The Social Tax: Redistributive Pressure and Labor Supply

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

September 30, 2021

Preliminary draft

Please click here for the most recent version

Abstract

In developing countries, redistributive transfers within kin and social networks are frequent. We test whether transfer requests dampen the incentive to work among full-time piece-rate factory workers in Côte d'Ivoire. We offer workers blocked savings accounts, into which they can deposit increased earnings over 3-9 months. Accounts vary in whether they are private or known to their networks, altering the likelihood of transfer requests. This changes the effective "social tax" on earning *increases*, allowing to isolate substitution effects on effort. When accounts are private, take-up is substantively higher: 60% vs. 14%. Furthermore, workers increase labor supply and effort: 9.2% higher attendance and 14.5% higher output and earnings. This implies a 18% social tax rate on earned income. Because our design leaves liquid earnings unchanged, we find no decline in outgoing transfers, indicating no loss in redistribution. Instead, our findings suggest that potential welfare benefits of redistributive arrangements may come at an efficiency cost.

JEL: J22, J24, H24, D61, D64, O12

---

*Carranza: World Bank (ecarranza@worldbank.org); Donald: World Bank (adonald@worldbank.org); Grosset: Columbia University (fg2429@columbia.edu); Kaur: University of California at Berkeley and NBER (supreet@berkeley.edu). This paper greatly benefited from comments by Michael Best, François Gerard, Jessica Goldberg, Sylvie Lambert, Guilherme Lichand, Karen Macours, Owen Ozier, Golvine de Rochambeau, Simone Schaner, Krzysztof Zaremba, and various seminar participants. Julia Buzan, Oumar Koné, Tiphaine Forzy, Chris Tullis, Ambika Sharma, Prathyush Parasuraman, Shelby Carvalho and Cécile Delcuvellerie provided superb research assistance. We gratefully acknowledge financial support from the World Bank's Umbrella Facility for Gender Equality, the World Bank's Jobs Umbrella Multidonor Trust Fund and the National Science Foundation (Kaur's CAREER award SES 1848452). We also thank Innovations for Policy Action (IPA), especially Nicolò Tomaselli, Henriette Hanicotte, Samuel Kembou Nzalé, Mireille Nuguhe Gbagbo and Augustin Kouadio for assistance in the implementation of this project. This paper is a product of the World Bank Africa Gender Innovation Lab and the World Bank Jobs Group. Research approved by IPA IRB. AEA RCT Registry ID: AEARCTR-0003821. All errors and omissions are our own.

1 Introduction

In developing countries, informal financial transfers within social and kin networks are ubiquitous and frequent (Banerjee and Duflo, 2007; Fafchamps, 2011). For example, in the setting of this study, Côte d'Ivoire, full-time factory workers report transferring 26% of their income to others outside their household on average, and 77% made at least one transfer in the past 3 months. The presence of frequent transfers has traditionally been understood as reflecting informal risk sharing, improving welfare by substituting for missing insurance markets.[1]

However, a broad body of work in social science—spanning economics, anthropology, and sociology—has discussed the possibility that informal redistributive arrangements could have distortionary effects (O'Brien, 2012; Portes, 1998). These literatures provide qualitative accounts that individuals face social pressure to share income gains, even if they arise from exerting costly effort. Such a possibility departs from efficient risk sharing models, in which transfers should only be triggered by idiosyncratic shocks, not by higher effort. This departure may arise if it is difficult for others to distinguish effort from shocks, or simply if cultural norms dictate that richer individuals should redistribute to poorer ones.[2] If the gains from effort increases are indeed redistributed to others, then this could distort labor supply—potentially lowering output and income (e.g., Alger and Weibull, 2008; Hoff and Sen, 2011). This would suggest that redistributive arrangements, while potentially welfare improving, may come at an important efficiency cost.

In this paper, we empirically examine this possibility. We conceptualize redistributive arrangements as a "social tax" on earnings. We focus on the domain of labor supply—the primary means through which the poor generate income. The main goal of our paper is to test whether redistributive arrangements distort worker labor supply, output, and earnings.

We work with full-time female factory workers in cashew processing plants in Côte d'Ivoire. The 473 workers in our sample are employed by OLAM, a large transnational agroprocessing firm, with an average tenure at the firm of 1.5 years. Workers are paid their wages twice a month in cash. 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 workers' effort and their income.

Workers in our setting report facing heavy demands to share income with extended family and others in their social network (Figure 2). For example, 77% state that income gains from higher effort would lead to more transfer requests. In addition, 73% state they have

[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), Rosenzweig (1988), among many others.

[2] For example, Platteau (2000) provides descriptive evidence on strong egalitarian norms for redistribution in Africa. Such norms may reflect how risk-sharing motives have become implemented in societies, or may simply constitute an additional (and non-mutually exclusive) reason for why transfers occur. Note that this does not rule out the possibility that some transfers are driven by altruism.

difficulty saving for large goals because if they "put money aside, someone else will ask for it". Workers perceive that it is socially costly to turn down requests for cash—unless they can credibly convey that they do not have the money available.[3] Consequently, workers engage in informal strategies to convert earnings to illiquid form—for example, by buying goods immediately after payday (e.g. Miracle et al., 1980; Goldberg, 2017). This suggests that methods that credibly lock away funds could make it easier to avoid transfer requests, leading workers to retain more of their earned income.

To construct our empirical test, we introduce a tool to lower redistributive pressure on income gains: a blocked savings account into which workers can transfer earnings *increases*. Workers who opt in choose a threshold, which must be higher than their baseline earnings. In each biweekly paycycle, any amount they earn *above* this threshold is automatically deposited into the blocked account; the remainder of their earnings is paid in cash as usual on payday.[4] The funds in the account cannot be accessed until the end of the blocked period (3-9 months). We develop this product in partnership with the largest savings bank in Côte d'Ivoire, Banque Populaire (BPCI). 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 approach offers two important benefits. 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 does not induce income effects, so that any labor supply response isolates substitution effects. Consequently, if there is a positive social tax, our intervention should unambiguously increase labor supply—generating a clean test for our hypothesis. Second, under our design, cash-on-hand is unchanged: expected take-home cash pay is not lower by construction. Consequently, workers should have similar levels of disposable income to redistribute—making it unlikely that our intervention makes others in the network worse off through a reduction in transfers; we validate this prediction using endline data. As we discuss in Section 2, in the presence of volatility, these predictions may not hold exactly. However, we view our design as largely mitigating two key concerns that would arise from simply enabling workers to move existing earnings into blocked accounts.

To test for the role of redistributive pressure, we vary whether the existence of the blocked accounts is private or would become known to the worker's network—changing the likelihood of transfer requests against money in the account. We randomize workers into three treatment conditions over the course of the experiment: 1) Control (no account); 2) Private (private blocked savings account); and 3) Non-private (blocked savings account whose existence is revealed to the network). In the private account condition, no one except the

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

[4]Workers choose one threshold that applies to all future paycycles. They can revise this threshold up to three times, and can opt out of a threshold (i.e. any future deposits) at any point.

worker knows of the account's existence. In the non-private condition, workers are told that if they take up the account and save in it, then members of their social network may receive two publicity text messages that advertise the blocked account product and mention that the worker had successfully saved in their account—as a promotional tool for the bank.[5]

If redistributive pressure exists, then workers in the Private treatment group would expect to receive and fulfill fewer undesirable transfer requests on income gains than workers in the Non-private and Control groups. Consequently, we predict that workers in the Private group will increase their labor supply—and therefore their total earnings—relative to those in the other treatment arms. Note that if there is no "social tax", then this prediction should not hold. For example, if all transfers reflect either altruistic giving (i.e. workers internalize the full utility gain from someone else's consumption of the transfer), or efficient risk sharing (i.e. transfers only respond to idiosyncratic shocks), then we would not expect differences between the Private and Non-private conditions. In addition, we use additional design features to further rule out potential confounds below, such as privacy concerns or self-control problems.

We conduct the experiment in two phases. In Phase 1, a subset of workers is randomized into either the Private or Control conditions. This enables us to compare the overall impact of the Private condition against having no account at all. Because our population of workers was largely unbanked at baseline, this "preparatory" phase was intended to instill trust in both the privacy of the Private condition and in the security of the accounts—e.g., workers saw that deposits occurred as expected, and funds could be withdrawn after the unblock date at the end of the Phase 1 treatment period.[6]

In Phase 2, we conduct our key test by randomizing workers into either the Private or Non-private account treatment conditions. We both add new workers to the sample for this phase, and also cross-randomize workers who participated in Phase 1. We conduct tests to verify that treatment effects are not sensitive to Phase 1 treatment status. We also leverage the cross-randomization to help test for confounds below.

In Phase 1, 43% of the Private group workers take up the blocked savings account. Their total earnings go up by 12% (p-val 0.022) relative to the Control group. This is driven in part by an attendance increase of 6.9 percentage points (10.3%)—suggesting increased labor supply on both the extensive and intensive margins. This indicates large output and

[5] All workers provided baseline details about network members. In the Non-private treatment arm, if the worker declined to take up the account or did not save in it, no information would be shared with network members. This was explained as a way for the bank—which was at the time actively trying to increase take-up of its blocked account product and was already advertising in the area—to advertise to community members as part of a publicity drive. Firms commonly use text messages and personal referral programs to advertise products in this setting.

[6] For example, in one of the factory plants, a former employee from an insurance company had collected deposits from workers and then disappeared. Once Phase 1 was over, we were able to credibly announce to workers that Private accounts had been offered and implemented successfully, providing reputational benefits for our intervention before launching our key test in Phase 2.

earnings benefits from simply offering workers the blocked accounts.

In Phase 2, we test more directly for our hypothesized mechanism. Under the Private condition, take-up of the blocked accounts is 60%. In contrast, under the Non-private condition, account take-up is only 14%. This indicates that blocked accounts are substantially less desirable as a savings vehicle if others network members would learn of their existence.

Consequently, the private accounts lead to substantially larger income gains. Compared to the Non-private condition, Private condition workers increase their output and earnings by a striking 14.5% (p-val < 0.001). As before, this is driven partly by an increase in attendance: Private workers are 5.5 percentage points (9.2%) more likely to show up to work (p-val 0.028), accounting for about two-thirds of the total treatment effect on output.[7] These magnitudes are similar if we restrict analysis to workers who did not receive account offers in Phase 1.

Overall, these results indicate that reducing the likelihood of transfer requests on savings has a marked impact on workers' willingness to supply labor, and consequently their total earnings. The magnitude of the average treatment effect from offering Private blocked accounts (instead of Non-private ones) is sizable—equivalent to how much earnings would rise if each worker worked an additional 1.04 days in every 2-week paycycle.[8]

While the Private accounts enable workers to accumulate savings, we find no evidence of reduced transfers. As expected, *cash* take-home pay levels in the Private and Non-private conditions are statistically indistinguishable—in line with the design of the blocked accounts. Consistent with the idea that the accounts did not alter redistributive pressure on the regular cash component of earnings, we find no discernible decline in the total amount of transfers from workers to other households. This indicates that our intervention increased income for the Private account holders without making others in their network worse off—suggesting a potential Pareto improvement.

A potential concern with the interpretation of our results is that the Non-private treatment may operate through some channel other than redistributive pressure. For example, workers may have a hedonic preference for privacy, or be embarrassed that they need a blocked account to save. More generally, the mechanics through which the Non-private treatment was implemented—text messages sent to others mentioning a savings account held by the worker—could introduce other potential mechanisms.

To address such concerns, we construct a placebo test with workers who took up Private ac-

[7]Using an endline survey, we document that these earnings gains are not simply reflecting substitution with other income-generating activities; rather they reflect an increase in total income. This is consistent with the fact that, at baseline, only 89% of workers have no income outside the factory.

[8]The average treatment effect on daily earnings is 228 FCFA, corresponding to a 2,736 FCFA increase per paycycle (comprised of 12 workdays). Mean daily earnings among the Non-private group is 2,635. This gives $2736/2635 = 1.04$ workdays per paycycle.

counts in *Phase 1* but were offered Non-private accounts in *Phase 2*. Three months after the start of Phase 2, we asked these workers for permission to send promotional text messages to their network members advertising that they had saved in a blocked account through the bank in the past (i.e. in Phase 1), in exchange for a token compensation.[9] This therefore mimics the Non-private treatment in its mechanics—but for accounts where money would already be spent long ago—providing little scope for triggering transfer requests. 85% of workers accepted this offer. This indicates that simply having one's network discover that one has used a blocked account to save is not the reason take-up plummets in the Non-private treatment. This further supports our interpretation of the role of redistributive pressure.[10]

Blocked accounts can also boost effort if workers have self-control problems in consumption. However, under present-focus alone, effects should be similar between the Private and Non-private blocked account treatments. Redistributive pressure is therefore necessary to explain our results—so that self-control is not a confound per se. Conditional on the Private accounts reducing transfer requests, the existence of self-control problems could still help contribute to our observed treatment effects. To gauge the potential relevance of this, we surprise a random subset of workers with the option to forego depositing extra earnings in one upcoming paycheck cycle—varying whether this offer is made 4 days before the worker's payday versus on the day of the payday. Counter to the predictions of basic present-focus models, workers are not relatively more likely to want to unblock savings on the day of the payday itself.

To estimate the effective social tax rate that our results imply, we partner with the factory to randomly vary workers' piece rate wages within the context of their normal work for 6 days at the end of the experiment. Combining the resultant elasticity from this piece rate exercise with the 14.5% output distortion from our reduced form estimates, we can construct an estimate of the social tax rate in our context. Our most conservative estimate of the social tax rate faced by the average worker is 18%—consistent with prior work that pressures to redistribute are large.

Our study contributes to a growing literature on redistributive pressure and its impacts on economic behavior. A long tradition of qualitative work documents strong social pressure for individuals to share their income with others (Scott, 1976; Kennedy, 1988; Platteau, 2000, 2014). Numerous studies argue that such pressure can rationalize savings and consumption behaviors in observational data, such as the propensity to hold illiquid savings (Baland et al., 2011; Di Falco and Bulte, 2011; Dillon et al., 2020; Grimm et al., 2013; Alby et al., 2020; Baland et al., 2016). In addition, a robust body of lab-in-the-field experiments—pioneered by Jakiela and Ozier (2016)—offer cash windfalls to individuals and show that they will take costly actions to hide this income from their network (Beekman et al., 2015; Goldberg, 2017;

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

[10] Fear of theft also cannot explain low take-up of Non-private accounts. Workers walk home from the factory with their entire cash earnings in their pockets twice each month, on days that are publicly known. The release and withdrawal of savings from the bank is considerably more private than this.

Squires, 2018; Fiala, 2018; Boltz et al., 2019). For example, Squires (2018) uses estimates from such a windfall exercise to structurally estimate sizable productivity implications for Kenyan micro-entrepreneurs. Finally, heterogeneity analysis in field studies indicates that the benefits of improved savings technologies are higher for those who report facing more redistributive pressure at baseline (Dupas and Robinson, 2013; Riley, 2020).

We advance this literature by providing direct causal evidence that redistributive pressure distorts field behavior. We focus on an important, natural, high-stakes field setting: labor supply among full-time workers, within the context of existing long-run employment. We offer the first causal evidence that redistributive pressure distorts labor supply decisions—and consequently productivity and earnings. This provides empirical support for the idea that redistributive arrangements may come at an efficiency cost.

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

Our study has potential implications for understanding one set of barriers to productivity growth in developing countries, particularly in Sub-Saharan Africa. If redistributive pressures distort labor supply, they could also hamper other economic activity—from technology adoption to human capital accumulation. This mirrors concerns about informal redistribution expressed in some of the earliest development literature (Lewis, 1955; Tam et al., 1957). In addition, our findings raise the question of whether improving safety nets for the poor may not only benefit recipients, but also affect the productivity of non-recipients.[11] While only speculative, these possibilities point to potential directions for additional research.

The rest of the paper proceeds as follows. Section 2 presents a brief conceptual framework that motivates our empirical test. Section 3 describes the context, including background on redistributive pressure. We detail the design in Section 4 and protocols in Section 5. Section 6 outlines the empirical strategy and Section 7 presents the results. Section 8 evaluates potential confounds, and Section 9 estimates the social tax rate implied by our estimates. Section 10 concludes.

2 Conceptual Framework

The goal of our experiment is to lower the social tax rate workers face on earned income – and to study labor supply responses. The key challenge is that the income and substitution

[11]Studies document that formal safety nets may serve as a substitute for informal insurance arrangements (Dupas et al., 2017; Mobarak and Rosenzweig, 2012).

effects of a tax cut push in opposite directions, leading to potentially ambiguous effects on labor supply. Our intervention addresses this challenge by dampening the informal tax rate on earnings *increases* only. This allows for a sharper test of our hypothesis that redistributive pressures dampen labor supply.

We use a simple theoretical framework of labor supply decisions under taxation to ground those key features of our experimental design. Take a worker deriving utility from her consumption c (positively) and from the effort she puts at work e (negatively). We represent her preferences with a standard utility function $u(c, e)$, where $u_c(c, e) > 0$, $u_e(c, e) < 0$, $u_{cc}(c, e) < 0$ and $u_{ee}(c, e) < 0$.

Output, over which the worker earns a piece-rate w, is a deterministic function of effort. She then purchases her private consumption c by using her labor earnings ew and non-labor income y. Absent taxation, she chooses her preferred level of effort ${e_0}^*(w, y)$ such that the marginal product from her effort equals the marginal cost from her effort: $w \times u_c\left({c_0}^*, {e_0}^*\right) = -u_e\left({c_0}^*, {e_0}^*\right)$.

We conceptualize redistributive pressures as a tax on earnings, applied at a marginal rate τ.[12] The net amount of labor income available to the worker is therefore $(1 - \tau)we$. This pivots down her budget constraint, as shown in Figure 1a. The trade-off between exerting costly effort and consuming is altered, distorting the private incentive to work. Under a baseline social tax level of τ_1, the preferred level of effort is now ${e_1}^*((1 - \tau_1)w, y)$.

Figure 1: Tax Rate

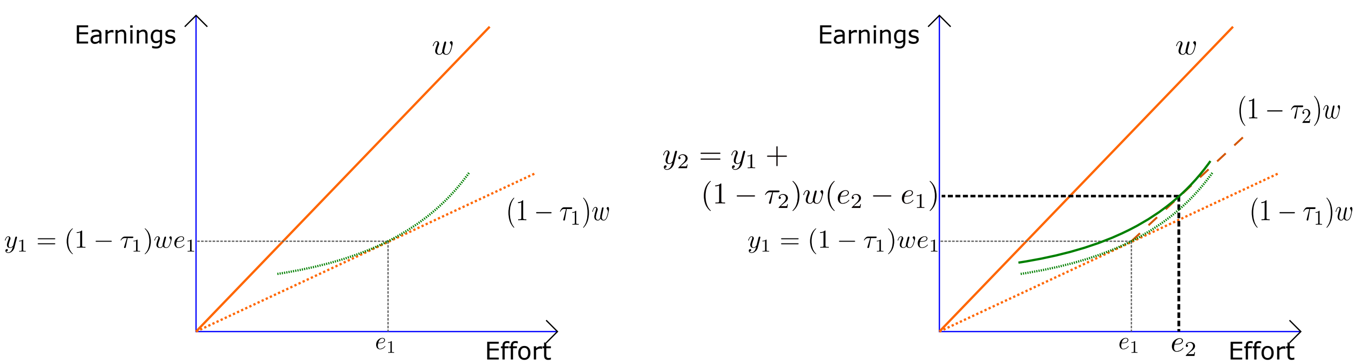

(a) Baseline choice of effort   **(b)** Effect of the intervention

Notes: Panel A shows the choice of effort in a piece-rate setting with a social tax at rate τ_1. Panel B shows the change in optimal effort when the social tax rate is dampened on earnings *increases* above the baseline level (from the rate τ_1 to the rate τ_2).

[12]The tax rate need not be constant across all earnings levels for our qualitative predictions to hold.

A simple test for our hypothesis that redistributive pressures dampen labor supply would be to reduce the social tax rate by shielding all labor earnings from redistributive pressures, and to observe the response of labor supply. Nonetheless, this simple approach bears with it two difficulties.

First, as is the case with all taxes, the net effect of reducing the social tax is ambiguous: effort could increase or decrease (Hausman, 1985). On the one hand, the payoff to work is now higher, increasing the incentive to supply labor (substitution effect). On the other hand, the gain of earnings available for consumption decreases the incentives to work more (income effect). Estimating the overall effect only would make it harder to identify the economic magnitude of the distortion, and to estimate the prevailing social tax rate. Second, enabling workers to shield their earnings from redistributive pressures could have the potentially undesirable effect of reducing existing transfers to kin. [13]

One way to overcome those concerns is to lower redistributive pressures on earnings *increases* only. Specifically, consider reducing the Social tax rate to $\tau_2 < \tau_1$, for all $e > {e_1}^*$, while keeping the tax rate at τ_1 for $e \leq {e_1}^*$ (Figure 1b).

This has three appealing features.

First, it produces a sharper test for our hypothesis that redistributive pressure dampens effort: a lower tax rate for additional effort, $e - {e_1}^*$, will produce an effect on labor supply only if ${e_1}^* < {e_0}^*$ (i.e if labor supply had previously been below the first best). In this case, labor supply and output will increase under the lower tax rate: ${e_2}^* > {e_1}^*$, where ${e_2}^*$ is the optimal labor supply in Figure 1b.

Second, if the level of effort is unchanged (at ${e_1}^*$), the worker's disposable earnings are also unchanged (at $\tau_1 w {e_1}^*$). As such, as demonstrated in Appendix A.1, our test does not induce an income effect, and identifies a pure substitution effect.[14] As demonstrated in section 9, this will allow us to back-out the prevailing social tax rate faced by workers.

Third, our intervention has the potential to be Pareto-improving. By construction, it cannot decrease the worker's average earnings. A revealed preference argument further indicates than the worker expects being weakly better off with the private illiquid savings account (and associated potential increased effort and earnings). A welfare loss may be experienced, though, if the worker experiences a sudden consumption need: she might not be able to adjust her cash-on-hand through increased labor supply and fully cover this temporary expense

[13]For example, suppose that workers are in a risk sharing network where it is possible to renege ex-post on commitments made to their network. No one will want to engage in risk sharing ex-ante if it is not possible to enforce it ex-post, which can lead to an unraveling of risk sharing arrangements.

[14]Because earnings are stochastic, there is some potential for income effects—e.g., in the weeks that the cost of effort is low. This should be minor such that, to a first order, effort changes are largely from substitution effects.

shock. Members of the worker's redistributive network can also be weakly better off, since the specific kink we introduce in the social tax schedule limits the possibilities for workers to reduce their transfers below baseline levels. We verify it by testing for the impact of our intervention on transfers made by the workers.[15]

3 Context: Redistributive Pressures

We work with full-time piece rate workers, employed in cashew processing plants in Cote D'Ivoire. The majority of workers in our sample is female, with an average tenure at the firm of 1.5 years (lower quartile of 1 year, upper quartile of 2 years).

In our study setting, transfers are common and frequent. Workers transfer a significant share of their earnings (26% on average) to others outside of their household.[16] Of this, 66% is redistributed to extended family, and the remainder to non-family members in workers' networks. In addition, the amount redistributed is positively correlated with individual earnings.

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

Workers indicate that increasing income through higher labor supply will lead to more transfer requests. Figure 2 documents survey responses to questions about redistributive pressure. 77% of workers in our sample 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 B). In addition, workers state that redistributive pressure hinders their ability to accumulate savings. 73% state they have difficulty saving over time for large goals because if they put money aside, someone else will ask for it (Panel

[15]From a non-financial perspective, members of the worker's redistributive network might be negatively affected by the intervention if they value the time that the worker spend with them, and if that time is reduced due to her increased labor supply. This need not be the case, especially if workers increase their efforts on the intensive margin rather than on the extensive margin; or if they substitute time away from other income-generating activities that didn't benefit their redistributive network members. We provide treatment effect estimates along those specific margins.

[16]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.

A).

Figure 2: Motivational Evidence: Redistributive Pressure

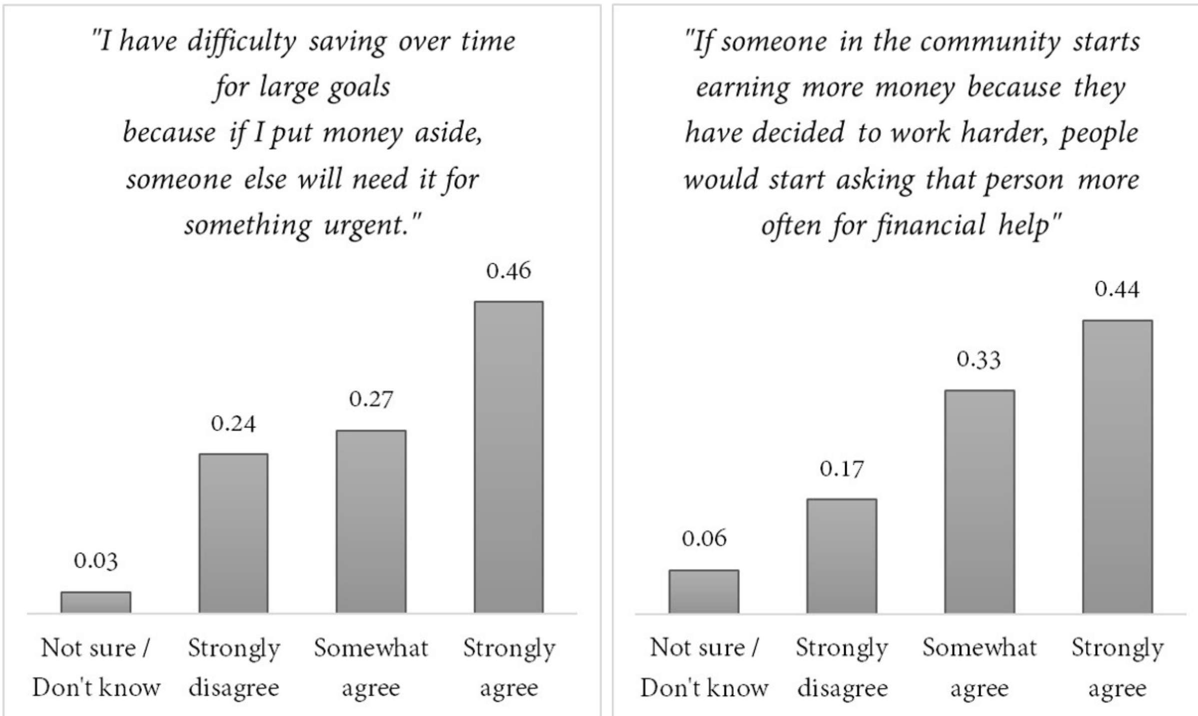

(a) Savings **(b)** Labor income

Notes: N=420 cashew factory workers in Cote D'Ivoire.

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

Consistent with this idea, workers employ a variety of strategies to make their funds inaccessible for redistribution. For example, workers report buying household goods immediately after payday, storing money with others, and participating in ROSCAs. Such strategies were described by workers during qualitative research conducted as part of the preparatory fieldwork for our study (McNeill and Pierotti, 2020), 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., 2020). However, workers perceive the efficacy of such strategies to be limited, as indicated by Figure 2a. In our study, we draw on these existing informal strategies to design a blocked savings account to help workers accumulate long-run savings.

4 Experimental Design

We design an experiment to test whether redistributive pressures distort labor supply—and therefore output and earnings—among full-time factory workers in Cote D'Ivoire. Workers are paid piece rates for output, and receive their wages twice a month in cash. We provide additional details about the factories and production task in Section 5.1 below.

4.1 Blocked savings accounts

To construct our experimental design, we design a tool to lower redistributive pressure on income gains: a blocked savings account into which workers can transfer earnings *increases*. Workers who opt in choose a threshold, which must be higher than their baseline earnings. In each biweekly paycycle, any amount they earn *above* this threshold is automatically deposited by the factory into the blocked account; the remainder of their earnings is paid in cash as usual on payday. We develop this product in partnership with the largest savings bank in Côte d'Ivoire, Banque Populaire (BPCI).

Specifically, using administrative data from the factories, we compute baseline earnings to be the worker's average earnings per paycycle in the past 3 months. Denote this baseline level ${y_{i1}}^*$ for worker i. If a treatment worker takes-up the account, she chooses a threshold level for deposits T_i, which must be at least as large as her baseline earnings: $T_i \geq {y_{i1}}^*$. For every pay period, any earnings *above* T_i are automatically and privately deposited into the savings account. In other words, the deposit amount is $max\{0, y_{it} - T_i\}$, where y_{it} is earnings for worker i in pay period t. The remainder of the earnings (up to T_i) is paid out in cash on the worker's regular payday.

The funds in the account cannot be accessed until the end of the blocked period (3-9 months).[17] Once the account is unblocked, it converts to a regular savings account with the bank. At that time, workers can withdraw the saved funds, or leave them in the accounts.

We conceptualize this product as reducing the effective social tax rate on earnings increases. Because the funds deposited into the blocked savings account cannot be accessed, they can-

[17]For logistical reasons, the initial set of blocked accounts offered were for 9 months. In later enrollment waves, they were 3 months long in duration.

not be used to fulfill transfer requests—at least not until cash is disbursed after the unblock period. This enables workers to potentially accumulate a large chunk of savings. Consequently, if workers decide to increase their effort, 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. We use the treatment arms below to further vary the likelihood that workers can retain the income they save in the blocked accounts.

4.2 Treatments

Our primary manipulation varies the extent to which the existence of the blocked accounts is known to others in a worker's network. This, in turn, alters whether workers are likely to receive transfer requests against the money saved in the blocked accounts.

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

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

In the Non-private condition, workers are told that if they take up the account and save in it, members of their social network may receive 2 text messages that reveal that the worker has a blocked account and has achieved savings in it.Each of the advertisement text 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 with the bank. The first text message would be sent the week after the worker achieved savings in the account for the first time. The second text message would be sent just before the unblock date, and mention that the worker's savings would be available soon for their use. If the worker declined to take up the account or did not save in it, no information would be shared with network members. Workers were given this information before deciding whether to take up the account in the Non-private condition.

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

In the presence of redistributive pressure, the Non-private treatment increases the probability of transfer requests against savings in the blocked accounts. Specifically, if network members know that the accounts exist and the funds will be unblocked soon, then this will prompt them to ask workers for transfers against the saved amount. In contrast, under the Private treatment, the funds are held without anyone's knowledge, and workers are more likely to be able to retain them for their own use.

Consequently, if there is redistributive pressure, then by lowering the probability of transfer requests, the Private treatment lowers the extent to which income gains will be redistributed others. As a result, we predict that workers in the Private group will increase their labor supply—and therefore their total earnings—relative to those in the other two treatment arms. As we discuss below, we use additional design features to further rule out potential confounds below, such as privacy concerns or self-control problems.

4.3 Randomization

We conduct the experiment in two phases. In Phase 1, a subset of workers is randomized into either the Private or Control conditions. This enables us to compare the overall impact of the Private condition against having no account at all. In Phase 2, we conduct our key test by randomizing workers into either the Private or Non-private account treatment conditions. Table 1 provides an overview of the randomization design.

Table 1: Experiment Design Overview

Phase 1	Phase 2
Private treatment (50%)	**Private** (50%)
	Non-private (50%)
Control (50%)	**Private** (50%)
	Non-private (50%)
Not in Phase 1	**Private** (50%)
	Non-private (50%)

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

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

The randomization design in Phase 2 maximizes statistical power for our primary test: comparing effects of the Private treatment vs. the Non-private treatment on worker output and earnings. First, in addition to adding new workers to the sample for this phase, we also cross-randomize workers who participated in Phase 1—stratifying Phase 2 treatment status by Phase 1 treatment assignment. In the analysis, we conduct tests to verify that treatment effects are not sensitive to Phase 1 treatment status—for example, by excluding workers who were assigned to the Private condition in Phase 1. In addition, we also leverage the cross-randomization to help test for confounds below. Second, we did not include a pure Control condition in Phase 2. Consequently, we cannot directly compare the Non-private condition with the Control condition. Since this comparison is not core to our predictions, our chosen design maximizes our ability to test our key hypothesis.

4.4 Auxiliary tests

We supplement our design with three additional sets of tests.

4.4.1 Placebo test for confounds

One potential concern with our design is that the mechanics through which the Non-private treatment was implemented—text messages sent to others mentioning a savings account held by the worker—could introduce other potential mechanisms that could confound the interpretation of our results. For example, workers might not want their family members and friends to be bothered by a publicity SMS; they might prefer that financial matters stay private; or might be ashamed of revealing that they need to use a formal device in order to achieve savings.

To address such concerns, we undertake two supplementary placebo tests. First, we ask workers for permission to send a publicity SMS about the bank's savings product—similar to the one under the Non-private treatment, but without mentioning the worker's details—to their network members. We do not offer workers with any compensation in exchange. We use this to assess the extent to which a desire to avoid having family members receive an

SMS advertisement drives low take-up in the Non-private condition.

Second, we exploit the cross-randomization in our design to construct a stronger placebo test. We undertake the test with workers who took up Private accounts in *Phase 1* but were offered Non-private accounts in *Phase 2*. Three months after the start of Phase 2, we asked these workers for permission to send promotional text messages to their network members advertising that they had saved in a blocked account through the bank in the past (i.e. in Phase 1). In exchange, workers were offered a small token compensation of 1,000 FCFA.[18] This therefore closely mimics the Non-private treatment in its mechanics—but for accounts where money would already be spent long ago—providing little scope for triggering transfer requests. If take-up rates of this offer are high, this would suggest that simply having others know that you used a blocked account per se is not a major barrier to take-up. Under our hypothesis that the primary reason to prefer privacy is a desire to avoid transfer requests, we predict that take-up of this offer will be substantively higher than take-up of the Non-private accounts.

4.4.2 Self-control in consumption

As we discuss in detail Section 8.2 below, self-control problems alone cannot explain our predicted effects—and are therefore not a confound per se. Specifically, under present-focus alone, effects should be similar between the Private and Non-private blocked account treatments. However, conditional on the Private accounts reducing transfer requests, the existence of self-control problems could still help contribute to our observed treatment effects. Consequently, while the potential presence of present focus does not impact our qualitative test for the existence of redistributive pressure, it has implications for estimating the tax rate implied by our treatment effects.

To gauge the potential relevance of this, we test a core prediction of basic time inconsistency models. Specifically, if agents want to lock away savings so that their future self does not frivolously spend it, then demand for such a commitment device should be higher earlier in the paycycle relative to the day of the payday itself. To test this prediction, we surprise workers with the option to opt-out of the direct deposit of their earnings increases for one pay cycle only—randomly varying whether the option is provided four days before the payday, or on the payday itself. Because workers are always paid a minimum of several days after the end of the paycycle, this is after the effort decision for that paycycle has already been made.

Because of shocks, some degree of opt-out is expected even for workers without self-control problems in consumption. Nonetheless, if workers are present-focused sophisticates, they are more likely to want blocked earnings when offered the option further away from the payday,

[18] As a benchmark, this amount is less than half a day of earnings, and ends up corresponding to less than 5% of the Phase 2 average treatment effect.

knowing that their future self will splurge. When the payday arrives however, they may be tempted to keep all of their earnings in cash, just this one time. We view this as only a suggestive test, which we use to gauge the empirical relevance of self-control in consumption for understanding the benefits of the blocked accounts in our specific setting.[19]

4.4.3 Piece rate variation

We experimentally estimate the workers' effort elasticity to the piece-rate wage. We partner with the factories to randomly vary the piece rate they face for that day's output over the course of 6 days. The piece rate takes 4 possible values: 15% below the usual piece rate, the usual piece rate, 15% above the usual rate, and 30% above the usual rate. During this 6-day period, each morning after arriving to the factory, the worker draws a piece rate level for that day in an i.i.d. fashion. This randomization is done by drawing a slip out of a bag.[20] We use this as an input for estimating the prevailing marginal social tax rate implied by our experimental estimates.

5 Implementation and Protocols

5.1 Job features: Factory setting, production task, and payment

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

Côte d'Ivoire is the world's second-largest producer of raw cashew nuts but only processes 7% of them domestically, with current processing capacity at less than 50% utilization (World Bank 2018, 2020). The government and industry perceive that two major constraints to increasing domestic processing, a national development priority, are low labor productivity and unmet labor demand despite relatively high wages. Our hypothesized mechanism—due to its implications for labor supply and productivity—has relevance for both these constraints.

Production task. Workers in the experiment are engaged in manually peeling cashew nuts. This entails gently rubbing off with the fingers or a knife those parts of the peel that are

[19]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 supported by previous work—for example, Kaur et al. (2015) and Augenblick et al. (2015). Under hyperbolic time preferences, this test is valid regardless of the length of time periods in the utility function.

[20]Out of fairness concerns, workers know that they cannot lose from the activity, even if they repeatedly draw the -15% slip. Specifically, they know that they will earn, at the end of the paycycle, at least as much as they would have earned if the usual rate had been applied to their output.

still attached to the cashew after it undergoes mechanized peeling. Workers fill up buckets with cashews and return to their workstation for peeling.

Workers complete a set workday from 8am to 5pm, Mondays through Saturdays, with a one-hour break for lunch. 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.

Payment. Each worker receives a linear piece rate for her output.[21] The entirety of workers' earnings are comprised of their piece rate wages. Consequently, changes in effort translate directly into changes in worker earnings. Workers are paid their earnings twice a month in cash.

5.2 Experiment Protocols

Sample. We enroll 473 full-time workers, of which 464 are women, in the manual peeling sections of the two Olam factory plants.[22] We enroll all eligible workers in the manual peeling sections across both factory sites. To be eligible, workers must have regularly attended work and have been present at the factory for a sufficiently long period. We impose tenure requirements to minimize attrition from the experiment, since newly joined workers have high turnover rates. It is unclear ex-ante whether treatment effects for more short-term workers would be higher or lower. All workers in our sample report their factory income as their primary source of earnings, and 89% report having no other source of income.

Blocked accounts. The blocked savings accounts are designed and implemented jointly with Banque Populaire (BPCI), and represent a financial innovation relative to what was available on the market. They are free of charge for those are offered the accounts through the experiment, and have no minimum deposit requirements. In this setting, opening a formal bank account is an administrative hassle: individuals must obtain formal documents and travel to the bank to fill out the application and submit them. To help minimize this hurdle for workers who were offered the blocked accounts in the experiment, a bank employee is stationed in the factories during the enrollment period to help collect applications. In addition, by directly depositing worker earnings above the threshold into the accounts, the factory makes making deposits much easier—since otherwise deposits need to be done manually by traveling to a bank branch. Finally, savings accounts typically charge fees unless workers maintain a minimum savings balance—a requirement that is also waived for those offered blocked accounts through our experiment. Consequently, while all workers can open a nor-

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

[22] Across all cashew-processing facilities managed by Olam, roughly 95 per cent of the workers in peeling sections are women. Women are typically more likely than men to sign up for (and be hired for) manual peeling jobs.

mal savings account with the bank, in practice, our intervention greatly lowers the barriers to obtaining and using a formal bank account. Consistent with this, only 1% of Control group workers have a formal bank account at endline.

Workers who take up the blocked account choose one threshold that applies to all future pay-cycles. They can revise this threshold as many times as they want in the first month, and up to three times after that. 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 against mistakes, and allow workers the ability to re-optimize thresholds after experiencing the accounts if they wanted.

After the end of the blocked period, the account converts to a regular savings account. At that time, workers have the option to withdraw all or part of their accumulated savings, or to lock their existing savings for another blocked period.[23] The length of the blocked accounts is 9 months in Phase 1, and shortened to 3 months in Phase 2 in order to complete the experiment on the timetable agreed upon with implementing partners.[24]

Announcement and Training. To ensure workers understand how the blocked accounts worked, we undertake training sessions within the factory with all workers who were assigned to the Private or Non-private conditions. Workers attend the sessions in small groups of about 5 workers each. These sessions are attended by a bank staff member to answer questions, and led by a moderator from the research team. The sessions explain the rules of the accounts, including choosing thresholds, and work through examples. At the end of the session, each worker takes a quiz to verify comprehension of the account rules. If the worker scores below 80% on the quiz, they are retrained one-on-one by a moderator.

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 they have 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. In addition, workers can use an SMS to check their savings balance with the bank at any time. Consequently, our protocols ensure privacy of savings deposits—while leaving the process of obtaining payment of the cash component of earnings on paydays

[23] If workers re-block their existing savings for another blocked period, this only applies to prior savings already in the account. No future earnings are directly deposited—unless dictated by the worker's treatment status in a subsequent phase of the experiment. Consequently, once the blocked period ends, there is no expected impact in the social tax rate on future earnings.

[24] We were scheduled to complete the experiment in the summer of 2019. Delays in Phase 1 implementation reduced the time remaining to complete Phase 2. To stick to our timetable, we therefore reduced the time length of the Phase 2 accounts to be 3 months. This does not affect the internal validity of our estimates, since we only compare worker outcomes within each phase.

unchanged. These payment protocols were the same for both the Private and Non-private treatments.

Enrollment and roll-out. In Phase 1, workers were enrolled into the experiment in three waves between 2017-2019. The launch in 2017 involved only one of the factory plants and a subset of workers—to establish operations and implementation joint with the bank and Olam. The primary waves of Phase 1 occurred from July 2018-March 2019 in both factory locations simultaneously. Phase 2 was then conducted in both factory locations simultaneously from April 2019-July 2019.

6 Empirical strategy

6.1 Data sources

Our primary data source is Olam's detailed daily administrative data at the worker level. This data includes individual daily records of attendance and output (the quantity of nuts processed), and earnings–our main labor supply outcomes. These data are used by Olam to compute earnings, and the amount to be deposited into workers' accounts in each paycycle.

We supplement this with data collected through three sets of phone surveys conducted during the experiment and a more detailed in-person endline survey. These include information on perceptions about redistributive pressure, financial transfers with network members, and earnings outside the factory. Respondents received a small compensation for answering the surveys, with compensation for time set to be commensurate with the hourly wage a typical worker would earn at Olam.

Table 2 provides descriptive statistics for our study sample, and checks for balance across control and treatment arms, separately for Phase 1 and Phase 2. Balance tests for equal means of baseline measures confirm the randomization generated balanced treatment assignments. We find some evidence for imbalance in baseline attendance rates, and show robustness to this in the analysis below.[25]

6.2 Estimation

Our primary outcome is workers' daily earnings (in FCFA). Given the linear piece rate incentive scheme, this is similar to examining effects on output, and captures workers' labor

[25] The imbalance in attendance is limited to the sample of workers from site 1. In site 2, there is perfect balance on earnings and attendance (Appendix Table A1). To verify our results–particularly the effects on the extensive margin–are not driven by this imbalance, we show that our treatment effect estimates on both earnings and on attendance are robust to estimations using the sample of workers from site 2 only. In addition, since we run difference in differences regressions to estimate treatment effects, our estimates should not be sensitive to some baseline imbalance in levels.

Table 2: Balance

Variable	C. mean	T. mean	Diff.	p-value	N
		Phase 1			
Baseline earnings	1756	1700.63	-56	.56	4455
Baseline attendance	.68	.63	-.05	.03	4455
Leaves before announce	.12	.11	-.01	.76	400
Has valid ID	.51	.5	0	.95	353
Speaks Dioula	.49	.52	.03	.64	353
Speaks Baoule	.28	.31	.03	.53	353
		Phase 2			
Baseline earnings	1694	1604	-89	.28	3804
Baseline attendance	.77	.7	-.07	.03	3804
Leaves before announce	.14	.15	.01	.82	370
Has valid ID	.55	.61	.06	.31	317
Speaks Dioula	.41	.39	-.03	.65	317
Speaks Baoule	.26	.3	.04	.42	317

Notes: This table presents the average in the control group, the average in the treatment group, the difference (treatment mean - control mean), the p-value associated with this difference, and the number of observations in each regression. Observations are at the worker level. Dioula and Baoule are the most-spoken languages in our sample. Baseline earnings are average daily earnings in the 15 days before the treatment announcement. Baseline attendance are average days of presence in the 15 days before the treatment announcement. As explained above, Olam's piece-rates are set to reach Côte d'Ivoire's minimum wage for full-time workers, but workers may earn less either through lower attendance or productivity. Leaves before announce denotes workers who were eligible for the study but left the factory before treatment assignment was announced. Standard errors are clustered at the worker level.

supply. We also study effects on attendance to disentangle the intensive (effort while at work) and the extensive (attendance) margins of labor supply.

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

$$y_{it} = \beta PrivateAcct_{it} + \gamma_i + \delta_t + \epsilon_{it}, \tag{1}$$

where y_{it} is the outcome of interest for worker i on date t. The $PrivateAcct_{it}$ indicator equals one if worker i was assigned to receive the Private treatment on date t (i.e. during the days comprising the blocked period). In addition, γ_i and δ_t are worker and paycycle fixed effects, respectively. We also show robustness to including strata$\times$paycycle fixed effects in the regressions, as well as robustness to other specifications.[26] Standard errors are clustered

[26] In Phase 1, strata are factory x treatment waves. In Phase 2, strata are treatment assignment in Phase 1, separately by wave in each factory—and are therefore collinear with the worker fixed effects. The gap

at the individual-level to allow for serial correlation of ϵ_{it} (Bertrand et al., 2004; Abadie et al., 2017).

7 Results

7.1 Take-up of blocked accounts

Figure 3: Phase 1 and Phase 2: Account Take-up

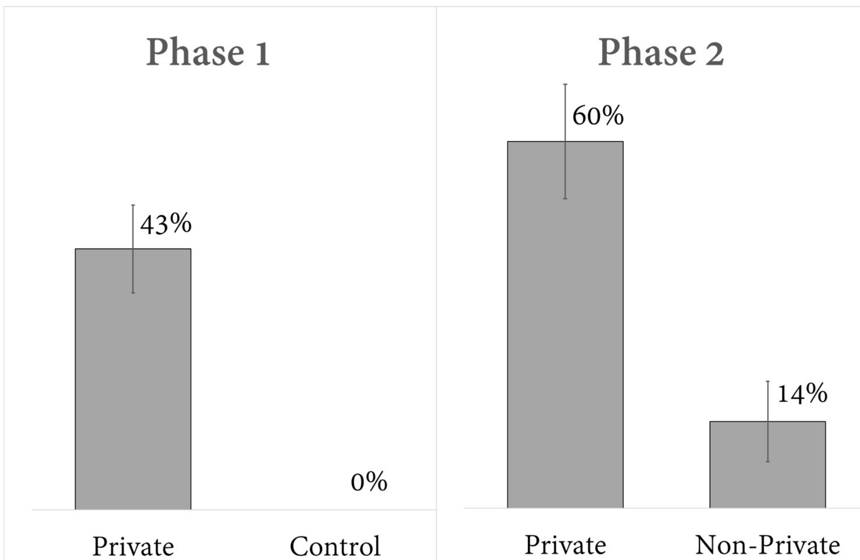

Notes: Means and 95% CIs. SEs clustered at the worker-by-phase level.
Phase 1: N = 409 worker-waves; 354 workers. Phase 2: N = 317 workers.

In Phase 1, the proportion of Private group workers who decide to take up the blocked savings account is 43%. Our key test is demand for the blocked accounts in Phase 2. In Phase 2, take-up rate for Private blocked accounts rises to 60%. In contrast, under the Non-private condition, account take-up is only 14%. This indicates that the blocked accounts are substantially less desirable as a savings vehicle if others' in the workers' social network would learn of their existence.

Those accounts were actively used by workers. In Phase 1, workers who took up the accounts achieved savings in 28% of their paycyles on average.This translates into savings deposited into blocked accounts in 11% of worker-pay cycles overall. In Phase 2, commensurate with the relatively higher take-up rate of Private accounts, savings were deposited into worker accounts in 16% and 3% of pay cycles in the Private and Non-private conditions, respectively. Note that since thresholds had to be set weakly above baseline earnings, these numbers

between the end of Phase 1 and start of Phase 2 is two weeks. For consistency, we use a baseline period of two weeks in both the Phase 1 and Phase 2 treatment analyses; results are robust to other baseline periods.

under-estimate the fraction of paycycles in which workers earned more than their baseline earnings levels.

The take-up rates for the blocked accounts are reasonably high in comparison to previous studies. This may be driven by several features of our setting. First, we invested a lot of time up front in factories to build trust among workers in opening formal bank accounts; the role of trust is consistent with the large increase in Private account take-up between Phase 1 and Phase 2. In addition, we removed many logistical barriers to opening and making deposits into accounts, as discussed in Section 5.2 above. In addition, in contrast to traditional commitment device studies, there were no penalties in failing to make a target which lowered risk from participating in the accounts. Further consistent with the popularity of the accounts, among workers who took up a Private account when offered the first time, 100% chose to do so again if they were randomized to be offered it a subsequent time.

7.2 Effects on labor supply and earnings

The Private blocked accounts raise worker output and earnings. In Phase 1, workers assigned to the Private account treatment increase their total earnings by 187 FCFA per day (12%) on average relative to Control workers (Table 3, Panel A, Col. 3, p-val 0.022). These results indicate large output and earnings benefits from simply offering workers the blocked accounts.

Table 3, Panel A, Cols. (4)-(5) provide the key test of our hypothesis by comparing earnings under the Private vs. Non-private account conditions in Phase 2. Compared to the Non-private condition, Private condition workers increase their output and earnings by 227.9 FCFA (Col. 5, p-val < 0.001). This corresponds to a striking 14.5% average treatment effect on earnings. On average, this magnitude is equivalent to how much earnings would rise if each worker worked an additional 1.04 days in *every 2-week paycycle*.[27]

In Col. (6), we verify that our treatment effects are not driven by Phase 1 treatment status. Specifically, since we cross-randomized treatment assignment across the two Phases, one potential concern is that treatment effects may be heavily influenced by workers who were assigned to Private accounts in Phase 1 and then Non-private accounts in Phase 2. In Col. (6), we exclude from the analysis workers who were assigned to Private accounts in Phase 1—eliminating the scope for this concern. The results are nearly identical.

Our conceptual framework predicts an increase in labor supply, but is agnostic as to the specific channel: increased effort on the intensive margin (i.e. increased productivity while

[27] The average treatment effect on daily earnings is 228 FCFA, corresponding to a 2,736 FCFA increase per paycycle (comprised of 12 workdays). Mean daily earnings among the Non-private group is 2,635. This gives $2736/2635 = 1.04$ workdays per paycycle.

Table 3: Treatment effects on earnings and attendance

	Phase 1			Phase 2		
		Full Sample		Full Sample		Restricted
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Daily earnings						
Private (vs. Control)	169.5**	174.0**	187.3**			
	(78.06)	(87.49)	(81.31)			
Private (vs. Non-Private)				218.3***	227.9***	243.5***
				(60.35)	(60.39)	(73.85)
Control mean	1587	1587	1587	1570	1570	1639
Panel B: Attendance						
Private (vs. Control)	0.0571*	0.0613*	0.0688**			
	(0.0304)	(0.0338)	(0.0313)			
Private (vs. Non-private)				0.0513**	0.0553**	0.0547*
				(0.0252)	(0.0251)	(0.0301)
Control mean	0.67	0.67	0.67	0.60	0.60	0.61
Worker FE	Yes	Yes	Yes	Yes	Yes	Yes
Paycycle FE	Yes	Yes	No	Yes	No	No
Strata FE	No	Yes	No	Yes	No	No
Strata-paycycle FE	No	No	Yes	No	Yes	Yes
N: worker-days	99215	99215	99215	38222	38222	24465
N: workers	353	353	353	317	317	203

Notes: Dependent variable in panel A: daily earnings. Dependent variable in panel B: attendance. Attendance equals 1 if the worker is present at the factory that day, 0 otherwise. DiD specification, with 15 days of baseline. In Phase 1, strata are treatment waves by factory. In Phase 2, strata are treatment assignment in Phase 1 separately by wave in each factory. The strata FE in Phase 2 are thus collinear to the Worker FE and the treatment variable. Column (6) excludes workers who were offered a private account in Phase 1. Standard errors clustered at the worker level. *, **, *** indicate significance at the 10, 5 or 1% level.

at work) or on the extensive margin (i.e. increased attendance at the factory). To investigate this question, in Table 3, Panel B, we estimate the treatment effects on workers' attendance, obtained from Olam's administrative data.[28] Private workers increase their attendance at the factory by 6.9 percentage points (10.3%) compared to Control workers and by 5.5 percentage points (9.2%) relative to Non-private workers (Cols. 3 and 5, respectively). The Phase 2 effect on attendance accounts for about two-thirds of the total treatment effect on earnings.

In addition, we verify that these effects are not simply reflecting a substitution away from other income-generating activities outside of the factory jobs. Using earnings data from the endline surveys, we find no treatment effects on earnings outside the factory. This is consistent with the fact that, at baseline, 89% of workers report having no earnings outside of their factory job.

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

[28]Given the low flexibility of work hours (see Section 5.1), we focus on attendance as the extensive margin, and treat the rest of the treatment effects as intensive margin.

their total income.

7.3 Effects on transfers

As discussed in Section 2, we designed our treatment to be potentially Pareto-improving: enabling workers to increase their Olam earnings with increased effort, while leaving kin no worse off due to unchanged amounts available for redistribution. In Table 4, we examine what happens to the likelihood of transferring to kin–and amount of transfers made–by workers in Phase 2 of the study.[29]

Table 4: Transfers to Network Members

	Panel A: Likelihood of transfer			
	To anyone	To family	To non-family	From anyone
	(1)	(2)	(3)	(4)
Private	0.0419	0.0540	-0.00523	0.0779
	(0.0464)	(0.0582)	(0.0583)	(0.0598)
Control mean	0.784	0.510	0.405	0.484
	Panel B: Amount transferred			
	To anyone	To family	To non-family	Net amount
Private	4998.2**	3918.2	1080.0	3261.5
	(2532.5)	(2398.1)	(951.8)	(3244.8)
Control mean	13129	9967	3162	4848

Notes: N=298 workers. Endline survey data: transfers to and from others outside worker's household. Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Regressions include strata dummies.

We do not find significant changes in either the likelihood of having made a transfer in the last three months (Panel A, Cols. 1-3) or the likelihood of having received a transfer in the last three months (Panel A, Col. 4). Over the same time period, estimates of total transfer amounts given to family members and non-family members (Panel B, Cols. 2-3), as well as net transfers (Panel B, Col. 4) are not statistically significant. However, when we pool all transfers made by the worker in the last three months (Col. 1), we detect a significant increase as a result of treatment.

This positive effect could reflect the fact that while Private accounts lower the social tax on income gains, they do not drive the tax down to zero–especially since most workers set

[29] At endline, workers were asked how much they transferred to different categories of individuals over the previous three months. Transfers to parents, siblings, adult children, aunts/uncles and other family members are summed to create the 'family' category, while transfers to friends, acquaintances and others are summed to create the 'non-family' category. Neither the likelihood of receiving transfers nor net transfers (total transferred minus received) are dis-aggregated, since the endline survey did not collect an exhaustive list of transfers received by category of individual. We use Phase 2 assignment to provide a cleaner treatment impact estimate, since a portion of Phase 1 workers were re-offered a private account in Phase 2 while others were not, as shown in Table 1.

thresholds higher than their baseline earnings (as shown in Figure A1). Workers could have thus transferred some of the earnings gains that were not automatically deposited into the blocked account, redistributing some of their savings due to altruistic reasons or requests that happened to arrive around their withdrawal date. Regardless, overall, our results indicate that the income gains achieved by Private group workers did not come at the expense of lower redistribution to others, and may have led to aggregate welfare gains.

8 Confounds and Mechanisms

8.1 Non-private treatment: Threats to validity

As discussed in Section 4.4.1, we undertake two placebo tests to examine whether low take-up of the Non-private accounts may be driven by some channel other than an expected increase in transfer requests against the savings in the account. This includes not wanting family members to be spammed with a publicity SMS, a hedonic preference for keeping one's financial matters private, or embarrassment at others discovering that the worker needed a blocked account to save.

Note that any such other explanation must cause sufficient disutility that it is worth giving up the 14.5% average increase in income obtained by Private workers in Phase 2. In addition, note that such explanations cannot on their own explain why the Private accounts do have positive benefits (in both Phase 1 and Phase 2). Consequently, one would then need to consider an additional mechanism (aside from redistributive pressure) to explain the positive Private account treatment effects.

In our first placebo test, we ask workers for permission to send a publicity SMS about the bank's savings product—similar to the one under the Non-private treatment, but without mentioning the worker's details—to their network members. We do not offer workers with any compensation in exchange. 95% of workers in the Non-private treatment group give permission to have these messages sent to their network. Consequently, it is unlikely that a desire to avoid having family members receive an SMS advertisement drives low take-up in the Non-private condition.

In our second, sharper, test, we undertake an exercise with workers who had *taken up* a Private account in *Phase 1* and were assigned to the Non-private treatment in *Phase 2*. Three months after the end of Phase 1 (i.e. after the start of Phase 2), we ask these workers for permission to advertise their prior savings in a Private account to their network via a publicity SMS. Since the money would have already been spent by the time this SMS is sent, this placebo test mimics the Non-private treatment in terms of the type of information revealed—but does not have nearly the same scope to trigger transfer requests (see discussion in Section 3). Workers were offered a token 1,000 FCFA compensation in exchange for this—a magnitude that is less than 5% of the Private treatment effect in Phase 2. A striking 85% of workers agree to this offer.[30] This strongly indicates that simply having one's financial information revealed to others is not the reason workers decide to forego having a blocked account in the Non-private treatment.

Finally, fear of theft also cannot explain low take-up of Non-private accounts. Workers walk home from the factory with their entire cash earnings in their pockets twice each month, on days that are publicly known. The release and withdrawal of savings from the bank is considerably more private than this.

8.2 Self-control benefits of blocked accounts

Blocked accounts can also boost effort if workers have self-control problems in consumption. However, under present-focus alone, effects should be similar between the Private and Non-private blocked account treatments. Redistributive pressure is therefore necessary to explain our results. In other words, self-control alone is not sufficient to explain our Phase 2 treatment effects, and is therefore not a confound per se.

Conditional on the Private accounts reducing transfer requests, the existence of self-control problems could still help contribute to our observed treatment effects. We gauge the potential empirical relevance of present-focus through the opt-out test described in Section 4.4.2.

Counter to the predictions of basic present-focus models, workers are not relatively more likely to want to unblock savings on the day of the payday itself. The proportion of workers who decide to keep their earnings in the blocked account 4 days before payday is 86% , versus 94% on payday—a difference that is not statistically distinguishable from zero (Appendix Figure A3). This is counter to the predictions of basic present focus models—with the results actually going in the opposite direction. Note that a small proportion of workers opting out in a given paycycle is to be expected—even in the absence of any self-control problems—due, for example, to volatility which may raise the need for cash in some weeks.

Thus, while individuals may face self-control problems in consumption and savings in general, as has been documented in other settings (e.g., Ashraf et al., 2006; Schilbach, 2019), these patterns suggest this mechanism is unlikely to strongly drive effects of the blocked accounts in our specific experiment. Finally, note that the blocked accounts cannot help with self-control problems in effort provision, because the accounts push the receipt of earnings (i.e. the returns to effort) to be even further into the future (O'Donoghue and Rabin, 1999; Kaur et al., 2015; Augenblick et al., 2015)

[30] In contrast, among workers who took up Private blocked accounts in Phase 1, the fraction that take-up Private and Non-private accounts in Phase 2 is 90% and 31%, respectively (Appendix Table A4).

9 Estimation of Social Tax Rate

Methodology. Our experimental design can be used to estimate the prevailing rate of the social tax faced by workers in our context. To do so, we take advantage of the specific kink in the workers' budget constraint induced by our intervention.

By dampening the social tax rate on earnings *increases* only, our intervention allows us to identify a pure substitution effect – as explained in section 2. We can thereby formally relate the change in the social tax rate τ induced by the provision of a private blocked account to our estimated change in effort e through the compensated elasticity of effort to the net-of-tax wage rate ζ^c:

$$\frac{de_2{}^*}{e_2{}^*} = \zeta^c \frac{d(1-\tau_2)}{(1-\tau_2)} \tag{2}$$

A formal demonstration for Equation 2 is provided in Appendix section A.1, building upon the model introduced in section 2. 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:

$$\frac{1-\tau_1}{1-\tau_2} = \left(\frac{e_1{}^*}{e_2{}^*}\right)^{\frac{1}{\zeta^c}} \tag{3}$$

Estimate for $\frac{e_1{}^*}{e_2{}^*}$. We can recover an estimate of $\frac{e_1{}^*}{e_2{}^*}$ from our main treatment effect estimates. Given the linearity between earnings and production, the percentage increase in earnings induced by our intervention dampening the Social tax rate equals the related percentage increase in effort, which we measure by the quantity produced.

Estimate for ζ^c. To calibrate ζ^c, we experimentally varied the piece-rate faced by workers on their output of the day, as described in Section 4.4.3, and measured its effects on their output. This yields an estimated elasticity of 0.25 (Table A6, col 2, p-val 0.034).[31]

Note that ζ^c is a compensated elasticity. Our approach minimizes income effects associated with the piece-rate variation and therefore seems sensible. A worker keeping her level of effort unchanged would expect an increase in earnings due to the experiment of 3.75% at the next pay cycle[32]. That change in income is small, and becomes negligible when compared to a worker's lifetime earnings. It is therefore unlikely that the experiment would trigger major income effects – so the estimated elasticity can be used as measure of a compensated

[31] We find asymmetric labor supply responses around the usual piece rate, possibly due to the compensation scheme introduced to ensure that workers cannot lose earnings during the piece-rate variation exercise and that could incentivize workers to not reduce effort as much when drawing the rate lower than the usual one. To obtain a conservative tax rate estimate, we prefer to exclude those observations.

[32] Each of the 6 days of the experiment, the worker could face either -15%, +0%, +15% or +30% of the usual rate, each with probability 1/4. For a standard pay cycle covering 2 weeks of earnings, this amounts to $\frac{-15+0+15+30}{4} \times \frac{6}{12} + 0 \times \frac{6}{12} = 3.75$.

elasticity.

In this experimental piece-rate variation activity however, workers could only respond along the productivity margin – while they could respond along both productivity and attendance in our main intervention. As such, our estimate is not a perfect mapping for the elasticity of effort ζ^c.[33] Rather, we can recognize that ζ^c, the compensated elasticity of effort to the net-of-tax wage rate, can be decomposed into the sum of the compensated elasticity of productivity ζ_p^c and of the compensated elasticity of attendance ζ_a^c. We can thereby use the elasticity estimated from the piece-rate variation activity as an upper bound for the elasticity of productivity ζ_p^c: if both the productivity and attendance margins had been available during the piece-rate variation exercise, workers would have increased productivity at most as much as they did when only the productivity margin was available. This is a conservative choice, since an upper bound on the elasticity translates into a lower bound on the estimated social tax rate.

As detailed in Section 7.2, the 14.5% increase in earnings from providing private blocked savings accounts can be decomposed into a 9.2% increase in attendance and a 5.3% increase in productivity. If we assume that our intervention led to the same percentage change in the social tax rate faced by workers on their attendance than on their productivity, then we can infer that the workers' elasticity of attendance is about 74% higher than their elasticity of productivity. We then have $\zeta^c = \zeta_p^c + \zeta_a^c = \zeta_p^c + \zeta_p^c \times 1.74 = 0.25 + 0.25 \times 1.74 = 0.67$. We will therefore use 0.67 as our preferred estimate for the compensated elasticity of effort, used in the computation of the social tax rate.

Social tax rate estimates. Assuming that the provision of private blocked savings accounts brings the marginal Social tax rate on earnings increases to 0, and using our preferred conservative elasticity estimate of 0.67, we find that the prevailing marginal social tax rate prior our intervention was of 18%.

Since our intervention might not remove all redistributive pressures on labor earnings increases for the workers who choose to take-up the private blocked accounts, Appendix Figure A4 presents the estimated prevailing Social tax rate for a range of possible τ_2. It also displays the statistical uncertainty inherent to both the treatment effects and the elasticity estimates, by including 90% confidence intervals computed by bootstrap.

Our preferred elasticity estimate is likely an upper bound, and thus provides a conservative

[33] To allow workers to choose a response to the piece rate variation along margins more similar to those of our main intervention, we would have needed to assign workers with the same piece rate over multiple days. (Over at least 3 days, given the 60% attendance rate in the "Non-Private" group.) Doing so would however have induced income effects over the paycycle: the estimated elasticity would not clearly have been a compensated elasticity, which is what we need. Further, to satisfy the operational and ethical requirement that the workers all face a similar tax rate on average over the variation experiment, the implementation period would have been much longer than operationally possible.

lower bound for the prevailing social tax rate. On the other end of the spectrum, the low elasticity of attendance experimentally estimated by Goldberg (2016) in Malawi yields a social tax rate closer to 23%. Appendix Table A7 presents the prevailing tax rate implied by a range of reasonable elasticities of effort.

10 Conclusion

Informal transfers among kin groups and social networks are important for coping with risk and are typically viewed as unequivocally positive. This is especially the case when kin networks fill a space created by missing insurance markets. We argue that while informal transfer arrangements may enable important welfare benefits, such benefits likely come at the cost of efficiency. Efficient implementation of transfers relies on perfect information of shocks experienced by the agents participating in the transfer arrangement. In reality, shocks and effort are only imperfectly observable and transfers are enforced via norms of redistribution, with members who start earning more (for example, by taking up a job in a factory or other employment in the formal sector) expected to increase financial assistance to other members.

We provide evidence that this form of social insurance may turn into social taxation, dampening incentives to exert effort and accumulate wealth. We develop a financial product to mimic a decrease in social taxation: a private blocked savings account into which worker earnings are automatically transferred every fortnight, provided they clear a threshold at least as high as their average earnings pre-intervention. 75% of workers in our sample report having trouble saving because if they put money aside, someone will ask them it. When offered the possibility of opening the account, which allows workers to save without cash on hand and enables refusal of transfer requests without cost, 43% accept.

Consistent with kin taxation as an impediment to labor supply, accounts strongly raise individual labor supply, but only when others do not know of their existence. The increase in labor supply appears to be driven by both an increase in attendance and an increase in effort while on the job. The observed effects on attendance and earnings cannot be explained by privacy concerns, and we show the existence of redistributive pressure is a necessary mechanism to explain the results. We exogenously vary workers' piece-rate wages to compute the marginal social tax rate implied by our results, which we estimate to be 18%. We do not find evidence that the intervention leads workers to reduce the amount they transfer to kin, suggesting that reducing redistributive pressure may be Pareto-improving.

Our findings suggest that social norms around redistribution may constitute a barrier to economic development. New economic activities—from adopting cash crops in agriculture to gaining employment in prima facie higher-return nonfarm jobs—can be undermined if individuals do not capture the rents from their efforts. The effect of redistributive norms on labor supply may be particularly important for explaining obstacles to value-chain upgrading through a labor-intensive processing sector recruiting low-skill workers. In Côte d'Ivoire, as in other developing economies, two barriers to the development of a more formal sector in rural areas are the high unit-labor costs stemming from low labor productivity, and the unmet demand for low-skill labor (World Bank 2018). Both of these channels are consistent with workers facing an increase in social taxation if they take up a more formal job, thus disincentivizing labor supply.

Our findings speak to the importance of strengthening formal safety nets. Tackling the likely underlying cause for redistributive norms—the lack of formal consumption smoothing mechanisms—could improve output and growth. Government social safety net programs could not only improve the welfare of the unemployed, but could also boost the productivity of the employed by reducing their responsibility for redistribution. Similarly, improving access to formal health or livelihood insurance could have externality benefits by increasing the aggregate output of others in the same kin network.

Two additional points merit mentioning. First, since virtually our entire sample is composed of women, we cannot speak to gender differences in the prevalence or severity of redistributive pressure. Formative qualitative research among the cashew factory workers found that both men and women face redistributive pressure (McNeill and Pierotti, 2020), and existing evidence on redistributive pressure confirms this. That said, a vast body of international work has documented women's lower-decisionmaking power in the household. It is possible that men and women face different compositions of intra vs. inter-household pressure and ensuing effects on labor supply. Moreover, since women's labor force participation and earnings are generally lower compared to men's, improving access to formal safety nets and financial products that shield women from redistributive pressure may be important avenues for economic development. Both of these hypotheses merit exploration in future research.

Second, our financial innovation is a promising policy tool for improving earnings among the poor, regardless of the mechanism through which it operates. Our experiment offers some lessons for the design of blocked savings instruments. First, trust in institutions is a major determinant of account take-up. Many workers who did not take-up reported being swindled by past financial institutions. Take-up increased in each subsequent implementation wave, with individuals learning from others' experience with the instrument. Moreover, virtually everyone who took up account once did so again when offered to a 2nd or 3rd time. The blocked accounts continued to be offered in one of the factories once the research team left, upon completion of our study. However, the implications of such a financial tool for risk sharing are less clear. General implementation may not necessarily be Pareto-improving, as it could exacerbate ex-post reneging. This is also a fertile area for future research.

References

Abadie, A., S. Athey, G. W. Imbens, and J. Wooldridge (2017): "When should you adjust standard errors for clustering?" .

Alby, P., E. Auriol, and P. Nguimkeu (2020): "Does Social Pressure Hinder Entrepreneurship in Africa? The Forced Mutual Help Hypothesis," *Economica*, 87, 299–327.

Alger, I. and J. W. Weibull (2008): "Family Ties, Incentives and Development: A Model of Coerced Altruism," in *Arguments for a Better World: Essays in Honor of Amartya Sen, Volume 2*, Oxford: Oxford University Press.

Anderson, S. and J. M. Baland (2002): "The economics of roscas and intrahousehold resource allocation," *The quarterly journal of economics*.

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.

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.

Baland, J.-M., I. Bonjean, C. Guirkinger, and R. Ziparo (2016): "The economic consequences of mutual help in extended families," *Journal of Development Economics*, 123, 38–56.

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

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

Beekman, G., M. Gatto, and E. Nillesen (2015): "Family networks and income hiding: evidence from lab-in-the-field experiments in rural Liberia," *Journal of African Economies*, 24, 453–469.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Differences-In-Differences Estimates?" *The quarterly journal of economics*, 119, 249–275.

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.

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.

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.

Dillon, B., J. De Weerdt, and T. O'Donoghue (2020): "Paying More for Less: Why Don't Households in Tanzania Take Advantage of Bulk Discounts?" *The World Bank Economic Review*, forthcoming.

Dupas, P., A. Keats, and J. Robinson (2017): "The Effect of Savings Accounts on Interpersonal Financial Relationships: Evidence from a Field Experiment in Rural Kenya," *The Economic Journal*, 129, 273–310.

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

Efron, B. (1982): *The jackknife, the bootstrap and other resampling plans*, SIAM.

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.

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.

Hansen, B. (2020): *Econometrics*.

Hausman, J. A. (1985): "Taxes and labor supply," New York: North-Holland Publishers, vol. 1 of *Handbook of Public Economics*, chap. 4, 213–263.

Hoff, K. and A. Sen (2011): "The Kin System as a Poverty Trap?" in *Poverty Traps*, ed. by S. Bowles, S. N. Durlauf, and K. Hoff, Princeton, NJ: Princeton University Press, 95–115.

Jakiela, P. and O. Ozier (2016): "Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies," *The Review of Economic Studies*, 83, 231–268.

Karaivanov, A. and R. M. Townsend (2014): "Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes," *Econometrica: Journal of the Econometric Society*, 82, 887–959.

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

Kennedy, P. (1988): "African Capitalism: The Struggle for Ascendancy; and Dietz, J. and D. James (eds)(1990)," *Progress Toward Development in Latin America*.

Lewis, A. W. (1955): *Theory of Economic Growth*, Routledge, 1 edition ed.

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

McNeill, K. and R. Pierotti (2020): "Reason-Giving for Resistance: Relational Work and Obfuscation in Informal Financial Assistance," .

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

Mobarak, A. M. and M. Rosenzweig (2012): "Selling Formal Insurance to the Informally Insured," *mimeo, Yale University*.

O'Brien, R. (2012): "Depleting Capital? Race, Wealth and Informal Financial Assistance," *Social Forces*, 91, 375–396.

O'Donoghue, T. and M. Rabin (1999): "Doing It Now or Later," *The American Economic Review*, 89, 103–124.

Platteau, J.-P. (2000): *Institutions, social norms, and economic development*, vol. 1, Psychology Press.

——— (2014): "Redistributive Pressures in Sub-Saharan Africa: Causes, Consequences, and Coping Strategies," in *Africa's Development in Historical Perspective*, Cambridge University Press, 153–207.

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

Riley, E. (2020): "Resisting social pressure in the household using mobile money: Experimental evidence on microenterprise investment in Uganda," .

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.

Schilbach, F. (2019): "Alcohol and Self-Control: A Field Experiment in India," *The American Economic Review*, 109, 1290–1322.

Scott, J. C. (1976): *The Moral Economy of the Peasant: Rebellion and Subsistence in Southeast Asia.*, Yale University Press.

Somville, V. (2011): "Daily Collectors, Public Good Provision and Private Consumption: Theory and Evidence from Urban Benin," Tech. Rep. 1106.

Squires, M. (2018): "Kinship Taxation as an Impediment to Growth: Experimental Evidence from Kenyan Microenterprises," .

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.

World Bank (2018): "Cashew Value-Chain Competitiveness Project: Project Appraisal Document," Tech. rep., The World Bank Group, Washington, DC.

——— (2020): "World Development Report 2020: Trading for Development in the Age of Global Value Chains," Tech. rep., The World Bank Group, Washington, DC.

A Appendix

Figure A1: Distribution of baseline earnings relative to the chosen threshold

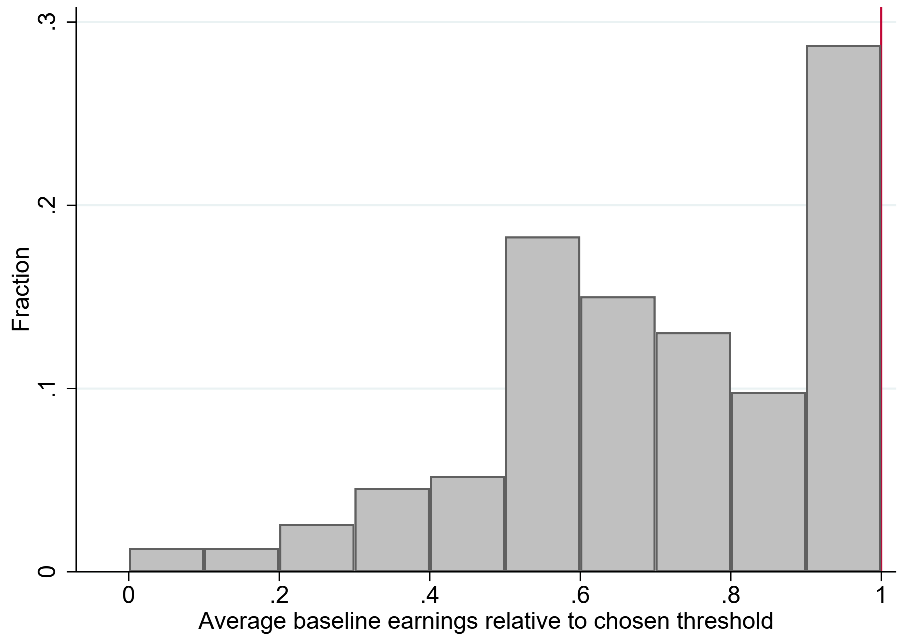

Notes: The figure presents the distribution of average baseline earnings relative to the threshold: $\frac{\text{average baseline paycycle earnings}}{\text{chosen threshold}}$.
Sample of the 153 treated workers holding an illiquid private savings account, pooling both Phases of the intervention.

Figure A2: Distribution of treatment period earnings relative to the chosen threshold

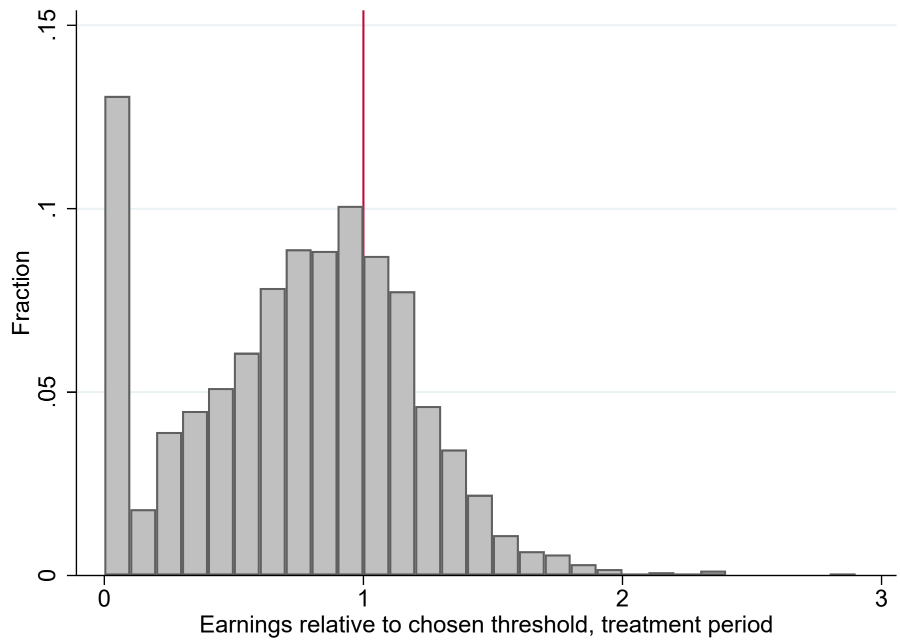

Notes: The figure presents the distribution of treatment period earnings relative to the threshold: $\frac{\text{earnings per paycycle}}{\text{chosen threshold}}$.
Sample of the 153 treated workers holding an illiquid private savings account, pooling both Phases of the intervention. N=2,272 worker-paycycles.

Figure A3: Refusal to Opt Out

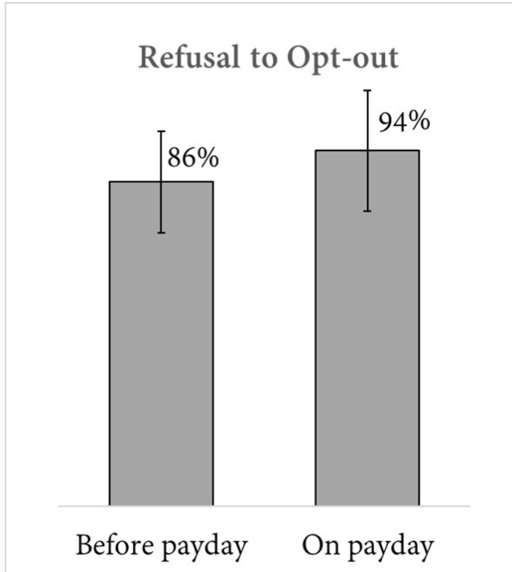

N = 61 factory workers, Côte d'Ivoire
Means and 95% confidence intervals

Figure A4: Social tax rate

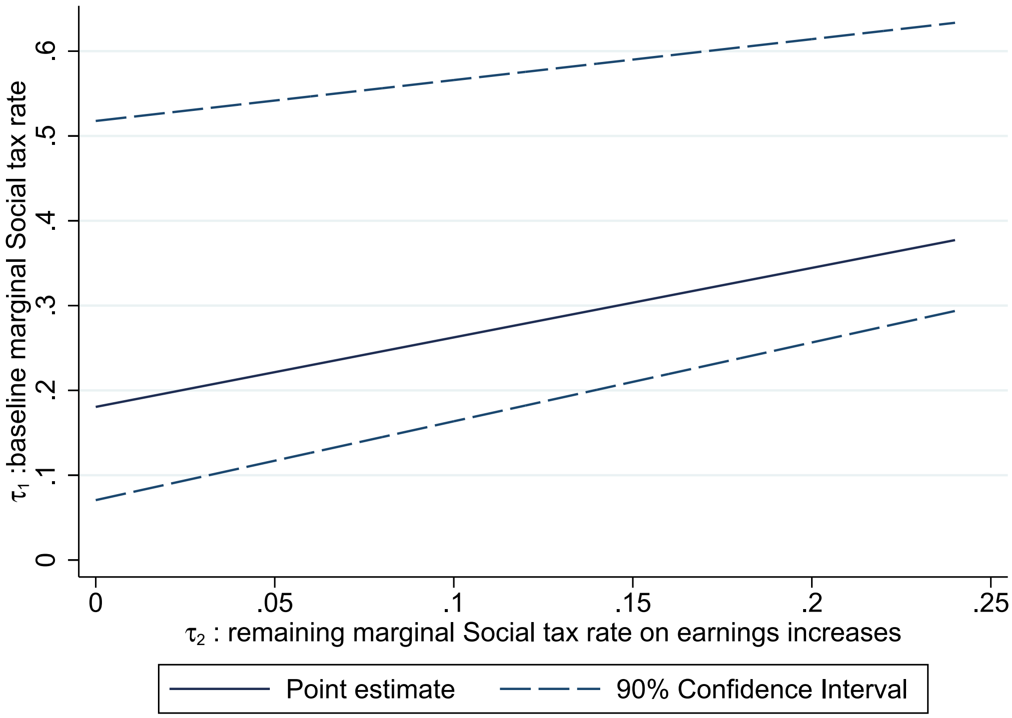

Notes: The figure presents the estimation of the baseline marginal Social tax rate τ_1 for various possible values of the marginal Social tax rate on earnings increases that remains when workers are offered an illiquid private savings account τ_2. It includes biased-corrected percentile intervals, based on 10,000 wild cluster bootstrap samples constructed using Rademacher random variables. The implementation follows the exposition from the Hansen 2020 textbook of this technique introduced by Efron (1982).

Table A1: Balance - Site 2

Variable	C. mean	T. mean	Diff.	p-value	N
Phase 1					
Baseline earnings	1965	1924.18	-40	.82	1365
Baseline attendance	.75	.72	-.03	.51	1365
Leaves before announce	.02	.11	.09	.05	112
Has valid ID	1	1	0	.	105
Speaks Dioula	.36	.46	.1	.32	105
Speaks Baoule	.4	.34	-.06	.53	105
Phase 2					
Baseline earnings	1871	1901	30	.79	1620
Baseline attendance	.69	.7	.01	.87	1620
Leaves before announce	.16	.11	-.05	.35	156
Has valid ID	1	1	0	.	135
Speaks Dioula	.26	.29	.03	.7	135
Speaks Baoule	.32	.37	.05	.57	135

Table A2: Take-up rates - All waves

Site	Wave	Dates	Treated workers	Compliance rate
Site 1	Wave 1	Jun 17 - Mar 18	51	.53
Site 1	Wave 2	Jul 18 - Mar 19	36	.47
Site 1	Wave 3	Nov 18 - Mar 19	44	.36
Site 2	Wave 1	Jul 18 - Mar 19	50	.36
Site 1	Phase 2 - Private	Apr 19 - Jul 19	83	.64
Site 1	Phase 2 - Non-Private	Apr 19 - Jul 19	99	.15
Site 2	Phase 2 - Private	Apr 19 - Jul 19	73	.56
Site 2	Phase 2 - Non-Private	Apr 19 - Jul 19	62	.13

Notes: This table presents, for each wave in each factory, the number of treated workers (excluding those who left the factory before the treatment announcement) and the rate of compliant workers (i.e. the share of treated workers with an open account).

Table A3: Take-up rates - Phase 2

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

Notes: This table disaggregates the Phase 2 take-up results. "Eligible Phase 1" denotes workers who have been eligible for at least one of the three waves of phase 1 (as opposed to workers newly eligible for phase 2). "Treated Phase 1" denotes workers who have been in the treatment group for at least one wave of Phase 1. "Compliant Phase 1" denotes workers who complied with the treatment assignment (i.e. opened an illiquid savings account with the BPCI) for at least one wave of Phase 1.

Table A4: Phase 2 Treatment effects on earnings and attendance – by Phase 1 treatment status

	(1)	(2)	(3)	(4)
Panel A: Daily earnings				
Private (vs. non-private)	227.9***	219.3**	281.2**	201.5*
	(60.39)	(88.08)	(130.1)	(104.8)
Non-private mean	1570	1539	1771	1452
Panel B: Attendance				
Private (vs. non-private)	0.0539**	0.0559	0.0503	0.0558
	(0.0250)	(0.0363)	(0.0523)	(0.0446)
Non-private mean	0.60	0.58	0.65	0.58
Phase 1 status	All	Not in study	Control	Private
N: worker-days	38222	14447	10018	13757
N: workers	317	120	83	114

Notes: Dependent variable in panel A: daily earnings. Dependent variable in panel B: attendance. Attendance equals 1 if the worker is present at the factory that day, 0 otherwise. DiD specification, with 15 days of baseline. All specifications include worker and paycycle FE. Standard errors clustered at the worker level. *, **, *** indicate significance at the 10, 5 or 1% level.

Table A5: Main results - By site

	All sites		**Site 1 only**		**Site 2 only**	
	Earnings	Attendance	Earnings	Attendance	Earnings	Attendance
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Private (vs. control)	176.407**	0.054*	168.898*	0.052	227.312*	0.069
	(82.246)	(0.032)	(91.085)	(0.036)	(132.149)	(0.046)
Control group mean	1586.65	0.67	1461.67	0.66	1871.86	0.69
N: worker-day	99215	99215	73280	73280	25935	25935
N: worker	353	353	248	248	105	105
Private (vs. non-private)	164.497**	0.073**	111.760	0.052	237.230**	0.103***
	(81.557)	(0.034)	(91.930)	(0.044)	(116.962)	(0.039)
Non-private mean	1570.27	0.60	1182.25	0.52	2195.58	0.72
N: worker-day	38222	38222	22022	22022	16200	16200
N: worker	317	317	182	182	135	135

Notes: DiD specification, with 15 days of baseline earnings. All regressions include worker and pay cycle FE. In (3)-(4), an interaction term (not displayed) between treatment status and "site 2" dummy is added; N and control group mean are for Site 1 only. In (5)-(6), an interaction term (not displayed) between treatment status and "site 2" dummy is added; N and control group mean are for factory 2 only. S.e. clustered at the worker level. *, **, *** indicate significance at the 10, 5 or 1% level.

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

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

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

Table A7: Baseline social tax rate, as a function of effort elasticity and post-intervention social tax rate

Effort elasticity	0.4	0.6	0.8	1.0
Endline tax rate				
0%	29%	20%	16%	13%
2%	30%	22%	17%	14%
4%	32%	23%	19%	16%
6%	33%	25%	21%	18%

Notes: This table presents the baseline social tax rate τ_1 estimated for various values of the social tax rate on earnings increases remaining after the intervention τ_2 and of the effort elasticity ζ^c. Note that our piece-rate variation exercise estimates the productivity elasticity to be $\zeta^c_p = 0.25$. Our estimate for the attendance elasticity is $\zeta^c_a = 0.44$ (so ζ^c=0.67). A reasonable lower bound is from Goldberg (2016) at $\zeta^c_a = 0.15$ (so ζ^c=0.5). Equation (3) is used for the computation of τ_1.

A.1 Demonstrations: tax rate, effort, and compensated elasticity

Baseline labor supply decision

At baseline, a worker solves $\max_{c,e} u(c,e)$ under the budget constraint $BC1 : (1-\tau_1)we+y = c$. Her optimal decision is thus ${e_1}^*\left((1-\tau_1)w, y\right)$. Let us denote the baseline choice made by the worker as e_1.

Labor supply decision under treatment

Our intervention modifies the budget constraint for workers in the treatment group. By dampening the social tax rate from τ_1 to τ_2 on earnings increases only (i.e., for all $e \geq e_1$), a kink is introduced in her budget constraint. Specifically, it now is:

$$BC2 : (1-\tau_1)we_1 + \mathbb{1}_{e\geq e_1}\left\{(1-\tau_2)w(e-e_1)\right\} + y = c$$

Note the trivial result: ${e_2}^* \geq e_1$, where ${e_2}^*$ is the worker's choice under $BC2$.[34] In what follows, we thus use the budget constraint $BC2^* : (1-\tau_1)we_1 + (1-\tau_2)w(e-e_1) + y = c$.

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

Slutsky equation.

Here, we derive Slutsky's equation applied to our model.
Let us define $\tilde{e}\left((1-\tau_2)w, u\right)$ as the Hicksian (compensated) supply of effort and $\gamma\left((1-\tau_2)w, u\right)$ as the expenditure function. By duality of utility maximization and expenditure minimization, we have:

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

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

$$\frac{\partial \tilde{e}\left((1-\tau_2)w, u\right)}{\partial\left[(1-\tau_2)w\right]} = \frac{\partial {e_2}^*\left((1-\tau_2)w, \mathbb{Y}\right)}{\partial\left[(1-\tau_2)w\right]} + \frac{\partial {e_2}^*\left((1-\tau_2)w, \mathbb{Y}\right)}{\partial \mathbb{Y}}\frac{\partial \gamma\left((1-\tau_2)w, u\right)}{\partial\left[(1-\tau_2)w\right]} \tag{4}$$

By Shephard's Lemma,

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

(Note the minus sign. It comes from the budget constraint being $\mathbb{Y} = c - (1-\tau_2)we$ and thus from effort as being a "bad" instead of a "good".)

[34] Proof by contradiction: Suppose that ${e_2}^* < e_1$. Then $BC2$ becomes $(1-\tau_1)we + y = c$, which is $BC1$. Since e_1 is the optimal choice under $BC1$, we must have ${e_2}^* = e_1$. This contradicts ${e_2}^* < e_1$.

Combined with the starting equality, we get:

$$\frac{\partial \tilde{e}\left((1-\tau_2)w,u\right)}{\partial\left[(1-\tau_2)w\right]}=\frac{\partial {e_2}^{*}((1-\tau_2)w,\mathbb{Y})}{\partial\left[(1-\tau_2)w\right]}-\frac{\partial {e_2}^{*}((1-\tau_2)w,\mathbb{Y})}{\partial \mathbb{Y}}{e_2}^{*}((1-\tau_2)w,\mathbb{Y}) \tag{5}$$

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

Income and substitution effects.

If $\tau_2=\tau_1$, then $\mathbb{Y}=y$ and ${e_2}^{*}\left((1-\tau_2)w,\mathbb{Y}\right)={e_1}^{*}\left((1-\tau_1)w,y\right)=e_1$.

Starting from this baseline situation, what happens when we dampen the tax rate τ_2 applied above the kink e_1?
In other words, we want to compute:

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

Applying Slutsky's equation, where $\tilde{e}\left((1-\tau_2)w,u\right)$ is the Hicksian (compensated) supply of effort,

$$\frac{d{e_2}^{*}\left((1-\tau_2)w,\mathbb{Y}\right)}{d\tau_2}=-w\frac{\partial \tilde{e}\left((1-\tau_2)w,u\right)}{\partial\left[(1-\tau_2)w\right]}+w(e_1-{e_2}^{*})\frac{\partial {e_2}^{*}\left((1-\tau_2)w,\mathbb{Y}\right)}{\partial \mathbb{Y}} \tag{6}$$

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

We thus have:

$$\frac{d{e_2}^{*}\left((1-\tau_2)w,\mathbb{Y}\right)}{d\tau_2}=-w\frac{\partial \tilde{e}\left((1-\tau_2)w,u\right)}{\partial\left[(1-\tau_2)w\right]} \tag{7}$$

There is no income effect; only a substitution effect.

We can further observe that $-\frac{d{e_2}^{*}((1-\tau_2)w,\mathbb{Y})}{d\tau_2}>0$, which corresponds to our prediction.

Taking Equation 7 to the data.

We have demonstrated that the effect of dampening the social tax rate on worker's effort increases, when starting from a baseline tax schedule with no kink, is given by Equation 7. Re-arranging it, by recognizing that ${e_2}^{*}\left((1-\tau_2)w,\mathbb{Y}\right)=\tilde{e}\left((1-\tau_2)w,u\right)$ (by definition) and that $-d\tau_2=d\left[(1-\tau_2)\right]$, we obtain:

$$
\begin{aligned}
(7) \Leftrightarrow \frac{d{e_2}^* \left((1-\tau_2)w, \mathbb{Y}\right)}{{e_2}^* \left((1-\tau_2)w, \mathbb{Y}\right)} &= \frac{\partial \tilde{e} \left((1-\tau_2)w, u\right)}{\partial (1-\tau_2)} \frac{(1-\tau_2)}{\tilde{e} \left((1-\tau_2)w, u\right)} \frac{d(1-\tau_2)}{(1-\tau_2)} \\
&\equiv \zeta^c \frac{d(1-\tau_2)}{(1-\tau_2)}
\end{aligned} \tag{8}
$$

Where ζ^c is the compensated elasticity of effort to the net-of-tax wage rate.

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

$$
\begin{aligned}
\log \left(\frac{{e_2}^*}{{e_1}^*} \right) &= \zeta^c \log \left(\frac{1-\tau_2}{1-\tau_1} \right) \\
\Leftrightarrow \frac{1-\tau_1}{1-\tau_2} &= \left(\frac{{e_1}^*}{{e_2}^*} \right)^{\frac{1}{\zeta^c}}
\end{aligned} \tag{9}
$$

We now have obtained the equation that we will use to estimate the social tax rate faced by workers at baseline.