causing these residual output wedges.

One key source of wedges in the literature on misallocation is government policy which distorts incentives. Progressive (formal) taxes or scale-dependent regulations, for example, are likely to cause misallocation. In this context, however, firms are largely outside the reach of government, as none pay direct formal taxes. However, 14% of the firms in my sample are registered with the county government. This registration alerts the government to their presence, but (at least in practice) does not carry fiscal or regulatory responsibilities. It may however increase the extent to which government officials interact with or pressure the firm and its owner. Indeed, I find that registered firms are more likely to face positive output wedges (53% versus 22% for firms that are not registered). Of course large firms are also far more likely to be registered, and so this difference may be related to other firm characteristics.

⁵³See section 5.4.3 for more details on how this is calculated.

Table 14: Household characteristics of CT participants

	(1) Non-participant (lab)	(2) Participant (lab)	(3) <i>p</i> -value of diff.
Age of beneficiary	43.347	39.506	0.000
Female beneficiary	0.739	0.790	0.014
Female head of household	0.356	0.316	0.089
Total household members	6.125	6.426	0.004
HH members above 16	2.552	2.493	0.304
Somali	1.000	0.996	0.173
Income source: Business	0.105	0.154	0.006
Income source: Livestock	0.683	0.630	0.029
Income source: Casual labour	0.115	0.151	0.043
Income source: Farming	0.006	0.018	0.048
Household income (monthly)	66.95	68.19	0.582
Education expenditure	4.547	4.664	0.748
Number of assets	0.962	1.145	0.001
Distance from school (hrs)	0.377	0.327	0.000
Safe walk to school	0.941	0.954	0.240
<i>N</i>	505	1805	
Proportion	0.219	0.781	

Notes: All data from baseline survey of the cash transfer RCT, mid-2013, before being randomized into control or treatment. Column 1 includes CT participants who did not participate in the lab experiment. Column 2 are all those CT participants who did take part in the lab experiment. ‘Beneficiary’ refers to the household member assigned to receive the cash transfers (if randomized into treatment group). Each household could only have one beneficiary. Income source values are 1 if the category is the main source of income. (‘Other’ is excluded.) Household income and expenditure are in USD, converted at the rate of 100 KES to 1 USD. Number of assets is from a list of six potential assets: mobile phone, radio, wheelbarrow, plough, donkey cart, and bicycle. ‘Distance from school’ is number of hours walking to nearest primary school. ‘Safe walk to school’ is equal to 1 if the respondent considers it safe for children to walk to school alone. ‘Somali’ is equal to 1 if the participant is Somali, 0 otherwise.

Other potential sources of output wedges in this setting include gender discrimination, political connections, safety from theft or predation, optimization errors, and heterogeneity in risk aversion and impatience. Discrimination may for example bar some entrepreneurs from using certain types of equipment, or engaging in certain types of economic activities. It is common for women, for example, to be restricted to a limited set of business types. Fear of theft would lead to consequences very similar to kinship taxation, and also show up as output wedges. Indeed they are conceptually very similar—both imply that some fraction of output (in expectation) will be taken from the entrepreneur. Optimization errors as in Kremer et al. (2013), which could result in either positive or negative output wedges, are likely common in a setting with widespread illiteracy and very limited record-keeping. Under non-separability of the firm and the household’s problem because of imperfect financial markets, both risk aversion and impatience affect the firm’s choice of inputs. Heterogeneity in these individual characteristics would also show up as output wedges.

Some part of the residual output wedges, however, probably does not reflect underlying economic distortions. One issue is that, in the absence of panel data, this model abstracts from uncertainty. In a world with productivity shocks, inputs may be chosen optimally *ex ante*, without being optimal *ex post*. Since we can only infer *ex post* productivity from the available data, any shocks which were realized after inputs were chosen will contribute to the residual dispersion of output wedges.

Another source of output wedge dispersion is likely to be measurement error. The informal firms in this sample do not keep records on sales or expenditures, which may make this problem even more important than for typical studies in the misallocation literature that use data from formal manufacturing firms. For example, Hsieh and Klenow (2009) use the US level of misallocation in manufacturing firms as a benchmark, and consider their reallocation exercise as one where Indian and Chinese firms match the US level of misallocation. Using zero misallocation as a benchmark of efficiency would be misleading. I present the results from column 2 in table 12 because achieving a 69% gain in aggregate efficiency is realistic, but simply as an exercise to contrast output wedges (τ^y) from capital-labor wedges (τ^k), and to show that kinship taxation accounts for a large share of what I think of as an upper bound on the costs of output wedges. What I do need is that the measurement error in the relevant firm statistics are not systematically different for firms which face kinship taxation from those that do not. If this were the case, it would presumably take the form of high kinship tax firms understating their profits and overstating their use of inputs. This would systematically lead me to underestimate the cost of kinship taxation.

Unfortunately my data are not sufficient to quantify the importance of each of these potential sources of distortions from output wedges. The only wedge I can directly identify

is from kinship taxation, where I use responses from the lab experiment to illuminate the results from this analysis of firm-level characteristics.

Appendix D Experimental instructions

Text in italics are instructions for the enumerator that were not read aloud. Bold text indicates skipping logic.

As a reward for answering my questions and cleaning the beans, I will now give you two chances to receive some money. The first time you might receive airtime or some coins, the second time you might receive a few hundred shillings.

D.1 Airtime

Read exactly from script. Do not add information that is not in script unless the participant asks you a question.

You will play a Lottery where you may win some Safaricom airtime or some coins. As you can see, here are some coloured pencils. Soon we will do a lottery. I will ask you to decide on one colour then put all of them in this basket and choose one without looking. If you draw same colour as you decided, you will win the 50 Sh airtime card or coins as a prize. If you draw any other colour you win nothing. But first, we will play the Ticket Game to see whether the prize of the lottery will be coins or the airtime card.

So, we will do two games. First we will play the Ticket Game to decide whether the prize is coins or airtime. Second we will do a Lottery with the coloured pencils to see whether you win the prize

D.1.1 Ticket game

Here is what we will do for the Ticket Game: As you can see, here are eight (8) tickets with different numbers on them. Each ticket gives you a choice between the airtime card and some value of coins you might win, for example 40 Sh. For each ticket, you must tell me whether you prefer receiving the value on the ticket in cash, or the airtime card. After you make this decision for each of the tickets, I will put each ticket in this basket and ask you to draw one without looking.

The prize in the lottery will be either the airtime card or the value in cash on the ticket you pick, depending on the decision you made for that ticket. You cannot change your choice after you draw from the basket your decisions are final. The best strategy for you is to make

your decision for each ticket based on whether you really would prefer airtime or the cash, without thinking about the other tickets. Do you have any questions?

One of the following two sets of questions was chosen randomly by the data entry software.

Alternative 1 *Start with the 50 Sh ticket, and show it to the participant*

- If the prize was either this 50 Sh airtime card or 50 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 45 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 40 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 35 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 30 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 25 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 20 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 15 Sh coins, which would you choose?

Alternative 2 *Start with the 15 Sh ticket, and show it to the participant*

- If the prize was either this 50 Sh airtime card or 15 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 20 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 25 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 30 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 35 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 40 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 45 Sh coins, which would you choose?
- If the prize was either this 50 Sh airtime card or 50 Sh coins, which would you choose?

If the data entry software detects an inconsistent decision: One of your decisions does not make sense, and means that your participant prefers a smaller amount of cash to a larger amount of cash. *Go back and confirm each of the eight decisions with the participant.*

If the data entry software detects decisions are still inconsistent: *You have not fixed the mistake. Confirm that you have asked the questions a second time and you are sure that this is what the participant wants.*

If participant chose 15 Sh coins instead of 50 Sh airtime card: Why did you decide 15 Cash instead of the Airtime card?

If participant chose 40 Sh coins instead of 50 Sh airtime card, but not 15 Sh coins instead of 50 Sh airtime card: Why did you decide 40 Cash instead of the Airtime card?

If participant chose 50 Sh airtime card instead of 40 Sh coins: Why did you decide the Airtime card instead of 40 Cash?

Now you will pick one ticket from the basket. Please draw one ticket, quickly and without looking.

Which ticket did they pick from the basket? According to the eight decisions you selected, their choice was [...] for the ticket you drew.

Based on the ticket you drew, the prize could either be [...] Sh in coins, or the airtime card. Earlier you decided for this ticket you would rather have the [coins/airtime] and that is what we will use. You cannot change your decision now.

D.1.2 Lottery

Now we have completed the Ticket Game, and are ready to do the lottery. Here is what we will do for the Lottery: As you can see, here are eleven coloured pencils. They will all go in this basket. These six pencils have bright colours, and you can decide one of these to be the winning pencil.

If chose airtime over coins for ticket that was drawn: Remember that during the Ticket Game, you decided you would rather get the 50 Sh airtime card than [...] Sh. If you draw the colour you decided, you will win the 50 Sh airtime card as a prize. If you draw another colour you win nothing.

If chose coins over airtime for ticket that was drawn: Remember that during the Ticket Game, you decided you would rather get the [...] Sh coins than the 50 Sh airtime card. If you draw the colour you decided, you will win [...] Sh coins as a prize. If you draw another colour you win nothing.

Now, choose one of the six bright colour pencils. You cannot change your mind later. *Add the 11 pencils to the basket, shake it, and ask the participant to draw only one. Did they win or lose the lottery?*

If drew the winning colour, and prize was airtime: *Give them the prize, which is an airtime card*

If drew the winning colour, and prize was coins: *Give them the prize, which is [...] Sh in coins*

We are going to play this game again, with a larger prize. It is important you understand how the game is played. Do you have any questions about the game?

D.2 Hiding

Read exactly from script. Do not add information that is not in script unless the participant asks you a question.

You will play a Lottery where you may win some money. Everything we do now will be the same as we just did, but with different prizes. Soon we will do another lottery. Again you will choose one colour then choose one from the basket without looking. If you draw same colour as you decided, you will receive up to 500 Sh and this will be announced to the other participants in your group. If you draw any other colour you win nothing. But first, I would like to offer you the chance to receive money without announcing it, using the Ticket Game.

So, we will do two games. First we will play the Ticket Game to decide whether you might receive money that is announced or money that is not announced. Second we will do a Lottery with the coloured pencils to see whether you win the prize.

D.2.1 Ticket game

Here is what we will do for the Ticket Game: As you can see, here are eight (8) tickets with different numbers on them. Each ticket gives you a choice between the 500 Sh ANNOUNCED and some money NOT ANNOUNCED, for example 400 Sh. After you make this decision for each of the tickets, I will put each ticket in this basket and ask you to draw one without looking.

The prize in the lottery will be either 500 Sh ANNOUNCED or the amount NOT ANNOUNCED on the ticket you pick, depending on the decision you made for that ticket. You cannot change your choice after you draw from the basket—your decisions are final. The best strategy for you is to make your decision for each ticket based on whether you really would prefer ANNOUNCED or NOT ANNOUNCED, without thinking about the other tickets. Do you have any questions?

One of the following two sets of questions was chosen randomly by the data entry software.

Alternative 1 *Start with the 500 Sh ticket, and show it to the participant*

- If the prize was either 500 Sh ANNOUNCED or 500 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 450 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 400 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 350 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 300 Sh NOT ANNOUNCED, which would you choose?

- If the prize was either 500 Sh ANNOUNCED or 250 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 200 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 150 Sh NOT ANNOUNCED, which would you choose?

Alternative 2 *Start with the 150 Sh ticket, and show it to the participant*

- If the prize was either 500 Sh ANNOUNCED or 150 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 200 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 250 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 300 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 350 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 400 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 450 Sh NOT ANNOUNCED, which would you choose?
- If the prize was either 500 Sh ANNOUNCED or 500 Sh NOT ANNOUNCED, which would you choose?

If the data entry software detects an inconsistent decision: One of your decisions does not make sense, and means that your participant prefers a smaller amount of NOT ANNOUNCED money to a larger amount of NOT ANNOUNCED money. *Go back and confirm each of the eight decisions with the participant.*

If the data entry software detects decisions are still inconsistent: *You have not fixed the mistake. Confirm that you have asked the questions a second time and you are sure that this is what the participant wants.*

If participant chose 150 Sh NOT ANNOUNCED instead of 500 Sh ANNOUNCED: Why did you decide 150 Sh NOT ANNOUNCED instead of 500 Sh ANNOUNCED?

If participant chose 400 Sh NOT ANNOUNCED instead of 500 Sh ANNOUNCED, but not 150 Sh NOT ANNOUNCED instead of 500 Sh ANNOUNCED: Why did you decide 400 Sh NOT ANNOUNCED instead of 500 Sh ANNOUNCED?

If participant chose 500 Sh ANNOUNCED instead of 400 Sh NOT ANNOUNCED: Why did you decide 500 Sh ANNOUNCED instead of 400 Sh NOT ANNOUNCED?

Now you will pick one ticket from the basket. Please draw one ticket, quickly and without looking.

Which ticket did they pick from the basket? According to the eight decisions you selected, their choice was [...] for the ticket you drew.

Based on the ticket you drew, the prize could either be [...] not announced, or [...] announced. Earlier you decided for this ticket you would rather have the [*announced/notannounced*] and that is what we will use. You cannot change your decision now.

D.2.2 Lottery

Now we have completed the Ticket Game, and are ready to do the lottery. Here is what we will do for the Lottery: As before, here are eleven coloured pencils. They will all go in this basket. These six pencils have bright colours, and you can decide one of these to be the winning pencil.

If chose 500 Sh ANNOUNCED over NOT ANNOUNCED for ticket that was drawn: If you draw the colour you decided, you will win the 500 Sh ANNOUNCED as a prize. If you draw another colour you win nothing. Remember that during the Ticket Game, you decided you would rather get the 500 Sh ANNOUNCED than [...] Sh NOT ANNOUNCED.

If chose NOT ANNOUNCED over 500 Sh ANNOUNCED for ticket that was drawn: If you draw the colour you decided, you will win the [...] NOT ANNOUNCED as a prize. If you draw another colour you win nothing. Remember that during the Ticket Game, you decided you would rather get the [...] Sh NOT ANNOUNCED than the 500 Sh ANNOUNCED.

If someone in this group were to receive money without us announcing it and they did not tell anyone, would others find out? **If answer yes:** Why do you think it would not stay secret?

Now, choose one of the six bright colour pencils. You cannot change your mind later. *Add the 11 pencils to the basket, shake it, and ask the participant to draw only one. Did they win or lose the lottery?*

If drew the winning colour, and prize was airtime: *Give them the prize, which is 500 Sh ANNOUNCED*

If drew the winning colour, and prize was coins: *Give them the prize, which is [...] Sh NOT ANNOUNCED*

Ask this only after giving the secret money: Do you think anyone has seen me give you the money? If yes please explain.