The Social Tax: Redistributive Pressure and Labor Supply

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

May 2025

Abstract

In low-income communities in both rich and poor countries, redistributive transfers within kin and social networks are frequent. Such arrangements may distort labor supply—acting as a "social tax" that dampens the incentive to work. We document that across countries, from the United States to Côte d'Ivoire, low-income groups report strong pressure to share earned income with others; in addition, 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 9 months. Workers may only deposit earnings increases, relative to baseline, mitigating income effects on labor supply. Offering Private accounts raises work attendance by 6.5% and earnings by 9.4%. These treatment effects are concentrated among workers who report higher redistributive pressure at baseline. To obtain further suggestive evidence on mechanisms, in a supplementary experiment, we vary whether blocked accounts are private or known to the worker's network. When accounts are private, take-up is substantively higher (60% vs. 14%), with a resultant 8.8% higher earnings. Outgoing transfers do not decline, indicating no loss in redistribution. 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-Touba: CREST, Institut Polytechnique de Paris (florian.grosset@ensae.fr); Kaur: University of California at Berkeley and NBER (supreet@berkeley.edu). This paper greatly benefited from comments by the handling editor Oriana Bandiera, anonymous referees, 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, Cécile Delcuvellerie, Kenza Chaabouni, and Peirong Shi 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 Nuguhe 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 21% of their income to others outside their household on average, and 79% made at least one transfer in the past 3 months. Similarly, in the United States, data from the PSID indicates that among Black Americans, high earners 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 United States. In each plot, the unit of observation is a geographic sub-location × ethnic group (or race) × 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 correla-

[1] See Karaivanov and Townsend (2014), De Weerdt and Fafchamps (2011), De Weerdt 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.

[2] A 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

tions and are 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

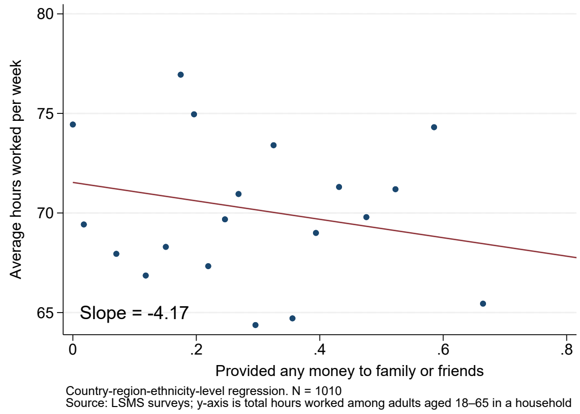

(A) West Africa

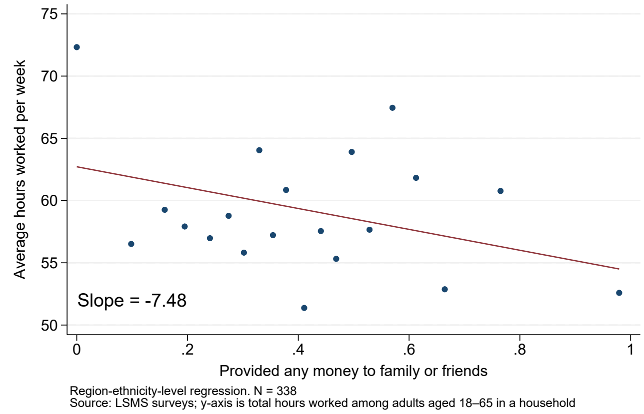

(B) Côte d'Ivoire

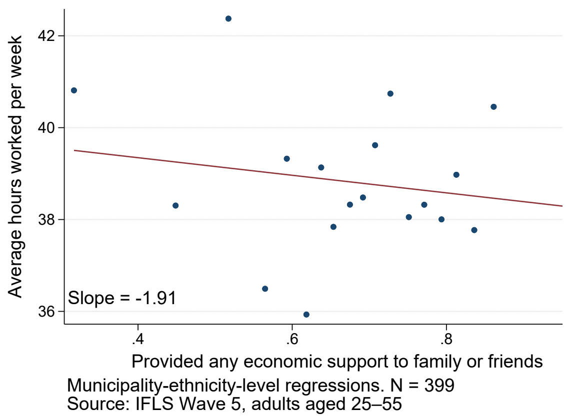

(C) Indonesia

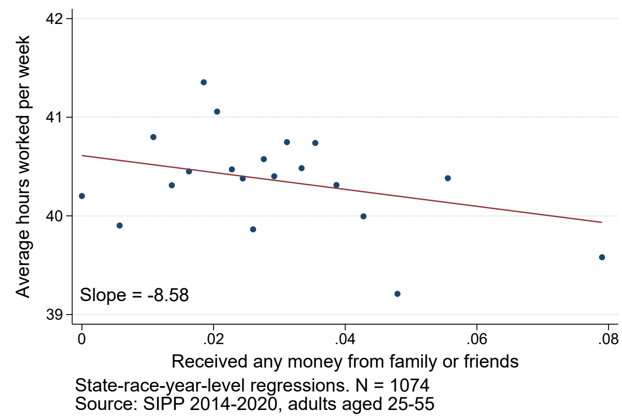

(D) United States

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. LSMS data from ANSD (2022); INS-CIV (2022); INS-Niger (2022); INSAE (2022); INSD (2022); INE (2022); INSEED (2022); INSTAT (2022), IFLS data from Strauss et al. (2016), SIPP data from U.S. Census Bureau (2014, 2018, 2019, 2020).

In this paper, we develop a causal test to empirically examine this possibility. We conceptualize redistributive pressure as a "social tax" on earnings. This offers a parallel between the impacts of redistributive pressure and that of formal taxes.[3] We focus on

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.

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

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

We work with 474 full-time 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 1.5 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. 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: 76% 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 credibly convey that they do not have the money available. 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 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 Private 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 remaining earnings are 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 tax on earnings gains alone does not induce income effects, only substitution effects.[4] Con-

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".

[4]Denote the worker's baseline labor supply as e_0. Under our intervention, if she continues to supply e_0, then her tax rate (and therefore net earnings) remain unchanged—i.e. there is no income effect. The rate is

sequently, 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. However, note that these conditions 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 tax rate implied by our effects would be a lower bound.

In our main field experiment, we randomize some workers to receive an offer to sign up for a Private blocked savings account (treatment), while the remaining Control workers continue with the status quo of no account. There is substantial demand for the Private accounts: 45% of workers take them up. Being offered a Private account leads to sizable increases in labor supply and earnings: workers' total output, and consequently earnings, increases by 9.4% (p=0.025). This is accompanied by an increase in attendance at work of 4.37 percentage points, or 6.5% (p=0.092). Because almost all workers (85%) have no earnings outside the factory, these treatment effects constitute increases in workers' total income. Note that these treatment effects are large—equivalent to how much earnings would rise if each worker worked an additional 0.76 days in every 2-week paycycle.

To assess the potential role of redistributive pressure as a mechanism, we undertake two suggestive tests. First, we design a "mechanism experiment" which creates a more targeted "proof of concept" test for redistributive pressure. After the conclusion of the main experiment, we again offer blocked accounts to a sample of workers, but vary whether network members would have knowledge of the accounts. Specifically, as before, a random half of workers in the sample is offered Private blocked accounts, whose details are known only to the worker. The other half of workers is offered a Non-private account, which adds the feature that some of the workers' network members would learn that the worker held savings in the account as well as the unblock date.[5] This holds constant the internal benefits of the accounts (e.g., any self-control or goal setting effects), but varies whether the accounts actually shield workers from redistributive requests.

only lower for supply above e_0. Consequently, starting from the baseline of e_0, a worker who switches from the status quo to our intervention faces a pure substitution effect only.

[5]If workers take up and save in a Non-private account, then network members would receive up to two text messages in which the bank would advertise the accounts, explain the worker had saved in it, and encourage the recipient to also sign up for such an account. The second text message would relay the savings level and be timed to let the recipient know that funds would be unblocked the following week. 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. In the text we present additional tests to argue that a general desire for privacy or risk of theft cannot explain our findings in the mechanism experiment.

We find that the blocked accounts are substantially less desirable as a savings vehicle when they do not shield workers from redistributive requests. While 60% of workers take up the Private accounts in the mechanism experiment, only 14% do so for the Non-private ones—a 77% reduction (p<0.001). We ask workers who decline the Non-private accounts their primary reasons for doing so: 96% of workers state that the account would cause net transfer requests from others outside of their household to go up. These differences in take-up lead to substantial differences in labor supply and earnings. Relative to the Non-private account, being offered the Private account increases earnings by 8.8% (p=0.063) and attendance by 6.6 percentage points or 8.4% (p=0.045). The magnitude of the effects on earnings is quite similar to that in the main experiment.

Second, we examine heterogeneity in impacts. Consistent with our hypothesized mechanism, treatment effects are concentrated among workers who face more redistributive pressure at baseline. For example, among those who receive frequent transfer requests, being offered a Private blocked account increases earned income by 15% (p=0.008). In stark contrast, among those who infrequently make transfers, the Private account has no discernible effect. Moreover, this pattern of heterogeneity by baseline redistributive pressure holds in each of the main experiment and mechanism experiment.

Together, the mechanism experiment and heterogeneity results offer suggestive evidence for the relevance of redistributive pressure in driving (at least some of) the effects of the blocked accounts. In contrast, we do not find direct evidence in support of alternate channels. For example, we find that our effects are not primarily driven by intra-household pressure to share earnings with one's spouse: we see no evidence that effects are stronger among the 59% of women who have a partner or spouse. In addition, we find no change in labor supply during the announcement period—the weeks after workers have learned their treatment status but before the accounts become active—inconsistent with a fairness or morale effects story. Finally, inconsistent with a self-control mechanism, in a supplementary exercise, we find no evidence that workers are more likely to demand access to blocked savings accounts when they are four days away from their payday, relative to on the payday itself. In addition, such internal mechanisms could not explain the findings in the mechanism experiment, or why our treatment effects are concentrated among those who report high redistributive pressure at baseline. However, while these arguments indicate that redistributive pressure is likely necessary to explain our results, internal benefits of the blocked accounts might still contribute to the observed magnitude of treatment effects.

Finally, we examine treatment effects on 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 improved output and welfare among treated workers, without reducing financial redistribution to the network.

We conclude with a discussion of the social tax rate implied by our experimental results, and the associated parallel with formal taxation. Existing labor supply elasticity estimates, which are primarily from richer countries, indicate small elasticities; this in turn implies a large social tax rate of about 40% for the average worker in our experiment. However, a body of evidence indicates much larger labor supply elasticities for low-income workers, from Uber drivers in the US to casual workers in India. These larger elasticities imply a much smaller social tax rate, of about 9-13% for the average worker and 17-21% for those who take up the Private accounts. In either case—because our reduced form treatment effects are a composite of the social tax rate and labor supply elasticity—our experimental results indicate that the scope for labor supply distortions is large, at least in our specific context.

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 Weerdt 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, 2024).[6] 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, 2013b; Riley, 2024; Squires, 2024).[7]

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

[6]For 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 (2024) uses the willingness to pay to keep a cash windfall private in the lab to structurally estimate sizable productivity implications among micro-entrepreneurs.

[7]A large 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; Zhou and Mahadeshwar, 2024), 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, 2024). More generally, our study relates to the growing literature on the impediments to female labor force participation (e.g. Jayachandran, 2015).

to work—altering labor supply, with implications for productivity and income.[8]

Our study also contributes to the literature on the impact of savings on labor supply. We offer novel evidence that blocked accounts (or accounts with some illiquid feature) can directly increase labor supply.[9] We do so within the context of a highly scalable policy intervention (i.e. direct deposits made by the firm). Moreover, the literature on savings interventions has found mixed evidence on whether improved savings access affects labor supply, with some studies finding increases (Callen et al., 2019; Horn et al., 2021; Field et al., 2021), and others finding no changes in labor supply (e.g. Somville and Vandewalle, 2023; Dupas and Robinson, 2013a). We directly tie the impacts of savings interventions to redistributive pressure—helping provide guidance on one set of conditions under which savings improvements can be expected to increase labor supply.[10]

Our study has implications for understanding one potential set of impediments to labor supply and productivity in low-income settings, particularly Sub-Saharan Africa (Lewis, 1955; Tam et al., 1957). If redistributive pressures distort work incentives, it could also hamper other costly actions that increase future income, from human capital investment to technology adoption. 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 Motivation: Redistributive Pressure

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 (21% on average) to others outside of their household.[11] Of this, 61% is redistributed within the extended family outside the household, 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

[8]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, or being able to lock away existing earnings. 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.

[9]Our intervention has a natural parallel with work on commitment savings, which has primarily focused on savings levels as an outcome (Ashraf et al., 2006; Dupas and Robinson, 2013c).

[10]For example, both Dupas and Robinson (2013b) and Pomeranz and Kast (2024) suggest that savings interventions lead to greater savings increases among individuals who face higher redistributive pressure.

[11]We 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.

average earnings redistribute more to their networks on average (Appendix Figure A1).[12]

Transfer requests occur for diverse reasons—including unexpected shocks (illness), expected expenditures (school fees), investments (housing improvements), and consumption (people showing up at meal times). Respondents express a desire to avoid many 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.

Figure 2 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 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: 76% 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 in our study setting of Côte d'Ivoire perceive the cost of refusing requests to include social stigma or isolation—making it unpleasant, for example, to attend extended family or network gatherings, which are an important source of socialization and utility. 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 A2, 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.

Consistent with this idea, workers employ a variety of strategies to make their funds inaccessible for redistribution. For example, workers report buying household goods im-

[12] An increase in average earnings of 1 FCFA is correlated with an increase in reported transfers of 0.041 FCFA (p=0.011, 95% CI=[0.009,0.073]). Note that some young workers transfer large amounts (as a fraction of their income) to their parents.

Figure 2: Motivational Evidence: Redistributive Pressure in Côte d'Ivoire

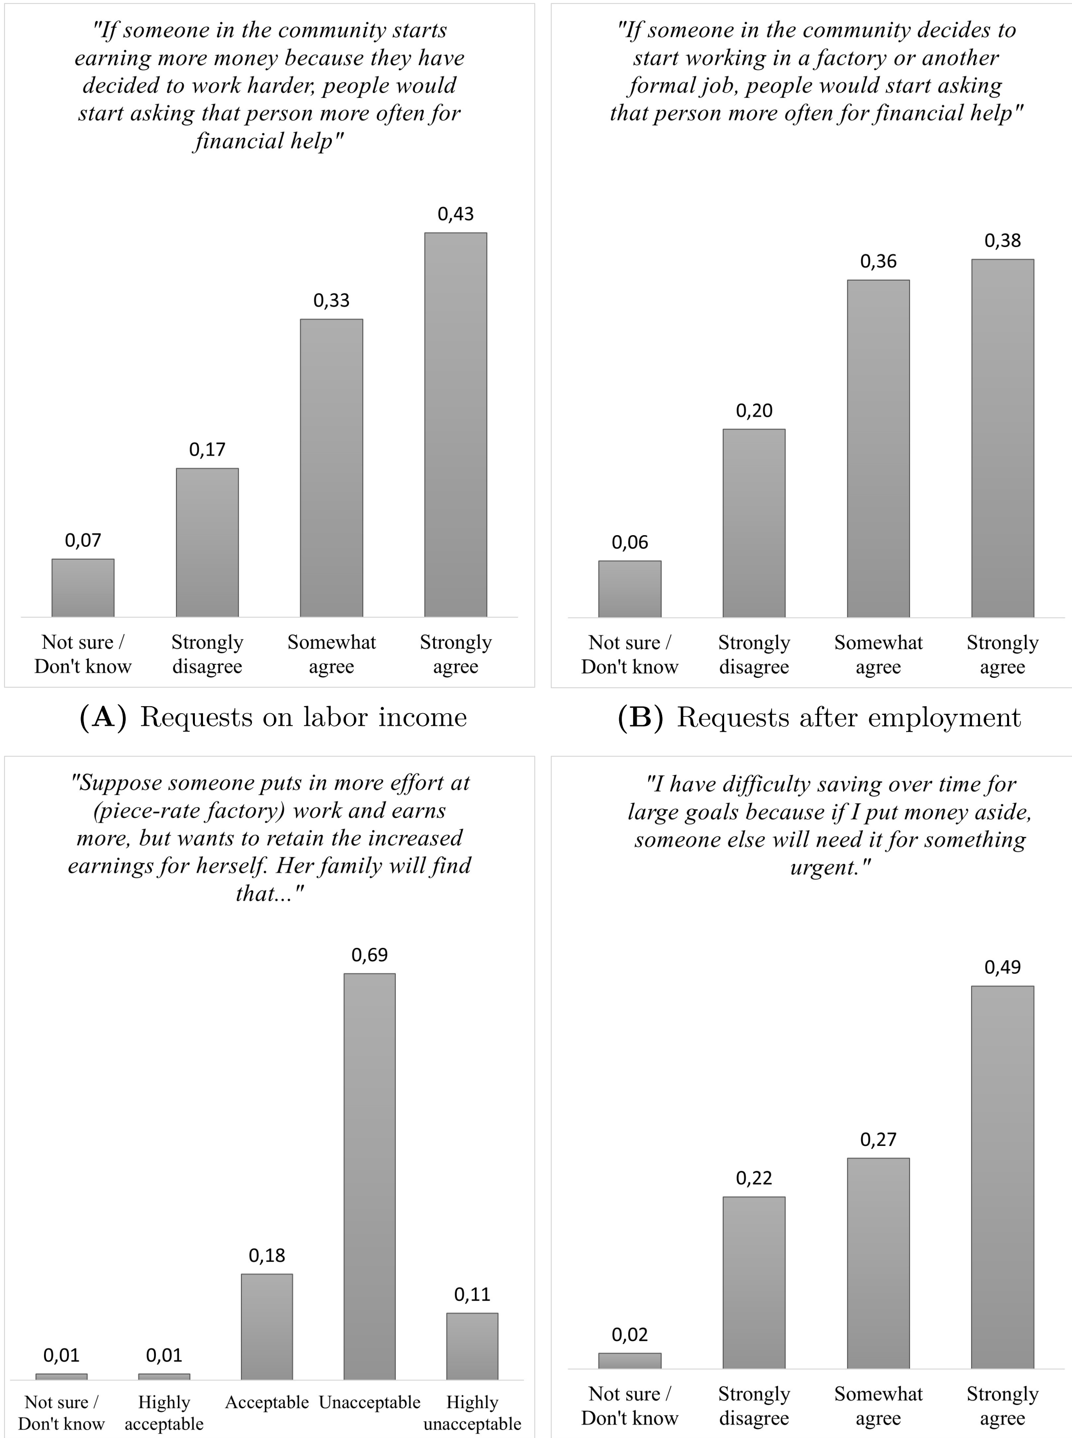

(A) Requests on labor income

(B) Requests after employment

(C) Transfer norms

(D) Inability to save

Notes: N=412 (Panels A and B), N=486 (Panel C) and N=429 (Panel D) cashew factory workers in Côte d'Ivoire.

mediately 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 2D. In our study, we draw on these existing strategies to design a blocked savings account to help shield savings from redistributive pressure.

3 Context: Factory Workers in Côte d'Ivoire

Workers. The workers in our study are full-time laborers employed in cashew-processing plants run by Olam, a large multinational agro-processing company that controls 80% of the processed volume of cashews in the country. We work in two factory plants in central Côte d'Ivoire, about 230 km away from the economic capital of Abidjan. We enroll 474 full-time workers in our study, of which 464 are women.[13] In our sample, the average worker has worked at the firm for 1.5 years (25th percentile of 1 year, 75th percentile of 2.1 years).

Production task. Workers 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.

Wages and publicity of payments. Each worker receives a linear piece rate for her output.[14] The entirety of workers' earnings are comprised of their piece rate wages. Consequently, changes in effort or attendance translate directly into changes in worker earnings. Workers are paid their earnings twice a month in cash.[15] The pay process is visible among co-workers, such that co-workers have awareness about the earnings of others. The timing of payment is generally common knowledge in workers' local communities. In our sample, workers' factory income is their primary source of earnings, and 85% report having no other source of income.

[13]The 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.

[14]The specific piece rate changes based on the quality of the nuts, which fluctuates over time and is exogenous to the worker.

[15]Earnings at the factory are set so as to exceed Côte d'Ivoire's minimum wage for full-time attendance.

4 Conceptual Framework

We use a simple model of labor supply under taxation to motivate our research design. Following Mirrlees (1971) and Gruber and Saez (2002), 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_0. 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.[16] 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 $y = (1 - \tau_0)we$. We use e_0 to denote her utility-maximizing level of labor supplied under tax rate τ_0 (see Figure 3A).

To test for the presence of a tax, and its subsequent effects on labor supply, we seek to lower τ_0 by enabling the worker to avoid transfer requests on (some portion of) her earned income. However, simply lowering τ_0 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_0 (substitution effect). On the other hand, the increase in net earnings under labor supply levels $e \leq e_0$ would decrease the incentive to work (income effect). This would make it challenging to use the worker's labor supply response to diagnose 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*

[16] Note that if factory workers have "won the lottery" by obtaining a factory job, then transfers to those without factory jobs could simply reflect what one would expect under informal insurance. However, the factory plants in our sample typically have excess demand for labor: at least some network members who make transfer requests could also obtain full-time jobs there. Consequently, the presence of transfers in our setting would require a more subtle model of informal insurance.

Figure 3: Tax Rate

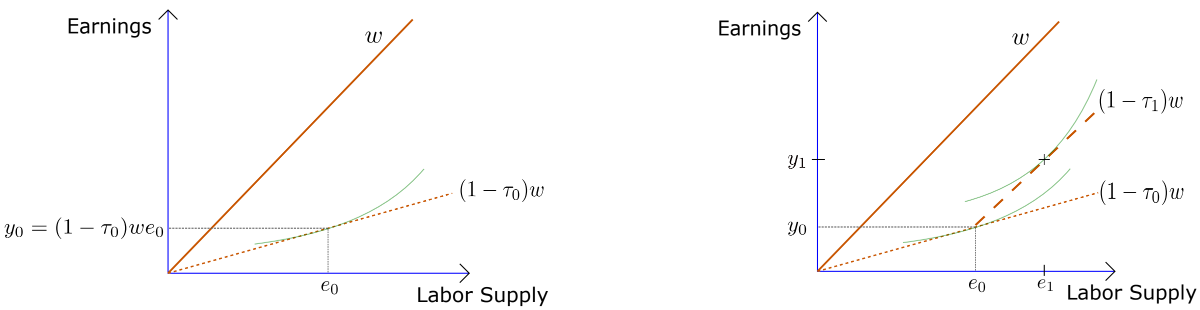

(A) Baseline tax schedule   **(B)** New schedule: reduced tax on income gains

Notes: Panel A: labor supply (e_0) under a linear piece rate with social tax rate τ_0. Panel B: change in optimal labor supply when the social tax rate is reduced to τ_1 on earnings above e_0.

only. Specifically, consider reducing the social tax rate to $\tau_1 < \tau_0$ only for $e > e_0$, while keeping the tax rate at τ_0 for $e \leq e_0$ (Figure 3B).

This new tax rate induces a kink in the worker's budget constraint. The worker consequently chooses her labor supply level e_1 to solve $\max_{c,e} u(c, e)$ under the budget constraint:

$$c \leq \mathbb{1}_{e \leq e_0} \{(1-\tau_0)we\} + \mathbb{1}_{e > e_0} \{(1-\tau_0)we_0 + (1-\tau_1)w(e-e_0)\}.$$

First, note the trivial result: $e_1 \geq e_0$.[17] This already rules out the possibility of a labor supply decrease. We can therefore rewrite the budget constraint as: $(1-\tau_1)we + \mathbb{Y} = c$, where $\mathbb{Y} \equiv (\tau_1 - \tau_0)we_0$. The worker's optimal choice of effort under the new tax schedule is $e_1((1-\tau_1)w, \mathbb{Y})$.

We can derive the change in e_1 induced by a change in τ_1. We provide an overview of the key results here; see Appendix A.3 for details. Applying the Slutsky equation, we obtain:

$$\frac{de_1((1-\tau_1)w, \mathbb{Y})}{d\tau_1} = -w \frac{\partial \tilde{e}((1-\tau_1)w, u)}{\partial [(1-\tau_1)w]} + w(e_0 - e_1) \frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial \mathbb{Y}}, \qquad (1)$$

where $\tilde{e}((1-\tau_1)w, u)$ is the Hicksian (compensated) labor supply. On the right hand side 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 3A to that in Figure 3B. We begin with the baseline situation of no kink in the budget constraint, so $\tau_1 = \tau_0$. Then, the effect of introducing the new tax schedule—i.e.,

[17] Proof by contradiction: Suppose that $e_1 < e_0$. Then the budget constraint becomes $(1-\tau_0)we + y = c$, which is the budget constraint under which e_0 is the optimal choice. This contradicts $e_1 < e_0$.

decreasing τ_1 above e_0 when starting from the baseline of $e = e_0$—is given by:

$$-\frac{de_1\left((1-\tau_1)w, \mathbb{Y}\right)}{d\tau_1} = w\frac{\partial \tilde{e}\left((1-\tau_1)w, u\right)}{\partial\left[(1-\tau_1)w\right]} > 0. \tag{2}$$

The income effect term drops out: the increase in effort from the change in tax rate above the kink is only driven by the substitution effect term. Intuitively, because the tax rate on earnings up until labor supply level e_0 has not changed, the net earnings up until e_0 are unchanged, eliminating the income effect from the response.[18]

The result in Equation (2) indicates that moving from the tax schedule in Figure 3A to that in Figure 3B delivers an unambiguous prediction on labor supply: *labor supply will increase.* In other words, if workers face an initial social tax rate $\tau_0 > 0$, and our intervention lowers this tax rate to some $\tau_1 < \tau_0$ only for effort levels above e_0, then workers will increase their labor supply. Further, the amount paid out as a social tax is weakly higher: $\tau_0 we_0 + \tau_1 w(e_1 - e_0) \geq \tau_0 we_0$. As a result, redistribution to the network should not decline, and may even increase if $\tau_1 > 0$.

Discussion. The tax schedule in Figure 3B helps resolve the two key challenges that would arise if we simply lowered τ_0, 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_0. 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.[19] However, we see no evidence for this condition: the fraction redistributed does not increase with paycheck-to-paycheck fluctuations in income (Appendix Figure A3). While suggestive, this indicates that income effects are unlikely to play a meaningful role in workers' labor supply reactions in our experiment. Moreover, the presence of income effects would only make it harder for us to detect effects on labor supply, leading us to underestimate 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

[18] To see this intuition in more detail: Suppose we lowered the tax rate on all earnings. Then if the worker remains at $e = e_0$, her tax rate and therefore net earnings are higher—this is the income effect. In contrast, under the tax schedule in Figure 3B, if the worker remains at $e = e_0$, her tax rate and therefore net earnings are unchanged—there is no income effect. Consequently, starting from the status quo tax schedule—where $\tau_1 = \tau_0$ and the worker exerts $e = e_0$—reducing τ_1 produces only a pure substitution effect.

[19] For example, suppose earnings in each period are $we_0 + \epsilon$, where ϵ is a random variable with mean zero, so that average earnings are $\overline{we_0}$. 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 $\overline{we_0}$, then in periods with $\epsilon > 0$, workers will face a lower tax rate even though their labor supply remains the same ($e = e_0$). This could both generate some income effects on labor supply, and also reduce redistribution to the network. However, the scope for both these effects would still be less under the approach in Figure 3B than if we allowed workers to hide any part of their existing earnings.

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. The potential benefits of mitigating the two concerns above motivates our experimental intervention, which seeks to vary the extent to which earnings *increases* are sheltered from transfer requests.

5 Blocked Savings Accounts

5.1 Overview of Accounts

We seek to construct a design that mimics the approach in the tax schedule in Figure 3B 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.*

First, using administrative data from the factories, we compute baseline earnings to be the worker's average earnings per paycycle in the past 3 months. Workers who opt in to the account then choose a threshold, which must be at least as high as their computed baseline earnings.

In each biweekly paycycle, any amount workers earn *above* their chosen threshold is automatically deposited by the factory into the blocked account; the remainder of their earnings is paid in cash as usual on payday. Deposits and savings levels are private, known only to the worker. The funds in the account cannot be accessed until the end of the blocked period (3-9 months).[20] 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.

Because the funds deposited into the blocked savings account cannot be accessed, they cannot be used to fulfill transfer requests. 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.

5.2 Implementation Details

Account opening. The savings product is administered by one of the largest banks in Côte d'Ivoire, Banque Populaire (BPCI). BPCI previously offered goal-based blocked savings ac-

[20] It is possible that a regular formal savings account, rather than a blocked account, could also be an effective tool. We used blocked accounts to ensure workers could credibly say they did not have cash when asked for transfers, obviating the need to lie.

counts. We worked with the bank to offer date-based accounts, with no minimum balance or monthly fees, in exchange for removing the interest rate on savings in the account. We also eased the logistical barriers to opening and using an account by having a bank employee stationed at the factories to help with initial paperwork, and having the factory directly deposit earnings into the account. Consistent with such barriers, less than 2% of Control group workers have any type of formal bank account at endline.

Announcement. Treatment status is chosen 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 deciding 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. The gap between when workers learn their treatment status and when accounts take effect for most workers is 1-2 weeks.

Training. We undertake training sessions within the factory with all workers offered a blocked account. 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 comprehension quiz. If the worker scores below 80% on the quiz, they are retrained one-on-one by a moderator.

While 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 (with the restriction that the threshold must always remain above their baseline earnings). 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 to re-optimize thresholds after experiencing the accounts if they want.

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.

6 Main Experiment: Private Blocked Savings Accounts

We begin by testing for the impact of the Private blocked account on labor supply and earnings. In Section 7, we supplement this with mechanism tests for the role of redistributive pressure.

6.1 Research design

Randomization. We randomize workers in both factory plants to either receive a 9-month Private blocked account (Treatment) or no account (Control). Workers were enrolled in the experiment in three staggered waves, with data collection ending in March 2019. Randomization is at the individual level, stratified by factory plant, baseline attendance, enrollment wave, and ID card availability. Additional information about the timing and implementation of each wave is provided in Appendix A.4.

Data. Our primary data source is Olam's detailed administrative data on each worker's daily attendance, output (the quantity of nuts processed), and earnings.[21] These administrative data are used by Olam to compute workers' cash payments and the amount deposited into the savings accounts (when relevant) in each paycycle. We supplement this with survey data, which includes measures of transfers to others (see Appendix A.5 for details on survey data collection).

Appendix Table A1 provides descriptive statistics and checks for balance in baseline covariates across treatment arms (Cols. 1-3). Across treatments, the average worker in our sample has a baseline attendance rate of 66% and unconditional daily earnings of 1736 FCFA (3.1 USD).

In pairwise t-tests between treatment assignment and each of these variables taken individually, 3 out of 12 variables are imbalanced at baseline. However, we cannot reject the null hypothesis that these baseline covariates do not jointly explain assignment to the Treatment group (F-test p-value = 0.151). Moreover, all our key outcome variables are balanced in trends at baseline.

Our primary outcome measure, earnings, is unfortunately imbalanced in levels at baseline for the main experiment. However, this baseline imbalance is not driving our estimated effects. First, our primary empirical specification controls for any baseline differences in levels using a differences-in-differences approach that includes individual fixed effects. The identifying assumption is of parallel trends—whose validity is supported by balance in the earnings trends during the baseline period (Appendix Table A1 and Figure 4). Second, the baseline imbalance is driven by one specific randomization wave, and removing this wave from our sample recovers balance on earnings while keeping the estimated treatment effects similar (Appendix Table A2, Col. 6). Third, we find similar treatment effects using only endline data and controlling for baseline covariates (Appendix Table A3).

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

[21] Note that this data does not include information on worker hours. However, given the fixed factory timings of 8am-5pm, it is reasonable to use attendance as a meaningful measure of the extensive margin of labor supply.

our measure of workers' labor supply. We also study effects on attendance (i.e. the extensive margin decision of deciding to come to work).

Given the staggered waves of enrollment and treatment assignment, a natural approach is to estimate a stacked difference-in-differences regression (Cengiz et al., 2019; Baker et al., 2022). This method involves creating a "clean 2x2" dataset for each randomization event e, stacking these datasets together, and estimating:[22]

$$y_{ite} = \beta PrivateAccount_{ie} \times Post_{te} + \gamma_{ie} + \delta_{te} + \epsilon_{ite}, \tag{3}$$

where y_{ite} is the outcome of interest for worker i on date t within the randomization event e. The $PrivateAccount_{ie}$ indicator equals one if worker i was assigned to receive the Private account treatment in randomization event e. $Post_{te}$ is an indicator that equals one on dates when the blocked savings account is active, within randomization event e. The γ_{ie} and δ_{te} are worker-by-event and day-by-event fixed effects, respectively. Standard errors are clustered at the worker level. Throughout our analyses, we use a baseline period that includes the 12 working days before the treatment announcement date, as well as the period between the announcement date and the first day in which earnings count towards the savings account. We show robustness to alternate baseline periods and empirical specifications.[23] This specification naturally extends to an event-study specification—to estimate the dynamic effects of the treatment—and directly maps to our experimental design.

Our coefficient of interest is β, which captures the treatment effect of being offered a Private blocked account, relative to the omitted category of being offered no account (Control). We predict that $\hat{\beta} > 0$.

6.2 Results: Impact of Private blocked accounts

Account take-up. 45% of workers who are offered the blocked savings accounts take them up. This is a robust take-up rate, and likely reflects several features of our setting: the large time investment we undertook (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.

[22] We consider 8 randomization events in the Main Experiment, defined in Appendix A.4.2. For each, the corresponding clean 2x2 dataset is a balanced panel that includes worker-day observations such that: no worker was already treated at baseline; each worker is randomly assigned to Treatment or Control during the randomization event; and each worker retains the same assignment for the full post-treatment period.

[23] The gap between the end of the main experiment and start of the mechanism experiment (see below) is two weeks, or 12 working days. We consequently use a 12-day baseline period across all tables.

Among workers who take up a blocked account, 15% choose a threshold close to the minimum (i.e., a round number within 10% above their baseline earnings), while the remaining 85% select a threshold that is higher (see Appendix Figure A4). Overall, the median threshold chosen corresponds to 155% 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.[24]

Table 1: Treatment Effects

	Take-Up	Earnings	Attendance	Retention
	(1)	(2)	(3)	(4)
Panel A: Main Experiment				
Private (vs. Control)	0.453	162.4	0.0437	-0.0275
	(0.0396)	(71.92)	(0.0258)	(0.0479)
	[0.000]	[0.025]	[0.092]	[0.567]
Sample mean in control	0.00	1720.67	0.67	0.72
N: worker-days	451	122916	122916	451
N: workers	354	354	354	354
Panel B: Mechanism Experiment				
Private (vs. Non-Private)	0.460	152.9	0.0655	0.0379
	(0.0481)	(82.08)	(0.0326)	(0.0551)
	[0.000]	[0.063]	[0.045]	[0.492]
Sample mean in control	0.14	1730.90	0.78	0.58
N: worker-days	317	38222	38222	317
N: workers	317	317	317	317
Panel C: Pooled Experiments				
Private (vs. Control+Non-Private)	0.452	157.6	0.0547	0.00179
	(0.0321)	(54.88)	(0.0206)	(0.0370)
	[0.000]	[0.004]	[0.008]	[0.961]
Sample mean in control	0.06	1725.01	0.72	0.66
N: worker-days	768	161138	161138	768
N: workers	474	474	474	474

Notes: Dependent variable is account take-up in Col. (1), daily earnings (in FCFA) in Col. (2), daily attendance in Col. (3), and an indicator variable equal to 1 if the worker remained at the factory until the last week of the wave in Col. (4). Stacked difference-in-differences specification. All regressions include day-by-wave and worker-by-wave fixed effects. Standard errors clustered by worker.

Effects on labor supply and earnings. Being offered a Private account substantially

[24]Because 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.

increases labor supply and earnings. Table 1, Panel A reports treatment effects using administrative data from the firm. Relative to the Control group, workers who are offered the Private blocked accounts increase their output and earnings by 162.4 FCFA or 9.4% (p=0.025, Col. 2). Figure 4 plots the corresponding event study estimates.

Figure 4: Treatment Effects over Time, by Week

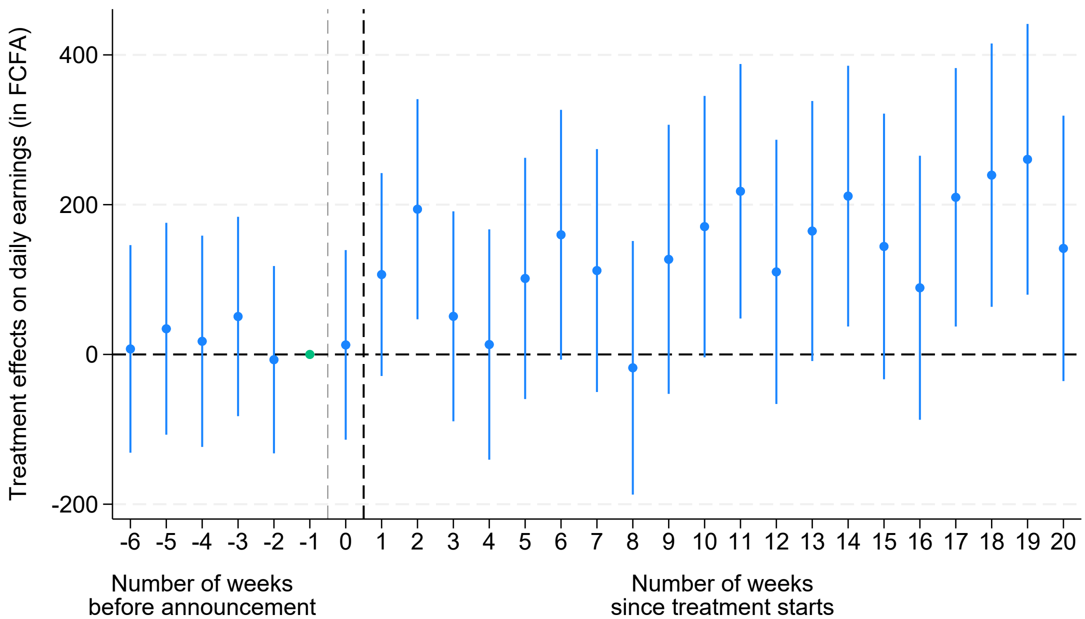

Notes: Effects of offering Private blocked savings account (vs. being assigned to Control) on daily earnings. Effects reported separately by week (defined as 6 working days). Week number 0 corresponds to the first week of the period between the treatment announcement (denoted with a gray dashed line) and the effective start of the treatment (first day in which earnings count towards the savings accounts, denoted with a black dashed line).

These effects are driven, in part, by a suggestive 4.37 percentage point (pp) or 6.5% increase in attendance at work (Panel A, Col. 3, p=0.092).[25] This reflects a reduction in absenteeism; we find no discernible changes in turnover (Col. 4). In Appendix Tables A2 and A3, we verify that our treatment effects are robust to alternate empirical specifications, including using only intervention period data to estimate effects.

These magnitudes are economically meaningful. For example, the treatment effect on earnings is equivalent to how much earnings would rise if the average worker worked an additional 0.76 day in *every 2-week paycycle.*[26] In addition, these effects are not simply reflecting

[25]Using a simple back of the envelope calculation, attendance accounts for 69% of the overall earnings effect; however, the confidence intervals around this estimate are wide.

[26]The average treatment effect on daily earnings is 162.3 FCFA, corresponding to a 1,948 FCFA increase per paycycle (comprised of 12 workdays). Mean daily earnings conditional on working among the Control group is 2,569. This gives 1948/2569 = 0.76 workdays per paycycle.

a substitution away from other income-generating activities: using data from our endline surveys, we find no treatment effects on earnings outside the factory. This is consistent with the fact that, at baseline, 85% of workers report having no earnings outside of their factory job, and on average, 96% of total income comes from factory earnings. Consequently, the effects in Table 1 reflect an increase in total earned income.

7 Mechanism: Redistributive Pressure

While we designed the blocked accounts to mitigate *external* redistributive pressure, they could also have *internal* benefits, such as helping with self-control problems or goal setting. In Section 7.1, we use a supplementary experiment to create a suggestive test for the role of redistributive pressure as a mechanism. In Section 7.2. we test heterogeneity in treatment effects by proxies of baseline redistributive pressure. Both these sections therefore seek to provide distinct, and complementary, pieces of positive evidence for the role of redistributive pressure. In Section 7.3, we explicitly examine alternate channels, such as self-control.

7.1 Mechanism experiment: Non-private accounts

To construct a more targeted test for redistributive pressure, we undertake a mechanism experiment. We offer a blocked account to all workers, but vary the extent to which the account—and most importantly, its unblock date—is known to others in the network. This enables us to hold the internal benefits fixed, while varying the extent to which the accounts actually shield workers from redistributive requests. This allows us to test directly for the role of redistributive pressure, but at the expense of a less naturalistic intervention.

Design of Non-private accounts. 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 others are more likely to request transfers immediately after expected pay dates. 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 shortly after the unblock date—reducing the usefulness of the blocked accounts as a way to avoid redistributive pressure on income gains.

Specifically, in the Non-private account, workers are told they are being offered a "Publicity program" by the bank. In exchange for receiving the "free" blocked savings account, they consent to enroll in a program where their achievements would be conveyed to some of their network members via two text message advertisements. This was couched as the bank's attempt to increase its clientele—a credible claim both because the bank was indeed actively advertising, and because personal referrals and SMS advertising are extremely common in this setting. To enable this, at baseline, we asked all workers to give us the names and phone numbers of up to five network members.

Workers were told the two SMS messages 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. In addition, the second SMS, sent shortly before the unblock date, would relay the savings level and that the account would soon be unblocked: "«Worker's name» 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).

In the mechanism experiment, we randomize 317 workers to receive either a three-month Non-private blocked account, or a three-month Private blocked account; the design of the Private accounts is the same as in the main experiment, except for the shorter duration of the blocked period. This sample is comprised of 120 newly added workers who did not participate in the main experiment, and 197 workers from the main experiment (with randomization stratified by their main experiment treatment status) (see Appendix Table A4).

Effects of Private vs. Non-private accounts. Table 1, Panel B presents the impact of being offered the Private accounts relative to the omitted category of the Non-private accounts. While a robust 60% of workers take up the Private accounts, only 14% do so for the Non-private ones—a 77% reduction (p<0.001, Col. 1). This is consistent with our hypothesis that the blocked savings account is substantially less desirable as a savings vehicle when it does not shield workers from redistributive requests.

This sharp decline in take-up reduces the scope for Non-private accounts to impact labor supply. Consistent with this, being offered the Private account leads to 152.9 FCFA or 8.8% higher earnings per week (Col. 2, p=0.063), and 6.6 p.p. or 8.4% higher attendance (Col. 3, p=0.045).[27] The magnitude of these effects on earnings is quite similar to that from the main experiment. Consequently, in Panel C, we pool data from both the main and mechanism experiments, estimating the impact of being offered a Private account relative to the pooled comparison group of being assigned to either the Control or Non-private arms. Effects are similar in magnitude, and more precise given the larger sample size.

[27]Note that we do not interpret these effects as the direct productivity effects of the Private account, since they are mediated through take-up. Rather, these effects indicate that blocked accounts have large earnings benefits for the compliers (i.e. those who are affected by redistributive pressure, taking up blocked accounts in Private but not Non-private), but lose this value when other network members would know of their existence.

Reasons for low Non-private take-up. To get direct evidence on why Non-private accounts are perceived as less attractive, we asked workers who declined to take up a Non-private account their primary reasons for doing so (Figure 5). Consistent with our hypothesis, 96% of workers said that the account would cause net transfer requests from outside the home to go up. In addition, 46% said that their spouse's contribution to household spending could go down. Only 5% reported other reasons, primarily citing concerns like jealousy from network members. This role of redistributive pressure as the main reason for not taking up the Non-private account is also directly visible from the open answers that workers provided when asked why they didn't want to take up (Appendix Figure A5).[28]

Interpretational concerns. We examine three sets of interpretational concerns with the Non-private accounts. The first is that the Non-private accounts may lead to more redistributive requests than the status quo by making the worker's cash-on-hand especially salient (i.e. lead to a tax rate above τ_0)—leading us to overestimate the treatment effects of reducing redistributive pressure. Note that this requires distortionary effects from redistributive pressure, and so is not qualitatively a confound per se, but would have quantitative implications for our results.

However, in our setting, the factory's paydays are known publicly and a salient trigger for redistributive requests. In contrast, in the Non-private accounts, for the cash to be accessible, the worker must physically withdraw it from the bank (which could happen weeks after the unblock date). 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 Private accounts in the mechanism experiment, their total savings at the end of the 3-month blocked period is roughly comparable to (i.e. equivalent to 105% of) their average take home cash pay in a given paycycle. 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.

Second, 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 publicly known days. However, not a single worker in our sample reports ever having been robbed on the way home from work, and only 0.9% report ever having faced theft in relation to payday.

[28]In Figure 5, workers were provided a set of options and asked to select the ones that were the primary reason for declining the account, or supply their own reason. Before giving them this list, we also elicited open ended qualitative responses on their reason for declining. We code the answers to the open ended responses by category in Appendix Figure A5. Among those that replied, 23% report not wanting a blocked account (of any kind). Many others (30%) did not give an informative response, saying something like "I don't want a Non-private account." Redistributive pressure is by far the most common unprompted reason given, both overall (42%) and among those giving a specific reason(91%).

Figure 5: Non-private Take-up: Drivers of Decision

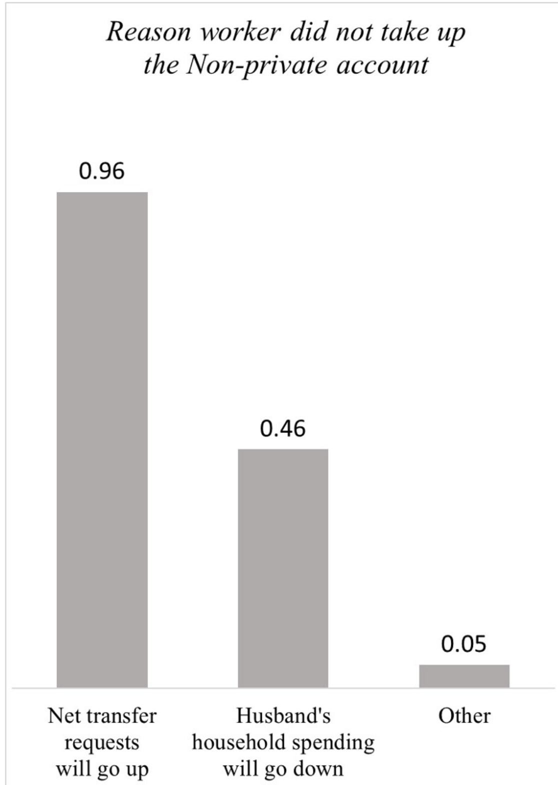

Notes: Reasons for not taking up Non-private account. Elicited from workers assigned to the Non-private treatment who refused to take-up the Non-private blocked accounts; collected when workers report their take-up decision to research staff. Workers were asked what were major factors that drive their 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.

Third, a general desire for privacy (rather than redistributive pressure per se) could drive the unpopularity of the Non-private accounts. We construct a test for whether workers generally object to others knowing about their financial lives. We undertake the test with workers who took up *Private* accounts in the main experiment, but were offered *Non-private accounts* in the mechanism experiment. Three months after the end of the main experiment, we ask these workers for permission to send text messages to their network members advertising that they had saved in a blocked account through the bank in the past. 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 earnings gain for workers who take up Private accounts in the mechanism experiment. This placebo 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.[29] A striking 88% of workers agree to this offer.

Of course, this is not exactly equivalent to the Non-private treatment. However, it indicates 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 to explain why workers leave so much money on the table (i.e. 9% of their full-time earnings) by refusing Non-private accounts (relative to Private accounts).

This is perhaps unsurprising, given the cultural context. As we discussed above, workers are paid the cash portion of their earnings in front of each other by the factory. Women frequently participate in ROSCAs, so that their payments, loans, and payout dates are common knowledge in their community. Workers are therefore used to many aspects of their financial lives being shared with others.

7.2 Heterogeneity in treatment effects by redistributive pressure

As a further test of mechanisms, we examine whether treatment effects are higher for workers who report more redistributive pressure at baseline.

In Table 2, we examine three baseline proxies for redistributive pressure (Cols. 4-6); in each case, we find that treatment effects of the Private accounts are concentrated among those who face higher pressure. For example, in Col. (4), among workers who make frequent transfers to others at baseline, being offered the Private account increases earnings by 272 FCFA (14.7%, p =0.008). However, among those who infrequently make transfers, we cannot reject that the Private accounts have no treatment effects on earnings. We find qualitatively similar patterns in Col. (5), where we examine heterogeneity by whether workers state that they cannot accumulate savings due to redistributive pressure.[30] In Col. (6), we examine another proxy for the severity of 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

[29] Recall from Section 2 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.

[30] Among the questions in Figure 2, at baseline in both the main and mechanism experiments, we only asked the question in Panel D: "I have difficulty saving over time for large goals, because if I put money aside, someone else will need it for something urgent". The covariate in Col. (5) equals zero for workers who strongly agree with this statement.

altruism. In qualitative work, workers expressed particular frustration about such transfers. As in Col. (4), treatment effects are markedly large among those facing such pressure, and substantively smaller and insignificant among those facing less pressure.

Table 2: Heterogeneous Treatment Effects: Redistributive Pressure

	Baseline covariate:						
	Redistributive Pressure						**Intra-HH**
	Low Pressure (PCA)			Infrequ. share money	Savings not taxed	Not taxed by acquaint.	Has a partner
	Pooled	Main exper.	Mechan. exper.				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Private account	255.2	277.0	243.6	272.1	148.2	334.5	160.0
	(98.57)	(158.6)	(126.8)	(102.0)	(78.69)	(155.5)	(117.5)
	[0.010]	[0.082]	[0.056]	[0.008]	[0.060]	[0.032]	[0.174]
Private account	-295.0	-579.4	-221.1	-235.8	-131.1	-234.6	-39.62
$\times$ Baseline covariate	(136.7)	(302.3)	(167.7)	(136.1)	(176.2)	(184.6)	(144.4)
	[0.031]	[0.056]	[0.188]	[0.084]	[0.457]	[0.205]	[0.784]
Sample (experiments)	Pooled	Main	Mechan.	Pooled	Pooled	Pooled	Pooled
Dep var mean, Covar.=0	1828	1749	1888	1853	1721	1971	1764
Dep var mean, Covar.=1	1662	1606	1681	1644	1808	1677	1711
Share: covariate = 1	0.39	0.27	0.50	0.56	0.22	0.74	0.59
P-val: sum = 0	0.689	0.247	0.838	0.702	0.914	0.316	0.174
N: observations	161138	122916	38222	161138	161138	161138	161138
N: workers	474	354	317	474	474	474	474

Notes: Unit of observation is worker-day. Dependent variable is daily earnings (in FCFA). Covariate in each of columns (1)-(6) is an indicator variable equal to one if the worker reports facing little redistributive pressure. Individual proxies for (facing little) redistributive pressure in Cols. (4)-(6) are: providing financial assistance to others less than once per month (Col. 4); not strongly agreeing 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. 5); reporting never giving financial assistance to acquaintances (Col. 6). Covariate in Cols. (1-3) is an indicator for whether the first principal component from a PCA on the variables in Cols. (4-6) is above median (hence equal to 1 if baseline redistributive pressures are low). Last, in Col. (7), covariate equals one if the worker reports having a partner. For each observation, we use as covariate the most recent baseline value. Cols. (2) and (3) use data only from the main and mechanism experiments, respectively; remaining columns pool data from both experiments. The Col. (6) covariate was only asked during the mechanism experiment, so all variation is estimated off the mechanism experiment in this column. Each column shows results from a stacked difference-in-differences specification. 12 working days of baseline data. All regressions include day-by-wave and worker-by-wave fixed effects. Standard errors clustered at the worker level.

In Col. (1) of Table 2, we examine heterogeneity by whether the first principal component of these three proxies is above the sample median (equal to 1 if baseline redistributive pressures are low). The results are similar to those in Cols. (4)-(6) but more precise. The above analysis looks at the pooled sample of all workers (comparing the Private accounts against the pooled omitted category of Control and Non-private Accounts, as in Panel C of Table 1). In Cols. 2 and 3, we examine this heterogeneity separately for the main and mechanism experiments, respectively. We find that in each experiment, effects are concentrated among those who face higher redistributive pressure; for those who face low pressure, we cannot reject that there are no treatment effects in either experiment. Overall, these findings are consistent with previous work indicating the relevance of heterogeneity in redistributive pressure (e.g. Dupas and Robinson, 2013b; Riley, 2024; Squires, 2024).

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. In Col. (7), we examine heterogeneity by whether the worker has a partner; while effects are noisier, we see no evidence that impacts are larger among women who have a partner. The estimated effect for single workers is 160 FCFA (9.1%, p=0.174), and the estimated effect for workers with a partner is slightly smaller at 120 FCFA (7%, p=0.174). This pattern is similar across both the main and mechanism experiments (Appendix Table A5). This is also consistent with the results in Col. (6) indicating that effects are concentrated among those who want to avoid requests outside their household and close family and friends. Finally, our data indicate there were no changes in average reported intra-household bargaining power (Donald, forthcoming)—possibly because the average woman in our sample had been working full-time for years. 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 above, some women who refuse Non-private accounts cite concerns about changes in their partner's contribution to household expenses. However, our findings in Table 2 suggest that redistributive pressure *outside* 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, 2024).

7.3 Confounds: Internal benefits

Sections 7.1 and 7.2 present positive evidence for the role of redistributive pressure. In this section, we examine evidence for whether the Private accounts affected output through *internal* psychological benefits: fairness concerns, self-control, and goal setting.

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. Treatment assignment occurred by selecting (confidential) worker IDs via a lottery in the factory, conducted by us, with the bank present. Relatedly, the marketing of the blocked accounts, the lottery, and the paperwork in the signup process conveyed that the employer, Olam, had no role in picking who received accounts. 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 A6). 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=0.079).

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 accounts: in both cases, sophisticates should see value in the accounts and choose them. Redistributive pressure is therefore necessary to explain our results.

To gauge the potential relevance of present focus, we test a core prediction of basic time inconsistency models. In one paycycle, we surprise workers with the option to opt out of depositing earnings into their blocked accounts for just that paycycle—randomly varying whether this option is provided four days before the payday, or on the payday itself.[31] Note that we would expect some opt out even in the absence of any self-control problems—for example, due to shocks that may increase cash needs in some weeks. 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,

[31] 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.

[32] Present-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.

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 A6). 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 present focus.

Similarly, we do not find evidence for heterogeneity in treatment effects using proxies for present focus. Following Kaur et al. (2015), we document that on average, workers are substantively more likely to come to work the closer they are to the end of the paycycle (i.e. the last set of days where their labor supply will matter for their next paycheck) (Appendix Table A7). There is substantial heterogeneity across workers in these paycycle effects. However, we see no evidence that workers with more pronounced paycycle effects have larger treatment effects from being offered the blocked accounts (Appendix Table A8).

Thus, while individuals may face self-control problems in general (e.g., Ashraf et al., 2006; Brune et al., 2021), the above two tests suggest this mechanism is unlikely to strongly drive the effects of the blocked accounts in our specific experiment. This may be because commitment devices become ineffective if individuals are partially naive or the threshold is not set exactly at the right value given a person's specific discount function (Heidhues and Kőszegi, 2009; Bai et al., 2021). Moreover, the blocked accounts push the receipt of earnings (i.e. the returns to effort) even further into the future—a force that could actually *decrease* the effort of present-focused agents (O'Donoghue and Rabin, 1999; Kaur et al., 2015).

Goal setting. Related to the above, because the blocked accounts require selecting a threshold, they may motivate workers through goal-setting or soft commitment. Again, 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 right after their threshold: the threshold creates a "goal" and workers become motivated to reach the goal; thus, once the goal is reached, motivation to continue working harder should go down. In contrast, our hypothesized mechanism of redistributive pressure implies no such bunching: the benefits of a lower tax rate kick in after the worker crosses the threshold, and so the accounts are only useful if workers overshoot the threshold. Consistent with this latter prediction, we see substantial overshooting in the data (Appendix Figure A7). Among workers with blocked accounts, earnings are more than 10% away from the chosen threshold in 82% of paycycles. Conditional on earning weekly above the threshold, earnings are on average 24% higher than the worker's chosen threshold level; they overshoot the chosen threshold by over 5% in 84% of these paycycles, and by over 10% in 70% of them.

Discussion. The blocked accounts could in principle offer a range of internal benefits.[33]

[33] In addition to the prominent internal mechanisms discussed above, the accounts could reduce stress or worry about meeting future expenses (Kaur et al., 2025), potentially amplifying the effects of reducing the social tax.

While we offer some suggestive evidence against such concerns above, it is not necessarily conclusive. Our view is that any alternate explanation for our findings must explain why demand for the blocked accounts plummets when they are Non-private, and why workers cite redistributive pressure as the main reason for not taking up the accounts. Moreover, it must explain why treatment effects are concentrated entirely among those who report high amounts of redistributive pressure at baseline; among those with low redistributive pressure, we find little evidence for treatment effects. Consequently, we argue that redistributive pressure is likely necessarily to explain our findings. That said, conditional on taking up the accounts, workers' labor supply might be affected by internal benefits of the accounts. This implies that, while our paper offers qualitative evidence for redistributive pressure, caution is needed to interpret the specific magnitude of the treatment effects.

8 Effects on transfers

We designed the blocked accounts so that the expected *cash* component of biweekly take-home pay did not decline. Consequently, we should see no decline in transfers to others. In fact, transfers may weakly go up for two reasons. First, since the thresholds chosen by workers were often higher than average baseline earnings, treated workers would have seen a weakly positive effect on take-home cash earnings—some of which could have been taxed by their network. Second, when the accounts were unblocked, some of the savings could have been redistributed to others.

In Table 3, we examine treatment effects on transfers to individuals outside of the worker's household. In Panel A, we examine the extensive margin of making or receiving any transfers. We do not find significant changes in the likelihood of having made a transfer to anyone in the last three months in either the main or mechanism experiment (Cols. 2 and 3, respectively). In the mechanism experiment, we also asked more detailed information about transfers to subgroups: we find no discernible impact on making a transfer to family (Col. 4), making a transfer to non-family (Col. 5), or in the net transfer position (Col. 6). In Panel B, we examine the total amount of transfers. Point estimates are generally positive, but largely insignificant—indicating that transfers to the network did not go down, and may even have increased.

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—making workers better off, without reducing financial support to the network.[34] This offers a proof of concept that it may be

[34]In addition, we find positive (but statistically insignificant) improvements for workers in a measure of subjective well-being.

Table 3: Transfers to Network Members

	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Likelihood of transfer					
	To anyone			To family	To non-family	Net position
	Pooled	Main exper.	Mechan. exper.	Mechan. exper.	Mechan. exper.	Mechan. exper.
Private account	0.0453	0.0380	0.0513	0.0587	-0.00797	-0.0535
	(0.0399)	(0.0694)	(0.0454)	(0.0578)	(0.0569)	(0.100)
	[0.257]	[0.585]	[0.259]	[0.310]	[0.889]	[0.595]
Dep var mean	0.652	0.504	0.784	0.510	0.405	0.314
	Panel B: Amount transferred					
	To anyone			To family	To non-family	Net amount
	Pooled	Main exper.	Mechan. exper.	Mechan. exper.	Mechan. exper.	Mechan. exper.
Private account	3338.1	650.8	5524.7	4435.4	1089.3	3535.2
	(2088.3)	(3360.7)	(2599.0)	(2450.6)	(973.4)	(3288.3)
	[0.111]	[0.847]	[0.034]	[0.071]	[0.264]	[0.283]
Dep var mean	11698	10100	13129	9967	3162	4848

Notes: Data from endline surveys, covering transfers in past 3 months. In the phone endline survey to the main experiment, data is only available for aggregate transfers to individuals outside the household. This phone survey did not include an exhaustive questionnaire to prompt recall of all transfers, and so the reported levels may be underestimates. In the in-person endline survey to the mechanism experiment, disaggregated data for specific categories of individuals is available. 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)-(5). In Col. (6), dependent variable equals 1 if the respondent transfers more than they receive, -1 if they receive more than they transfer, and 0 if they give and receive the same amount. Panel B dependent variable is the continuous transfer amount sent in Cols. (1)-(5) and the net amount of transfers sent (transfers sent minus transfers received) in Col. (6). N=235 main experiment workers (Col. 2), N=299 mechanism experiment workers (Cols. 3-6), and N=381 workers in pooled sample of mechanism and main experiments (Col. 1). Standard errors clustered at the worker level.

possible to design mechanisms that undo distortions from redistributive pressure, without making others in the network worse off.

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 draw from the literature on formal taxation, with an overview here and detailed derivations in Appendix A.3.

Building on the model introduced in Section 4, we can rewrite Equation 2 to obtain the following expression:

$$\frac{de_1}{e_1} = \zeta \frac{d(1-\tau_1)}{(1-\tau_1)}, \tag{4}$$

where ζ is the compensated elasticity of labor supply. This expression describes how e_1 changes with τ_1—starting from the case where τ_1 equals τ_0, hence e_1 equals e_0 (i.e. moving from Figure 3A to 3B). 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. Equation 4 thus indicates that:

$$\frac{1-\tau_0}{1-\tau_1} \approx \left(\frac{e_0}{e_1}\right)^{\frac{1}{\zeta}} \tag{5}$$

We can recover an estimate of $\frac{e_0}{e_1}$ from our main treatment effect estimate of 9.43% (Table 1, Col. 2, Panel A). 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 use a range of estimates from the literature. In addition, we also randomized piece rates in the factory (see Appendix A.4.5 for details).

Finally, expression 5 requires an estimate of τ_1, the social tax rate faced by treated workers under the Private account. The most conservative assumption (yielding the smallest tax rate estimate) is that $\tau_1 = 0$: the accounts fully eliminate any social tax. In reality, the effects on transfers in Table 3, along with the fact that many workers set thresholds above baseline earnings, suggest that τ_1 is likely greater than zero. We consequently also present estimates for a range of values of τ_1.

We present estimates of the social tax rate implied by our results in Table 4. 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 Local Average Treatment Effect (LATE) for the subset of workers induced to choose the accounts.

Table 4: Social tax rate estimates

	Labor supply elasticity							
	0.10	0.18	0.25	0.50	0.75	0.98	1.00	1.25
Private account tax rate (τ_1)	**Panel A: Baseline tax rate τ_0, Sample average (ITT)**							
$\tau_1 = 0\%$	59%	39%	30%	17%	11%	9%	9%	7%
$\tau_1 = 2.5\%$	60%	41%	32%	19%	14%	11%	11%	9%
$\tau_1 = 5\%$	61%	42%	34%	21%	16%	13%	13%	12%
Private account tax rate (τ_1)	**Panel B: Baseline tax rate τ_0, among Compliers (LATE)**							
$\tau_1 = 0\%$	85%	65%	53%	32%	22%	18%	17%	14%
$\tau_1 = 2.5\%$	85%	66%	54%	33%	24%	20%	19%	16%
$\tau_1 = 5\%$	86%	67%	55%	35%	26%	22%	21%	18%

Notes: This table presents the baseline social tax rate, τ_0, faced by workers, estimated for various values of τ_1 (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 $\zeta = 0.98$. Expression (5) is used for the computation of τ_0.

The current literature, drawing on estimates primarily from richer countries, tends to find that labor supply elasticities are usually fairly small (Bargain and Peichl, 2016). Similarly, using estimates from a field experiment in Malawi, Goldberg (2016) finds a low attendance elasticity (i.e. 0.15); this, when combined with our treatment effect estimates, suggests a total labor supply elasticity of 0.18.[35] Table 4 indicates that such low elasticities imply large social tax rates: 39%-42% for the average worker, and 65%-67% for workers who take up the Private blocked accounts.

However, the literature suggests that the labor supply elasticity tends to be higher for lower-income workers and gig economy workers, who may have more ability to adjust hours and have lower labor force attachment (e.g. Meghir and Phillips, 2010; Bargain and Peichl, 2016; Zidar, 2019; Chen et al., 2019).[36] Consistent with this, using rainfall shocks as demand shifters, Kaur (2019) finds a labor supply elasticity between 0.8-1.3 among casual agricultural workers in India.[37] This would imply a much lower tax rate: 9-13% for the average worker, and 17-21% for those who take up the accounts.

The above discussion highlights the fact that 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 elasticity—which in turn would imply that even modest tax rates can lead to large changes in labor supply behavior. Alternately, if the elasticity is low, this necessarily implies an extremely large tax rate. In light of this, the large treatment effects of the Private accounts in Table 1 can be viewed as a useful and transparent signal on the potential for meaningful distortions on labor supply.

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 suggest 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

[35]The full labor supply elasticity is the sum of the extensive margin (i.e. attendance) elasticity plus the intensive margin (i.e. output conditional on attendance) elasticity. We use the ratio of the attendance to productivity treatment effects in our experiment to proxy for the ratio of attendance to productivity elasticities and recover a total labor supply elasticity estimate. See Appendix A.3.6 for details and a discussion.

[36]In Appendix A.3.7, we theoretically examine how the labor supply elasticity changes with individuals' level of consumption, for the simple case of utility that is separable in consumption and effort. We document that, all else equal, the labor supply elasticity will be increasing in the marginal utility of consumption—implying that poorer workers would be more responsive to tax rate changes.

[37]Positive rainfall shocks shift labor demand, with associated wage and employment changes tracing the slope of the labor supply curve. Positive shocks lead to a 6.3% or 7.2% change in the wage (depending on specification), with an associated estimated employment change of 8.3% or 5.7%, respectively.

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 affect earnings and productivity.

It is worth noting that redistributive pressure is not specific to the study setting of Côte d'Ivoire. A long tradition of work in the social sciences argues that such pressure is relevant in a wide range of settings, including in rich countries such as the United States (Stack, 1974; O'Brien, 2012; Wherry et al., 2019). To provide motivational evidence on the potential generality of our mechanism, we collected data from residents in the United States via Prolific (Figure 6). The findings indicate that a meaningful share of Americans—albeit smaller in absolute terms than in Côte d'Ivoire—similarly report facing prevalent redistributive pressure. In this sample, 43% of low-income Black and Hispanic individuals in the US deem it likely or very likely that, if someone in their community decides to make more money by working longer hours, their extended family or friends would start asking that person more often for financial help (Panel A). Consistent with evidence from Côte d'Ivoire, individuals who start a new job also face redistributive pressure. For instance, 41.5% of low-income Black and Hispanic individuals in the US deem it likely or very likely that, if someone in their community gets hired for a full-time job, their extended family or friends would start asking that person more often for financial help (Panel C).

The magnitude of our findings is potentially surprising given that, in rich countries, labor supply does not react strongly to changes in formal taxation; most studies find elasticities close to zero (Bargain and Peichl, 2016). On one hand, as we discuss in Section 9, there is evidence that this may be a consequence of the high labor force attachment, coupled with inability to flexibly change hours conditional on being employed, in rich countries. This is consistent with the higher labor supply elasticities among workers with less attachment, lower-income workers, and gig economy workers. For example, Chen et al. (2019) find extremely large elasticities among Uber drivers, typically ranging between 1.8 to 2. On the other hand, if the labor supply elasticity is indeed low, then our estimates imply extremely high social tax rates. This is consistent with Squires (2024), who estimates that 30% of the micro-entrepreneurs in his sample face a social tax rate of 50%. It is impossible to distinguish between these two potential interpretations without reliable estimates of the labor supply elasticity for workers in developing countries—estimates which are sorely lacking. For example, existing elasticity estimates—0.15 in Goldberg (2016) and up to 1.3 in Kaur (2019)—are vastly different in value, and each have their limitations. Given the importance of this parameter as an input into labor market policy, this is an important gap in the development labor literature.

Regardless, as we discuss at the end of Section 9, the magnitude of our experimental impacts suggests the potential for meaningful distortions on labor supply. This has potential

Figure 6: Redistributive Pressure in the United States

(A) Requests on labor income low-income households

(B) Requests on labor income high-income households

(C) Requests after employment low-income households

(D) Requests after employment high-income households

Notes: Redistributive pressure reported by respondents in the US, for three racial groups (black, hispanic, white) and two income groups (low-income: annual income below USD 60,000; high-income: annual income above USD 60,000). In each panel, N=150 respondents. Data collected on Prolific.

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 2, Panel C). Among those who do hold these jobs, our experimental findings suggest the potential for social taxation to 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; 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? Moreover, could the presence of a social tax lead to complementarities in labor supply within the network, generating the possibility of multiple equilibria (Hoff and Sen, 2011; Donald and Grosset-Touba, 2025)? 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.[38] This is in line with strategies, such as ROSCA participation, that are already commonly employed in this setting. 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., 2023).

Finally, our findings suggest an additional route through which improving formal safety nets could boost productivity: by reducing demands for redistribution on others in bene-

[38]We 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 increasing their trust in the offered accounts. Moreover, virtually everyone who took up the account once did so again when offered a subsequent time.

ficiaries' 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.

Almås, I., A. Armand, O. Attanasio, and P. Carneiro (2018): "Measuring and 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.

ANSD (2022): "Enquête Harmonisée sur les Conditions de Vie des Ménages 2018-2019 Sénégal," Data set from the Agence National de la Statistique et de la Démographie (ANSD). Provided by World Bank, Development Data Group, ref. SEN_2018_EHCVM_v01_M `https://doi.org/10.48529/HHHX-J012`.

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.

Baker, A. C., D. F. Larcker, and C. C. Y. Wang (2022): "How much should we trust staggered difference-in-differences estimates?" *J. Financ. Econ.*, 144, 370–395.

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 (2023): "Changes in Social Network Structure in Response to Exposure to Formal Credit Markets," *The Review of Economic Studies*.

Banerjee, A. V. and E. Duflo (2007): "The Economic Lives of the Poor," *The Journal of Economic Perspectives*, 21, 141–167.

Bargain, O. and A. Peichl (2016): "Own-wage labor supply elasticities: variation across time and estimation methods," *IZA Journal of Labor Economics*, 5, 1–31.

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.

Callen, M., S. De Mel, C. McIntosh, and C. Woodruff (2019): "What are the headwaters of formal savings? Experimental evidence from Sri Lanka," *The Review of Economic Studies*, 86, 2491–2529.

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.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019): "The Effect of Minimum Wages on Low-wage Jobs," *The Quarterly Journal of Economics*, 134, 1405–1454.

Chen, M. K., P. E. Rossi, J. A. Chevalier, and E. Oehlsen (2019): "The value of flexible work: Evidence from Uber drivers," *Journal of political economy*, 127, 2735–2794.

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*.

Donald, A. and F. Grosset-Touba (2025): "Complementarities in Labor Supply: How Joint Commuting Shapes Work Decisions," .

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 (2013a): "Savings constraints and microenterprise development: Evidence from a field experiment in Kenya," *American Economic Journal: Applied Economics*, 5, 163–192.

——— (2013b): "Why Don't the Poor Save More? Evidence from Health Savings Experiments," *The American Economic Review*, 103, 1138–1171.

——— (2013c): "Why don't the poor save more? Evidence from health savings experiments," *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, 2342–2375.

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.

Gruber, J. and E. Saez (2002): "The elasticity of taxable income: evidence and implications," *Journal of Public Economics*, 84, 1–32.

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.

Horn, S., J. Jamison, D. Karlan, and J. Zinman (2021): "Does lasting behavior change require knowledge change? Evidence from savings interventions for young adults," *Evidence from Savings Interventions for Young Adults (January 2021). Global Poverty Research Lab Working Paper.*

INE (2022): "Inquérito Harmonizado sobre as Condições de vide dos Agreagados Familiares 2018-2019 Guinea-Buissau," Data set from the Instituto Nacional de Estatística (INE). Provided by World Bank, Development Data Group, ref. GNB_2018_EHCVM_v01_M `https://doi.org/10.48529/1EKB-M086`.

INS-CIV (2022): "Enquête Harmonisée sur les Conditions de Vie des Ménages 2018-2019 Côte d'Ivoire," Data set from the Institut National de la Statistique (INS). Provided by World Bank, Development Data Group, ref. CIV_2018_EHCVM_v01_M `https://doi.org/10.48529/8WH3-BF40`.

INS-Niger (2022): "Enquête Harmonisée sur les Conditions de Vie des Ménages 2018-2019 Niger," Data set from the Institut National de la Statistique (INS). Provided by World Bank, Development Data Group, ref. NER_2018_EHCVM_v01_M `https://doi.org/10.48529/GGAM-AX39`.

INSAE (2022): "Enquête Harmonisée sur les Conditions de Vie des Ménages 2018-2019 Benin," Data set from the Institut National de la Statistique et de l'Analyse Économique (INSAE). Provided by World Bank, Development Data Group, ref. BEN_2018_EHCVM_v01_M `https://doi.org/10.48529/RN3K-Z374`.

INSD (2022): "Enquête Harmonisée sur les Conditions de Vie des Ménages 2018-2019 Burkina Faso," Data set from the Institut National de la Statistique et de la Démographie (INSD). Provided by World Bank, Development Data Group, ref. BFA_2018_EHCVM_v01_M `https://doi.org/10.48529/WV88-J486`.

INSEED (2022): "Enquête Harmonisée sur les Conditions de Vie des Ménages 2018-2019 Togo," Data set from the Institut National de la Statistique et des Etudes Economiques et Démographiques (INSEED). Provided by World Bank, Development Data Group, ref. TGO_2018_EHCVM_v01_M `https://doi.org/10.48529/WW9Z-D865`.

INSTAT (2022): "Enquête Harmonisée sur les Conditions de Vie des Ménages 2018-2019 Mali," Data set from the Institut National de la Statistique (INSTAT). Provided by World Bank, Development Data Group, ref. MLI_2018_EHCVM_v01_M `https://doi.org/10.48529/90E9-4E91`.

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.

Jayachandran, S. (2015): "The roots of gender inequality in developing countries," *Annual review of economics*, 7, 63–88.

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. (2019): "Nominal wage rigidity in village labor markets," *American Economic Review*, 109, 3585–3616.

Kaur, S., M. Kremer, and S. Mullainathan (2015): "Self-Control at Work," *The Journal of Political Economy*, 123, 1227–1277.

Kaur, S., S. Mullainathan, S. Oh, and F. Schilbach (2025): "Do financial concerns make workers less productive?" *The Quarterly Journal of Economics*, 140, 635–689.

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*.

Meghir, C. and D. Phillips (2010): "Labour supply and taxes," *Dimensions of tax design: The Mirrlees review*, 202–74.

Miracle, M. P., D. S. Miracle, and L. Cohen (1980): "Informal savings mobilization in Africa," *Economic Development and Cultural Change*, 28, 701–724.

Mirrlees, J. A. (1971): "An Exploration in the Theory of Optimum Income Taxation," *The Review of Economic Studies*, 38, 175–208.

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.

Pomeranz, D. and F. Kast (2024): "Savings accounts to borrow less: experimental evidence from Chile," *Journal of Human Resources*, 59, 70–108.

Portes, A. (1998): "Social Capital: Its Origins and Applications in Modern Sociology," *Annual Review of Sociology*, 24, 1–24.

Riley, E. (2024): "Resisting social pressure in the household using mobile money: Experimental evidence on microenterprise investment in Uganda," *American Economic Review*, 114, 1415–1447.

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.

Somville, V. and L. Vandewalle (2023): "Access to banking, savings and consumption smoothing in rural India," *Journal of Public Economics*, 223, 104900.

Squires, M. (2024): "Kinship taxation as an impediment to growth: experimental evidence from Kenyan microenterprises," *The Economic Journal*, 134, 2558—2579.

Stack, C. B. (1974): *All Our Kin: Strategies for Survival in a Black Community*, Harper and Row.

Strauss, J., F. Witoelar, and B. Sikoki (2016): "The Fifth Wave of the Indonesia Family Life Survey (IFLS5): Overview and Field Report," Data set, ref. WR-1143/1-NIA/NICHD, downloaded from `https://www.rand.org/well-being/social-and-behavioral-policy/data/FLS/IFLS/download.html`.

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.

U.S. Census Bureau (2014): "2014 Survey of Income and Program Participation (SIPP)," Dataset downloaded from https://www.census.gov/programs-surveys/sipp/data/datasets.2014.html.

——— (2018): "2018 Survey of Income and Program Participation (SIPP)," Dataset downloaded from https://www.census.gov/programs-surveys/sipp/data/datasets/2018-data/2018.html.

——— (2019): "2019 Survey of Income and Program Participation (SIPP)," Dataset downloaded from https://www.census.gov/programs-surveys/sipp/data/datasets/2019-data/2019.html.

——— (2020): "2020 Survey of Income and Program Participation (SIPP)," Dataset downloaded from https://www.census.gov/programs-surveys/sipp/data/datasets/2020-data/2020.html.

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.

Zhou, A. and R. Mahadeshwar (2024): "The Impact of Intra-Household Income Hiding on Labor Productivity," .

Zidar, O. (2019): "Tax cuts for whom? Heterogeneous effects of income tax changes on growth and employment," *Journal of Political Economy*, 127, 1437–1472.

A Online Appendix

A.1 Appendix Figures

Figure A1: Cross-sectional average earnings and transfers

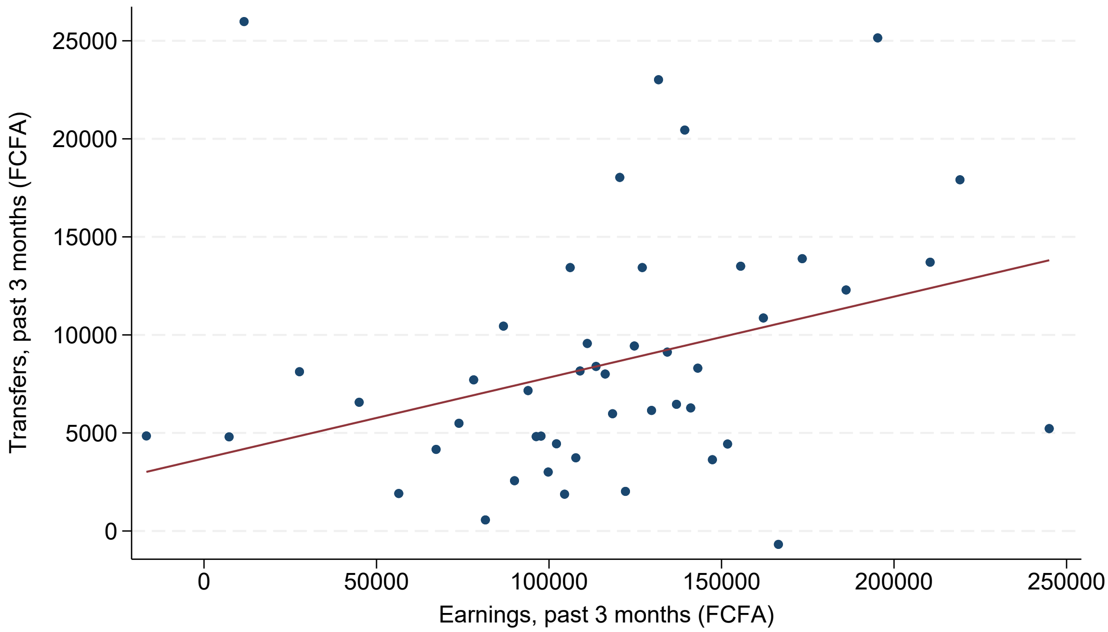

Notes: Relationship between average amount transferred and average factory earnings, both in the 3 months prior to the survey (in FCFA). Earnings from factory administrative data. Transfers from 2 rounds of worker phone surveys. Observations residualized from survey fixed effects. Transfers top-coded at the 99th percentile. Line of best linear fit reported. 356 observations, from workers not offered a Private savings account.

Figure A2: Psychological cost of lying about earnings' availability

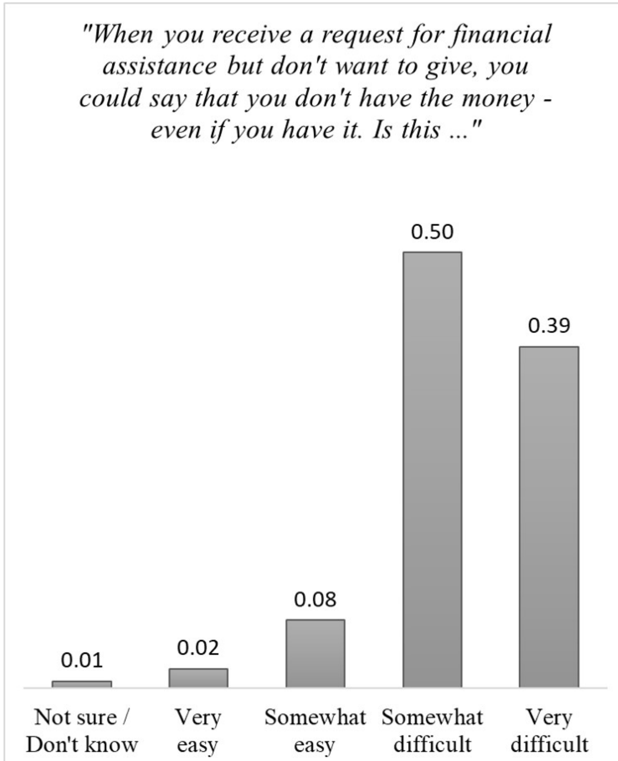

Notes: In-person survey with 488 cashew factory workers in Côte d'Ivoire.

Figure A3: Earnings shocks and transfers

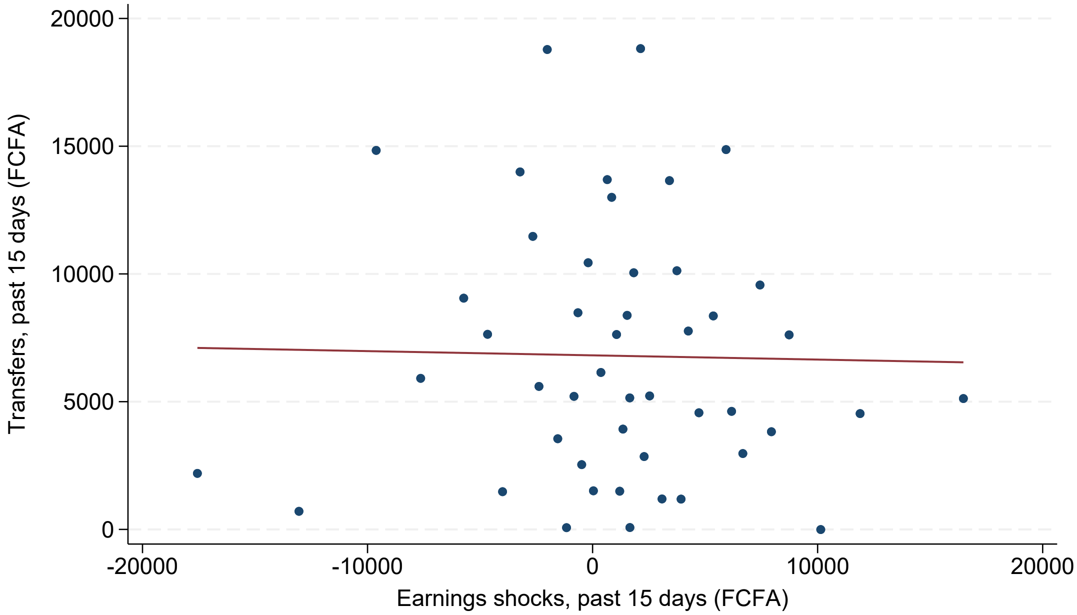

Notes: Relationship between transfers and earnings shocks. Earnings from factory administrative data. Earnings shocks defined as the difference between earnings in the 15 days prior to the survey relative to average earnings in past 3 months excluding that paycycle (i.e. the 75 days before). Transfers 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. 365 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 A4: Distribution of chosen threshold relative to baseline earnings

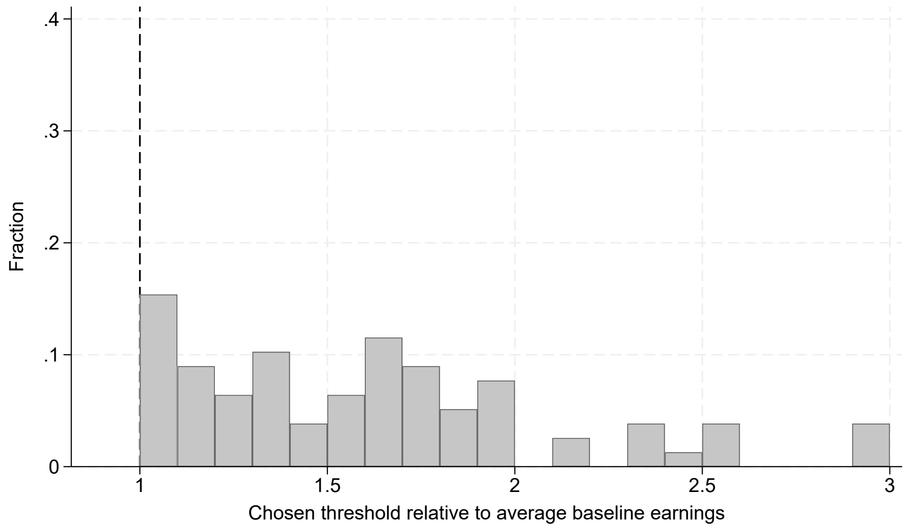

Panel A: Main Experiment

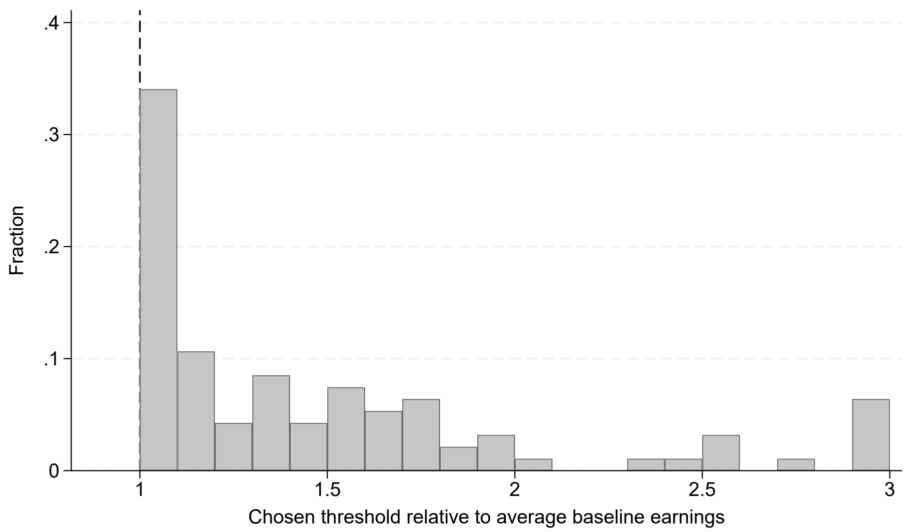

Panel B: Mechanism Experiment

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 = 78 treated workers holding a Private savings account in the main experiment, N=94 in the mechanism experiment.

Figure A5: Non-private Take-up: Open-Ended Motives

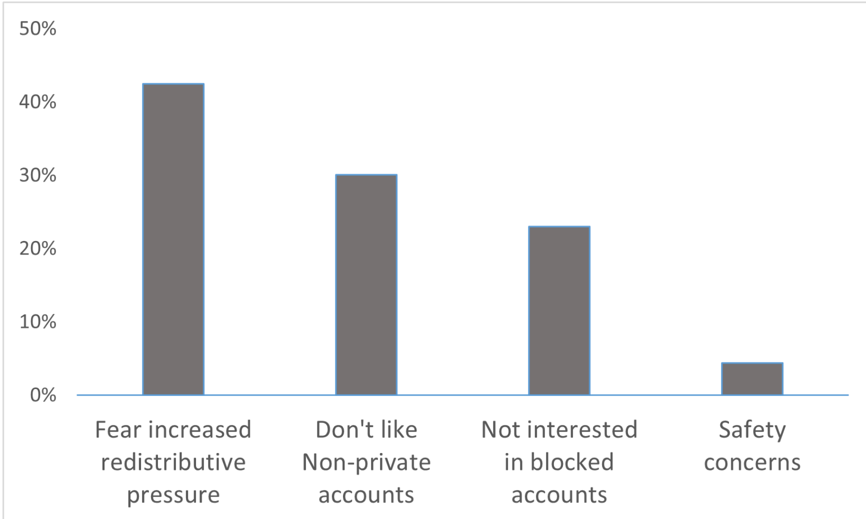

Notes: Drivers of take-up decision for the Non-private account. Elicited from workers assigned to the Non-private treatment who refused to take-up the Non-private blocked accounts; collected when workers report their take-up decision to research staff. Workers were ask the main reasons why they declined to take-up the offered account, and the research team coded the reasons given. N=113 workers.

Figure A6: Refusal to Opt Out

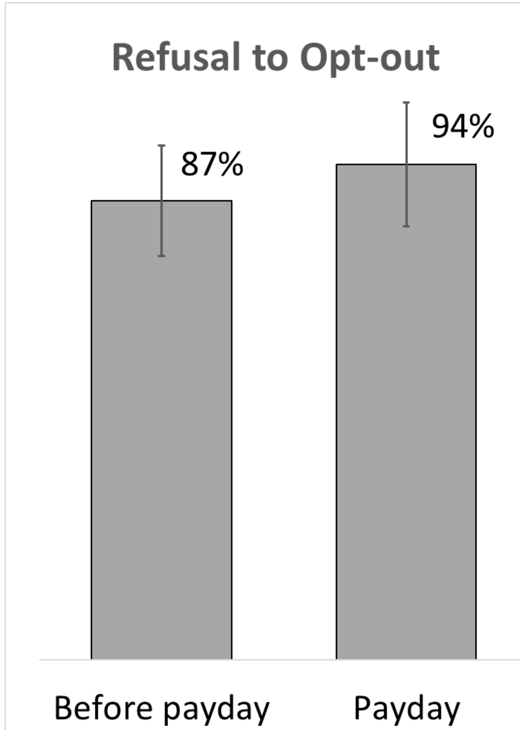

Notes: A subset of workers who took up blocked accounts was surprised with the option to opt out of having to deposit earnings increases into the accounts for one upcoming cycle. The figure shows the proportion who refused this offer to opt out, with 90% confidence intervals, separately for workers offered this option 4 days before the payday (left bar) and on the payday (right bar). N=92 workers chosen for exercise and who had blocked accounts in the mechanism experiment.

Figure A7: Distribution of treatment period earnings relative to the chosen threshold

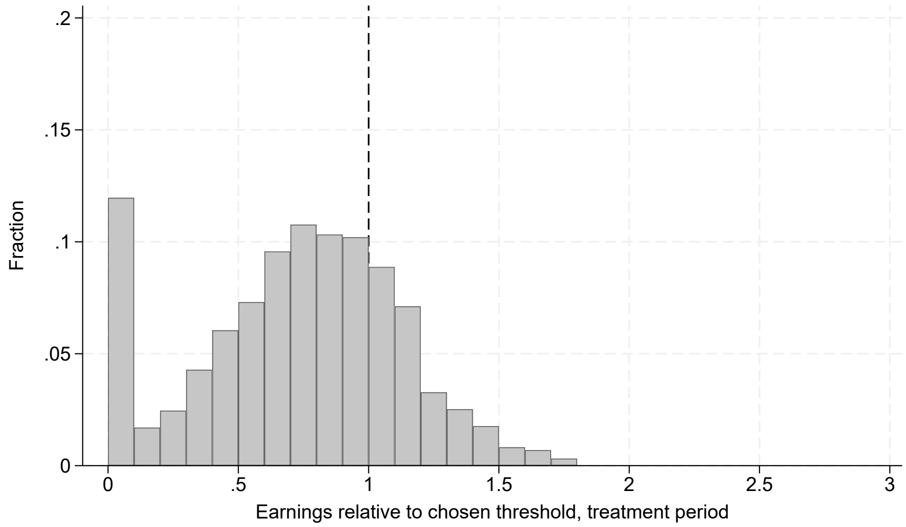

Panel A: Main Experiment

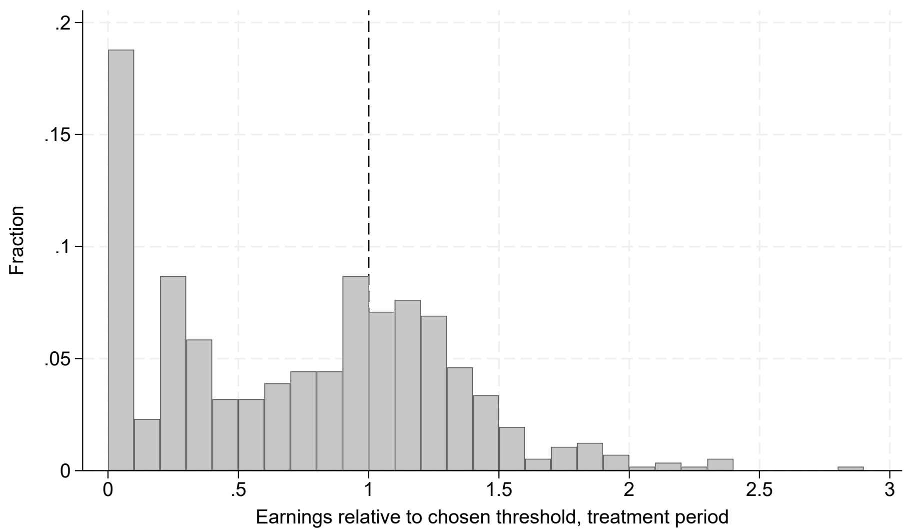

Panel B: Mechanism Experiment

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: $\frac{\text{earnings per paycycle}}{\text{chosen threshold}}$. Main experiment: N=1,588 worker-paycycles, among the 78 treated workers holding an active Private savings account. Mechanism experiment: N=564 worker-paycycles, among the 94 treated workers holding an active Private savings account.

A.2 Appendix Tables

Table A1: Balance

Variable	Main Experiment			Mechanism Experiment		
	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	259	-2	.769	466	-2	.934
	[169]	(8)		[234]	(25)	
Daily earnings (level)	1800	-165	.053	1783	-91	.295
	[790]	(85)		[745]	(87)	
Daily earnings (trend)	4.18	-1.33	.551	1.1	-2.65	.125
	[19.78]	(2.23)		[14.38]	(1.72)	
Daily attendance (level)	.697	-.052	.020	.727	-.055	.039
	[.209]	(.022)		[.203]	(.026)	
Daily attendance (trend)	-.001	-.001	.476	-.001	-.001	.108
	[.008]	(.001)		[.005]	(.001)	
Workers' characteristics						
Has an ID	.620	-.005	.907	.761	.044	.411
	[.487]	(.045)		[.429]	(.054)	
Is a woman	.986	.005	.688	.994	-.013	.305
	[.118]	(.013)		[.079]	(.013)	
Speaks Dioula	.556	-.032	.529	.410	-.025	.646
	[.498]	(.05)		[.493]	(.055)	
Speaks Baoule	.257	.059	.192	.255	.04	.424
	[.438]	(.045)		[.437]	(.05)	
Heterogeneity variables						
Infrequently share money	.500	.012	.894	.642	-.043	.453
	[.503]	(.091)		[.481]	(.058)	
Savings not taxed	.243	-.176	.009	.259	.006	.910
	[.432]	(.067)		[.439]	(.053)	
Not taxed by acquaintances	.	.	.	.710	.068	.202
	[.]	(.)		[.455]	(.053)	
Has a partner	.473	.118	.202	.687	-.101	.075
	[.503]	(.092)		[.465]	(.057)	
Omnibus F-test p-value			.151			.671

Notes: Summary statistics and tests for baseline balance by treatment group. Cols. (1)-(3) use data from the main experiment, and Cols. (4)-(6) use data from the mechanism experiment. 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 intervention waves, 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 2. Note that the "Not taxed by acquaintances" was not collected in the main experiment 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.

Table A2: Treatment effects – Robustness to alternative specifications

	(1)	(2)	(3)	(4)	(5)	(6)
Private account	162.4	162.4	162.4	164.0	177.7	162.4
	(71.92)	(71.50)	(71.46)	(75.19)	(77.14)	(71.92)
	[0.025]	[0.024]	[0.024]	[0.030]	[0.022]	[0.025]
Time FE	Day	Week	Paycycle	Day	Day	Day
Estimation sample	Main	Main	Main	Exclude first active paycycle	Exclude announcement period from baseline period	Exclude imbalanced randomization event
N: observations	122916	122916	122916	122916	122916	122916
N: workers	354	354	354	354	354	354

Notes: Results from the main experiment. Dependent variable: daily earnings. The omitted category is not being offered an account. Col. (1) is the same specification as Col. (1) of Table 1. As compared to our main specification, Col. (2) has week fixed effects; Col. (3) has paycycle 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 excluding the announcement period from the baseline period. Col. (6) shows robustness to excluding the newly eligible workers for the 3rd wave of the first study site, which is the randomization event that drives the imbalance in the level of baseline earnings reported in Table A1. All regressions include worker fixed effects. Standard errors clustered by worker.

Table A3: Treatment effects – Intervention period data only

	DiD	Intervention period data only		
	(1)	(2)	(3)	(4)
Private account	162.4	114.0	125.9	114.9
	(71.92)	(64.53)	(64.71)	(64.60)
	[0.025]	[0.078]	[0.052]	[0.075]
Controls	Worker FE, day FE	Earnings (linear)	Earnings (linear) + heterogeneity table covariates	Lasso (double-selection)
N: observations	122916	113999	113999	113999
N: workers	354	354	354	354

Notes: Results from the main experiment. Dependent variable: daily earnings. The omitted category is not being offered a Private account. Col. (1) is the same specification as Col. (2) of Table 1. Cols. (2)-(4) use endline data only, and control for intervention wave fixed effects. In addition, Col. (2) controls for baseline earnings (linearly), Col. (3) for linear baseline earnings plus the baseline heterogeneity variables used in Table 2, and Col. (4) for controls selected using double-selection lasso. Standard errors clustered by worker.

Table A4: Take-up rates - Mechanism Experiment

Mechanism Expe.	Eligible Main	Treated Main	Compliant Main	Take-up rate	N.
Private	All	All	All	.6	156
Private	Yes	No	-	.72	39
Private	Yes	Yes	All	.62	55
Private	Yes	Yes	Yes	.9	30
Private	No	-	-	.52	62
Non-Private	All	All	All	.14	161
Non-Private	Yes	No	-	.18	44
Non-Private	Yes	Yes	All	.19	59
Non-Private	Yes	Yes	Yes	.32	34
Non-Private	No	-	-	.07	58

Notes: This table disaggregates the mechanism experiment take-up results. "Eligible Main" denotes workers who have been eligible at any point in the main experiment (as opposed to workers newly eligible only in mechanism experiment). Overall, among the workers eligible for the mechanism experiment, 120 are newly eligible (62 assigned to Private + 58 assigned to Non-private) while the others (317-120=197) were already eligible for the main experiment. "Treated Main" denotes workers who were offered a Private account during the main experiment. "Compliant Main" denotes workers who took up a Private blocked accounts in the main experiment.

Table A5: Heterogeneous Treatment Effects: Having a Partner

	Pooled (1)	Main experiment (2)	Mechanism experiment (3)
Private account	160.0	224.9	162.9
	(117.5)	(165.6)	(146.6)
	[0.174]	[0.175]	[0.267]
Private account × Has a Partner	-39.62	-161.7	-23.57
	(144.4)	(260.5)	(177.8)
	[0.784]	[0.535]	[0.895]
Dep var mean if no partner	1764	1724	1798
Dep var mean if doesn't have partner	1711	1683	1720
Share: has partner	0.59	0.53	0.64
P-val: sum = 0	0.174	0.756	0.168
N: observations	161138	122916	38222
N: workers	474	354	317

Notes: Unit of observation is worker-day. Dependent variable is daily earnings (in FCFA). Col. (1) pools both experiments; Col. (2) restricts the sample to the main experiment; Col. (3) restricts the same to the mechanism experiment. Each column shows results from a stacked difference-in-differences specification. 12 working days of baseline data. All regressions include day-by-wave and worker-by-wave fixed effects. Standard errors clustered at the worker level.

Table A6: Treatment effects during announcement period and first paycycle

	Main experiment		Mechanism expe.	Pooled
	(1)	(2)	(3)	(4)
Private X Announcement Period	23.67	41.79	-3.458	16.34
	(66.45)	(69.45)	(72.40)	(52.58)
	[0.722]	[0.548]	[0.962]	[0.756]
Private X First Treatment Paycycle	151.4	151.4	19.88	95.89
	(69.69)	(69.69)	(94.71)	(57.90)
	[0.031]	[0.031]	[0.834]	[0.098]
Private X Remaining Paycycles	179.4	179.2	169.9	172.7
	(79.96)	(80.02)	(100.9)	(63.03)
	[0.025]	[0.026]	[0.093]	[0.006]
Announcement Period	First week	Full	Full	Full
Sample mean in control	1721	1721	1731	1725
N: worker-days	122145	122916	38222	161138
N: workers	354	354	317	474

Notes: Unit of observation is worker-day. Dependent variable is daily earnings (in FCFA). For each observation, we use as covariate the most recent baseline value. Each column shows results from a stacked difference-in-differences specification. Relative to the baseline period of 12 working days before the treatment announcement, the table reports the effects of being offered a Private blocked savings accounts separately during the announcement period, during the first treatment paycycle, and during the remaining treatment paycycles. Since workers might have imperfect knowledge of the specific start of the first treatment paycycle, Col.1 only consider the first week of the announcement period. Cols.2-4 use the full announcement period. All regressions include day-by-wave and worker-by-wave fixed effects. Standard errors clustered at the worker level.

Table A7: Paycycle Effects

	Attendance (pp)		Earnings (FCFA)	
	(1)	(2)	(3)	(4)
Days since paycycle starts	0.379	0.304	18.07	13.83
	(0.0577)	(0.0559)	(1.558)	(1.380)
	[0.000]	[0.000]	[0.000]	[0.000]
Controls	No	Yes	No	Yes
Sample mean at paycycle start	69	69	1545	1545
N: observations	70553	70553	70553	70553
N: workers	461	461	461	461

Notes: Unit of observation is worker-day. Dependent variable is attendance (in percentage points (pp): 100 if present, 0 if absent) in Cols. (1)-(2) and daily earnings (in FCFA) in Cols. (3)-(4). For example, Col. (1) indicates that for each day closer to the end of the paycycle, attendance increases by 0.379pp. Sample of the 8 full paycycle closest to the announcement date, for each randomization event. Each specification controls for randomization wave fixed effects. In addition, the even columns control for worker-by-wave, day of week-by-wave and paycycle-by-wave fixed effects. Standard errors clustered at the worker level.

Table A8: Heterogeneous Treatment Effects: Self-Control

Low self-control measured by:	Paycycle effects in attendance			in earnings
	Continuous	Positive	Above Median	Continuous
	(1)	(2)	(3)	(4)
Private account	150.6	153.1	139.8	131.3
	(55.59)	(88.20)	(72.79)	(56.96)
	[0.007]	[0.083]	[0.055]	[0.022]
Private account X Low self-control	13.02	11.47	40.98	1.090
	(30.84)	(113.4)	(110.4)	(1.239)
	[0.673]	[0.919]	[0.711]	[0.379]
N: observations	161138	161138	161138	161138
N: workers	474	474	474	474

Notes: Unit of observation is worker-day. Dependent variable is daily earnings (in FCFA). We measure the baseline self-control in labor supply of each worker by estimating its paycycle effects in labor supply. For each worker, and within each randomization event, we regress its attendance in pp and earnings in FCFA (over the 8 full paycycles closest to the treatment announcement date) on the number of days since the start of the fortnight, including worker, day-of-week and fortnight fixed effects. The aggregate paycycle effects are reported in Table A7. We report the heterogeneity in treatment effects across both our experiments, by the continuous paycycle effect in attendance (Col. 1), an indicator for having a positive paycycle effect in attendance, which reflects having potential self-control issues (Col. 2), an indicator for having above median paycycle effect in attendance (Col. 3), and the continuous paycycle effect in earnings (Col. 4). Each column shows results from a stacked difference-in-differences specification. 12 working days of baseline data. All regressions include day-by-wave and worker-by-wave fixed effects. Standard errors clustered at the worker level.

Table A9: Effort elasticity estimates from experimental piece-rate variation

	log(output)				
	(1)	(2)	(3)	(4)	(5)
log(Piece-rate)	0.166	0.175	0.168	0.159	0.246
	(0.0703)	(0.0704)	(0.0708)	(0.0707)	(0.115)
	[0.019]	[0.013]	[0.018]	[0.025]	[0.034]
Day FE	Yes	No	No	No	Yes
Linear time trend	No	No	Yes	Yes	No
Quadratic time trend	No	No	No	Yes	No
Lowest rate excl.	No	No	No	No	Yes
N: worker-days	1528	1528	1528	1528	1164
N: workers	303	303	303	303	301

Notes: Unit of observation is worker-day. Workers learned their piece rate for the day after arriving to work. Dependent variable is log output. Cols. (1) and (5) controls for day fixed effects, Col. (2) has no time controls, Cols. (3)-(4) control for a linear and quadratic time-trend, respectively. Col. (5) excludes the worker-days with the lowest piece-rate (lower than usual) drawn. Standard errors clustered by worker.

A.3 Model Appendix

A.3.1 Baseline labor supply decision

As described in Section 4, at baseline, a worker solves $\max_{c,e} u(c,e)$ under the budget constraint $BC1 : (1-\tau_0)we = c$. τ_0 denotes the linear marginal tax rate faced on income, as show in Figure 3A. Her optimal labor supply decision is thus $e_0\left((1-\tau_0)w\right)$. Denote the baseline choice made by the worker as e_0.

A.3.2 Labor supply decision under new tax schedule

Suppose we introduce an alternate tax schedule, as shown in Figure 3B. By dampening the social tax rate from τ_0 to τ_1 on earnings increases only (i.e., for all $e \geq e_0$), a kink is introduced in her budget constraint. Specifically, it now is:

$$BC2 : c = \mathbb{1}_{e \leq e_0}\left\{(1-\tau_0)we\right\} + \mathbb{1}_{e > e_0}\left\{(1-\tau_0)we_0 + (1-\tau_1)w(e-e_0)\right\}.$$

Note the trivial result: $e_1 \geq e_0$, where e_1 is the worker's choice under $BC2$.[39] In what follows, we thus use the budget constraint $BC2^* : (1-\tau_0)we_0 + (1-\tau_1)w(e-e_0) = c$.

We introduce $\mathbb{Y}$ as the lump-sum income shift due to the tax schedule kink, defined as $\mathbb{Y} \equiv (\tau_1 - \tau_0)we_0$. Since the choice variable e does not enter $\mathbb{Y}$, we can re-write the worker's labor supply decision under Treatment as $\max_{c,e} u(c,e)$ under $BC2^* : (1-\tau_1)we + \mathbb{Y} = c$. Her optimal decision is thus $e_1\left((1-\tau_1)w, \mathbb{Y}\right)$. Note the similitude in form with the baseline labor supply decision.

A.3.3 Slutsky equation.

To study how the level of effort e_1 responds to a change in the tax rate τ_1 above the kink, it is useful to derive the Slutsky equation applied to our model. This allows us to separate the effort response to the change in the tax schedule into a substitution and an income effect.

We define $\tilde{e}\left((1-\tau_1)w, u\right)$ as the Hicksian (compensated) labor supply and $\gamma\left((1-\tau_1)w, u\right)$ as the expenditure function. By duality of utility maximization and expenditure minimization, we have:

$$\tilde{e}\left((1-\tau_1)w, u\right) = e_1((1-\tau_1)w, \gamma\left((1-\tau_1)w, u\right))$$

Taking the derivative on both sides with respect to $(1-\tau_1)w$, we have:

$$\frac{\partial \tilde{e}\left((1-\tau_1)w, u\right)}{\partial\left[(1-\tau_1)w\right]} = \frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial\left[(1-\tau_1)w\right]} + \frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial \mathbb{Y}} \frac{\partial \gamma\left((1-\tau_1)w, u\right)}{\partial\left[(1-\tau_1)w\right]} \tag{A.1}$$

[39] Proof by contradiction: Suppose that $e_1 < e_0$. Then $BC2$ becomes $(1-\tau_0)we = c$, which is $BC1$. Since e_0 is the optimal choice under $BC1$, we must have $e_1 = e_0$. This contradicts $e_1 < e_0$.

Note that by Shephard's Lemma,

$$\frac{\partial \gamma\left((1-\tau_1)w, u\right)}{\partial\left[(1-\tau_1)w\right]} = -\tilde{e}\left((1-\tau_1)w, u\right)$$

(Note the minus sign. It comes from the budget constraint being $\mathbb{Y} = c - (1-\tau_1)we$ and thus from labor supply, e, being a "bad" instead of a "good".)

Substituting this to the previous result:

$$\frac{\partial \tilde{e}\left((1-\tau_1)w, u\right)}{\partial\left[(1-\tau_1)w\right]} = \frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial\left[(1-\tau_1)w\right]} - \frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial \mathbb{Y}} e_1((1-\tau_1)w, \mathbb{Y}) \tag{A.2}$$

This is the Slutsky equation applied to our model, which relates the compensated and uncompensated labor supply responses.

A.3.4 Income and substitution effects.

If $\tau_1 = \tau_0$, then $\mathbb{Y} = 0$ and $e_1\left((1-\tau_1)w, \mathbb{Y}\right) = e_0\left((1-\tau_0)w\right) = e_0$.

Starting from this baseline situation, what happens when we dampen the tax rate τ_1 applied above the kink e_0?

Using the chain rule, applied to $e_1\left((1-\tau_1)w, \mathbb{Y}\right)$, we obtain:

$$\begin{aligned}\frac{de_1\left((1-\tau_1)w, \mathbb{Y}\right)}{d\tau_1} &= \frac{\partial e_1\left((1-\tau_1)w, \mathbb{Y}\right)}{\partial\left[(1-\tau_1)w\right]}\frac{\partial\left[(1-\tau_1)w\right]}{\partial \tau_1} + \frac{\partial e_1\left((1-\tau_1)w, \mathbb{Y}\right)}{\partial \mathbb{Y}}\frac{\partial \mathbb{Y}}{\partial \tau_1} \\ &= -w\frac{\partial e_1\left((1-\tau_1)w, \mathbb{Y}\right)}{\partial\left[(1-\tau_1)w\right]} + we_0\frac{\partial e_1\left((1-\tau_1)w, \mathbb{Y}\right)}{\partial \mathbb{Y}}\end{aligned}$$

We have derived above the following Slutsky's equation:

$$\frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial[(1-\tau_1)w]} = \frac{\partial \tilde{e}((1-\tau_1)w, u)}{\partial[(1-\tau_1)w]} + \frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial \mathbb{Y}} e_1((1-\tau_1)w, \mathbb{Y})$$

Thus,

$$\frac{de_1((1-\tau_1)w, \mathbb{Y})}{d\tau_1} = -w\frac{\partial \tilde{e}((1-\tau_1)w, u)}{\partial[(1-\tau_1)w]} + w(e_0 - e_1)\frac{\partial e_1((1-\tau_1)w, \mathbb{Y})}{\partial \mathbb{Y}} \tag{A.3}$$

We start from the baseline where $\tau_1 = \tau_0$, and so we start from the situation where $e_0 - e_1 = 0$. It follows that:

$$\frac{de_1\left((1-\tau_1)w, \mathbb{Y}\right)}{d\tau_1} = -w\frac{\partial \tilde{e}\left((1-\tau_1)w, u\right)}{\partial\left[(1-\tau_1)w\right]} \tag{A.4}$$

There is no income effect in the change in effort induced by moving from the baseline situation to the new budget constraint. The change in effort is only determined by a substitution effect.

We can further observe that $-\frac{de_1((1-\tau_1)w,\mathbb{Y})}{d\tau_1} > 0$, 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.

We know that $d[(1-\tau_1)] = -d\tau_1$, therefore:

$$\frac{de_1((1-\tau_1)w,\mathbb{Y})}{d(1-\tau_1)} = w\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial[(1-\tau_1)w]}$$

$$\frac{de_1((1-\tau_1)w,\mathbb{Y})}{d(1-\tau_1)}\frac{1}{e_1((1-\tau_1)w,\mathbb{Y})} = \frac{w}{e_1((1-\tau_1)w,\mathbb{Y})}\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial[(1-\tau_1)w]}$$

By duality of utility maximization and expenditure minimization: $e_1((1-\tau_1)w,\mathbb{Y}) = \tilde{e}((1-\tau_1)w,u)$ for $\mathbb{Y} = \gamma((1-\tau_1)w,u)$, so:

$$\frac{de_1((1-\tau_1)w,\mathbb{Y})}{d(1-\tau_1)}\frac{1}{e_1((1-\tau_1)w,\mathbb{Y})} = \frac{w}{\tilde{e}((1-\tau_1)w,u)}\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial[(1-\tau_1)w]}$$

We can rewrite the last term:

$$\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial[(1-\tau_1)w]} = \frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial(1-\tau_1)}\frac{\partial(1-\tau_1)}{\partial[(1-\tau_1)w]} = \frac{1}{w}\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial(1-\tau_1)}$$

Substituting this into the previous equation, we get:

$$\frac{de_1((1-\tau_1)w,\mathbb{Y})}{d(1-\tau_1)}\frac{1}{e_1((1-\tau_1)w,\mathbb{Y})} = \frac{1}{\tilde{e}((1-\tau_1)w,u)}\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial(1-\tau_1)}$$

Assuming that $d(1-\tau_1) \propto (1-\tau_1)$:

$$\frac{de_1((1-\tau_1)w,\mathbb{Y})}{e_1((1-\tau_1)w,\mathbb{Y})}\frac{1}{d(1-\tau_1)} = \frac{1-\tau_1}{\tilde{e}((1-\tau_1)w,u)}\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial(1-\tau_1)}\frac{1}{1-\tau_1}$$

$$\frac{de_1((1-\tau_1)w,\mathbb{Y})}{e_1((1-\tau_1)w,\mathbb{Y})} = \frac{1-\tau_1}{\tilde{e}((1-\tau_1)w,u)}\frac{\partial\tilde{e}((1-\tau_1)w,u)}{\partial(1-\tau_1)}\frac{d(1-\tau_1)}{1-\tau_1}$$

Hence,

$$\frac{de_1((1-\tau_1)w, \mathbb{Y})}{e_1((1-\tau_1)w, \mathbb{Y})} = \zeta \frac{d(1-\tau_1)}{1-\tau_1} \tag{A.5}$$

with $\zeta = \frac{1-\tau_1}{\tilde{e}((1-\tau_1)w,u)} \frac{\partial \tilde{e}((1-\tau_1)w,u)}{\partial(1-\tau_1)}$ the compensated elasticity of labor supply to the net-of-tax wage rate.

Equation A.5 describes how e_1 varies when τ_1 changes—starting from the situation where τ_1 equal to τ_0, hence e_1 equals e_0. 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:

$$\begin{aligned} &\log\left(\frac{e_1}{e_0}\right) \approx \zeta \log\left(\frac{1-\tau_1}{1-\tau_0}\right) \\ &\Rightarrow \frac{1-\tau_0}{1-\tau_1} \approx \left(\frac{e_0}{e_1}\right)^{\frac{1}{\zeta}} \end{aligned} \tag{A.6}$$

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

A.3.6 Elasticity decomposition.

To obtain an estimate for ζ, the compensated elasticity of labor supply to the net-of-tax wage rate, it can be useful to 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,

$$e = p \times a \tag{A.7}$$

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

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

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

$$\begin{aligned}(A.8) \Leftrightarrow & \frac{de}{e}\frac{(1-\tau)}{d(1-\tau)} = \frac{da}{a}\frac{(1-\tau)}{d(1-\tau)} + \frac{\times dp}{p}\frac{(1-\tau)}{d(1-\tau)} \\ \Leftrightarrow & \zeta = \zeta_a + \zeta_e\end{aligned} \tag{A.9}$$

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 daily at the worker level, over 4 possible values. See Appendix A.4.5 for details. 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 0.166 (Appendix Table A9, Col. 1).[40]

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 labor supply—moved both the extensive and intensive margins. Specifically, we find that the extensive margin of attendance accounts for 83% of the total effect of the Private treatment on output (Table 1, Panel C, Cols. 2 and 3). Assuming that the ratio of effects on the extensive and intensive margins reflects the ratio of elasticities along the extensive and intensive margins implies an attendance elasticity 4.88 times as large as the productivity elasticity.

Our piece-rate variation exercise thus implies an estimate of ζ_a of 0.81. Since $\zeta = \zeta_a + \zeta_e$, the total estimated value of ζ is 0.98.

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.

A similar reasoning can be applied to estimate ζ when we only have an estimate of the elasticity of attendance ζ_a. For instance, Goldberg (2016) estimates an attendance elasticity

[40] 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.

of 0.15 among daily laborers in Malawi—implying a total labor supply elasticity of 0.18.

A.3.7 Determinants of elasticity.

There is a strong parallel between the social tax and formal taxation. The literature on formal taxation is largely based on rich countries, and emphasized differences in behavior (in terms of both labor supply and taxable income elasticities) by income level within those countries. In this appendix, we examine the question of whether we might expect the labor supply elasticity to the income tax rate to differ among higher vs. lower income individuals.

To enable this investigation, we examine a simple case where utility is separable between consumption and effort. Formally, the individual chooses its effort e and consumption c to maximize its utility $u(c) - k(e)$, subject to its budget constraint $we + y = c$ where$w$ is the net-of-tax wage rate (i.e., $w = (1-\tau)\overline{w}$, with τ the marginal income tax rate and $\overline{w}$ the gross wage rate) and y non-labor income.

The choice of effort is characterized by the first-order condition:

$$wu'(we + y) = k'(e)$$

Deriving again with respect to effort e, we obtain:

$$u'(we(w) + y) + wu''(we(w) + y)(e(w) + we'(w)) - k''(e)e'(w) = 0$$

Re-arranging yields:

$$\frac{w}{e}\frac{u'(c) + we(w)u''(c)}{k''(e) - w^2u''(c)} = \frac{\partial e}{\partial w}\frac{w}{e} \tag{A.10}$$

Where the right-hand side is the elasticity of labor supply (effort) to the net-of-tax wage rate.

Recall that $w = (1-\tau)\overline{w}$, so $\frac{\partial e}{\partial w}\frac{w}{e} = \frac{\partial e}{\partial(1-\tau)}\frac{(1-\tau)}{e}$: studying how the labor supply elasticity to w differs for high versus low income individuals is equivalent to studying how the labor supply elasticity to $(1-\tau)$ differs.

We see from Equation (A.10) that:

1. The elasticity of effort is higher for individuals with higher marginal utility from consumption, *ceteris paribus.*

2. The elasticity of effort is higher for individuals with higher curvature of their utility function, *ceteris paribus.*

3. The elasticity of effort is lower for individuals with higher curvature of their cost function, *ceteris paribus.*

If the utility from consumption $u(c)$ is increasing and concave, as is commonly assumed, this implies that individuals with lower consumption levels have *higher* elasticity of effort.

A.4 Protocols Appendix

A.4.1 Timeline

Overview. The intervention was implemented over the course of three years, from 2017 to 2019. The main experiment was conducted in three waves, over two sites. We started by a first site, with a 9-month blocked period for the accounts. Building on this promising trial, we expanded the intervention with a new 9-month wave in the same site – maintaining workers from the first wave in the program, and enrolling newly eligible workers. Three months after the start of this second wave, we enrolled once again newly eligible workers at this site – this time with a 5-month blocked period. Along with the second 9-month wave in the first site, we expanded the intervention to a second site with a 9-month wave. The mechanism experiment was then implemented in both factories, over three months.

Main experiment, Wave 1, Site 1. The main experiment was first launched in one factory, with treatment assignments announced on June 7th, 2017. For workers who were offered the Private blocked savings accounts and chose to enroll, the first day of earnings counting towards the accounts was June 12th, 2017. In practice, 75% of the workers who chose to enroll had their accounts active within a fortnight of that date, and 100% had their accounts active within a month. For this wave, the last day of earnings counting towards the account was March 31st, 2018. This is also the date in which the accounts were unblocked and workers could access their savings.

Main experiment, Wave 2, Site 1. The intervention was expanded with a new 9-month wave in the same site. First, all workers already offered a blocked savings account in wave 1 were offered the possibility to keep being enrolled in wave 2, with accounts being re-blocked for the new wave. It would have been difficult to stop offering them an account they enjoyed, while offering it to other workers at the factory at the same time. Second, all workers who had not been offered the savings accounts in wave 1 but satisfied the eligibility criteria for wave 2 were randomly assigned into Treatment (being offered the blocked savings account) or Control (not being offered the accounts) for wave 2. This includes (i) workers who were already eligible for wave 1 but were randomly assigned to Control (Given the high level of enthusiasm and pressing demand among workers for the accounts, we made the ethically-motivated decision to give wave 1 Control workers the chance to be offered the account in wave 2 – at the expense of some statistical power for our analysis.); (ii) workers who had been flagged by factory management as deserving an account in wave 1 but did not satisfy the eligibility criteria then and were not offered an account; and (iii) workers newly eligible for wave 2. For all these workers, treatment assignments were announced on June 26th, 2018. Their earnings counted towards the savings accounts from July 1st, 2018 to March 31st, 2019.

Main experiment, Wave 3, Site 1. To accommodate the recurrent influx of new workers at the factory and the high turnaround, we launched a last wave of the experiment a few months after the start of wave 2. Specifically, treatment assignments were announced on October 22rd, 2018, with earnings counting towards the savings accounts starting November 1st, 2018. The unblocking date for these accounts was the same as for wave 2: March 31st, 2019. The total blocked period was thus shorter for these accounts, at 5 months. The workers randomly assigned to Treatment or Control under these waves were (i) workers who had been eligible and assigned to Control under wave 1, but were no longer eligible for wave 2 so had not been randomly assigned to Control or Treatment then, and (ii) workers newly eligible for wave 3, who had not been eligible for either wave 1 or 2.

Main experiment, Site 2. The intervention was expanded to a second site, at the same time as the second wave in the first site. There, a single wave was implemented, with savings accounts blocked for 9 months. (Unlike Site 1, not enough workers became eligible for the intervention in the few months after the launch of this wave to justify launching another wave there.) Treatment assignments were announced on June 19th, 2019, with earnings counting towards the savings accounts from July 1st 2018 to March 31, 2019.

Mechanism experiment. The mechanism experiment was implemented in both factories, at the same time. Treatment assignments were announced on March 27th, 2019 at the first site and on March 25th, 2019 at the second site. In both sites, the first and last days of earnings counting towards the savings accounts were April 16th, 2019 and July 15th, 2019.

Additional experimental activities. To avoid any influence on the launch of the mechanism experiment (and the associated take-up decisions), we conduct the two auxiliary experimental activities after the start of this mechanism experiment. The experimental opt-out test was conducted over two consecutive pay cycles: the last fortnight of May, and the first fortnight of June 2019; while the experimental piece-rate variation activity was implemented in the first site over the course of a week, from June 17 to June 22, 2019.

A.4.2 Stacked datasets

The implementation process of the experiment in multiple waves of randomization naturally lends itself to a stacked difference-in-differences empirical specification for the analysis. As described in Cengiz et al. (2019) and Baker et al. (2022), this involves creating a "clean 2 x 2" dataset for each randomization event, and stacking these datasets together before running the estimation. Each "clean 2 x 2" dataset is a balanced panel at the worker-day level, including pre- and post-treatment assignment data for workers assigned to either Treatment or Control in this randomization event. Importantly, no workers should be treated at baseline: the dataset is "clean" in the sense that all workers start without having been offered a blocked savings accounts, and then some workers are randomly assigned to be offered the accounts while the others remain without.

We consider 8 randomization events when analyzing our main experiment (hence 8 "clean 2 x 2" datasets), which directly reflect the implementation process described above. These randomization events have been defined to maximize the statistical power of the analysis by efficiently using all available data. However, the results are robust to considering alternative definitions of these randomization events, as long as we maintain a set of clean 2 x 2 datasets to satisfy the requirements from the stacked difference-in-differences specification.

We start from Wave 1 in Site 1. As described above, all workers assigned to Treatment in wave 1 remained assigned to Treatment throughout wave 1 and wave 2 (with accounts unblocked for a short period between the two waves), while workers assigned to Control in wave 1 were subsequently re-randomly assigned to either Control or Treatment (with that re-randomization happening in either wave 2 or wave 3, depending on the workers).

1. Our first clean 2 x 2 dataset includes the workers assigned to Treatment in wave 1 as well as the workers assigned to Control in wave 1 who will be assigned to Control in waves 2/3 too. We include data from the baseline of wave 1 until the end of wave 2 (while excluding the data in the unblocking period between the end of wave 1 and the start of wave 2).

2. Our second clean 2 x 2 dataset includes the workers assigned to Treatment in wave 1 as well as the workers assigned to Control in wave 1 who will be assigned to Treatment in waves 2/3. To keep it a clean dataset in which treatment assignment remains constant throughout, we only include data from the baseline of wave 1 until the end of wave 1.[41]

We then consider the randomization events from Wave 2 in Site 1:

3. Our third clean 2 x 2 dataset includes workers who were eligible in wave 1, had been assigned to Control in wave 1, were eligible for wave 2, and were randomized to either Treatment or Control in wave 2. It includes data from the baseline of wave 2 until the end of wave 2.

4. Our fourth clean 2 x 2 dataset includes workers who were not eligible in wave 1, had been flagged by the factory management in wave 1 as workers who should be offered the account, were not offered an account in wave 1, were eligible for wave 2, and thus were randomized to either Treatment or Control in wave 2. It includes data from the baseline of wave 2 until the end of wave 2.

[41] We do include the same workers (those assigned to Treatment in wave 1) in multiple datasets, but account for this in the inference by clustering standard errors at the worker level.

5. Our fifth clean 2 x 2 dataset includes workers who were not eligible in wave 1, became eligible in wave 2, and were randomized to either Treatment or Control in wave 2. It includes data from the baseline of wave 2 until the end of wave 2.

We similarly consider the randomization events from Wave 3 in site 1:

6. Our sixth clean 2 x 2 dataset includes workers who were eligible in wave 1, had been assigned to Control in wave 1, but were not eligible in wave 2. They were randomized to either Treatment or Control in wave 3. It includes data from the baseline of wave 3 until the end of wave 3.

7. Our seventh clean 2 x 2 dataset includes workers who were not eligible in wave 1 or wave 2, but became eligible in wave 3. They were randomized to either Treatment or Control in wave 3. It includes data from the baseline of wave 3 until the end of wave 3.

Finally, we consider the randomization in site 2:

8. Our eight clean 2 x 2 dataset includes workers who were randomized into Control or Treatment in site 2. It includes data from the baseline of this wave until the end of this wave.

When pooling the two experiments to estimate the effects of being offered a Private blocked savings account relative to being offered either a Non-private blocked savings account or nothing, we consider the mechanism experiment as one additional randomization event.

A.4.3 Offer and Implementation of Private blocked accounts (Both experiments)

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 experiment using a lottery, where ID numbers are drawn by the research team to assign treatment status. In the main experiment, there is one lottery drawing for each intervention wave. In the mechanism experiment, 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.[42]

3. The account will be blocked for a period of 9 months in the main experiment[43], and 3 months in the mechanism experiment, after which the worker may withdraw all of her money free of charge. During 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.[44]

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.

[42] In exchange for waiving the fees, the BPCI did not pay out interest on the savings.

[43] 5 months for the third wave in the first site.

[44] In practice, this happened for one worker over the course of the study.

5. Workers are allowed to revise their threshold during an initial probation period of two months.[45]

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 your 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.[46]

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 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 subcon-

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

[46]Certificates 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.

tractor, instead of an individual account.[47] 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.[48] 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.

A.4.4 Non-private accounts (Mechanism Experiment)

The implementation protocol for the mechanism experiment 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 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 the main experiment — 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 did not 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

[47] The 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.

[48] Our 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.

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 the main experiment), 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 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.

Note that in practice, we had obtained the contact information of workers’ network members through the baseline survey for the mechanism experiment. During that baseline survey, we asked workers to “share with us the name and phone number of five people you know" — we remained general, and prompted workers to think about their friends, family, and acquaintances. We told them that those individuals could be contacted by us and offered opportunities. Consequently, the workers could have named anyone and we did not confirm those numbers at the time of the baseline survey. Consequently, Thus, there was the potential for workers to name contacts who would have had no impact on the real social tax rate of these workers. However, given workers' perception of how information spreads in networks, even if the person named was not a source of the social tax, they would have the ability to pass on information to others in the network that was a source of transfer requests. This helped ensure power for our mechanism experiment design.

A.4.5 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.[49]

A.5 Data Appendix

Human Resources records. Our primary data source is Olam's detailed administrative data on each worker's daily attendance, output, and earnings. Output is measured by weighting at the end of each day the quantity of nuts processed by each worker during the day. Workers are fully paid piece-rate, so there is a direct correspondence between their output and earnings. These administrative data are used by Olam to compute workers' payments, and can thus be deemed as high quality.

Surveys. We supplement the administrative records with data collected through two sets of surveys. First, we conducted phone surveys prior to the launch of each wave of the main experiment, and prior to the launch of the mechanism experiment, to obtain baseline data for these experiments. These surveys include information on perceptions about redistributive pressure, and coarse information about financial transfers with network members.

Specifically, for the main experiment, we conducted a phone survey in April 2018 among a random subset of the workers eligible for the intervention at that time. This survey serves as baseline for waves 2 and 3 in Site 1.[50] Due to operational reasons, we unfortunately were not able to collect survey data ahead of the launch of Wave 1 in Site 1 for the main experiment, and therefore lack survey baseline data for that wave. In addition, we conducted a phone survey in June 2018 serving as baseline survey for Site 2.

For the mechanism experiment, we conducted phone surveys in December 2018 and March 2019, serving as baseline for both sites. For workers surveyed twice, we use the latest

[49] While 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 paycycle 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 paycycle. A few weeks later, workers received the remainder of the due amount: $\sum_{d=1}^{6} [\text{output}_d \times (\text{experimental rate}_d - \text{usual rate})]$.

[50] Except for the workers who became newly eligible for treatment for wave 3 only, as they were not included in the April 2018 sampling frame.

answer available as their baseline measure.

Second, we conducted a more detailed in-person endline survey in August 2019, shortly after the end of the mechanism experiment, which includes details about financial transfers with network members.