

The Social Tax: Redistributive Pressure and Labor Supply

Eliana Carranza, Aletheia Donald, Florian Grosset, and Supreet Kaur*

August 18, 2022

Abstract

In low-income communities, pressure to share income with others may disincentivize work, distorting labor supply. We document that across countries, social groups that undertake more interpersonal transfers work fewer hours. Using a field experiment, we enable piece-rate factory workers in Côte d’Ivoire to shield income using blocked savings accounts over 3-9 months. Workers may only deposit earnings increases, relative to baseline, mitigating income effects on labor supply. We vary whether the offered account is private or known to the worker’s network, altering the likelihood of transfer requests against saved income. When accounts are private, take-up is substantively higher (60% vs. 14%). Offering private accounts sharply increases labor supply—raising work attendance by 10% and earnings by 11%. Outgoing transfers do not decline, indicating no loss in redistribution. Our estimates imply a 9-14% social tax rate. The welfare benefits of informal redistribution may come at a cost, depressing labor supply and productivity.

*Carranza: World Bank (ecarranza@worldbank.org); Donald: World Bank (adonald@worldbank.org); Grosset: Columbia University (fg2429@columbia.edu); Kaur: University of California at Berkeley and NBER (supreet@berkeley.edu). This paper greatly benefited from comments by Michael Best, François Gerard, Jessica Goldberg, Pamela Jakiela, Sylvie Lambert, Guilherme Lichand, Karen Macours, Owen Ozier, Léa Rouanet, Golvine de Rochambeau, Simone Schaner, Krzysztof Zaremba, and various seminar participants. Julia Buzan, Oumar Koné, Tiphaine Forzy, Chris Tullis, Ambika Sharma, Prathyush Parasuraman, Shelby Carvalho and Cécile Delcuvellerie provided superb research assistance. We gratefully acknowledge financial support from the World Bank’s Umbrella Facility for Gender Equality, the World Bank’s Jobs Umbrella Multidonor Trust Fund, and the National Science Foundation (Kaur’s CAREER award SES 1848452). We thank Innovations for Policy Action (IPA), especially Nicolò Tomaselli, Henriette Hanicotte, Samuel Kembou Nzalé, Mireille Nguhe Gbagbo and Augustin Kouadio for assistance with implementation. This project is a product of the World Bank’s Africa Gender Innovation Lab and Jobs Group. The findings, interpretations, and conclusions expressed in this paper do not necessarily represent the views of the World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent. Research approved by IPA IRB. AEA RCT Registry ID: AEARCTR-0003821. All errors and omissions are our own.

1 Introduction

“I am tired of giving [people] money...I am working to pay for my expenses, and [they] just come asking me for it all the time.”

— Interview with factory worker, Côte d’Ivoire (2016)

“When Magnolia and Calvin Waters inherited a sum of money, the information spread quickly to every member of their domestic network. Within a month and a half, all of the money was absorbed by participants in their network whose demands and needs could not be refused.”

— Carol Stack, *All Our Kin: Strategies for Survival in A Black Community* (1974)

In low-income communities, informal financial transfers within social and kin networks are ubiquitous and frequent (Banerjee and Duflo, 2007; Fafchamps, 2011). For example, full-time factory workers in Côte d’Ivoire report transferring 25-35% of their income to others outside their household on average, and 77% made at least one transfer in the past 3 months. Similarly, in the US, high earners among Black Americans share a substantial portion of their wealth with their network (O’Brien, 2012; Wherry et al., 2019). Frequent transfers have traditionally been understood as reflecting informal risk sharing, improving welfare by substituting for missing insurance markets.¹

However, work in the social sciences—spanning economics, anthropology, and sociology—has discussed the possibility that, despite these potential benefits, informal redistributive arrangements may also have distortionary effects (e.g., Lewis, 1955; Stack, 1974; Portes, 1998; Platteau, 2000). These literatures provide qualitative accounts that individuals face social pressure to share earned income. This, in turn, could disincentivize work—depressing labor supply levels and consequently earnings in lower-income populations.


To motivate this idea, Figure 1 indicates that work hours tend to be negatively correlated with the prevalence of transfers within social groups across a diverse range of settings: West Africa, Côte d’Ivoire, Indonesia, and the U.S. In each plot, the unit of observation is a geographic sub-location $\times$ ethnic group or race $\times$ year. The figure plots the average frequency of transfers in all years except the current year t (x-axis), against the average number of work hours in year t (y-axis).² Of course, these patterns simply reflect correlations and are

¹See Karaivanov and Townsend (2014), De Weerd and Fafchamps (2011), De Weerd and Dercon (2006), Fafchamps and Lund (2003), Ligon et al. (2002), Grimard (1997), Townsend (1994), Coate and Ravallion (1993), Rosenzweig and Stark (1989) and Rosenzweig (1988), among many others.

²A negative income shock in year t could increase transfers in year t . By using this leave-one-out approach, we attempt to avoid this direct correlation, and rather capture the correlation between employment

not necessarily causal. However, they motivate the possibility that informal redistributive arrangements could dampen labor supply levels.

Figure 1: Motivational Evidence: Redistribution and Hours Worked



Notes: Binscatter plots of the relationship between transfers in all years except the current year t (x-axis) and hours worked per week in wage employment plus self-employment in year t (y-axis). Panels A-B (C-D) display work hours for all adults in the household (for the individual respondent). Unit of observation is geographic sub-unit (determined by data availability) $\times$ ethnicity/race $\times$ year. The line of best linear fit and its slope are reported. Patterns are robust to dropping outliers.

In this paper, we develop a causal test to empirically examine this possibility. We conceptualize redistributive pressure as a “social tax” on earnings.³ We focus on the domain

levels and the general (time-invariant) tendency of a group to engage in transfers. In each panel, the negative correlations hold if we simply examine the contemporaneous correlations, and also if we remove outliers.

³We do not take a stance on the underlying microfoundation for redistributive pressure—such as second-

of labor supply, as it is the primary means through which the poor generate income. We test whether redistributive arrangements distort labor supply, output, and earnings.

We work with 474 full-time female factory workers in Côte d’Ivoire. The workers are employed in cashew processing plants run by Olam, a large transnational agro-processing firm, with an average tenure at the firm of 1.7 years. Workers are paid their wages twice a month in cash. Their labor supply is a function of both attendance at work and effort intensity while working. The entirety of their earnings is based on piece rates for output—the amount of peeled nuts—so that there is a direct mapping between labor supply and income.

Workers report frequent transfer requests from individuals outside of their households—both relatives and non-relatives.⁴ Consistent with our hypothesis, in our setting, 77% of workers believe that if they increased labor supply to earn more money, they would be subject to more transfer requests. Moreover, workers view redistributive pressure as hampering their ability to accumulate savings: 74% state they have difficulty saving for large goals because if they “put money aside, someone else will ask for it”. Turning down requests for cash is perceived as socially costly—unless workers can convey that they do not have the money available.⁵ Consequently, they engage in informal strategies to convert earnings to illiquid form—for example, by buying goods immediately after payday (e.g. Miracle et al., 1980; Goldberg, 2017). This suggests that methods that credibly lock away earnings could make it easier to retain more of one’s income—potentially increasing the incentive to work.

Drawing on this idea, we introduce a tool to lower redistributive pressure on income gains: a blocked savings account into which workers transfer earnings *increases*. Workers who opt in choose a threshold, which must be weakly higher than their baseline earnings. In each biweekly paycycle, any amount earned *above* this threshold is automatically deposited into the account; the remainder of their earnings is paid in cash as usual. The funds in the account cannot be accessed until the blocked period ends (3-9 months). We develop this product in partnership with one of the largest banks in Côte d’Ivoire, the Banque Populaire.

We conceptualize this product as reducing the effective social tax rate on earnings increases, while leaving the tax rate on preexisting levels of cash earnings unchanged. This design offers two important benefits—relative, for example, to an approach that lets workers move any of their existing earnings into blocked accounts. First, tax rate reductions usually generate opposing income and substitution effects—making it difficult to use labor supply responses to diagnose the existence or magnitude of a distortion. In contrast, lowering the

best risk-sharing with unobservable effort, or cultural sharing norms. This does not rule out the possibility that some transfers are driven by altruism; this portion of transfers would not constitute a “tax”.

⁴Reported pressure is generally from individuals who do not work at the factory.

⁵This indicates a psychological cost of lying about not having money and/or a real perceived social cost if one’s lie is discovered.

tax on earnings gains alone does not induce income effects, only substitution effects.⁶ Consequently, if there is a positive social tax, our intervention should unambiguously increase labor supply. Second, under our design, cash-on-hand is unchanged: by construction, expected take-home cash pay is not lower, so workers should have similar levels of disposable income to redistribute. This makes it unlikely that our intervention makes others in the network worse off through a reduction in transfers. In addition, while there is volatility in our setting, this is unlikely to undermine these predictions.⁷ However, note that while these two benefits help with the quantitative interpretation of our labor supply estimates, they need not hold for our test to be valid. Any increase in labor supply from our intervention would provide positive evidence for a labor supply distortion. Income effects would bias our results towards zero, so that the implied tax rate is a lower bound.

To test for the role of redistributive pressure, we vary the likelihood of transfer requests against saved income in the account. We randomize workers into three treatment conditions over the course of the experiment: 1) Control (no account); 2) Private (offer of a private blocked savings account); and 3) Non-private (offer of a blocked savings account whose existence and unblock date is revealed to the network). In the Private condition, no one except the worker is informed of the account’s existence. In the Non-private condition, workers are told that if they take up the account and save in it, then as a promotional tool for the bank, members of their social network may receive two publicity text messages that reveal that the worker is saving in a blocked account. The second publicity SMS would be sent shortly before the unblock date, stating that the funds would be released in the next week.⁸ This increases the probability that workers would receive transfer requests around the unblock date—mimicking, for example, the increase in requests that workers regularly experience around their biweekly paydays, the dates of which are known within the community.

If redistributive pressure exists, then workers with Private accounts should expect to retain more of any income gains than workers with Non-private or no accounts. Consequently, we predict that workers in the Private group will increase their labor supply—and therefore

⁶Denote the worker’s baseline labor supply as e_1 . Under our intervention, if she continues to supply e_1 , then her tax rate (and therefore net earnings) remain unchanged—i.e. there is no income effect. The rate is only lower for supply above e_1 . Consequently, starting from the baseline of e_1 , a worker who switches from the status quo to our intervention faces a pure substitution effect only.

⁷Volatility could introduce the scope for income effects if the social tax faced by a worker is higher when she has a positive income shock. We find no evidence for this in the data, suggesting that income effects are unlikely to play a meaningful role in workers’ labor supply reactions in our experiment. Moreover, note that the presence of income effects would only make it harder for us to detect effects on labor supply.

⁸This treatment was explained as a way for the bank—which was at the time actively trying to increase take-up of its blocked account product and was already advertising in the area—to advertise to community members as part of a publicity drive. Firms commonly use text messages and personal referral programs to advertise products in this setting.

their total earnings—relative to those in the Non-private or Control arms.

We conduct the experiment in two phases. In Phase 1, a subset of workers is randomized into either the Private or Control conditions. This enables us to compare the overall impact of the Private condition against having no account at all. Because our population of workers was largely unbanked at baseline, this “preparatory” phase was intended to instill trust in both the privacy of the Private condition and in the security of the accounts—e.g., workers saw that deposits occurred as expected, and funds could be withdrawn after the unblock date at the end of the Phase 1 treatment period.⁹ In Phase 2, we conduct our key test by randomizing workers into either the Private or Non-private account treatment conditions. We both add new workers to the sample for this phase, and also cross-randomize the workers from Phase 1. We conduct tests to verify that treatment effects are not sensitive to Phase 1 treatment status. We also leverage the cross-randomization to help test for confounds.

There is substantial demand for the Private accounts. When they are first offered in Phase 1, 43% of workers take them up. In Phase 2, take-up rises to 60%. In contrast, take-up of Non-private accounts in Phase 2 is only 14%—indicating that the blocked accounts are much less desirable as a savings vehicle when their existence is known to one’s network.

Being offered a Private account leads to sizable increases in labor supply and earnings. Overall, relative to the Control or Non-private conditions, the Private arm increases workers’ total output, and consequently earnings, by 11.4% ($p=0.012$). This is accompanied by an increase in attendance at work of 6.2 percentage points, or 9.7% ($p=0.023$). The magnitude of the treatment effects on earnings is similar when comparing the Private arm separately to either the Control condition (effect of 11.3%, $p=0.032$) or the Non-private condition (effect of 11.5%, $p=0.043$). Moreover, if we restrict our Phase 2 analysis to workers who were untreated in Phase 1, the results remain similar: being offered a Private account (relative to a Non-private one) increases earned income by 12.8% ($p=0.034$). This, along with other robustness checks, indicates that our effects are not sensitive to our cross-randomized design. Finally, because almost all workers (89%) have no earnings outside the factory, these treatment effects constitute increases in workers’ total income.

Note that the magnitude of these treatment effects is large—equivalent to how much earnings would rise if each worker worked an additional 1.19 days in every 2-week paycycle. This indicates that reducing the likelihood of transfer requests has a marked impact on workers’ willingness to supply labor, and consequently their total earned income.

Consistent with our hypothesized mechanism, treatment effects are concentrated among workers who face more redistributive pressure at baseline. For example, among workers who

⁹For example, in one of the factory plants, workers had previously been swindled out of savings deposits. Once Phase 1 was over, we were able to credibly announce to workers that Private accounts had been offered and implemented successfully, providing reputational benefits for our intervention before launching our key test in Phase 2.

report difficulty saving due to redistributive pressure, the Private treatment increases earned income by 18% ($p=0.018$); in contrast, among those who do not perceive such difficulty, we cannot reject that the Private treatment has no effect. Moreover, we find that our effects are not primarily driven by pressure to share earnings within the household. For example, the effects are suggestively stronger among those who do not have a partner or spouse; such workers see a 16% increase in earnings from being offered a Private account ($p=0.017$).

While the Private accounts enable workers to accumulate savings, we find no evidence of reduced transfers to others. In line with the design of the blocked accounts, there is no decline in *cash* take-home pay levels in the Private arm. Consequently, as expected, we find no discernible decline in transfers from workers to other households. These findings suggest that our intervention did not make others in workers’ networks worse off.

We argue that our findings cannot be explained by morale effects or fairness concerns. For example, we find no treatment effects during the 2-4 week period when workers know their treatment status but accounts are not yet activated; however, effects arise immediately in the first paycycle once savings are actually shielded from transfer requests. In addition, self-control problems alone cannot drive our findings, since they cannot explain the difference in take-up between the Private and Non-private treatments. Moreover, to gauge the potential relevance of self-control, we surprise a subset of workers with the option to forego depositing earnings into their account for one upcoming paycycle—varying whether this offer is made 4 days before the payday versus on the payday. Counter to the predictions of basic present-focus models, workers are not relatively more likely to want to unblock savings on the day of the payday itself.

In addition, because the Non-private treatment was implemented via text messages sent to others, this raises potential confounds related to privacy. To address such concerns, we construct a placebo test with workers who took up Private accounts in *Phase 1* but were offered Non-private accounts in *Phase 2*. Three months after the start of Phase 2, we ask these workers for permission to send promotional text messages to their network members advertising that they had saved in a blocked account through the bank in the past (i.e., in Phase 1), in exchange for a token compensation.¹⁰ This mimics some features of the Non-private treatment, but for accounts where the money would already be spent long ago—providing little scope for triggering transfer requests. 85% of workers accepted this offer. This indicates that simply having one’s network discover that one has used a blocked account to save is not the reason take-up plummets in the Non-private treatment. Similarly, fear of theft also cannot explain the low take-up of Non-private accounts: workers walk home from the factory with their entire cash earnings in their pockets twice each month on days that are publicly known, whereas savings withdrawal is considerably more private.

¹⁰We offered workers 1,000 FCFA—less than 5% of the Phase 2 average treatment effect.

To estimate the social tax rate implied by our results, we partner with the factory to randomize piece rate wages at the end of the experiment. Combining the resultant elasticity from this exercise with the 11.4% treatment effect on output from our experiment, we estimate that the social tax rate faced by the average worker in our sample is 9-14%, and is 19-23% for the subset of workers who actually take up the Private accounts.

Our study advances the literature on redistributive pressure and its impacts on economic behavior. A long tradition of qualitative work documents strong social pressure to share income with others in both developing countries (Scott, 1976; Kennedy, 1988; Platteau, 2000, 2014), and in low-income communities of color in rich countries (e.g. Stack, 1974; O’Brien, 2012; Wherry et al., 2019). Numerous studies using observational data argue that such pressure can rationalize behaviors such as the propensity to hold illiquid savings, as well as consumption, borrowing, transfer, entrepreneurship, and labor supply patterns (Di Falco and Bulte, 2011; Dillon et al., 2021; Baland et al., 2011; De Weerd et al., 2019; Grimm et al., 2013; Alby et al., 2020; Baland et al., 2016). In addition, a robust body of work—pioneered using lab-in-the-field experiments by Jakiela and Ozier (2016)—shows that individuals will take costly actions to keep income windfalls from their network (Beekman et al., 2015; Goldberg, 2017; Di Falco et al., 2018; Fiala, 2018; Boltz et al., 2019; Squires, 2021).¹¹ Finally, heterogeneity analysis in field studies indicates that the impacts of improved savings technologies or cash grants correlate with baseline levels of redistributive pressure (Dupas and Robinson, 2013; Riley, 2022; Squires, 2021).¹²

We build on and complement prior work by examining impacts in a high-stakes field setting: labor supply among full-time workers, within the context of long-run employment. We offer the first piece of direct evidence that redistributive pressure creates a disincentive to work—altering labor supply, with potentially large implications for productivity and income.

Note that because our specific intervention is designed to minimize income effects, it does not directly speak to the total impact of reducing existing redistributive pressures. In addition, while the blocked accounts are a proof of concept that it may be possible to boost individual earnings without decreasing redistribution, we do not view them as necessarily the most scalable policy approach. Rather, our intervention serves primarily as a tool to test whether redistributive pressure distorts labor supply, with utility consequences for workers.

Our study has potential implications for understanding one set of impediments to labor

¹¹For example, Goldberg (2017) varies whether a large windfall lottery payment is made in public or private, and finds that public recipients spend the money down more quickly. Squires (2021) uses the willingness to pay to keep a cash windfall private in the lab to structurally estimate sizable productivity implications among micro-entrepreneurs.

¹²A large related literature on intra-household bargaining indicates that women face pressure to share income with their spouses (e.g. Castilla and Walker, 2013; Bernhardt et al., 2019), and that separate accounts to hold savings can affect women’s bargaining power and outcomes (e.g. Ashraf, 2009; Ashraf et al., 2010; Schaner, 2015; Almås et al., 2018; Fiala, 2018; Field et al., 2021; Riley, 2022).

supply and productivity in low-income settings, particularly Sub-Saharan Africa. If redistributive pressures distort work incentives, they could also hamper other costly actions that increase future income, from human capital investment to technology adoption. This mirrors concerns expressed in some of the earliest development literature (Lewis, 1955; Tam et al., 1957). In addition, our findings raise the question of whether improved safety nets could affect the productivity of non-recipients by reducing their responsibility for redistribution (Dupas et al., 2017; Mobarak and Rosenzweig, 2012). While only speculative, these possibilities point to potential directions for additional research.

2 Conceptual Framework

We use a simple model of labor supply under taxation to motivate our experimental design. We model a worker who chooses consumption, c , and labor supply, e , to maximize utility $u(c, e)$. Her utility function represents standard preferences, with $u_c(c, e) > 0$, $u_e(c, e) < 0$, $u_{cc}(c, e) < 0$ and $u_{ee}(c, e) < 0$. She earns a piece rate, w , for each unit of effort supplied, so that gross earnings are we . We normalize the price of consumption to 1.


We conceptualize redistributive pressure as a “social tax” on wage earnings, which we denote as τ_1 . Note that we do not take a stance on the underlying microfoundation for this social tax. For example, it may reflect second-best risk-sharing arrangements, where effort is unobservable and difficult to distinguish from shocks. An alternate (and not mutually exclusive) possibility is that redistributive arrangements stem from cultural norms that entitle poorer individuals to seek support from richer ones (Platteau, 2000). In this section, we simply take the presence of such a “tax” as given. Note that this does not rule out the possibility that some transfers are driven by altruism; this portion of transfers would not constitute a “tax”. Our field experiment diagnoses whether such a tax does indeed exist, and attempts to quantify whether it is economically meaningful.

In the presence of a social tax, for any chosen level of labor supply, e , the worker’s take-home post-tax income (and consequently consumption) is $y = (1 - \tau_1)we$. We use e_1 to denote her utility-maximizing level of labor supplied under tax rate τ_1 (see Figure 2A).

To test for the presence of a tax, and its subsequent effects on labor supply, we seek to lower τ_1 by enabling the worker to avoid transfer requests on (some portion of) her earned income. However, simply lowering τ_1 presents two sets of important challenges.

First, as is the case with all taxes, the net effect on labor supply would be ambiguous (Hausman, 1985). On the one hand, the payoff to work would now be higher, increasing the incentive to supply labor above e_1 (substitution effect). On the other hand, the increase in net earnings under labor supply levels $e \leq e_1$ would decrease the incentive to work (income effect). This would make it challenging to use the worker’s labor supply response to diagnose

Figure 2: Tax Rate



Notes: Panel A: labor supply (e_1) under a linear piece rate with social tax rate τ_1 . Panel B: change in optimal labor supply when the social tax rate is reduced to τ_2 on earnings above e_1 .

the existence of a distortion, or to estimate the magnitude of the tax. Second, enabling workers to shield their earnings from redistributive pressures could have the potentially undesirable effect of reducing existing transfers to kin—posing ethical challenges.

One way to mitigate both these concerns is to lower the social tax on earnings *increases* only. Specifically, consider reducing the social tax rate to $\tau_2 < \tau_1$ only for $e > e_1$, while keeping the tax rate at τ_1 for $e \leq e_1$ (Figure 2B).

This new tax rate induces a kink in the worker's budget constraint. The worker consequently chooses her labor supply level e_2 to solve $\max_{c,e} u(c, e)$ under the budget constraint:

$$c \leq \mathbb{1}_{e \leq e_1} \{(1 - \tau_1)we\} + \mathbb{1}_{e > e_1} \{(1 - \tau_1)we_1 + (1 - \tau_2)w(e - e_1)\}.$$

First, note the trivial result: $e_2 \geq e_1$.¹³ Note that this already rules out the possibility of a labor supply decrease. We can therefore rewrite the budget constraint as: $(1 - \tau_2)we + \mathbb{Y} = c$, where $\mathbb{Y} \equiv (\tau_2 - \tau_1)we_1$. The worker's optimal choice of effort under the new tax schedule is $e_2((1 - \tau_2)w, \mathbb{Y})$.

We can derive the change in e_2 induced by a change in τ_2 . We provide an overview of the key results here; see Appendix A.3 for details. Applying the Slutsky equation, we obtain:

$$\frac{de_2((1 - \tau_2)w, \mathbb{Y})}{d\tau_2} = -w \frac{\partial \tilde{e}((1 - \tau_2)w, u)}{\partial [(1 - \tau_2)w]} + w(e_1 - e_2) \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial \mathbb{Y}}, \quad (1)$$

where $\tilde{e}((1 - \tau_2)w, u)$ is the Hicksian (compensated) labor supply. On the right hand side

¹³Proof by contradiction: Suppose that $e_2 < e_1$. Then the budget constraint becomes $(1 - \tau_1)we + y = c$, which is the budget constraint under which e_1 is the optimal choice. This contradicts $e_2 < e_1$.

of Equation 1, the first term is the substitution effect and the second is the income effect.

We can use Equation 1 to study how labor supply responds when moving from the tax schedule in Figure 2A to that in Figure 2B. We begin with the baseline situation of no kink in the budget constraint, so $\tau_2 = \tau_1$. Then, the effect of introducing the new tax schedule—i.e., decreasing τ_2 above e_1 when starting from the baseline of $e = e_1$ —is given by:

$$-\frac{de_2((1-\tau_2)w, \mathbb{Y})}{d\tau_2} = w \frac{\partial \tilde{e}((1-\tau_2)w, u)}{\partial [(1-\tau_2)w]} > 0. \quad (2)$$

The income effect term drops out, and there is only a pure substitution effect remaining. Intuitively, because the tax rate on earnings up until labor supply level e_1 has not changed, the net earnings up until e_1 are unchanged, eliminating the income effect.¹⁴

The result in Equation (2) indicates that moving from the tax schedule in Figure 2A to that in Figure 2B delivers an unambiguous prediction on labor supply: *labor supply will increase*. In other words, if workers face an initial social tax rate $\tau_1 > 0$, and our intervention lowers this tax rate to some $\tau_2 < \tau_1$ only for effort levels above e_1 , then workers will increase their labor supply. Further, the amount paid out as a social tax is weakly higher: $\tau_1 w e_1 + \tau_2 w (e_2 - e_1) \geq \tau_1 w e_1$. As a result, redistribution to the network should not decline, and may even increase if $\tau_2 > 0$.

Discussion. The tax schedule in Figure 2B helps resolve the two key challenges that would arise if we simply lowered τ_1 , for example, by enabling workers to hide any of their preexisting earnings. However, our model is deterministic. In our experimental setting, there is some volatility, and we use workers' baseline average output as our measure of e_1 . Volatility could reintroduce some scope for income effects under the following condition: if the tax *rate* is higher when workers experience a positive income shock.¹⁵ However, we see no evidence for this condition: the fraction redistributed does not increase with paycheck-to-paycheck fluctuations in income (Appendix Figure A1). While suggestive, this indicates that income effects are unlikely to play a meaningful role in workers' labor supply reactions

¹⁴To see this intuition in more detail: Suppose we lowered the tax rate on all earnings. Then if the worker remains at $e = e_1$, her tax rate and therefore net earnings are higher—this is the income effect. In contrast, under the tax schedule in Figure 2B, if the worker remains at $e = e_1$, her tax rate and therefore net earnings are unchanged—there is no income effect. Consequently, starting from the status quo tax schedule—where $\tau_2 = \tau_1$ and the worker exerts $e = e_1$ —reducing τ_2 produces only a pure substitution effect.

¹⁵For example, suppose earnings in each period are $w e_1 + \epsilon$, where ϵ is a random variable with mean zero, so that average earnings are $\overline{w e_1}$. Suppose that, in periods where $\epsilon > 0$, the worker gets more transfer requests and gives away a larger share of her gross income. If our intervention reduces the tax rate on earnings above $\overline{w e_1}$, then in periods with $\epsilon > 0$, workers will face a lower tax rate even though their labor supply remains the same ($e = e_1$). This could both generate some income effects on labor supply, and also reduce redistribution to the network. The scope for both these effects would still be less under the approach in Figure 2B than if we allowed workers to hide any part of their existing earnings, but would still nonetheless be present.

in our experiment. Moreover, note that the presence of income effects would only make it harder for us to detect effects on labor supply, and would lead us to underestimate the resultant level of the social tax rate.

More broadly, while the absence of income effects helps with the interpretation of our results, it is not necessary for our test to be valid. If labor supply increases under our intervention, this could only arise because the substitution effect dominates any potential income effects. In other words, an increase in labor supply would provide positive evidence that the social tax distorts labor supply. Consequently, while this approach offers an improvement over alternative approaches by mitigating the two concerns above, it does not need to eliminate them. The potential for these benefits motivates our experimental intervention, which seeks to vary the extent to which earnings *increases* are sheltered from transfer requests.

3 Context: Redistributive Pressures

We work with full-time piece rate workers, employed in cashew processing plants in Côte d’Ivoire. In our study setting, transfers are common and frequent. Workers transfer a significant share of their earnings (25-35% on average) to others outside of their household.¹⁶ Of this, 72% is redistributed within the extended family, and the remainder to non-family members. While workers within the factory may make loans to each other, transfer requests tend to arise from individuals outside of the factory. Workers with higher average earnings redistribute more to their networks on average (Appendix Figure A2).¹⁷

Transfer requests occur for diverse reasons—including unexpected shocks (illness), expected expenditures (school fees), investments (housing improvements), and consumption (people showing up at mealtimes). Respondents express a desire to avoid a large subset but not all of these requests. Moreover, requests often occur on or shortly after paydays, which are generally known to network members, when workers are more likely to have cash on hand. As we discuss in Section 2, we do not take a stance on whether transfer requests are part of risk-sharing arrangements, or the result of cultural sharing norms. We simply note that such requests are prevalent and that workers desire to avoid many of them, as has been documented in previous work. In addition, the factory plants employing the workers in our sample typically have excess demand for labor, particularly for long-term workers with regular attendance—suggesting that network members who make transfer requests could also obtain full-time jobs at the factory.

¹⁶We sum up total transfers recalled by workers in a survey, and divide this by total income. This may be an underestimate if individuals do not remember all financial and in-kind transfers.

¹⁷An increase in average earnings of 1 FCFA is correlated with an increase in reported transfers of 0.048 FCFA ($p=0.006$, 95% CI=[0.014,0.082]). Note that some young workers transfer large amounts (as a fraction of their income) to their parents.


Figure 3 documents that workers believe that if they increase their income by increasing labor supply, they would be subject to more redistributive requests (Panels A and B). For example, 77% agree or strongly agree with the statement “If someone...starts earning more because they have decided to work harder, people would start asking that person more often for financial help” (Panel A). These responses match beliefs and anecdotes from baseline qualitative interviews. For example, consistent with Panel A, one worker said, “I left my village to come work, because of the money. But if I earn a little more, [they] will come to ask me [for money], so that I give [it to them]. And...at home, I have nothing.” Consistent with Panel B, another worker said, “I can say that [requests for transfers] have increased [since I started working at the factory] because before it was only my mother who came to ask me for money or my older sister who is married in a village near here, but now almost everyone calls me to ask me for money.”

In addition, 80% of workers believe that refusing to share such income gains with others would result in social disapprobation (Panel C). Workers also state that redistributive pressure hinders their ability to accumulate savings: 74% state they have difficulty saving over time for large goals because if they put money aside, someone else will ask for it (Panel D).

Turning down transfer requests is often deemed socially unacceptable if the worker has cash on hand. Workers perceive the costs for refusing requests to include social stigma or isolation—making, for example, it unpleasant to attend extended family or network gatherings, which are an important source of socialization and utility in this setting. However, workers can turn down transfer requests with no (or muted) consequences if they can credibly claim to have insufficient funds to share. Note that this indicates either an expected social cost if one is found to have lied, and/or a psychological cost of lying. In qualitative interviews, workers expressed the presence of both costs. In addition, in Appendix Figure A3, we document that workers find it psychologically difficult to refuse requests, consistent with evidence from the behavioral economics literature on the utility cost of lying (Gneezy, 2005; Feldhaus and Mans, 2014)). Overall, this suggests that enabling workers to lock away earnings so that they are inaccessible would effectively lower their (perceived) social tax. We draw on this idea in constructing our design.

Consistent with this idea, workers employ a variety of strategies to make their funds inaccessible for redistribution. For example, workers report buying household goods immediately after payday, storing money with others, and participating in ROSCAs. Such strategies were described by workers during qualitative research conducted as part of the preparatory fieldwork for our study (McNeill and Pierotti, 2021), and have also been documented in the prior literature (e.g. Anderson and Baland, 2002; Somville, 2011; Boltz and Villar, 2013; Goldberg, 2017; Dillon et al., 2021). However, workers perceive the efficacy of such informal strategies to be limited, as indicated by Figure 3D. In our study, we draw

Figure 3: Motivational Evidence: Redistributive Pressure



Notes: N=420 (Panels A and B), N=488 (Panel C) and N=459 (Panel D) cashew factory workers in Côte d'Ivoire.

on these existing strategies to design a blocked savings account to help shield savings from redistributive pressure.

4 Experimental Design

We design an experiment to test whether redistributive pressures distort labor supply—and therefore output and earnings—among full-time factory workers in Côte d’Ivoire. Workers are paid piece rates for output, and receive their wages twice a month in cash. We provide an overview of the design in this section, summarize the implementation protocols in Section 5, and provide more details in Appendix A.4.

4.1 Blocked savings accounts

We seek to construct a design that mimics the approach in the tax schedule in Figure 2B above: a tool to lower redistributive pressure on income gains. To approximate this approach, we introduce a blocked savings account into which workers can transfer earnings *increases*. Workers who opt in choose a threshold, which must be weakly higher than their baseline earnings. In each biweekly paycycle, any amount they earn *above* this threshold is automatically deposited by the factory into the blocked account; the remainder of their earnings is paid in cash as usual on payday.

Using administrative data from the factories, we compute baseline earnings to be the worker’s average earnings per paycycle in the past 3 months. The funds in the account cannot be accessed until the end of the blocked period (3-9 months). The savings product is administered by one of the largest banks in Côte d’Ivoire, Banque Populaire (BPCI).

Because the funds deposited into the blocked savings account cannot be accessed, they cannot be used to fulfill transfer requests—at least not until cash is available after the unblock period. This enables workers to potentially accumulate a large lump sum of savings. Consequently, if workers decide to increase their labor supply, the incremental earnings from that effort (which then get deposited into the account) will potentially be “taxed” at a lower rate than if the account were not available. In other words, the accounts are designed to make it more likely that any increases in productivity are retained by workers for their own future use.

4.2 Treatments

Our goal is to develop a treatment design that isolates the role of redistributive pressure in labor supply decisions. While the blocked accounts reduce the scope for such pressure, they

may also introduce other benefits, such as helping with self-control problems. Consequently, to construct a cleaner test, we seek to vary—among workers who are offered a blocked account—the extent to which the product actually shields savings from transfer requests.

To achieve this, we draw on the way in which transfer requests arise in our context. Because workers are paid on a regular schedule by the factory, their network members know roughly on which days they will walk home with a large amount of cash in their pockets. As discussed above, workers report that network members are more likely to request transfers immediately after expected pay dates—when they anticipate that workers will have cash on hand. We mimic this by designing a “non-private” version of the blocked accounts: members of the worker’s network learn, shortly before the unblock date, that the worker is about to have access to a meaningful amount of liquid cash. Under our hypothesized mechanism, this would lead some network members to make transfer requests against the savings in the account, just as they do on payday—reducing the usefulness of the blocked accounts as a way to avoid redistributive pressure on income gains.

Specifically, over the course of the experiment, we randomize workers into three treatment conditions:

1. *Private account*: Workers are offered a blocked savings account as described above. No one else is informed of the account’s existence except for the worker.
2. *Non-private account*: Workers are offered a blocked savings account. However, the existence of the account and the approximate unblock date are known to others in the worker’s network.
3. *Control*: Workers are not offered a blocked savings account.

The Private account is implemented as described in Section 4.1 above. Specifically, all details of the account remain private to the worker. Consequently, earnings increases—saved in the blocked accounts—are relatively shielded from transfer requests.

To implement the Non-private condition, workers are informed that they are being offered an account as part of a new “Publicity program” run by the bank to encourage others to open savings accounts. Under this program, if the worker takes up the account and saves in it, members of their social network may receive up to two publicity text messages. Each message would let the recipient know that the bank is offering blocked account products, mention that the worker has one of these accounts and has successfully saved in it, and encourage the recipient to also open a similar account with the bank. In addition, the second SMS, sent shortly before the unblock date, would relay the savings level and that the account would soon be unblocked: “«The worker» will already be able to access her savings in the next week!”. If the worker declines to take up the account or does not save in it, no

information would be shared with network members (see Appendix A.4). Consequently, if an individual takes up a Non-private account and saves in it, her network members would know in advance that she will soon have a sizable amount of liquid cash available. This approach builds on the design of previous lab-in-the-field experiments, which offer individuals a large cash windfall, and test whether they are willing to pay to prevent others in their network from learning about the windfall (e.g., Jakiela and Ozier, 2016).

The Non-private treatment was explained as a way for the bank—which was at the time actively trying to increase take-up of its blocked account product and was already advertising in the area—to advertise to community members as part of a publicity drive. Firms commonly use text messages and personal referral programs to advertise products in this setting.

The Non-private treatment raises two sets of possible interpretational concerns. The first is whether it may increase the tax rate above the baseline level, τ_1 , by prompting more requests than under the status quo. Note that this requires the presence of distortionary effects from redistributive pressure, and so qualitatively is not a confound per se, but may have quantitative implications for our results. In Section 8.2, we argue that the salience of the Non-private treatment is not necessarily larger than under workers’ publicly known regular paydays. In addition, our design enables us to use the effects of Private relative to Control to obtain one suggestive benchmark around this potential concern.

The second potential concern is that workers may turn down Non-private accounts not due to redistributive pressure, but due to privacy considerations. In Section 8.4, we offer evidence against this confound using several pieces of evidence: a placebo SMS exercise, data from workers on reasons for turning down Non-private accounts, and heterogeneity analysis.

The Private and Control conditions enable us to examine the effect of a Private blocked account against the status quo. In addition, the Private and Non-private conditions enable us to compare workers who all have access to the same blocked account product, but differ in their expectation of whether they will be able to retain the resultant savings for their own use. We should only observe differences under this variation in privacy if external pressures play a role in labor supply decisions. Consequently, we predict that workers in the Private group will increase their labor supply—and therefore their total earnings—relative to those in both the Control and Non-private treatment arms. In contrast, if redistributive pressure causes no labor supply distortions—for example, if all transfers reflect altruism or efficient risk sharing (where transfers respond only to shocks, but not to changes in labor supply)—then we should not expect differences across treatment arms.

4.3 Randomization

We conduct the experiment in two phases. In Phase 1, a subset of workers is randomized into either the Private or Control conditions. In a subsequent Phase 2, we randomize workers into either the Private or Non-private account treatment conditions. Table 1 provides an overview of the randomization design.

Table 1: Experiment Design Overview

Phase 1	Phase 2
Private treatment (50%)	Private (50%) Non-private (50%)
Control (50%)	Private (50%) Non-private (50%)
<i>Not in Phase 1</i>	Private (50%) Non-private (50%)

Notes: Phase 2 randomization is stratified by the Phase 1 status of the study participants: private treatment arm in Phase 1; control arm in Phase 1; or not in the study in Phase 1. The latter are workers eligible at the time of the enrollment for Phase 2, who were not yet at the factory or were not yet eligible at the time of the enrollment for Phase 1; they represent 38% of the Phase 2 sample.

Because our population of workers was largely unbanked at baseline, the “preparatory” Phase 1 was intended to instill trust in both the privacy of the Private condition and in the security of the accounts. This was especially important in one of the factory plants, where workers had previously been swindled: a former employee from an insurance company had collected deposits from workers and then disappeared. In addition, introducing the Non-private treatment to workers immediately may have cast doubt on whether privacy would be maintained even in the Private condition—given the low trust environment in our setting. Overall, ex-ante, since we were unsure whether such various trust issues would lead to low take-up rates for even the Private accounts, we began with Phase 1 to ensure the design did not fail for such operational reasons. Once Phase 1 was over, we were able to credibly announce to workers that Private accounts had been offered and implemented successfully in their factory in the past. Workers knew that if there had been issues with the accounts, their coworkers would speak up—providing reputational benefits for our intervention before launching our key test in Phase 2.

The randomization design in Phase 2 maximizes statistical power for our primary test—comparing effects of the Private treatment vs. the Non-private treatment—in two ways. First, in addition to adding new workers to the sample for this phase, we also cross-randomize workers who participated in Phase 1—stratifying Phase 2 treatment status by Phase 1 treat-

ment assignment. In the analysis, we conduct tests to verify that treatment effects are not sensitive to Phase 1 treatment status—for example, by restricting the analysis to workers who were not in the Phase 1 sample. Second, we do not include a pure Control condition in Phase 2. Consequently, we cannot directly compare the Non-private condition with the Control condition. Since this comparison is not core to our predictions, our chosen design maximizes our ability to test our key hypothesis. In addition, we leverage our cross-randomized design to test for confounds below.

4.4 Auxiliary tests: privacy, self-control and piece-rate variation

We supplement our design with three additional sources of variation. We provide an overview here, and explain each test in more detail in Section 8. First, towards the end of the experiment, we undertake placebo tests to examine workers’ willingness to share information about their savings behavior with their network members. We again offer to send publicity SMS messages, but which advertise information about *past* savings behavior under blocked accounts, where funds would already have been spent—obviating the scope for transfer requests against the savings. This enables us to examine the value of privacy on its own, when redistributive pressure is not a relevant force. A high willingness to share such information would suggest that privacy concerns are unlikely to drive low take-up of Non-private accounts. In addition, we supplement our analysis with several pieces of data collected through surveys, in order to better understand the mechanisms underlying our results.

Second, self-control problems in consumption could increase the value of blocked savings accounts, but this channel could not alone explain differences in Private and Non-private take-up. However, to more fully examine the relevance of self-control, we test a core prediction of time inconsistency models. Among workers who take up blocked accounts, we surprise them with the option to opt out of the direct deposit for the current paycycle only. We randomly vary whether the option is provided four days before the payday or on the day of the payday itself.

Third, at the end of the experiment, we partner with the factory to randomize piece rates. We use this to compute a labor supply elasticity, which we can use as an input into our estimates of the social tax rate implied by the results of our main experiment.

5 Implementation and Protocols

5.1 Job features

Background: cashew processing factories. The workers in our study are full-time laborers employed in cashew-processing factories run by Olam, a large multinational agro-processing company that controls 80% of the processed volume of cashews in Côte d’Ivoire. The two factory plants with which we work are located in central Côte d’Ivoire, about 230 km away from the national capital, Abidjan.

Côte d’Ivoire is the world’s second-largest producer of raw cashew nuts but only processes 7% of them domestically (World Bank 2018, 2020). While increasing domestic processing is a national development priority, the government and industry view low labor productivity and unmet labor demand despite relatively high wages as two primary impediments. Our hypothesized mechanism—due to its implications for labor supply and productivity—has relevance for both these constraints.

Production task. Workers in the experiment are engaged in manually peeling cashew nuts. This entails gently rubbing off with the fingers or a knife those parts of the peel that are still attached to the cashew after it undergoes mechanized peeling. Workers fill up buckets with cashews and return to their workstation for peeling. Production is strictly an individual activity, with no joint production of any kind. Workers’ daily output is determined by the weight (in kg) of how many cashews they have peeled that day. Workers complete a set workday from 8 am to 5 pm, Mondays through Saturdays, with a one-hour break for lunch. Consequently, while workers have latitude to choose attendance, there is less flexibility in choosing work hours within a day.

Payment. Each worker receives a linear piece rate for her output.¹⁸ The entirety of workers’ earnings are comprised of their piece rate wages. Consequently, changes in effort translate directly into changes in worker earnings. Workers are paid their earnings twice a month in cash.¹⁹

5.2 Experiment Protocols

Sample. Over the course of the experiment, we enroll 474 full-time workers, of which 464 are women, in the manual peeling sections of the two Olam factory plants.²⁰ We enroll all eligible workers in the manual peeling sections across both factory sites. To be eligible, workers must

¹⁸The specific piece rate changes based on the quality of the nuts, which fluctuates over time and is exogenous to the worker. Nuts that are more difficult to peel are paid higher piece rates.

¹⁹Earnings at the factory are set so as to exceed Côte d’Ivoire’s minimum wage for full-time attendance.

²⁰The factory also employs men, but they are usually not employed in tasks where payment is a piece rate based on individual production, making such tasks incompatible with our research design.

be “regular” workers at the factory, proxied by having baseline work attendance of at least about 50% in the three months before the launch of the intervention, and have started the job at least 2 months before launch. We impose tenure requirements to minimize attrition from the experiment, since newly joined workers have high turnover rates. It is unclear ex-ante whether treatment effects for more short-term workers would be higher or lower. In our sample, the average worker has worked at the firm for 1.7 years (25th percentile of 1.2 years, 75th percentile of 2.4 years). For workers in our sample, their factory income is their primary source of earnings, and 89% report having no other source of income.

Blocked accounts. The blocked savings accounts are designed and implemented jointly with the Banque Populaire de Côte d’Ivoire (BPCI). While blocked savings accounts previously existed in this setting, we worked with the bank to offer date-based (rather than only goal-based) accounts, with no minimum balance or monthly fee requirements, in exchange for removing the interest rate on savings in the account. In this setting, opening a formal bank account is an administrative hassle: individuals must obtain formal documents and travel to the bank to fill out the application and submit them. To help minimize this hurdle for workers who were offered the blocked accounts in the experiment, a bank employee is stationed in the factories during the enrollment period to help collect applications. In addition, by enabling the factory to directly deposit worker earnings above the threshold into the accounts, we make it much easier to make deposits—since otherwise, deposits need to be done manually by traveling to a bank branch. Consequently, while all workers can open a savings account with the bank, in practice, our intervention greatly lowers the barriers to obtaining and using a formal account. Consistent with this, only 1% of Control group workers have any type of formal bank account at endline.

Workers who take up the blocked account choose one threshold that applies to all future paycycles. They can revise this threshold up to three times. In addition, workers can opt out of having a threshold at any point, which would halt any additional future deposits from being made into the account. These provisions prevent mistakes, and allow workers the ability to re-optimize thresholds after experiencing the accounts if they want.

After the end of the blocked period, the account converts to a regular savings account: workers may withdraw all or part of their accumulated savings. Alternately, they may lock their existing savings for another blocked period.²¹ The length of the blocked accounts is 9 months in Phase 1, and shortened to 3 months in Phase 2 in order to complete the experiment on the timetable agreed upon with implementing partners.²²

²¹If workers re-block their existing savings for another blocked period, this only applies to prior savings already in the account. No future earnings are directly deposited—unless dictated by the worker’s treatment status in a subsequent phase of the experiment.

²²Delays in Phase 1 implementation reduced the time remaining to complete Phase 2. To stick to our timetable, we therefore reduced the time length of the Phase 2 accounts to 3 months. This does not affect

Announcement and Training. Treatment status is chosen in each phase using a lottery, where ID numbers are drawn by the research team to assign treatment status. This helps promote feelings of fairness, and also makes it clear that Olam, the employer, is in no way making the decision of who receives accounts. Drawing ID numbers (rather than names) enables workers to know if they themselves are chosen, while maintaining the privacy of selected workers.

To ensure workers understand how the blocked accounts work, we undertake training sessions within the factory with all workers assigned to the Private or Non-private conditions. Workers attend the sessions in small groups of about 5 workers each. These sessions are attended by a bank staff member to answer questions, and led by a moderator from the research team. The sessions explain the rules of the accounts, including choosing thresholds, and work through examples. At the end of the session, each worker takes a quiz to verify comprehension of the account rules. If the worker scores below 80% on the quiz, they are retrained one-on-one by a moderator. The gap between when workers learn their treatment status and when accounts take effect for most workers is 3-4 weeks: 1-2 weeks of pure announcement period, plus about 2 weeks to get the paperwork in place.

Privacy of deposits. Workers who enroll in the blocked accounts continue to be paid the take-home portion of their earnings (i.e., any amount earned less than the threshold) in cash, in the same way as before. Any amount earned above the threshold is directly deposited by the factory to the bank, and this amount is not discussed when payments are distributed to help maintain privacy. Instead, workers enrolled in the accounts are given a small receipt discreetly at a different time that verifies how much was deposited into their savings account. Consequently, our protocols ensure the privacy of savings deposits—while leaving the process of obtaining payment of the cash component of earnings on paydays unchanged. These payment protocols were the same for both the Private and Non-private treatments.

Enrollment and roll-out. In Phase 1, workers were enrolled in the experiment in three staggered waves, with data collection for Phase 1 ending in March 2019. Phase 2 was then conducted from April 2019-July 2019.

6 Empirical strategy

6.1 Data sources

Our primary data source is Olam’s detailed daily administrative data at the worker level. This includes individual daily attendance, output (the quantity of nuts processed), and

the internal validity of our estimates, since treatment assignment is stratified within a given phase.

earnings. Note that this data does not include information on worker hours; however, given the low flexibility of work hours (see Section 5.1), it is reasonable to use attendance as a meaningful measure of the extensive margin of labor supply. These administrative data are used by Olam to compute earnings, and the amount to be deposited into workers’ accounts in each paycycle.

We supplement this with data collected through two sets of phone surveys, which form the basis for our baseline data for Phases 1 and 2. These include information on perceptions about redistributive pressure. In addition, we use a more detailed in-person endline survey, which includes details about financial transfers with network members.

Table 2 provides descriptive statistics for our study sample, and checks for balance in baseline covariates across treatment arms separately for each Phase. Across treatments, the average worker in our sample has worked at the factory for 620 days, and has a baseline attendance rate of 68%. Almost all workers (98%) are women.

Our primary outcome measure, earnings, is balanced at baseline. Of the 25 comparisons presented, only the comparisons on attendance are statistically distinguishable across treatment arms at the 5% level. Note that our primary empirical specification controls for any baseline differences in levels using a differences-in-differences approach that includes individual fixed effects. In addition, as we show below, we find similar treatment effects using only endline data controlling for baseline covariates.

6.2 Estimation

Our primary outcome is workers’ daily earnings (in FCFA). Given the linear piece rate incentive scheme, this is equivalent to examining effects on output, and serves as our measure of workers’ total labor supply. We also study effects on attendance to specifically examine the extensive margin decision of deciding to come to work.

We estimate effects for Phase 1 (comparing the Private treatment with Control) and for Phase 2 (comparing the Private treatment with the Non-Private treatment). To measure treatment effects, we use the administrative panel data from the factories to estimate difference-in-differences regressions at the worker-day level:

$$y_{it} = \beta PrivateAcct_i \times Post_t + \alpha PrivateAcct_i \times Announcement_t + \gamma_i + \delta_t + \epsilon_{it}, \quad (3)$$

where y_{it} is the outcome of interest for worker i on date t . The $PrivateAcct_i$ indicator equals one if worker i was assigned to receive the Private account treatment. $Post_t$ is an indicator that equals one on dates when the blocked savings account is active, and $Announcement_t$ is an indicator that equals one during the announcement period (after workers are told their

Table 2: Balance

Variable	Phase 1			Phase 2		
	Control Mean/SD (1)	Private - Control Difference/SE (2)	P-value of difference (3)	Non-Private Mean/SD (4)	Private-Non Private Difference/SE (5)	P-value of difference (6)
Baseline labor supply						
Tenure at factory	651 [241]	-1 (27)	.961	576 [244]	0 (26)	.997
Earnings	1822 [863]	-106 (101)	.296	1726 [670]	-105 (82)	.200
Earnings above median	.52 [.501]	-.042 (.054)	.439	.478 [.501]	.041 (.056)	.467
Attendance	.666 [.207]	-.051 (.024)	.034	.771 [.223]	-.066 (.03)	.028
Attendance above median	.414 [.494]	-.049 (.051)	.340	.472 [.501]	-.094 (.055)	.092
Workers' characteristics						
Has an ID	.618 [.487]	-.031 (.058)	.594	.765 [.426]	.04 (.054)	.451
Is a woman	.986 [.116]	-.004 (.012)	.738	.994 [.079]	-.013 (.013)	.305
Speaks Dioula	.554 [.498]	-.05 (.054)	.356	.41 [.493]	-.025 (.055)	.646
Speaks Baoule	.261 [.44]	.052 (.049)	.287	.255 [.437]	.04 (.05)	.424
Heterogeneity variables						
Savings not taxed	.24 [.43]	-.108 (.063)	.088	.266 [.444]	.004 (.054)	.934
Not taxed by acquaintances	. [.]	. (.)	.	.713 [.454]	.063 (.053)	.238
Has a partner	.467 [.502]	.141 (.08)	.081	.683 [.467]	-.093 (.057)	.105
Limited control over earnings	.231 [.425]	-.066 (.068)	.334	.17 [.377]	-.026 (.044)	.546

Notes: Summary statistics and tests for baseline balance by treatment group. Cols. (1)-(3) use Phase 1 data, and Cols. (4)-(6) use Phase 2 data. Cols. (1) and (4) present the sample mean, with standard deviations in brackets, for workers in Control and Non-Private, respectively. Cols. (2) and (5) report the coefficient from a regression of the baseline covariate on an indicator for being assigned to the Private arm, controlling for strata fixed effects, with the standard error in parentheses; Cols. (3) and (6) report the associated p-values. Standard errors clustered at the worker level. Definitions of heterogeneity variables provided in notes to Table 4. Note that the “Not taxed by acquaintances” was not collected in the Phase 1 baseline. Worker earnings reflect the fact that attendance is on average 68% at baseline in our sample. Earnings at the factory are set so as to exceed Côte d’Ivoire’s minimum wage for full-time attendance.

treatment status, but before workers submit paperwork and blocked accounts take effect),²³ these time indicators are mutually exclusive. The γ_i and δ_t are worker and paycycle fixed

²³This announcement period covers a 3-4 week gap between when treatment is announced and accounts take effect. It is comprised of 1-2 weeks of a pure announcement period after workers are informed of their treatment status, and an additional 2 weeks to get paperwork in place.


effects, respectively. The gap between the end of Phase 1 and start of Phase 2 is two weeks. For consistency, we use a baseline period of two weeks in both the Phase 1 and Phase 2 treatment analyses. We test for robustness to other baseline periods, and also to alternate empirical specifications. We also show robustness to restricting analysis using endline data (i.e. the period after accounts are in place) only. Standard errors are clustered at the worker level.

Our coefficient of interest is β , which captures the treatment effect of being offered a Private account, relative to the omitted category of being offered no account (Control) or a Non-private account. We predict that $\hat{\beta} > 0$. In addition, α captures the treatment effect of knowing one's treatment status in the announcement period before the blocked accounts actually take effect. Under our hypothesized mechanism of redistributive pressure, we predict that $\hat{\alpha}$ will be indistinguishable from zero.

7 Results

7.1 Take-up of blocked accounts

Figure 4: Take-up of Blocked Accounts



Notes: Means and 95% CIs. SEs clustered at the worker-by-phase level.
Phase 1: N = 354 workers. Phase 2: N = 317 workers.

In Phase 1, the proportion of Private group workers who take up the blocked savings account is 43%. In Phase 2, this take-up rate for the Private blocked accounts rises to 60%—potentially reflecting increased trust.²⁴ In contrast, under the Non-private condition, account

²⁴Further consistent with the high perceived value of the Private accounts, among those who select into

take-up is only 14%—a 77% decline in take-up relative to Private accounts ($p < 0.001$). This indicates that the blocked accounts are substantially less desirable when their details would be revealed to the worker’s network.

We use survey data to examine workers’ self-reported reasons for the low take-up of Non-private accounts (Appendix Figure A4). Among those who declined a Non-private account, a striking 96% say that an important factor in their decision was that network members’ knowledge of the account would lead to an increase in net transfer requests. In addition, 46% say that they worried this would decrease their partner’s contribution to household expenses—also a component of redistributive pressure in our setting. In contrast, only 5% of workers cite any other reason as motivating their decision not to take up a Non-private account. We return to these ideas when we examine heterogeneous effects in Section 7.3, where we also directly examine the potential role of intra-household bargaining for our results. Overall, while only suggestive, these responses indicate that concerns about redistributive pressure were prominent in workers’ minds when refusing the Non-private accounts.²⁵

Among workers who take up a blocked account, 26% choose a threshold close to the minimum (i.e., a round number within 10% above their baseline earnings), while the remaining 74% select a threshold that is higher (see Appendix Figure A5). Overall, the median threshold chosen corresponds to 139% of mean baseline earnings. This is consistent with some desire to maintain flexibility in liquid earnings among workers—leaving room for workers to adjust labor supply to respond to cash needs in a given paycycle, while still allowing them to make use of the blocked accounts to save a subset of their earnings increases.²⁶

Despite the choice of fairly high thresholds, the accounts were actively used by workers. In Phase 1, workers who took up the accounts achieved savings in 28% of their paycycles on average. In Phase 2, among those who opened Private accounts, savings were deposited in 38% of pay cycles. Note that since thresholds were usually set weakly above baseline earnings, these numbers under-estimate the fraction of paycycles in which workers earned more than their baseline earnings levels.

the accounts in Phase 1, 90% of those who are offered them again in Phase 2 take them up. Appendix Table A1 provides a full a tabulation of Phase 2 take-up by Phase 1 treatment status. Note that the high take-up rates of the Private accounts likely reflect several features of our setting: the large time investment (about 9 months) to establish trust in the factories before launching the experiment; the removal of many logistical barriers to opening and making deposits into accounts; and the lack of penalties in failing to make a target which lowered risk from participating in the accounts.

²⁵We did not ask such questions of workers who were offered the Private accounts, so we cannot compare across the treatments. This is because we did not want to confuse the Private account workers by potentially making them worry that knowledge of the Private accounts would become known to their network.

²⁶Because the decrease in the social tax is still above baseline earnings, note that this still mitigates the scope for income effects. However, this potentially decreases the potency of the blocked accounts as a way to reduce redistributive pressure, potentially dampening labor supply responses. Note that it would not have been operationally feasible to allow workers to select a new threshold in each paycycle.

7.2 Effects on labor supply and earnings

Being offered a Private account substantially increases labor supply and earnings. In Table 3, we examine treatment effects on earnings in Panel A, and on attendance in Panel B. Overall, relative to the Control or Non-private arms, the Private arm raises average worker earnings by 175.9 FCFA, corresponding to an 11.4% increase in income (Col. 1, Panel A, $p=0.012$).²⁷ This is driven, in part, by a 6.2 percentage point (pp) or 9.7% increase in attendance at work (Col. 1, Panel B, $p=0.023$).²⁸ This extensive margin treatment effect is driven entirely by a reduction in absenteeism; we find no changes in turnover (Appendix Table A2).

Table 3: Treatment Effects on Earnings and Attendance

	Pooled	Phase 1	Phase 2		
	All workers (1)	All workers (2)	All workers (3)	Not Treated Ph1 (4)	Not in Ph1 (5)
Panel A: Daily earnings					
Private account	175.9 (69.68) [0.012]	173.2 (80.72) [0.032]	184.4 (90.90) [0.043]	214.0 (100.9) [0.034]	276.5 (118.4) [0.020]
Sample mean in control	1547	1535	1602	1670	1598
Panel B: Attendance					
Private account	0.0622 (0.0273) [0.023]	0.0557 (0.0313) [0.076]	0.0821 (0.0368) [0.026]	0.0862 (0.0402) [0.033]	0.111 (0.0470) [0.018]
Sample mean in control	0.64	0.65	0.61	0.61	0.59
N: worker-days	137678	99638	38040	24360	14400
N: workers	474	354	317	203	120

Notes: Unit of observation is worker-day. Dependent variable is earnings in Panel A and attendance in Panel B. Each column shows results from the difference-in-differences specification in Equation 3. Col. (1) pools across Phase 1 and 2; the omitted category is being assigned to the Control or Non-private conditions. Col. 2 sample is Phase 1; the omitted category is being assigned to the Control condition. Cols. (3)-(5) sample is Phase 2; the omitted category is being assigned to Non-private condition. Col. (4) restricts analysis to Phase 2 sample workers who were not assigned to Private treatment in Phase 1. Col. 5 restricts to Phase 2 sample workers who were not included in the Phase 1 sample. All regressions include day and worker fixed effects. Standard errors clustered by worker-phase in Col. (1) and by worker in Cols. (2)-(5).

These patterns are similar when we examine effects separately by phase. In Phase 1, compared to the Control group, the Private arm increases earnings by 11.3% (Col. 2, $p=0.032$). Similarly, in Phase 2, compared to the Non-private group, those offered Private accounts increase earnings by 11.5% (Col. 3, $p=0.043$).

These magnitudes are economically meaningful. For example, on average, the size of the Phase 2 treatment effect is equivalent to how much earnings would rise if each worker worked

²⁷In each panel of Col. (1), we estimate a pooled version of Equation 3, where we combine the data from both phases. The coefficient on Private gives the average treatment effect relative to both the other treatment arms.

²⁸Using a simple back of the envelope calculation, attendance accounts for 85% of the overall earnings effect; however, the confidence intervals around this estimate are wide.

an additional 1.19 days in *every 2-week paycycle*.²⁹ In addition, these effects are not simply reflecting a substitution away from other income-generating activities outside of the factory jobs: we find no treatment effects on earnings outside the factory. This is consistent with the fact that, at baseline, 89% of workers report having no earnings outside of their factory job, and on average, 93% of total income comes from factory earnings. Consequently, these treatment effects reflect an increase in total earned income.

In Cols. (4)-(5), we verify that our Phase 2 treatment effects are not driven by Phase 1 treatment status. Specifically, since we cross-randomize treatment assignment across the two Phases, one potential concern is that treatment effects may be influenced by workers who were assigned to Private accounts in Phase 1 and then lost access to these accounts in Phase 2. In Col. (4), we exclude from the analysis workers who were assigned to Private accounts in Phase 1, so that the sample includes only workers who were in the Control group in Phase 1 or workers who were not in the Phase 1 sample at all. The results are nearly identical: among such workers, the Private treatment increases output by 214 FCFA or 12.8% ($p=0.034$). Moreover, even if we exclude all workers in the Phase 1 sample, and only examine effects on the 120 workers who only joined the study in Phase 2, the treatment effects are similar and remain significant (Col. 5, $p=0.020$). Moreover, we cannot reject that the treatment effects estimates across Cols. (3), (4), and (5) are the same ($p=0.427$); in other words, the effects across these three samples are statistically indistinguishable. In addition, we verify that our treatment effects are robust to alternate empirical specifications, including using only intervention period data to examine treatment effects (Appendix Tables A3 and A4).

These results indicate that reducing the likelihood of transfer requests on savings has a marked impact on workers' willingness to supply labor. This, in turn, has a large impact on their total earned income.

7.3 Heterogeneity in treatment effects

Consistent with our hypothesized mechanism, the Private accounts have larger labor supply impacts among workers who report more redistributive pressure at baseline. In Table 4, Col. (1), we examine heterogeneity by whether workers state that they cannot accumulate savings due to redistributive pressure—i.e. those that agree or strongly agree with the statement, “I have difficulty saving over time for large goals, because if I put money aside, someone else will need it for something urgent”. Among such workers, being offered the option to shield savings from transfer requests increases earnings by 249 FCFA or 15.0% ($p=0.018$).

²⁹The average treatment effect on daily earnings is 184.4 FCFA, corresponding to a 2,213 FCFA increase per paycycle (comprised of 12 workdays). Mean daily earnings conditional on working among the Non-private group is 2,626. This gives $2626/2213 = 1.19$ workdays per paycycle.

In contrast, among those who do not report facing a social tax on savings, the estimated effects are significantly lower, and we cannot reject that there is no impact of the Private treatment for these workers ($p=0.95$).

Table 4: Treatment effects heterogeneity

	Redistributive Pressure		Intra-household Control		Productivity
	Savings not taxed (1)	Not taxed by acquaintances (2)	Has a partner (3)	Limited control over earnings (4)	Baseline earnings above median (5)
Private account	249.4 (104.8) [0.018]	312.8 (126.6) [0.014]	281.7 (117.7) [0.017]	220.9 (106.3) [0.038]	274.6 (75.15) [0.000]
Private account X Covariate	-258.1 (138.9) [0.064]	-211.7 (130.2) [0.105]	-160.3 (115.9) [0.167]	-288.6 (179.5) [0.109]	-203.1 (78.75) [0.010]
Mean if covariate = 0	1665	1790	1778	1660	1254
Mean if covariate = 1	1782	1568	1603	1754	1811
Share: covariate = 1	0.22	0.74	0.59	0.17	0.50
P-val: sum = 0	0.95	0.32	0.27	0.71	0.40
N: observations	137678	137678	137678	137678	137678
N: workers	474	474	474	474	474

Notes: Dependent variable is earnings. Unit of observation is worker-day. Private account is an indicator that equals 1 if the worker is assigned to the Private treatment. The omitted category is being assigned to the Control or Non-private condition. Covariate in each column is a binary indicator. Col. (1) covariate equals 1 if the worker strongly disagrees with the statement “I have difficulty saving over time for large goals because if I put money aside, someone else will need it for something urgent.” Col. (2) covariate equals 1 if the worker answers ‘never’ to the question: “How often do you financially support acquaintances (not close friends or family)?”. Col. (4) covariate equals 1 if the worker does not answer “input into most or all decisions” to the question “How much input do you have in making decisions about the money you earn?”. Covariates in Cols. (1)-(4) are from baseline surveys conducted prior to the assignment of treatment status in each phase. Baseline covariates are missing for initial enrollees in Phase 1: any missing values are coded as 0, and a dummy indicating which observations have missing values as well as the interaction between that dummy and Private are included in the regressions. The bottom of the table reports the share of non-missing observations for which the covariate equals 1, as well as the p-value of the F-test for Private + PrivateXCovariate=0. Regressions build on the specification in Table 3, Col. (1). All regressions include day and worker fixed effects. Standard errors clustered by worker.

Due to the short length of the baseline phone surveys, we could not ask extensive questions about redistributive pressures or transfers.³⁰ However, we supplement this with another proxy for the severity of redistributive pressure: whether the individual reports making transfers to “acquaintances”, defined as individuals who the worker does not consider close family members or friends; such transfers are especially likely to reflect a social tax rather than altruism. In qualitative work, workers expressed particular frustration about such transfers. Among such individuals, gaining access to the Private accounts raises earnings by a striking 313 FCFA or 17.5% (Col. 2, $p=0.014$). Among workers who do not report such extreme social taxation, the treatment effects are substantively smaller and not statistically significant at

³⁰Specifically, among the questions in Figure 3, at baseline in both Phase 1 and Phase 2, we only asked the question in Panel D; this is the source of heterogeneity in Table 4, Col. (1).

conventional levels ($p=0.32$). Overall, these findings are consistent with previous work indicating the relevance of heterogeneity in the extent to which individuals face redistributive pressure (e.g. Dupas and Robinson, 2013; Riley, 2022; Squires, 2021).

In our setting, women report facing pressure to share income not only from individuals outside their household, but also from their partner or spouse. This raises the question of whether all our results are driven by intra-household pressures. However, we find sizable treatment effects of 15.8% among workers who report having no partner—so that there is no scope for an intra-household motive (Col. 3, $p=0.017$). Among women with a partner, our estimated effects are positive but smaller in magnitude and statistically insignificant ($p=0.27$). Moreover, we see similarly strong and significant treatment effects among workers who report having high levels of control over their earnings (Col. 4). This is also consistent with the results in Col. (2) indicating that effects are concentrated among those who want to avoid requests outside their household and close family and friends. However, these patterns do not rule out the possibility that some portion of our effects is driven by a desire to shield earnings within the household. For example, as discussed in Section 7.1, some women who refuse Non-private accounts cite concerns about changes in their partner’s contribution to household expenses. However, our findings in Table 4—especially the results in Col. (3)—indicate that redistributive pressure *outside* of the household plays an important role in driving the impact of the Private accounts.

Note that since virtually our entire sample is composed of women, we cannot speak to gender differences in treatment effects. However, qualitative research among the cashew factory workers in our setting indicates that both men and women face redistributive pressure (McNeill and Pierotti, 2021). This is consistent with previous studies, some (but not all) of which document redistributive pressure among both genders (e.g. Beekman et al., 2015; Boltz et al., 2019; Squires, 2021).

Finally, we find that treatment effects are concentrated among workers who have below median baseline earnings (Col. 5). While only suggestive, this is potentially consistent with the idea that, in the cross-section, workers with lower labor supply levels are those who face larger disincentives to work due to the social tax.

7.4 Effects on transfers

We designed the blocked accounts so that the expected *cash* component of biweekly take-home pay did not decline. Consistent with the fact that workers typically select thresholds higher than baseline earnings levels, the estimated effects on take-home cash pay are actually positive, though statistically insignificant. This indicates that treated workers had similar or slightly higher levels of liquid disposable income. Under our hypothesis that our intervention

does not lower the tax rate on cash-on-hand, we should see no decline in transfers to others.

Table 5: Transfers to Network Members

	(1)	(2)	(3)	(4)
Panel A: Likelihood of transfer	To anyone	To family	To non-family	From anyone
Private (vs. non-private)	0.0419 (0.0464) [0.37]	0.0540 (0.0582) [0.35]	-0.00523 (0.0583) [0.93]	-0.0780 (0.101) [0.44]
Non-private mean	0.784	0.510	0.405	0.314
Panel B: Amount transferred	To anyone	To family	To non-family	Net amount
Private (vs. non-private)	4998.2 (2532.5) [0.049]	3918.2 (2398.1) [0.10]	1080.0 (951.8) [0.26]	3261.5 (3244.8) [0.32]
Non-private mean	13128.8	9966.7	3162.1	4848.4

Notes: Data from in-person endline survey, covering transfers in past 3 months (in Phase 2). Observations at worker level. In Panel A, dependent variable is a binary indicator for providing any transfer outside the household to individuals in the given category in Cols. (1)-(3) and receiving any transfer in Col. (4). Panel B dependent variable is the continuous transfer amount sent in Cols. (1)-(3) and the net amount of transfers sent (transfers sent minus transfers received) in Col. (4). Regressions include strata dummies. N=298 Phase 2 workers. Robust standard errors in parentheses.

In Table 5, we examine treatment effects on transfers to individuals outside of the worker’s household. We focus on Phase 2, since the endline survey asked about transfers in the past 3 months to improve recall—covering only this phase of the study. In Panel A, we examine the extensive margin of making or receiving any transfers. We do not find significant changes in either the likelihood of having made a transfer in the last three months (Cols. 1-3) or the likelihood of having received a transfer in the last three months (Col. 4).

In Panel B, we find that workers increase total outgoing transfers to others (Col. 1, 38%, $p=0.049$). However, the impact on net transfers, while positive, is statistically insignificant (Col. 4, $p=0.32$). Given that the endline survey was conducted shortly after the Phase 2 unblock date, this suggestive increase in outgoing transfers may reflect a combination of increased transfers from higher cash-on-hand during the intervention period, as well as redistribution from the unblocked savings. This latter force could reflect both altruism, as well as the likely possibility that τ_2 —the social tax rate on the blocked savings account—is not zero. Overall, as intended by our design, our results indicate that the income gains achieved by Private group workers did not come at the expense of lower redistribution to others. Rather, they may have led to aggregate welfare gains.³¹

³¹In addition, we find positive (but statistically insignificant) improvements for workers in a measure of subjective well-being. Consistent with the results in Section 7.3, we find no changes in average reported intra-household bargaining power (Donald, forthcoming).

8 Confounds and Mechanisms

8.1 Fairness concerns

If workers who are not selected for a Private account feel unfairly treated or are disgruntled, they may lower their productivity (e.g. Breza et al., 2018)—biasing our treatment effects upwards. We designed our study to minimize the scope for such morale effects in two ways. First, in each phase, treatment assignment occurred by selecting (confidential) worker IDs via a lottery in the factory, conducted by us, with the bank present. Consequently, workers knew that selection was not an indication of favoritism but was rather a matter of chance, with each worker having the same probability of selection. Second, relatedly, both the marketing of the blocked accounts and the lottery conveyed that the employer, Olam, had no role in picking who received accounts. Similarly, the training for the accounts and the paperwork sign-up process were handled exclusively by research team (i.e. IPA) staff members and bank representatives. Consequently, even if workers felt disappointed that they did not receive a Private account, it is unclear why this should manifest as retaliation toward the firm. Moreover, unlike most previous morale effects studies, because 100% of wages are based on piece rates, any reduction in output hurts not only the firm but also the worker.

In addition, we directly test for this confound by leveraging the weeks between the announcement of treatment assignment and the beginning of the blocked period. If workers are disgruntled about not receiving the Private accounts, then we would expect to see some change in their output immediately once they learn their treatment status. However, we see no discernible effect on output during the announcement period (Appendix Table A5). Rather, the effects only arise after treatment workers’ savings are actually shielded from redistributive pressure. We can reject that, on average, the treatment effect of the Private accounts is the same in the announcement and post-treatment periods (Col. 1, $p=0.027$).

One potential concern with this analysis is that the effects of the blocked accounts may only become salient, and therefore only cause disgruntlement, once deposits actually start occurring. However, it is unclear why this should be the case from the perspective of the disgruntled workers (i.e. the Control or Non-private groups). The announcement (which is done via public lottery) was substantively more salient than deposits (which occur privately behind the scenes on occasional paycycles, with no public signal to other workers). Consequently, if morale effects were driving the large magnitude of our treatment effects, we might have expected to detect at least some responsiveness in the announcement period.

Moreover, in Phase 2, salience during the announcement period should be especially substantial: the accounts have already existed in the factory for months, and even the workers not offered Private accounts (i.e. the Non-private workers) receive an overview and training on how the accounts work. However, we see no evidence for announcement effects

even in Phase 2: the estimated coefficient is 12.9 FCFA, corresponding to 0.8% of the control mean (Appendix Table A5, Cols. 2-3, $p=0.859$). In contrast, effects arise immediately once the accounts are active for all workers: in the first paycycle when accounts are active, workers in Private increase earnings by 182 FCFA or 11.4% ($p=0.082$, Col. 3). This is despite the fact that, in this first active paycycle, workers would not have seen deposits yet occur in their blocked accounts. This effect in the first active paycycle is statistically distinguishable from the effect in the announcement period at the 10% level ($p=0.078$). This further indicates that the lack of discernible announcement effects does not stem simply from a lack of salience.

Finally, this concern cannot explain why effects are concentrated among workers who report more redistributive pressure at baseline. There is no a priori reason why such workers should exhibit stronger fairness concerns.

8.2 Interpretation: Effect of treatments on social tax

One potential concern with the Non-private accounts is that they may lead to a higher tax rate than the status quo (τ_1) by making the worker’s cash-on-hand especially salient—leading us to overestimate the treatment effects of reducing redistributive pressure. Note that this concern can only arise if redistributive pressure has distortionary effects, and so qualitatively does not undermine our test. However, it has implications for interpreting magnitudes.

In our setting, because paydays are publicly known to the network, the status quo level of visibility around when workers have cash on hand is quite publicly salient. In addition, it is not the case that workers are substantively more flush with liquidity after the unblock date relative to their normal paycycles: among workers who achieve savings in the Private accounts in Phase 2, their total savings at the end of the 3-month blocked period is roughly comparable to (i.e. equivalent to 114% of) their average take home cash pay in any one single biweekly paycycle. Moreover, in the Non-private treatment, the network learns that funds will be unblocked in the next week, but for the cash to be accessible, the worker must physically withdraw it from the bank (which could happen weeks in the future), and can even choose to block the funds again for another cycle so that they are not yet available. Consequently, while we cannot rule out that the Non-private messages would increase requests relative to the status quo, it is unclear whether they would necessarily do so.

Finally, we find that the magnitude of treatment effects when comparing Private to Control is similar to the effects when comparing Private to Non-private (Table 3, Cols. 2-3). If the tax rate under Non-private were much higher than under Control, we may have expected the latter comparison to yield substantively larger effects. While certainly not definitive, this provides one suggestive benchmark against the likelihood of this concern.

8.3 Self-control

Blocked accounts could boost effort if workers have self-control problems in consumption. Time inconsistent sophisticates may decide it is not worth working hard today because their future selves will be tempted to frivolously spend savings. However, under this mechanism alone, take-up should be similar between the Private and Non-private blocked account treatments: in both cases, sophisticates should see value in the accounts and choose them. Redistributive pressure is therefore necessary to explain our results. Consequently, self-control is not a confound per se: it does not undermine our qualitative test for the existence of redistributive pressure. However, conditional on the Private accounts reducing transfer requests, present-focus could still help contribute to our observed treatment effects.

To gauge the potential relevance of this, we test a core prediction of basic time inconsistency models. We surprise workers with the option to opt out of directly depositing earnings into their accounts for the upcoming payday—randomly varying whether this option is provided four days before the payday, or on the payday itself. Because workers are always paid several days or more after the end of the paycycle—to allow the factory time to tally earnings—this offer occurs after the effort decision for that paycycle has already been made. Note that we would expect a small proportion of workers to opt out in a given paycycle even in the absence of any self-control problems—for example, due to shocks that may increase cash needs in some weeks. Rather, under basic time inconsistency models, the key prediction is that workers should be more likely to opt out on the payday itself, relative to further from the payday.³²

In contrast to this prediction, the proportion of workers who decide to keep their earnings in the blocked account 4 days before payday is 86%, versus 94% on payday (Appendix Figure A7). These means are not statistically distinguishable from each other, and the relative magnitudes actually go in the opposite direction from what one would expect under time inconsistency.

Thus, while individuals may face self-control problems in consumption and savings in general, as has been documented in other settings (e.g., Ashraf et al., 2006; Brune et al., 2021), these patterns suggest this mechanism is unlikely to strongly drive the effects of the

³²Present-focused sophisticates will seek to tie their hands to avoid their future self from splurging. However, when payday arrives, such workers may be tempted to keep all their earnings that day, just this one time. Under quasi-hyperbolic time preferences, this test relies on appropriately defining time periods: it is valid if the “self” on the payday demonstrates present focus. This is a common assumption in the literature, and is supported by previous work (Kaur et al., 2015; Augenblick et al., 2015; Augenblick, 2018). Under hyperbolic time preferences, this test is valid regardless of the length of time periods in the utility function. Note that we did not collect any baseline measures of self-control across workers. In addition, because there is a lag between the last day of a paycycle and when the funds for that paycycle are disbursed, we cannot use changes in output over the paycycle as a proxy for present focus as in Kaur et al. (2015). Consequently, we cannot examine heterogeneous effects by self-control.

blocked accounts in our specific experiment. This may be because, for example, optimal commitment device design hinges on fitting incentives to individual-level parameters, which can be difficult (Heidhues and Kőszegi, 2009; Bai et al., 2021). Finally, note that the blocked accounts cannot help with self-control problems in effort provision, because the accounts push the receipt of earnings (i.e. the returns to effort) even further into the future, which would actually *decrease* the effort of present-focused agents (O’Donoghue and Rabin, 1999; Kaur et al., 2015).

8.4 Privacy concerns

Another potential concern with our design is that the mechanics through which the Non-private treatment was implemented—text messages sent to others about the worker’s savings account—could introduce potential confounds. For example, workers might not want their network members to be spammed by a publicity SMS, they might prefer that their financial matters stay private, they may be ashamed of needing to use a blocked account to save, or they may fear theft. We offer four sets of arguments against privacy concerns as being the primary driver of our treatment effects.

First, note that privacy concerns cannot on their own explain all our results. They could undermine the interpretation of the Phase 2 effects (where we introduce the Non-private arm), but could not explain the Phase 1 treatment effects, which are very similar in magnitude.

Second, recall from Section 7.1 that we asked workers what factors drive their decision to refuse the Non-private account offer. While 96% of workers cite an increase in expected transfer requests, no workers volunteer the above privacy concerns as a driving factor (Appendix Figure A4). Third, such concerns cannot explain our heterogeneity effects: why treatment effects are concentrated among those who report facing more redistributive pressure at baseline. It is unclear why such workers should be more likely to value privacy, unless that value arises because of the redistributive pressure they face.

Fourth, we supplement these arguments with tests to assess whether workers could plausibly value privacy enough to explain why workers leave so much money on the table by refusing Non-private accounts (relative to Private accounts). Specifically, we undertake two supplementary placebo tests in Phase 2. In the first placebo test, we ask workers for permission to send two publicity text messages about the bank’s savings product to their network members. These messages would include the information, “The BPCI is a bank located in [factory site], which offers financial products to help workers save.” We do not offer any compensation in exchange. 95% of Non-private group workers give permission to have these messages sent to their network. Consequently, it is unlikely that a desire to

avoid having family members receive a “spam” SMS advertisement drives low take-up in the Non-private condition.

We construct a second, stronger test by leveraging the cross-randomization in our design. We undertake the test with workers who took up *Private* accounts in *Phase 1*, but were offered *Non-private* accounts in *Phase 2*. Three months after the end of Phase 1, we ask these workers for permission to send promotional text messages to their network members advertising that they had saved in a blocked account through the bank in the past (i.e. in Phase 1).³³ The text includes the language, “Last July, [worker name] used a blocked savings account with the BPCI that helped her save money”. In exchange, workers are offered a small token compensation of 1,000 FCFA—corresponding to 3 hours of work, and less than 4.5% of the estimated Phase 2 total earnings gain for workers who take up Private accounts. This SMS incorporates several features of the Non-private treatment: publicizing the bank’s blocked account product, giving the name of the worker, and stating that the worker had saved in a blocked account. However, this information is conveyed for accounts where the money would likely be spent long ago—making it easy for the worker to credibly state there are no longer funds available.³⁴ A striking 88% of workers agree to this offer.

The above two exercises are, of course, not exactly equivalent to the Non-private treatment. However, together, they indicate that workers do not inherently have an aversion to having some of their financial information revealed to network members. This suggests that, while privacy concerns may be present, they are unlikely to impact utility so severely that the Non-private treatment workers would give up 11.5% of their full-time earnings (i.e. the Phase 2 Intent to Treat effect) to avoid publicity. Consequently, we argue that this—combined with the other three arguments cited above—suggests that privacy concerns are unlikely to be the primary reason that the take-up of Non-private accounts plummets in Phase 2.

Finally, fear of theft also cannot explain the low take-up of Non-private accounts. Workers walk home from the factory with their entire cash earnings in their pockets twice each month, on days that are publicly known.³⁵ However, not a single worker in our sample reports ever having been robbed on the way home from work, and only 1.6% report ever having faced theft in relation to payday. In contrast, the withdrawal of savings from the

³³We undertake this with workers assigned only to Non-private in Phase 2 because we did not want workers with Private accounts in Phase 2 being asked by network members if they still had a blocked account.

³⁴Recall from Section 3 that holding savings itself does not violate redistributive norms; for example, workers regularly use informal illiquid savings technologies. Rather, it is considered unacceptable if a worker is known to have money and turns down a transfer request. Consequently, revealing that an account existed in the past would not trigger social disapprobation. In contrast, in the Non-private treatment, revealing that the worker has savings that will be unblocked next week would lead to pressure to make transfers.

³⁵Note that among workers with blocked accounts in Phase 2, the magnitude of savings at the end of the blocked period is slightly less than the average worker’s level of take-home cash pay in Phase 2.

bank (the timing of which can be chosen by the worker) is considerably more private than the factory’s payday cash payments.

8.5 Goal setting

Related to the above, because the blocked accounts require selecting a threshold, they may motivate workers to work harder due to a goal-setting or soft commitment motive. However, note that this cannot explain why take-up should be so drastically different under Private vs. Non-private accounts.

In addition, a goal-setting motive would imply that workers’ earnings should be bunched around the thresholds: once the goal is met, the incentive to work beyond it is much lower. In contrast, our hypothesized mechanism of redistributive pressure implies no such bunching: the accounts are only useful if workers overshoot the threshold, and we would expect heterogeneity in output beyond the threshold. Consistent with this latter prediction, we see no evidence of bunching in the data (Appendix Figure A6). Among workers with blocked accounts, earnings are within 10% of the chosen threshold in only 7.5% of paycycles. In addition, conditional on earning weakly above the threshold, earnings are on average 31% higher than the worker’s chosen threshold level.

9 Estimation of the Social Tax Rate

Our experimental design can be used to estimate the social tax rate faced by workers in our context. As before, we provide an overview here, with detailed derivations in Appendix A.3.

Building on the model introduced in Section 2, we can rewrite Equation 2 to obtain the following expression:

$$\frac{de_2}{e_2} = \zeta \frac{d(1 - \tau_2)}{(1 - \tau_2)}, \quad (4)$$

where ζ is the compensated elasticity of labor supply. This expression describes how e_2 changes with τ_2 —starting from the case where τ_2 equals τ_1 , hence e_2 equals e_1 (i.e. moving from Figure 2A to 2B). To bring this equation to the data, we apply the fact that a marginal relative change can be approximated by the natural logarithm of a percentage change. We can thus re-write Equation 4 as:

$$\frac{1 - \tau_1}{1 - \tau_2} = \left(\frac{e_1}{e_2} \right)^{\frac{1}{\zeta}} \quad (5)$$

We can recover an estimate of $\frac{e_1}{e_2}$ from our main treatment effect estimate of 11.4% (Table 3, Col. 1). Since earnings are a linear function of production due to the piece rate, this is the treatment effect on both earnings and output, which we use as our labor supply measure. To obtain an estimate for ζ , we exploit the fact that: $\zeta \equiv \zeta_a + \zeta_e$, where ζ_a is the extensive margin elasticity of attendance with respect to the net-of-tax wage, and ζ_e is the intensive margin elasticity of effort (conditional on attendance).

To estimate the labor supply elasticity, we partner with the factory to randomize workers' piece rate wages. At the end of the experiment, over the course of 6 days, we randomized the piece rate over 4 possible values: 15% below the usual piece rate, the usual piece rate, 15% above the usual rate, and 30% above the usual rate. We randomize daily at the worker level, implemented by each worker drawing a slip out of a bag that determines her piece rate for that day in an i.i.d. fashion.³⁶ Due to feasibility constraints in how this exercise could be implemented, workers learn their piece rate after arriving at the factory. Consequently, this variation allows us to estimate the intensive margin elasticity, ζ_e . This piece rate variation gives an estimate of ζ_e of 0.17 (Table 6, Col. 1).

Table 6: Labor Supply Elasticity

	log(Productivity)			
	(1)	(2)	(3)	(4)
log(Piece-rate)	0.166 (0.0703) [0.019]	0.175 (0.0704) [0.013]	0.168 (0.0708) [0.018]	0.159 (0.0707) [0.025]
Day FE	Yes	No	No	No
Linear time trend	No	No	Yes	Yes
Quadratic time trend	No	No	No	Yes
N: worker-days	1528	1528	1528	1528
N: workers	300	300	300	300

Notes: Unit of observation is worker-day. Workers learned their piece rate for the day after arriving to work. Dependent variable is log output. Col. (1) controls for day fixed effects, Col. (2) has no time controls, Cols. (3)-(4) control for a linear and quadratic time-trend, respectively. Standard errors clustered by worker.

To estimate the attendance elasticity, ζ_a , we exploit the fact that our intervention—which, by lowering the tax rate, mimics an increase in the effective piece rate per unit of

³⁶Note that in expectation, workers' average piece rate will be higher than the status quo over the 6-day period. To protect workers against earnings losses, we inform them that in the rare possibility that their average piece rate over the 6-day period is lower than the status quo rate, we will apply the status quo rate to their total production over those days. See Appendix A.4 for details. In addition, note that if there is inter-temporal substitution in labor effort across days—where working harder one day leads workers to work less hard the next day due to fatigue—then simply examining the change in effort with the daily change in piece rates would lead us to over-estimate the elasticity. However, we find no evidence for such effects: if anything, we find that working harder one day leads to increased effort the next day.

labor supply—moved both the extensive and intensive margin. Using the ratio of effects on attendance vs. the intensive margin in Table 3, we estimate an implied value of ζ_a of 0.94. Consequently, the total estimated value of ζ is 1.11. Note that using this decomposition approach likely is an upper bound on the elasticity, since if both the productivity and attendance margins had been available during the piece-rate variation exercise, workers would likely have increased productivity at most as much as they did when only the productivity margin was available. Using an upper bound for the elasticity yields a more conservative estimate of the tax rate. However, given the assumptions involved in going from the piece rate exercise to an elasticity estimate, we view this as only suggestive. In computing our tax rate estimate, we consequently present values for a range of elasticities.

Finally, expression 5 requires an estimate of τ_2 , the social tax rate faced by treated workers under the Private account. The most conservative assumption (yielding the smallest tax rate estimate) is that $\tau_2 = 0$: the accounts fully eliminate any social tax. In reality, the effects on transfers in Table 5, along with the fact that many workers set thresholds above baseline earnings, suggest that τ_2 is likely greater than zero. We consequently also present estimates for a range of values of τ_2 .

We present estimates of the social tax rate implied by our results in Table 7. Panel A provides estimates for the average worker based on the Intent to Treat (ITT) effects of the Private accounts, while Panel B provides estimates using the Treatment on the Treated (ToT) effects for the subset of workers who choose the accounts. Using our estimated elasticity of 1.11, we estimate that the average worker in our sample faces a baseline social tax rate of 9-14% (Panel A). Smaller values of the attendance elasticity, such as the estimate of 0.15 in Goldberg (2016), would imply a total labor supply elasticity of 0.32—yielding much larger estimates of the tax rate of at least 29%. In addition, the fact that only about half of workers take up the accounts indicates potentially substantial heterogeneity in the social tax rate. Among those who take up the Private accounts, revealing that they face a social tax, the estimated tax rate is at least 19% (Panel B). This is consistent with previous work indicating that some individuals face a high tax rate, while others face very little (e.g. Squires, 2021).

We present the tax rate estimates in Table 7 to provide some guidance on the range of likely values. Ultimately, our treatment effects are a composite of the tax rate and labor supply elasticity. Thus, even if the tax rate were modest, to account for our treatment effects, this would require a large effort elasticity—which in turn would imply that even modest tax rates can lead to large changes in labor supply behavior. Consequently, the large treatment effects of the Private accounts in Table 3 ultimately provide a useful and transparent signal on the potential for meaningful distortions on labor supply.

Table 7: Social tax rate estimates

	Labor supply elasticity					
	0.32	0.50	0.75	1.00	1.11	1.25
Panel A: Sample average (ITT)						
Endline rate: 0%	29%	19%	13%	10%	9%	8%
Endline rate: 2.5%	30%	21%	16%	12%	12%	11%
Endline rate: 5%	32%	23%	18%	15%	14%	13%
Panel B: Compliers (ToT)						
Endline rate: 0%	51%	37%	26%	20%	19%	17%
Endline rate: 2.5%	52%	38%	28%	22%	21%	19%
Endline rate: 5%	53%	40%	30%	24%	23%	21%

Notes: This table presents the baseline social tax rate, τ_1 , faced by workers, estimated for various values of τ_2 (the social tax rate on earnings increases for workers assigned to the Private arm) and of ζ (the labor supply elasticity). Note that our experimental estimates imply a labor supply elasticity $\zeta^c = 1.11$. Equation (5) is used for the computation of τ_1 .

10 Conclusion

Informal transfers among kin groups and social networks are important for coping with risk. Because they substitute for missing insurance markets, they are typically viewed as unequivocally positive. Our findings indicate that these important welfare benefits may come at a cost: social insurance can turn into social taxation, creating a disincentive to work. Because our intervention minimizes the scope for income effects, our results do not necessarily reflect the policy impact of reducing existing redistributive pressure on labor supply. However, the large magnitude of our treatment effects points to sizable distortions—raising the potential for the social tax to meaningfully lower aggregate earnings and productivity.

This has potential implications for understanding labor market malaise in developing countries. In many Sub-Saharan African countries, two major impediments to the growth and profitability of formal firms are difficulty in finding enough low-skilled workers to work regularly in formal jobs, as well as low labor productivity among those who do work (McMillan and Zeufack, 2022). For example, in Côte d’Ivoire, these two labor supply challenges, despite high wages, are cited as major obstacles to enabling domestic processing of cashews—considered an important policy priority for economic growth (World Bank, 2018, 2020). The presence of a social tax could contribute to both these labor supply challenges. For example, in our setting, 74% of workers state that taking a formal job would lead to increased transfer requests, despite the fact that such jobs are also readily available to those in workers’ networks (Figure 3, Panel C). Among those who do hold these jobs, our experimental findings suggest the potential for social taxation to substantively lower worker productivity.

More broadly, our results provide empirical grounding for long-held views expressing concerns about the distortionary effects of informal redistributive arrangements (Lewis, 1955;

Tam et al., 1957; Hoff and Sen, 2011; Platteau, 2014). If redistributive pressure affects the incentive to work, it may also affect the willingness to undertake other costly actions that are needed to generate future income. For example, could such pressures undermine the willingness to bear the risk to adopt new technologies, or undertake long-run investments such as in human capital? These possibilities point to interesting directions for future research.

While our intervention enables workers to increase their earnings without reducing redistribution to the network, we do not necessarily view it as a scalable policy solution. Rather, we view our intervention as primarily a tool to test for the relevance of the social tax for labor supply decisions. However, the success and popularity of our blocked account product speak to the potential of solutions that use illiquidity to help workers recoup returns from effort.³⁷ This is in line with strategies, such as ROSCA participation, that are already commonly employed in this setting. Consistent with the demand for such illiquid savings vehicles, the firm with which we worked has continued to offer workers the option to direct deposit earnings into blocked accounts, even after the conclusion of our study. However, the implications of such financial tools for risk sharing are less clear. General implementation may not necessarily be Pareto-improving, as it could reduce transfers to others. Our study suggests the importance of understanding these trade-offs, and developing scalable tools to lower social taxation without undermining risk-sharing arrangements (e.g. Dupas et al., 2017; Mobarak and Rosenzweig, 2012; Banerjee et al., 2021).

Finally, our findings suggest an additional route through which improving formal safety nets could boost productivity: by reducing demands for redistribution on others in beneficiaries’ networks. For example, could universal access to formal health or unemployment insurance have externality benefits due to decreased social taxation, amplifying their effects on investment and output? The possibilities above are of course only speculative. However, they suggest potentially broad implications of a social tax for economic behavior and policy.

References

- ALBY, P., E. AURIOL, AND P. NGUIMKEU (2020): “Does Social Pressure Hinder Entrepreneurship in Africa? The Forced Mutual Help Hypothesis,” *Economica*, 87, 299–327.
- ALMÁS, I., A. ARMAND, O. ATTANASIO, AND P. CARNEIRO (2018): “Measuring and

³⁷We find higher take-up rates of formal illiquid savings devices than many past studies. Both our qualitative work and earlier studies indicate that trust in institutions is a major determinant of account take-up (e.g. Bachas et al., 2021). Many workers who did not take up reported being swindled by past financial institutions. Take-up increased in each subsequent implementation wave, with individuals learning from others’ experiences with the instrument. Moreover, virtually everyone who took up the account once did so again when offered a subsequent time.

- changing control: Women’s empowerment and targeted transfers,” *The Economic Journal*, 128, F609–F639.
- ANDERSON, S. AND J. M. BALAND (2002): “The economics of roscas and intrahousehold resource allocation,” *The Quarterly Journal of Economics*, 117, 963–995.
- ASHRAF, N. (2009): “Spousal Control and Intra-household Decision Making: An Experimental Study in the Philippines,” *American Economic Review*, 99, 1245–77.
- ASHRAF, N., D. KARLAN, AND W. YIN (2006): “Tying Odysseus to the Mast: Evidence From a Commitment Savings Product in the Philippines,” *The Quarterly Journal of Economics*, 121, 635–672.
- (2010): “Female empowerment: Impact of a commitment savings product in the Philippines,” *World Development*, 38, 333–344.
- AUGENBLICK, N. (2018): “Short-term time discounting of unpleasant tasks,” *Unpublished Manuscript*.
- AUGENBLICK, N., M. NIEDERLE, AND C. SPRENGER (2015): “Working over Time: Dynamic Inconsistency in Real Effort Tasks,” *The Quarterly Journal of Economics*, 130, 1067–1115.
- BACHAS, P., P. GERTLER, S. HIGGINS, AND E. SEIRA (2021): “How debit cards enable the poor to save more,” *The Journal of finance*, 76, 1913–1957.
- BAI, L., B. HANDEL, E. MIGUEL, AND G. RAO (2021): “Self-control and demand for preventive health: Evidence from hypertension in India,” *Review of Economics and Statistics*, 103, 835–856.
- 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., E. BREZA, A. G. CHANDRASEKHAR, E. DUFLO, M. O. JACKSON, AND C. KINNAN (2021): “Changes in social network structure in response to exposure to formal credit markets,” Tech. rep., National Bureau of Economic Research.
- BANERJEE, A. V. AND E. DUFLO (2007): “The Economic Lives of the Poor,” *The Journal of Economic Perspectives*, 21, 141–167.
- BEEKMAN, G., M. GATTO, AND E. NILLESEN (2015): “Family networks and income hiding: evidence from lab-in-the-field experiments in rural Liberia,” *Journal of African Economies*, 24, 453–469.

- BERNHARDT, A., E. FIELD, R. PANDE, AND N. RIGOL (2019): “Household matters: Revisiting the returns to capital among female microentrepreneurs,” *American Economic Review: Insights*, 1, 141–60.
- BOLTZ, M., K. MARAZYAN, AND P. VILLAR (2019): “Income hiding and informal redistribution: A lab-in-the-field experiment in Senegal,” *Journal of Development Economics*, 137, 78–92.
- BOLTZ, M. AND P. VILLAR (2013): “Les liens des migrants internes et internationaux à leur ménage d’origine : portraits croisés de familles étendues sénégalaises,” *Autrepart*, N° 67-68, 103–119.
- BREZA, E., S. KAUR, AND Y. SHAMDASANI (2018): “The morale effects of pay inequality,” *The Quarterly Journal of Economics*, 133, 611–663.
- BRUNE, L., E. CHYN, AND J. KERWIN (2021): “Pay Me Later: Savings Constraints and the Demand for Deferred Payments,” *The American Economic Review*, 111, 2179–2212.
- CASTILLA, C. AND T. WALKER (2013): “Is ignorance bliss? The effect of asymmetric information between spouses on intra-household allocations,” *American Economic Review*, 103, 263–68.
- COATE, S. AND M. RAVALLION (1993): “Reciprocity without commitment: Characterization and performance of informal insurance arrangements,” *Journal of Development Economics*, 40, 1–24.
- DE WEERDT, J. AND S. DERCON (2006): “Risk-sharing networks and insurance against illness,” *Journal of Development Economics*, 81, 337–356.
- DE WEERDT, J. AND M. FAFCHAMPS (2011): “Social Identity and the Formation of Health Insurance Networks,” *The Journal of Development Studies*, 47, 1152–1177.
- DE WEERDT, J., G. GENICOT, AND A. MESNARD (2019): “Asymmetry of information within family networks,” *Journal of Human Resources*, 54, 225–254.
- DI FALCO, S. AND E. BULTE (2011): “A dark side of social capital? Kinship, consumption, and savings,” *The Journal of Development Studies*, 47, 1128–1151.
- DI FALCO, S., F. FERI, P. PIN, AND X. VOLLENWEIDER (2018): “Ties that bind: Network redistributive pressure and economic decisions in village economies,” *Journal of Development Economics*, 131, 123–131.
- DILLON, B., J. DE WEERDT, AND T. O’DONOGHUE (2021): “Paying More for Less: Why Don’t Households in Tanzania Take Advantage of Bulk Discounts?” *The World Bank Economic Review*, 35, 148–179.

- DONALD, A. (forthcoming): “Household-level Impacts of Women’s Financial Control: Experimental Evidence from Cote d’Ivoire,” *Unpublished Manuscript*.
- DUPAS, P., A. KEATS, AND J. ROBINSON (2017): “The Effect of Savings Accounts on Interpersonal Financial Relationships: Evidence from a Field Experiment in Rural Kenya,” *The Economic Journal*, 129, 273–310.
- DUPAS, P. AND J. ROBINSON (2013): “Why Don’t the Poor Save More? Evidence from Health Savings Experiments,” *The American Economic Review*, 103, 1138–1171.
- FAFCHAMPS, M. (2011): “Risk Sharing Between Households,” in *Handbook of Social Economics*, ed. by J. Benhabib, A. Bisin, and M. O. Jackson, North-Holland, vol. 1, 1255–1279.
- FAFCHAMPS, M. AND S. LUND (2003): “Risk-sharing networks in rural Philippines,” *Journal of Development Economics*, 71, 261–287.
- FELDHAUS, C. AND J. MANS (2014): “Who do you lie to? Social identity and the cost of lying,” Working Paper Series in Economics 76, University of Cologne, Department of Economics.
- FIALA, N. (2018): “Business Is Tough, but Family Is Worse: Household Bargaining and Investment Decisions in Uganda,” Tech. rep., Working paper, University of Connecticut.
- FIELD, E., R. PANDE, N. RIGOL, S. SCHANER, AND C. TROYER MOORE (2021): “On her own account: How strengthening women’s financial control impacts labor supply and gender norms,” *American Economic Review*, 111, 2342–75.
- GNEEZY, U. (2005): “Deception: The Role of Consequences,” *American Economic Review*, 95, 384–394.
- GOLDBERG, J. (2016): “Kwacha Gonna Do? Experimental Evidence about Labor Supply in Rural Malawi,” *American Economic Journal: Applied Economics*, 8, 129–149.
- (2017): “The effect of social pressure on expenditures in Malawi,” *Journal of Economic Behavior & Organization*, 143, 173–185.
- GRIMARD, F. (1997): “Household consumption smoothing through ethnic ties: evidence from Cote d’Ivoire,” *Journal of Development Economics*, 53, 391–422.
- GRIMM, M., F. GUBERT, O. KORIKO, J. LAY, AND C. J. NORDMAN (2013): “Kinship ties and entrepreneurship in Western Africa,” *International Journal of Entrepreneurship & Small Business*, 26, 125–150.
- HAUSMAN, J. A. (1985): “Taxes and labor supply,” New York: North-Holland Publishers, vol. 1 of *Handbook of Public Economics*, chap. 4, 213–263.
- HEIDHUES, P. AND B. KŐSZEGI (2009): “Futile attempts at self-control,” *Journal of the European Economic Association*, 7, 423–434.


- HOFF, K. AND A. SEN (2011): “The Kin System as a Poverty Trap?” in *Poverty Traps*, ed. by S. Bowles, S. N. Durlauf, and K. Hoff, Princeton, NJ: Princeton University Press, 95–115.
- 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.
- KARAIVANOV, A. AND R. M. TOWNSEND (2014): “Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes,” *Econometrica: Journal of the Econometric Society*, 82, 887–959.
- KAUR, S., M. KREMER, AND S. MULLAINATHAN (2015): “Self-Control at Work,” *The Journal of Political Economy*, 123, 1227–1277.
- KENNEDY, P. (1988): “African Capitalism: The Struggle for Ascendancy; and Dietz, J. and D. James (eds)(1990),” *Progress Toward Development in Latin America*.
- LEWIS, A. W. (1955): *Theory of Economic Growth*, Routledge, 1 edition ed.
- LIGON, E., J. P. THOMAS, AND T. WORRALL (2002): “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies,” *The Review of Economic Studies*, 69, 209–244.
- MCMILLAN, M. AND A. ZEUFACK (2022): “Labor Productivity Growth and Industrialization in Africa,” *Journal of Economic Perspectives*, 36, 3–32.
- MCNEILL, K. AND R. PIEROTTI (2021): “Reason-giving for resistance: obfuscation, justification and earmarking in resisting informal financial assistance,” *Socio-Economic Review*.
- MIRACLE, M. P., D. S. MIRACLE, AND L. COHEN (1980): “Informal savings mobilization in Africa,” *Economic Development and Cultural Change*, 28, 701–724.
- MOBARAK, A. M. AND M. ROSENZWEIG (2012): “Selling Formal Insurance to the Informally Insured,” *Unpublished Manuscript*.
- O’BRIEN, R. (2012): “Depleting Capital? Race, Wealth and Informal Financial Assistance,” *Social Forces*, 91, 375–396.
- O’DONOGHUE, T. AND M. RABIN (1999): “Doing It Now or Later,” *The American Economic Review*, 89, 103–124.
- PLATTEAU, J.-P. (2000): *Institutions, Social Norms, and Economic Development*, vol. 1, Psychology Press.
- (2014): “Redistributive Pressures in Sub-Saharan Africa: Causes, Consequences, and Coping Strategies,” in *Africa’s Development in Historical Perspective*, Cambridge University Press, 153–207.

- PORTES, A. (1998): “Social Capital: Its Origins and Applications in Modern Sociology,” *Annual Review of Sociology*, 24, 1–24.
- RILEY, E. (2022): “Resisting social pressure in the household using mobile money: Experimental evidence on microenterprise investment in Uganda,” Tech. rep., Working Paper, University of Oxford.
- ROSENZWEIG, M. R. (1988): “Risk, Implicit Contracts and the Family in Rural Areas of Low-Income Countries,” *The Economic Journal of Nepal*, 98, 1148–1170.
- ROSENZWEIG, M. R. AND O. STARK (1989): “Consumption Smoothing, Migration, and Marriage: Evidence from Rural India,” *The Journal of Political Economy*, 97, 905–926.
- SCHANER, S. (2015): “Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings,” *American Economic Journal: Applied Economics*, 7, 135–74.
- SCOTT, J. C. (1976): *The Moral Economy of the Peasant: Rebellion and Subsistence in Southeast Asia.*, Yale University Press.
- SOMVILLE, V. (2011): “Daily Collectors, Public Good Provision and Private Consumption: Theory and Evidence from Urban Benin,” Tech. Rep. 1106.
- SQUIRES, M. (2021): “Kinship Taxation as an Impediment to Growth: Experimental Evidence from Kenyan Microenterprises,” *Unpublished Manuscript*.
- STACK, C. B. (1974): *All Our Kin: Strategies for Survival in a Black Community*, Harper and Row.
- TAM, P., B. S. YAMEY, ET AL. (1957): *The Economics of Under-Developed Countries*, University of Chicago Press.
- TOWNSEND, R. M. (1994): “Risk and Insurance in Village India,” *Econometrica: Journal of the Econometric Society*, 62, 539–591.
- WHERRY, F. F., K. S. SEEFELDT, AND A. S. ALVAREZ (2019): “To Lend or Not to Lend to Friends and Kin: Awkwardness, Obfuscation, and Negative Reciprocity,” *Social Forces*, 98, 753–793.
- WORLD BANK (2018): “Cashew Value-Chain Competitiveness Project: Project Appraisal Document,” Tech. rep., The World Bank Group, Washington, DC.
- (2020): “World Development Report 2020: Trading for Development in the Age of Global Value Chains,” Tech. rep., The World Bank Group, Washington, DC.

A Online Appendix


A.1 Appendix Figures

Figure A1: Earnings shocks and transfers




Notes: Relationship between transfers and earnings shocks. Earnings from factory administrative data. Earnings shocks defined as the difference between earnings in the 15 days prior to the survey relative to average earnings in past 3 months excluding that paycycle (i.e. the 75 days before). Transfers from 2 rounds of worker phone surveys. Amount transferred in the 15 days prior to the survey. Observations residualized from survey fixed effects. Transfers top-coded at the 99th percentile. Line of best linear fit reported. 349 observations, from workers not offered a Private savings account. Note that earnings shocks may reflect shocks and/or paycycles where workers increase labor supply because they have increased cash needs, and can therefore spend the money quickly.

Figure A2: Cross-sectional average earnings and transfers




Notes: Relationship between average amount transferred and average factory earnings, both in the 3 months prior to the survey (in FCFA). Earnings from factory administrative data. Transfers from 2 rounds of worker phone surveys. Amount transferred in the 15 days prior to the survey. Observations residualized from survey fixed effects. Transfers top-coded at the 99th percentile. Line of best linear fit reported. 349 observations, from workers not offered a Private savings account.

Figure A3: Psychological cost of lying about earnings' availability




Notes: In-person survey with 488 cashew factory workers in Côte d'Ivoire.

Figure A4: Non-private Take-up : Drivers of Decision




Notes: Drivers of take-up decision for the Non-private account. Elicited from workers assigned to the Non-private treatment who refused to take-up the Non-private blocked accounts; collected when workers report their take-up decision to research staff. Workers were asked what were major factors that drive their decision not to take up the accounts. They could select as many options as they wanted, or provide their own. Figure shows the proportions of workers who select a given option as being important in their decision not to take-up. N=110 workers.

Figure A5: Distribution of chosen threshold relative to baseline earnings




Notes: The figure presents the distribution of the thresholds chosen by workers relative to the given worker's average baseline earnings in the 3 months before the intervention began: $\frac{\text{chosen threshold}}{\text{average baseline paycycle earnings}}$. The ratio is top-coded at 3. N = 153 treated workers holding a private savings account, pooling both Phases of the experiment.

Figure A6: Distribution of treatment period earnings relative to the chosen threshold



Notes: The figure presents the distribution of earnings per paycycle in the post-treatment period (after accounts are active), as a proportion of the worker's chosen threshold: $\frac{\text{earnings per paycycle}}{\text{chosen threshold}}$. N = the 153 treated workers holding a private savings account, pooling both Phases of the intervention. N=2,272 worker-paycycles.

Figure A7: Refusal to Opt Out



Notes: A subset of workers who took up blocked accounts was surprised with the option to opt out of having to deposit earnings increases into the accounts for one upcoming cycle. The figure shows the proportion who refused this offer to opt out, separately for workers offered this option 4 days before the payday (left bar) and on the payday (right bar). N=61 workers chosen for exercise and who had private blocked accounts in Phase 2.

A.2 Appendix Tables

Table A1: Take-up rates - Phase 2

Phase 2	Eligible Ph1	Treated Ph1	Compliant Ph1	Take-up rate	N.
Private	All	All	All	.60	156
Private	Yes	No	-	.72	39
Private	Yes	Yes	All	.62	55
Private	Yes	Yes	Yes	.90	30
Private	No	-	-	.52	62
Public	All	All	All	.14	161
Public	Yes	No	-	.18	44
Public	Yes	Yes	All	.19	59
Public	Yes	Yes	Yes	.31	35
Public	No	-	-	.07	58

Notes: This table disaggregates the Phase 2 take-up results. “Eligible Phase 1” denotes workers who have been eligible at any point in Phase 1 (as opposed to workers newly eligible only in Phase 2). “Treated Phase 1” denotes workers who were offered a Private account during Phase 1. “Compliant Phase 1” denotes workers who took up a private blocked accounts in Phase 1.

Table A2: Treatment effects on quits, during and at the end of the intervention

	Phase 1 & 2		Phase 1		Phase 2	
	During (1)	End (2)	During (3)	End (4)	During (5)	End (6)
Dep. variable: $\mathbb{1}\{quit\}$						
Private account	0.0114 (0.0325) [0.73]	0.00652 (0.0332) [0.84]	0.0180 (0.0414) [0.66]	0.0267 (0.0423) [0.53]	0.00195 (0.0495) [0.97]	-0.0225 (0.0510) [0.66]
Sample mean in control	0.22	0.25	0.20	0.21	0.26	0.30
N: worker-waves	774	774	457	457	317	317
N: workers	474	474	354	354	317	317

Notes: Dependent variable in cols. (1), (3) and (5): dummy indicating if the worker quit the factory during the treatment period (i.e. had left the factory at least one paycycle before the end of the period). Dependent variable in cols. (2), (4) and (6): dummy indicating if the worker had quit the factory as of the end of the treatment period. The omitted category is not being offered an account or offered a non-private account. All specifications include wave fixed effects. Standard errors clustered by worker.

Table A3: Treatment effects – Robustness to alternative specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Private account	175.9 (70.02) [0.012]	163.7 (69.11) [0.018]	178.4 (68.82) [0.0098]	182.8 (71.26) [0.011]	169.3 (68.72) [0.014]	162.6 (57.98) [0.0053]
Time FE	Day	Day- by-factory	Paycycle- by-factory	Day	Day	Day
Estimation sample	Main	Main	Main	Exclude first active paycycle	Include 2nd announcement paycycle in treatment period	Include first announcement paycycle in baseline
N: observations	137678	137678	137678	137678	137678	137678
N: workers	474	474	474	474	474	474

Notes: Dependent variable: daily earnings. The omitted category is not being offered an account or offered a non-private account. Col. (1) is the same specification as Col. (1) of Table 3. As compared to our main specification, Col. (2) has day-by-factory fixed effects; Col. (3) has paycycle-by-factory fixed effects. Col. (4) shows robustness to excluding the first active paycycle (when all accounts were active) from the estimation of treatment effects by interacting it out. Col. (5) shows robustness to including the 2nd post-announcement period paycycle (when workers had a grace period to get paperwork submitted for the accounts) as part of the treatment period. Col. (6) shows robustness to including the first paycycle post announcement (before any accounts were active) as part of the baseline period. All regressions include worker fixed effects. Standard errors clustered by worker.

Table A4: Treatment effects – Intervention period data only

	DiD	Intervention period data only	
	(1)	(2)	(3)
Private account	184.4 (90.90) [0.043]	166.3 (87.94) [0.060]	180.5 (88.58) [0.042]
Controls	Worker FE	Earnings (linear)	Earnings (linear) + heterogeneity table covariates
N: observations	38040	34054	34054
N: workers	317	317	317

Notes: Dependent variable: daily earnings. The omitted category is not being offered a private account. Col. (1) is the same specification as Col. (3) of Table 3. Cols. (2)-(3) use endline data only, and control for strata and day fixed effects. In addition, Col. (2) controls for baseline earnings (linearly) and Col. (3) for linear baseline earnings plus the baseline heterogeneity variables used in Table 4. In all specifications, the announcement period is interacted out from the treatment effect estimation. Sample includes all workers in Phase 2. Standard errors clustered by worker.

Table A5: Treatment effects during announcement period and first paycycle

	Pooled	Phase 2	
	(1)	(2)	(3)
Daily earnings			
Private X Announcement Period	38.79 (70.41) [0.582]	12.88 (72.55) [0.859]	12.88 (72.55) [0.859]
Private X Treatment Period	175.9 (69.68) [0.012]	180.8 (98.04) [0.066]	
Private X First Treatment Paycycle			182.1 (104.0) [0.081]
Private X Remaining Paycycles			180.4 (101.2) [0.075]
Pval: F-test of equal coefs.	0.027	0.063	0.078
N: worker-days	137678	38040	38040
N: workers	474	317	317

Notes: Dependent variable: daily earnings. Announcement Period takes value one from the date of the treatment assignment announcement until the day before the first accounts are active. In Col. (1)-(2), “Treatment Period” takes the value one from the first day of earnings when all accounts are active (i.e. in which earnings may be deposited into the blocked accounts). In Col. (3), “First Treatment Paycycle” takes the value one during the first paycycle when all accounts are active, and “Remaining Paycycles” takes the value one for all subsequent paycycles when earnings may be deposited into blocked accounts. The p-value of the F-test of equality of treatment effects during the announcement period and the remaining paycycles is reported. Col. (1) pools across Phases 1 and 2. The omitted category in Col. (1) is not being offered an account or offered a non-private account. The omitted category in Col. (2)-(3) is being offered a non-private account. DiD specification. Baseline: 15 days before the treatment assignment announcement. All specifications include day and worker fixed effects, as well as . treatment waves by factory fixed effects. Standard errors clustered by worker-phase.

Table A6: Effort elasticity estimates from experimental piece-rate variation

	log(output)	
	(1)	(2)
log(piece-rate)	0.166** (0.0703)	0.246** (0.115)
Lowest rate excl.	No	Yes
N: worker-days	1528	1164
N: workers	303	301

Notes: All specifications include day fixed-effects. “Lowest rate excl.” is “Yes” when worker-days with the lowest piece-rate (lower than usual) are excluded. Standard errors clustered at the worker level. *, **, *** indicate significance at the 10, 5, or 1% level.

A.3 Model Appendix

A.3.1 Baseline labor supply decision

As described in Section 2, at baseline, a worker solves $\max_{c,e} u(c,e)$ under the budget constraint $BC1 : (1 - \tau_1)we = c$. τ_1 denotes the linear marginal tax rate faced on income, as shown in Figure 2A. Her optimal labor supply decision is thus $e_1((1 - \tau_1)w)$. Denote the baseline choice made by the worker as e_1 .

A.3.2 Labor supply decision under new tax schedule

Suppose we introduce an alternate tax schedule, as shown in Figure 2B. By dampening the social tax rate from τ_1 to τ_2 on earnings increases only (i.e., for all $e \geq e_1$), a kink is introduced in her budget constraint. Specifically, it now is:

$$BC2 : c \leq \mathbb{1}_{e \leq e_1} \{(1 - \tau_1)we\} + \mathbb{1}_{e > e_1} \{(1 - \tau_1)we_1 + (1 - \tau_2)w(e - e_1)\}.$$

Note the trivial result: $e_2 \geq e_1$, where e_2 is the worker's choice under $BC2$.³⁸ In what follows, we thus use the budget constraint $BC2^* : (1 - \tau_1)we_1 + (1 - \tau_2)w(e - e_1) = c$.

Let us introduce $\mathbb{Y}$, defined as $\mathbb{Y} \equiv (\tau_2 - \tau_1)we_1$. Since the choice variable e does not enter in $\mathbb{Y}$, we can re-write the worker's labor supply decision under treatment as $\max_{c,e} u(c,e)$ under $BC2^* : (1 - \tau_2)we + \mathbb{Y} = c$. Her optimal decision is thus $e_2((1 - \tau_2)w, \mathbb{Y})$. Note the similitude in form with the baseline labor supply decision.

A.3.3 Slutsky equation.

Here, we derive Slutsky's equation applied to our model.

Let us define $\tilde{e}((1 - \tau_2)w, u)$ as the Hicksian (compensated) labor supply and $\gamma((1 - \tau_2)w, u)$ as the expenditure function. By duality of utility maximization and expenditure minimization, we have:

$$\tilde{e}((1 - \tau_2)w, u) = e_2((1 - \tau_2)w, \gamma((1 - \tau_2)w, u))$$

Taking the derivative on both sides with respect to $(1 - \tau_2)w$, we have:

$$\frac{\partial \tilde{e}((1 - \tau_2)w, u)}{\partial [(1 - \tau_2)w]} = \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial [(1 - \tau_2)w]} + \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial \mathbb{Y}} \frac{\partial \gamma((1 - \tau_2)w, u)}{\partial [(1 - \tau_2)w]} \quad (A.1)$$

By Shephard's Lemma,

$$\frac{\partial \gamma((1 - \tau_2)w, u)}{\partial [(1 - \tau_2)w]} = -\tilde{e}((1 - \tau_2)w, u)$$

³⁸Proof by contradiction: Suppose that $e_2 < e_1$. Then $BC2$ becomes $(1 - \tau_1)we = c$, which is $BC1$. Since e_1 is the optimal choice under $BC1$, we must have $e_2 = e_1$. This contradicts $e_2 < e_1$.

(Note the minus sign. It comes from the budget constraint being $\mathbb{Y} = c - (1 - \tau_2)we$ and thus from labor supply, e , being a “bad” instead of a “good”.)

Combined with the starting equality, we get:

$$\frac{\partial \tilde{e}((1 - \tau_2)w, u)}{\partial [(1 - \tau_2)w]} = \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial [(1 - \tau_2)w]} - \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial \mathbb{Y}} e_2((1 - \tau_2)w, \mathbb{Y}) \quad (\text{A.2})$$

This is the Slutsky equation that we use in the next paragraph.

A.3.4 Income and substitution effects.

If $\tau_2 = \tau_1$, then $\mathbb{Y} = 0$ and $e_2((1 - \tau_2)w, \mathbb{Y}) = e_1((1 - \tau_1)w) = e_1$.

Starting from this baseline situation, what happens when we dampen the tax rate τ_2 applied above the kink e_1 ?

In other words, we want to compute:

$$\begin{aligned} \frac{de_2((1 - \tau_2)w, \mathbb{Y})}{d\tau_2} &= \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial [(1 - \tau_2)w]} \frac{\partial [(1 - \tau_2)w]}{\partial \tau_2} + \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial \mathbb{Y}} \frac{\partial \mathbb{Y}}{\partial \tau_2} \\ &= -w \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial [(1 - \tau_2)w]} + we_1 \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial \mathbb{Y}} \end{aligned}$$

Applying Slutsky’s equation, where $\tilde{e}((1 - \tau_2)w, u)$ is the Hicksian (compensated) labor supply,

$$\frac{de_2((1 - \tau_2)w, \mathbb{Y})}{d\tau_2} = -w \frac{\partial \tilde{e}((1 - \tau_2)w, u)}{\partial [(1 - \tau_2)w]} + w(e_1 - e_2) \frac{\partial e_2((1 - \tau_2)w, \mathbb{Y})}{\partial \mathbb{Y}} \quad (\text{A.3})$$

Recall that we start from the baseline situation where $e_2((1 - \tau_2)w, \mathbb{Y}) = e_1((1 - \tau_1)w, y) = e_1$. As such $(e_1 - e_2) = 0$ and the second term of the equation drops out.

We thus have:

$$\frac{de_2((1 - \tau_2)w, \mathbb{Y})}{d\tau_2} = -w \frac{\partial \tilde{e}((1 - \tau_2)w, u)}{\partial [(1 - \tau_2)w]} \quad (\text{A.4})$$

There is no income effect; only a substitution effect.

We can further observe that $-\frac{de_2((1 - \tau_2)w, \mathbb{Y})}{d\tau_2} > 0$, which corresponds to our prediction.

A.3.5 Tax rate estimation.

We can use Equation A.4 to estimate the tax rate implied by our empirical estimates. Rearranging this equation, applying the fact that $e_2((1 - \tau_2)w, \mathbb{Y}) = \tilde{e}((1 - \tau_2)w, u)$ (by defi-

dition) and that $-d\tau_2 = d[(1 - \tau_2)]$, we obtain:

$$(A.4) \Leftrightarrow \frac{de_2((1 - \tau_2)w, \mathbb{Y})}{e_2((1 - \tau_2)w, \mathbb{Y})} = \frac{\partial \tilde{e}((1 - \tau_2)w, u)}{\partial(1 - \tau_2)} \frac{(1 - \tau_2)}{\tilde{e}((1 - \tau_2)w, u)} \frac{d(1 - \tau_2)}{(1 - \tau_2)} \equiv \zeta \frac{d(1 - \tau_2)}{(1 - \tau_2)} \quad (A.5)$$

Where ζ is the compensated elasticity of labor supply to the net-of-tax wage rate. This expression describes how e_2 varies when τ_2 changes—starting from the situation where τ_2 equal to τ_1 , hence e_2 equals e_1 .

To bring this equation to the data, we recognize that a marginal relative change can be approximated by the natural logarithm of a percentage change, and thus re-write it as:

$$\begin{aligned} \log\left(\frac{e_2}{e_1}\right) &= \zeta \log\left(\frac{1 - \tau_2}{1 - \tau_1}\right) \\ \Leftrightarrow \frac{1 - \tau_1}{1 - \tau_2} &= \left(\frac{e_1}{e_2}\right)^{\frac{1}{\zeta}} \end{aligned} \quad (A.6)$$

This gives us the equation we will use to estimate the social tax rate faced by workers at baseline.

A.3.6 Elasticity decomposition: Theory.

To obtain an estimate for ζ , the compensated elasticity of labor supply to the net-of-tax wage rate, we recognize that it can be decomposed into the sum of the compensated elasticity of attendance ζ_a and the compensated elasticity of effort (conditional on attendance) ζ_e .

A worker's labor supply (e) is indeed equal to its average effort while working (p) multiplied by its number of days present at work (a). Formally,

$$e = p \times a \quad (A.7)$$

Totally differentiating equation (A.7), we obtain:

$$\begin{aligned} de &= p \times da + a \times dp \\ \Leftrightarrow \frac{de}{e} &= \frac{p \times da}{p \times a} + \frac{a \times dp}{p \times a} \\ \Leftrightarrow \frac{de}{e} &= \frac{da}{a} + \frac{dp}{p} \end{aligned} \quad (A.8)$$

Multiplying each side of the equation by $\frac{(1-\tau)}{d(1-\tau)}$, we obtain the desired result:

$$(A.8) \Leftrightarrow \frac{de}{e} \frac{(1-\tau)}{d(1-\tau)} = \frac{da}{a} \frac{(1-\tau)}{d(1-\tau)} + \frac{\times dp}{p} \frac{(1-\tau)}{d(1-\tau)} \quad (A.9)$$

$$\Leftrightarrow \zeta = \zeta_a + \zeta_e$$

A.3.7 Elasticity decomposition - Practice.

From the piece rate variation implemented at the factory, we obtain an estimate of ζ_e of 0.166 (Table 6, Col. 1). We now need an estimate of ζ_a .

Attendance accounts for 85% of the total effect of the Private treatment on output (Table 3, Col. 1). This implies that the attendance elasticity is 5.67 times as large as the productivity elasticity. It therefore implies an estimate of ζ_a of 0.94.

Since $\zeta = \zeta_a + \zeta_e$, the total estimated value of ζ is 1.11.

A.4 Protocols Appendix

Offer and Implementation of Private blocked accounts (Phases 1-2)

Sample Eligibility. Eligible workers were required to either satisfy a minimum baseline attendance rate (in the 3 months before the start of the study), or be listed by the worker cooperative as a “permanent worker” in the factory plant. The minimum attendance rate was 60% in one factory site, and 45% in the other factory site (due to the second site being newer, with less established workers). In addition, workers were required to have been working at their factory for at least 2 months (to reduce attrition due to turnover). Finally, in one factory location, workers additionally needed a national ID card (‘CNI’) and a certificate of residence as per the bank’s documentation requirements.

Announcements. Representatives of the bank (BPCI) and research team (IPA), with support from Olam, jointly conduct brief announcements in the manual peeling sections of the factories, prior to the launch of each wave. Workers are informed that some of them will be offered a free product to help them save money but that, given that the product cannot be offered to all workers, the beneficiaries will be picked at random. Those selected workers will be invited to brief marketing sessions, in groups of 5-8 workers, covering the product’s key features. If selected workers are absent at the time of the announcement but come back to the factory prior to the launch of the program cycle, they are informed individually by the field staff.

Treatment status is chosen in each phase using a lottery, where ID numbers are drawn by the research team to assign treatment status. In Phase 1, there is one lottery drawing. In Phase 2, there is one drawing to select those receiving the Private accounts. A week later, we conduct a second drawing to select those receiving Non-private accounts—announced as a new Publicity program—in order to ensure there is no confusion among workers about their treatment assignment.

Marketing sessions. The sessions, conducted by IPA staff, last about 20 minutes. The sessions include a presentation of the key features of the accounts:

1. Participation in the program is fully voluntary, and workers will not face any consequences if they decline to take-up the product.
2. The offered product is a free savings account in a local partner bank. It has no fees during the program period, and has no minimum deposit requirement.³⁹
3. The account will be blocked for a period of 9 months in Phase 1, and 3 months in Phase 2, after which the worker may withdraw all of her money free of charge. During

³⁹In exchange for waiving the fees, the BPCI didn’t pay out interest on the savings.

the blocked period, savings can only be withdrawn if the worker were fired or unable to continue working at the factory — with an official letter from Olam or its hiring subcontractors as evidence. However, if the worker could prove the existence of a severe financial emergency (e.g. severe illness), then we allow workers to withdraw funds early.⁴⁰

4. Interested workers choose a threshold above which any earnings in each paycycle will be privately and automatically deposited into the savings account. If the worker earns an amount equal to or below the threshold, no money will be deposited in the account. This ensures that there is no risk that the accounts will squeeze them further in low-wage months. The threshold is constrained to be larger or equal to baseline average earnings per paycycle.
5. Workers are allowed to revise their threshold during an initial probation period of two months.⁴¹
6. After the end of the program, workers have the possibility to keep their accounts — converted into standard savings accounts, or potentially re-block the account.

To help workers better understand these features and make an informed take-up decision, they are then presented with a series of cases and asked how the accounts would operate in practice. Examples for those cases are “If you earn less than your set threshold during this fortnight, how much will be deposited into your account?” (the answer being nothing), and “If you chose a threshold of 15,000 FCFA and earn 18,000 FCFA this fortnight, how much will you save on your BPCI account?” (the answer being 3,000 FCFA).

Once workers understand how accounts operate, they are informed of the specific procedure to open up the account, including the required documentation. IPA would support them in gathering some of those documents, including by hiring a photographer and by paying for the issuance of a certificate of residency.⁴²

Finally, IPA staff individually administers a short quiz to workers to verify their understanding of the accounts. After answers to the quiz are recorded, any potential misperceptions are corrected.

Individual follow-up. IPA staff follow-up individually (and discreetly) with treatment workers in the subsequent days to answer any lingering questions. They then ask workers

⁴⁰In practice, this happened for one worker over the course of the study.

⁴¹In practice, the two months threshold was loosely applied: if a worker seemed to realized in good faith that the threshold was not appropriate, she would be able to revise it.

⁴²Certificates of residency are valid for a limited period of time, and are of limited use aside from opening up bank accounts. As such, it is highly unlikely that paying for certificates of residency could lead to any changes in labor supply by itself.

about their take-up decision. For workers who decide to open the offered private blocked account, IPA staff elicits their desired threshold. Workers are advised to choose a threshold allowing their cash earnings to cover their usual level of expenses (for consumption, transfers to kin, etc.), as well as the shocks they might incur over the course of the blocked period (e.g., illness).

Required documentation. To open a formal bank account in Côte d’Ivoire, individuals are legally bound to present a formal ID. This requirement represents a key impediment to financial inclusion. In one of the factory sites, workers receive no assistance with obtaining IDs. In the other factory site, where less than 30% of workers had a formal ID document, thanks to the dedication of the local Olam subcontractor in charge of hiring and payments, a solution was devised to lower the barriers to opening an account. The workers’ earnings above their chosen threshold were deposited into an account operated by that Olam subcontractor, instead of an individual account.⁴³ The subcontractor was responsible for handling the funds, and IPA monitored the process to ensure correct implementation. If during or at the conclusion of the program period the worker provided the bank with the correct documentation, an account was to be opened in the worker’s name and the total amount saved by that worker would be transferred by the bank. If the worker did not wish to continue with the program or did not obtain the necessary documentation upon the completion of the program period, the worker would be responsible for collecting all savings from the subcontractor, at the end of the project period. In the other site, there was no such mechanism, and presenting a valid ID was included as an eligibility criterion.

Receipts. During the program period, research staff privately deliver receipts individually to treatment workers who have opened an account, after each paycycle payment.⁴⁴ These receipts indicate the amount deposited that paycycle in the worker’s account, as well as its balance. These receipts are provided after each paycycle regardless of whether the worker exceeds her threshold and a deposit was actually made in the account.

Non-private accounts (Phase 2)

The implementation protocol for the second phase of the intervention was driven by a strong desire to avoid misperceptions and incorrect rumors regarding the features of the offered blocked savings accounts. In particular, to allow meaningful differences to emerge between the use of the Private blocked accounts and of the Non-private blocked accounts, it was

⁴³The subcontractor was a trusted local worker cooperative organization that helps Olam with hiring and maintains long-term relationships with workers, so that there is a high level of trust among workers in the subcontractor’s reliability in handling their funds.

⁴⁴Our partner BPCI computes the amount earned and to be saved, IPA staff checks the computation, and then details the exact way in which the deposits are made into the bank.

instrumental that workers in the Private treatment arm be convinced of the accounts' privacy.

We therefore started by offering accounts to those workers randomly assigned to the Private treatment arm. The announcements, marketing sessions, and individual follow-ups were identical to those in Phase 1 of the intervention — the only difference being that the accounts were now blocked for 3 months. The activities were implemented in both study sites over the course of the same week.

In the following week, we conducted these same activities (announcements, marketing sessions and individual follow-ups) for the Non-private treatment arm, labeled as a 'Publicity program'. To make sure that workers in the Private treatment arm understood that the publicity didn't apply to them, the factory-wide announcement specified that, "in an effort to further extend programs that help people to save, we will offer a new Publicity program to workers this coming week". At the same time as the marketing sessions and individual follow-ups for workers in the Non-private treatment arm were unfolding, IPA staff reached out privately to the workers in the Private treatment arm, to reassure them of the privacy of their accounts and savings. Due to these efforts, workers in our study sample appeared to understand well the specificity of their own treatment arm.

The content of the marketing sessions with workers in the Non-private treatment arm were identical to those in the Private treatment arm (and to those in Phase 1), except for description of the publicity feature of the account. We introduce the Non-private account to workers with the following script: "As you may know, we have been working in [study site 1] for over two years, and now in [study site 2] for almost a year. Due to the success of the previous programs, this week we are offering workers a new Publicity program. This is different than the the Private program offered in previous waves because this type of account is not private - it involves your network and community by letting them know about your savings account. Not all workers have been selected for this version of the program, and workers selected for this program were selected at random.

"This account is different than previous versions of the program in an important way: Because we're interested in publicizing to others in the community the importance of savings, if you choose to participate in the Publicity program, we will advertise to people in your community that you've accumulated savings through a new program. If you've already participated in a previous program with us, regardless of whether or not you decide to participate in this next phase, any past savings you've earned through the end of March will stay private, but details of your blocked savings after April 16th on may be shared."

We explained to workers that, by taking-up the account *and* achieving savings, they would give permission to advertise to people in their network/community that they have a savings account and some of the account features, via two SMS messages. To explain this, workers were told: "For example, if you were in the program, people in your network

would get this message in early July: ‘[Worker name] saved [amount] through a new program where she put aside some of her earnings with the BPCI. She’ll already be able to access her savings in the next week!’ Also, after the first time you achieve savings, we’ll send people in your network a message letting them know you’ve achieved some savings. We know that sometimes it is hard to save through the program for reasons outside of your control, so we won’t send any message to your friends and family if you don’t achieve savings.” Note that we decided to advertise the account only if the worker achieves savings in order to avoid confounds influencing take-up, threshold setting, or productivity due to potential shame in the event that a worker does not achieve savings.

Piece-rate variation

Announcement. Mid-June 2019, an announcement was made to workers: we will conduct a short-run program that would give them different wages for their work over different days. Specifically, over the course of a week, each worker would draw each morning a colored ball from a hat determining her piece-rate for the day. In case a worker repeatedly draws a low rate, she would be compensated so that, at the end of the week, she earns at least the amount she would have earned under the usual piece-rate. More details would be provided during short training sessions on the announcement’s day.

Specifically, workers had the possibility to draw one of four rates, with equal probability, to be applied to their production of the day: a rate 15% lower than the usual rate; the usual rate; a rate 15% higher than the usual rate; 30% higher than the usual rate. If workers did not adjust their production to the change in piece-rate, they could therefore expect an average increase in earnings of 7.5% that week. Nonetheless, out of bad luck, some workers might end up repeatedly drawing the rate lower than the rate applied absent our activity. To ensure that no worker could lose from the activity, earnings were computed as follows:

$$\text{earnings} = \max \left(\sum_{d=1}^6 \text{output}_d \times \text{usual rate}, \sum_{d=1}^6 \text{output}_d \times \text{experimental rate}_d \right)$$

Eligibility and training. All factory workers paid piece-rate based on their individual output were invited to participate in this activity in one of our two study sites. Workers participated in short training sessions in groups of 15-20 individuals. These sessions focused on explaining that workers would be offered, in the coming week, an opportunity to earn more money than they would usually earn, through a lottery. To make the piece-rate variation salient, workers were first reminded about their usual piece-rate pay structure (with examples and exercises to check understanding), before being walked through in more detail (again with examples and exercises) the idea and details of the piece-rate lottery.

Implementation. Upon arrival at the factory, and before collecting nuts to be processed during the day, each worker picks a ball out of a hat. The balls were identical plastic balls of four colors, each corresponding to a piece-rate. The correspondence between the colors and the rates was first introduced during the training sessions, and was made salient throughout the activity week with the use of posters. The worker’s name, unique ID, and applicable rate are recorded by IPA staff. At the end of the day, each worker’s production is collected, weighted and recorded as usual by Olam — with IPA staff sitting alongside them and noting the information as well. The data collected by IPA staff was then referenced against that of Olam’s administrative data to check for inconsistencies, prior to finalizing and making worker payments.⁴⁵

⁴⁵While the full payout from the activity was initially planned to occur at the next payday, operational constraints led to a disbursement in two tranches. At the payday associated with the pay cycle containing the piece-rate variation activity, the earnings for each worker were determined by applying the usual piece-rate to her total production over that pay cycle. A few weeks later, workers received the remainder of the due amount: $\sum_{d=1}^6 [\text{output}_d \times (\text{experimental rate}_d - \text{usual rate})]$.