Time vs. State in Insurance: Experimental Evidence from Contract Farming in Kenya*

Lorenzo Casaburi
University of Zurich

Jack Willis
Columbia University

September 28, 2017

Abstract

The gains from insurance arise from the transfer of income across states. Yet, by requiring that the premium be paid upfront, standard insurance products also transfer income across time. We show that this intertemporal transfer can help explain low insurance demand, especially among the poor, and in a randomized control trial in Kenya we test a crop insurance product which removes it. The product is interlinked with a contract farming scheme: as with other inputs, the buyer of the crop offers the insurance and deducts the premium from farmer revenues at harvest time. The take-up rate for pay-at-harvest insurance is 72%, compared to 5% for the standard pay-upfront contract, and the difference is largest among poorer farmers. Additional experiments and outcomes provide evidence on the role of liquidity constraints, present bias, and counterparty risk, and show that even a one month delay in premium payment increases demand by 21 percentage points.

*Lorenzo Casaburi, lorenzo.casaburi@econ.uzh.ch. Jack Willis, jack.j.willis@gmail.com. We wish to thank Vojta Bartos, Emily Breza, John Campbell, Michael Carter, Shawn Cole, Brian Dillon, Pascaline Dupas, Ram Fishman, Siddharth George, Ed Glaeser, Oliver Hart, Nathan Hendren, Cynthia Kinnan, Michael Kremer, David Laibson, Guilherme Lichand, Craig McIntosh, Mushfiq Mobarak, Sendhil Mullainathan, Nathan Nunn, Matthew Rabin, Gautam Rao, Martin Rotemberg, Frank Schilbach, Andrei Shleifer, Tavneet Suri, Alessandro Tarozzi, Argyris Tsiaras, Chris Udry, Andrew Weiss, and seminar audiences at ABCA conference, BASIS Technical Meeting, Barcelona Summer Forum, BREAD pre-conference, Cambridge, Columbia, Cornell (Dyson), CSAE Oxford conference, DIAL conference, ECBE conference, Harvard, IGC International Growth Week, IIES, INSEAD, Lindau meetings, NBER insurance meeting, NEUDC conference, NOVAFRICA conference, PACDEV conference, PSE, Princeton, Sciences Po, Stanford, TSE, UCL, and UCSD for useful comments. We are grateful to Noah Mambo, Mayara Silva and Rachel Steinacher for excellent research assistance and to the management of the contract farming company for their continued partnership. The project was funded by SIEPR, USAID BASIS and an anonymous donor, to whom we are very grateful. All errors are our own.

1 Introduction

In the textbook model of insurance, income is transferred across states of the world, from good states to bad. In practice, however, most insurance products also transfer income across time: the premium is paid upfront with certainty, and any payouts are made in the future, if a bad state occurs (Figure 1). As a result, the demand for insurance depends not just on risk aversion, but also on several additional factors, including liquidity constraints, intertemporal preferences, and trust. Since these same factors also make it harder to self-insure, charging the premium upfront may reduce demand for insurance precisely when the potential gains are largest, for example among the poor.[1]

This paper provides experimental evidence on the consequences of the transfer across time in insurance, by evaluating a crop insurance product which eliminates it. Crop insurance offers large potential welfare gains in developing countries, as farmers face risky incomes and have little savings to self-insure. Yet demand for crop insurance has remained persistently low, in spite of heavy subsidies, product innovation, and marketing campaigns (Cole and Xiong 2017). The transfer across time is a potential explanation. Farmers face highly cyclical incomes which they struggle to smooth across time, and insurance makes doing so harder: premiums are due at planting, when farmers are investing, while any payouts are made at harvest, when farmers receive their income.[2]

The insurance product we study eliminates the transfer across time by charging the premium at harvest, rather than upfront. We work in partnership with a Kenyan contract farming company, one of the largest agri-businesses in East Africa, which contracts small-holder farmers to grow sugarcane. As is standard in contract farming, the company provides inputs to the farmers on credit, deducting the costs from farmers' revenues at harvest time. We tie an insurance contract to the production contract and use the same mechanism to collect premium payments: the company offers the insurance product and deducts the premium (plus interest) at harvest.

Our first experiment shows that delaying the premium payment until harvest time results in a large increase in insurance demand. In the experiment we offered insurance to 605 of the farmers contracting with the company and randomized the timing of the premium payment.[3] Take-up of the standard, upfront insurance was 5%: low, but not out of line with results for other "actuarially-fair" insurance products in similar settings.[7] In contrast, when the premium was due at harvest

[1] For an example of the textbook model of insurance, see example 6.C.1 in Mas-Colell et al. (1995). Such purely cross-state insurance contracts do exist – examples include futures contracts and social security.

[2] Further, while farmers can often reduce their idiosyncratic income risk through informal risk sharing (Townsend 1994), similar mechanisms are less effective for reducing seasonal variation in income, since it is aggregate.

[3] The experiment was registered before baseline at the AEA RCT registry, ID AEARCTR-0000486, https://www.socialscienceregistry.org/trials/486

(including interest at 1% per month, the rate which the company charges on loans for inputs), take-up was 72%, substantially higher than results for other insurance products in similar settings. To benchmark this difference, in a third treatment arm we offered a 30% price discount on the upfront insurance premium. Take-up among this third group was 6%, not significantly different from take-up under the full-price upfront premium. Taken together, these results show that the farmers *do* have high demand for insurance, but they have a low willingness to pay for it upfront.

To help to identify the channels, we develop an intertemporal model of insurance demand, which shows that the transfer across time in insurance can help to explain why the poor demand so little of it. The model is based on a buffer-stock saving model (Deaton 1991) and includes a borrowing constraint, present-biased preferences, and imperfect contract enforcement. Liquidity constraints are central and play a dual role. First, they make paying the premium upfront more costly (if the borrowing constraint may bind, or almost bind, before harvest). Second, they make self-insurance (through consumption smoothing) harder, and thus increase the gains from risk reduction. As such, the transfer across time in insurance reduces demand precisely when the potential gains from insurance are largest – when liquidity constraints might bind. In the model, as in the real world, the poor are more susceptible to liquidity constraints, and thus are predicted to have both higher demand for pay-at-harvest insurance and a larger drop in demand when having to pay upfront. Heterogeneous treatment effects in the main experiment show that both predictions hold, both for the poor and for the liquidity constrained.

Two additional mechanism experiments (Ludwig et al. 2011) provide further evidence on channels. In the first, we test the most common reason farmers gave for not buying pay-upfront insurance in the main experiment: they did not have the cash. To test this, in this experiment we gave a subset of farmers cash, before offering them insurance later in the same meeting (similar to Cole et al. 2013a). The cash gift, being slightly larger than the cost of the premium, ensured that farmers did have money to purchase the insurance if they wished to. However, as acknowledged by Cole et al. (2013a), it may also have induced reciprocity. To address this, we cross-cut the cash treatment with a pay-upfront vs pay-at-harvest treatment. The difference-in-differences effect of the cash was 8%, small and not significant, showing that pay-upfront insurance was not the marginal expenditure, and ruling out the explanation that farmers simply did not have cash.[4] However, because the cash gift may not have completely removed liquidity constraints (farmers may still have wanted to borrow), two possible explanations remain: farmers were not liquidity constrained to begin with, or farmers were very liquidity constrained and hence had other uses for

[4]Barring reciprocity, if anything we would expect demand for pay-at-harvest insurance to fall slightly with a cash drop, hence the diff-in-diffs should be an upper bound on the effect on upfront take-up, net of reciprocity.

cash.

The second experiment on channels helps separate these two explanations. We consider the effect of a small delay in the premium payment, such that payment is not due immediately at sign-up.[5] The design targets present bias, of particular interest as it has implications for optimal contract design and for welfare interpretations. In the experiment, we compare insurance take-up across two groups. In both groups, during the visit, farmers had to choose between a cash payment, equal to the insurance premium, and free enrollment in the insurance. Farmers in the first group were told they would receive their choice immediately, whereas farmers in the second group were told they would receive their choice in one month's time.[6] Delaying delivery this way, by just one month, increased insurance take-up by 21 percentage points. We argue that the size of this effect is inconsistent with standard exponential discounting - if the discount rate was high, then why buy insurance in one month, given that it is a transfer across a subsequent eleven months – and that it is evidence of present bias (combined with liquidity constraints).

The final channel we consider is imperfect contract enforcement. If either party defaults before harvest time, then the farmer does not pay the premium at harvest, whereas the upfront premium is sunk. Tying the contracts together means that, for the farmer, defaulting on the insurance requires defaulting on the sales contract (side-selling), and vice versa. This has two implications. First, it reduces strategic default to avoid paying the harvest-time premium, the natural concern with removing the transfer across time, since farmers typically value the production relationship. In keeping with this, there was no significant difference in side-selling, or in yield conditional on not side-selling, across pay-upfront and pay-at-harvest treatment groups. Second, however, it can induce default: if the farmer defaults on the *sales* contract (for reasons unrelated to the insurance), he automatically defaults on the *insurance* contract. In our setting, before harvest, the company faced severe financial difficulties and temporarily shut-down their factory, resulting in long delays and uncertainty in harvesting. Because of this, twelve months after our experiment began, there was widespread side-selling: 52% of farmers side-sold or uprooted their crop, compared to a historical rate of less than 10%.

In spite of the large default rates ex-post, three arguments suggest that, ex-ante, any differential effect on take-up by the timing of the premium was limited. First, while survey measures of trust in the company are correlated with overall insurance take-up, their interactions with the timing of the premium are not, suggesting that the company defaulting on insurance payouts *after* harvesting

[5]Such delays have been shown to increase savings in other settings, such as Save More Tomorrow programs (Thaler and Benartzi 2004).

[6]Giving farmers the choice between the premium and insurance for free, rather than the choice of whether to buy insurance, eased liquidity constraints and enabled us to enforce payment in the one month treatment.

was more of a concern than potential side-selling. Second, assuming ex-ante expectations of side-selling are predictive of actual side-selling, then the correlation between take-up and actual side-selling should vary by premium timing. For both individual and local average side-selling, it does not. Third, using our model, we bound the differential effect by that of a proportional price cut on the take-up of upfront insurance – in particular, a price cut of the expected probability of side-selling, times the relative (expected) marginal utility of consumption conditional on side-selling, has a larger effect. But the main experiment showed that demand for upfront insurance is inelastic, so, to be important, one of these two terms would have had to be large; other results suggest they were not.

This paper adds to several strands of literature. First, many papers have investigated the demand for agricultural insurance and the factors which constrain it (Cole et al. 2013a; Karlan et al. 2014). Demand is generally found to be low, and interventions to increase it typically have small effects in percentage-point terms.[7] Many of the proposed explanations, such as risk preferences and basis risk (Mobarak and Rosenzweig 2012; Elabed et al. 2013; Clarke 2016), concern the transfer across states in insurance; we focus on the transfer across time.[8] Several studies have bundled insurance with credit (Gine and Yang 2009; Carter et al. 2011; Karlan et al. 2014; Banerjee et al. 2014; Ahmed et al. 2017), finding that take-up of credit increases little, and in some cases decreases. We effectively do the reverse, bundling credit with insurance. The closest paper to ours, Liu et al. (2016), finds that, for livestock mortality insurance in China, delaying premium payment increases take-up from 5% to 16%; Liu and Myers (2016) considers the theoretical implications. As far as we know, our paper is the first to provide experimental evidence on the effect of completely removing the intertemporal transfer from insurance contracts, and on the role of liquidity constraints, present bias and other channels. Additionally, we show theoretically and empirically, that the transfer across time is most costly for the poor, providing a potential explanation for their low insurance demand.[9]

Second, the transfer across time in insurance is studied implicitly in finance, but the focus is

[7] Karlan et al. (2014) find the highest take-up rates at actuarially fair prices among these studies, at around 40%; but most find significantly lower rates, for example, at around 50% of actuarially fair prices, Cole et al. (2013a) find 20-30% take-up, and at commercial price Mobarak and Rosenzweig (2012) find 15% take-up.

[8] The intertemporal transfer means a long line of work on investment decisions and financial market imperfections in developing country settings is also relevant (Rosenzweig and Wolpin 1993, Conning and Udry 2007). In particular, we add to evidence on liquidity constraints (Cohen and Dupas 2010) and present bias (Duflo et al. 2011) in similar settings.

[9] We add two further contributions relative to these existing papers. First, we work in a setting where contract enforcement is challenging, and consider a novel way to potentially improve it: tying the insurance contract to a production contract. This is important, since it is exactly in such settings where credit markets are likely to be inefficient, and hence paying the premium upfront will be costly. Second, we work with crop insurance, where seasonality increases the importance of the transfer across time.

on how insurance companies benefit by investing the premiums (Becker and Ivashina 2015), rather than on the cost for consumers, our focus. A recent exception is a largely theoretical literature (Rampini and Viswanathan 2010; Rampini et al. 2014) which argues that firms face a trade-off between financing and insurance. Rampini and Viswanathan (2016) apply similar reasoning to households.[10] These papers are part of a wide literature on how imperfect enforcement affects the set of financial contracts which exists (Bulow and Rogoff 1989; Ligon et al. 2002), to which we add by considering the implications of imperfect enforcement for the timing of insurance premiums.

Finally, our paper adds to a literature on the importance of interlinked contracts, i.e. contracts covering multiple markets, in developing country settings.[11] In particular, our work relates to research documenting informal insurance agreements in output and credit market contracts (Udry 1994; Minten et al. 2011), and to a recent line of empirical research on the emergence and impact of interlinked transactions (Macchiavello and Morjaria 2014, 2015; Blouin and Macchiavello 2017; Casaburi and Macchiavello 2016; Casaburi and Reed 2017; Ghani and Reed 2017).

The remainder of the paper is organized as follows. Section 2 describes the setting in which the experiment took place and discusses how tying an insurance contract to a production contract can affect enforcement. Section 3 presents the main experimental design and results. Section 4 develops an intertemporal model of insurance demand, which provides comparative statics and directs subsequent experiments. Section 5 presents evidence on channels, from the main experiment and from two additional experiments. Section 6 discusses the policy implications of our results, both for crop insurance and for insurance markets more generally, and presents ideas for future work. Finally, Section 7 concludes.

2 Setting, contract farming, and interlinked insurance

We work in Western Kenya with small-holder sugarcane farmers. Sugarcane is the main cash crop in the region, accounting for more than a quarter of total income for 80% of farmers in our sample. It has a long growing cycle (around sixteen months), leading to a long transfer across time in pay-upfront insurance, and it is not seasonal. Once planted, crops last upwards of three growing cycles; the first cycle, called the plant cycle, involves higher input costs and hence lower profits than the subsequent cycles, called the ratoon cycles. Crop failure is rare, but yields are subject to significant risks from rainfall, climate, pests and cane fires. Sugarcane farmers are typically poor, but not the poorest in the region: among our sample, 80% own at least one cow, the average total

[10] They show that limited liability results in poorer households facing greater income risk in equilibrium, even with a full set of state-contingent assets.

[11] See Bardhan (1980), Bardhan (1989), and Bell (1988) for summaries of this literature.

cultivated land is 2.9 acres, and the average sugarcane plot is 0.8 acres. Very few farmers in the study area have had experience with formal insurance.

2.1 Contract farming

We work in partnership with a contract farming company which has been working in the area since the 1970s. It is one of the largest agri-businesses in East Africa and contracts around 80,000 small-holder farmers. As is standard in contract farming - a production form of increasing prevalence (UNCTAD 2009) - farmers sign a contract with the company, at planting, which guarantees them a market and binds them to sell to the company, at harvest. The contract covers the life of the cane seed, meaning multiple harvests over at least four years. Each harvest, company contractors do the harvesting and transport the cane to the company factory, after which farmers are paid by weight, at a price set by the Kenyan Sugar Board (the regulatory body of the national sugar industry).

Interlinked credit A major benefit of contract farming is that buyers can supply productive inputs to farmers on credit, and then recover these input loans through deductions from harvest revenues. Such practice, often referred to as interlinking credit and production markets, is widespread. Our partner company provides numerous inputs in this way, such as land preparation, seedcane, fertilizer, and harvesting services, and charges 1% per month interest on the loans.[12]

Contract enforcement The supply of loans by the buyer raises the issue of contract enforcement, which will be important for considering insurance demand. In our setting, as is common in developing countries, the company must rely on self-enforcement of the contract - while illegal, farmers may side-sell (i.e. sell to another buyer, breaking the contract) with little risk of prosecution. By side-selling, farmers avoid repaying the input loan,[13] and are paid immediately upon harvesting, but are typically paid a significantly lower price for their cane (both because side-selling is illegal, and because sugarcane is a bulky crop, so that transport costs to other factories are high). While the company cannot directly penalize farmers for side-selling, it does collect any dues owed (plus interest) if the same plot is re-contracted in the future, or from other plots if the farmer contracts multiple plots.[14] Our administrative data does not tell us historical levels

[12] Inflation in Kenya was around 6% per annum during the study, so the real interest rate on input loans from the company was 6% per annum.

[13] Macchiavello and Morjaria (2014) show that, in the context of coffee in Rwanda, higher competition reduces input loans potentially for this reason.

[14] Debts remain on plots even if plots are sold, and are collected from future revenues regardless of who farms the plot. When we ran our experiment, debt collection was limited to the plot level: if a farmer defaulted on a loan on one plot, the company would not recover that loan from revenues from other plots farmed by the farmer. However the company changed this policy before harvest time for our farmers, so that defaulted loans on one plot could be recovered from other plots of the same farmer.

of side-selling, but does allow us to bound them. In the three years before the experiment, an average of 12% of plots which harvested in ratoon 1 did not harvest in ratoon 2 - an upper bound on side-selling / default, because it includes cases where farmers uproot the crop before inputs are applied (for example because of crop disease or poor yields). We could not ask farmers detailed questions about side-selling because it is illegal.

The company's main obligation under the contract is to harvest and purchase farmers' cane at a price set by the Kenyan Sugar Board. Farmers are well represented politically in the region, so serious breaches of the contract by the company are unlikely under normal circumstances. However, were the company to become insolvent, it would be unable to purchase the cane, in which case farmers may be forced sell to another buyer. This happened, temporarily, 12 months after the start of our experiment, affecting some of the farmers in our sample. In Section 5.4 we discuss in detail the implications for the interpretation of our results - to summarize, across multiple tests we find no evidence that ex-ante anticipation of this episode affected our main results, and we bound the size of the role it could have played.

Administration How the company coordinates with its farmers has two implications for our study. First, the company employs outreach workers to visit farmers in their homes and to monitor plots. These outreach workers market the insurance product we introduce, as detailed in the next section. Second, because of fixed costs in input provision, the outreach workers group neighboring plots into administrative units, called fields, which are provided inputs and harvested concurrently. As detailed in Section 3, we stratify treatment assignment at the field level in our experiments. Fields typically contain three to ten plots.

2.2 Interlinked insurance

In standard insurance contracts farmers pay the premium upfront and so bare all of the contract risk.[15] Pay-at-harvest insurance reduces the contract risk they face, as they do not pay the premium if the insurance company defaults before harvest time. However, it places significant contract risk on the insurer: the risk that farmers do not pay premiums when harvests are good. In contract farming settings, this risk may be reduced by using the same mechanism used to enforce repayment of input loans: the buyer can provide the insurance, and charge the premium as a deduction from farmers' harvest revenues.

Tying together the insurance and production contracts in this way, which we refer to as interlinking, will typically help enforce harvest-time premiums by increasing the cost to farmers of

[15]Consistent with this, trust has been shown to be an important issue in shaping insurance take-up in other settings (Dercon et al. 2011, Cole et al. 2013a).

defaulting on them. In an interlinked contract, the only way farmers can default on premiums is by defaulting on the sale contract. Doing so compromises all the gains from the relationship with the buyer, including the current and future purchase guarantees and future input provision. However, interlinking can also encourage default on the insurance contract, if a farmer wants to side-sell for some other reason (although, under the assumption that increasing the premium does not increase such side-selling, it can be priced into the insurance contract). In Section 4.3.1 we consider the question of contract enforcement theoretically.

Interlinking pay-at-harvest insurance with the production contract could increase side-selling in the latter, but there are two reasons to believe that this effect will be minimal in our setting. First, the insurance premium is small, and typically much smaller than the pre-existing input loans. Thus it is unlikely to be marginal in the strategic decision to default (a comparison between the static benefits of defaulting and the continuation value of the relationship).[16] Second, given the insurance design (detailed in the next section), the farmer has limited information about his likely payout when he has to decide whether to side-sell. In line with these arguments, Section 5.4 reports that interlinked insurance did not increase side-selling.

Finally, we note that in contract farming, since many of the inputs are provided by the company, the scope for insurance to affect productivity is reduced. In our setting farmers' only inputs are the use of their land and their labor for planting, weeding, and protecting the crop. This is a double-edged sword: insurance is less likely both to induce moral hazard, which would lower productivity, and to enable risky investments (Karlan et al. 2014), which would increase productivity.

3 Does the transfer across time affect insurance demand?

This section describes the main experiment of the paper, in which we compare take-up for insurance when the premium is paid upfront to take-up when the premium is paid at harvest time, thus removing the intertemporal transfer. Changing the timing of the premium increases take-up by 67 percentage points.

3.1 Experimental design

Treatment groups The experimental design randomized 605 farmers across three treatment groups (Figure 2). In all three treatments farmers were offered the same insurance product, described below; the only thing varied was the premium. In the first group (U1), farmers were offered the insurance product and had to pay the premium upfront, at "full price" (which across

[16] Further, if farmers value access to insurance in future years, insurance increases the continuation value of the relationship, and hence could actually reduce side-selling.

the study spanned between 85% and 100% of the actuarially fair price). In the second group (U2), premium payment was again upfront, but farmers received a 30% discount relative to the full price. In the third group (H), farmers could subscribe for the insurance and have the (full-price) premium deducted from their revenues at harvest time, including interest (charged at the same rate used for the inputs the company supplies on credit, 1% per month).[17] Randomization was at the farmer level and was stratified by field.

Insurance design The insurance was offered by the company and the payout design was the same across all experimental treatments. There was no intensive margin of insurance and farmers could only subscribe for their entire plot, not parts of it. The insurance had a double-trigger area yield design, preferred to a standard rainfall insurance because it had lower basis risk.[18] Under the design, a farmer received a payout if two conditions were met: first, if their plot yield was below 90% of its predicted level; and second, if the average yield in their field was below 90% of its predicted level. The design borrows from studies which used similar double-trigger products in other settings (Elabed et al. 2013), and its development relied on rich plot-level administrative panel data for predictions, simulations, and costing.[19] The product was very much a partial insurance product: in the states where payouts were triggered, it covered half of plot losses below the 90% trigger, up to a cap of 20% of predicted output. Finally, farmers would only receive any insurance payouts if they harvested with the company, as agreed under the production contract.

Insurance marketing The insurance was offered by company outreach officers during visits to the farmers. To reduce the risk of selecting farmers by their interest in insurance, the specific purpose of the visits was not announced in advance. 937 farmers were targeted, 638 (68%) of whom attended; the primary reason (75%) for non-attendance was that the farmers were busy somewhere far from the meeting location. To ensure that our sample consisted of farmers who were able to understand the insurance product, in an initial meeting outreach officers checked that target farmers mastered very basic related concepts (e.g. the concepts of tonnage and acres). A small number of farmers (5%), typically elderly, were deemed non-eligible at this stage. The resulting sample for randomization comprised 605 farmers. Compared to the 333 who did not

[17] The interest was added to the initial premium when marketing the insurance product.

[18] The company collects rainfall data through stations scattered across their catchment area. However, data quality is a concern and its predictive power is low.

[19] The data included production, plot size, plot location, and crop cycle, and was available for a subsample of contracted plots for 1985-2006 and for all contracted plots from 2008 onwards. The data was used to compute predicted yields at the plot and area level, which were needed for the double trigger insurance design. The historical data was also used to simulate past payouts and hence price the product, and to run simulations of alternative prediction models. Under the simulations the double-trigger product performed well on basis risk (Figure A.1) - 74% of farmers who would receive a payout with a single-trigger insurance continued to do so when the second area-level trigger was added – substantially better than an alternative based on rainfall.

enter the sample, they had slightly larger plots (0.81 vs. 0.75 acres; p-value=0.015) and similar yields (22.2 vs. 21.8 tons per acre; p-value=0.40).

After the initial meeting, the outreach officers explained the product in detail in one-to-one meetings with farmers, using plot-specific visual aids to describe the insurance triggers and payout scenarios. To ensure that farmers correctly understood the insurance product before being offered it, outreach officers verified that they could first answer basic questions about the product, e.g. the scenarios under which it would pay out, and would re-explain if not. Farmers then had three to five business days to subscribe, with premiums collected either immediately or during revisits at the end of this period.[20]

Sample selection Numerous farmer and plot criteria were used to select the sample, both to increase power and to improve the functioning of the insurance.[21] For example, the experiment targeted plots in the early stages of the ratoon cycles (in particular the first and second ratoons), i.e. plots which had already harvested at least once. This choice was made because the yield prediction model performs better for ratoon than for plant cycles.

Data collection We combine two sources of data for the analysis: survey data and administrative data. Our survey data comes from a short baseline survey, carried out by our survey team (before farmers were offered insurance) during the outreach-worker visits described;[22] 32 of 605 the farmers declined to be surveyed. As mentioned in section 2, the company keeps administrative data on all farmers in the scheme. It gives us previous yields, harvest dates, plot size, and growing cycle, and enables us to track whether the farmer sells cane to the company at the end of the cycle, and their yield conditional on doing so.

3.2 Balance

Table 1 provides descriptive statistics for the three treatment groups and balance tests. Since stratification occurred at the field level, we report p-values for the differences across the groups from regressions that include field fixed effects. Consistent with the specification we use for some of our analysis (and our pre-analysis plan) we also report p-values when bundling pay-upfront

[20] In practice, for a share of these farmers, revisits occurred one to two weeks after the first visit.

[21] The criteria used to select the sample were: plot size - large plots were removed from the sample, to minimize the insurer's financial exposure; plot yields - outliers were excluded, to improve the prediction of yield for the insurance contract; the number of plots in the field - fields with fewer than five plots were excluded, to improve power given the stratified design; the number of plots per farmer - the few farmers with multiple plots were only eligible for insurance for their smallest plot in the field; the number of farmers per plot - plots owned by multiple farmers were excluded; finally, while contracted farmers are usually subsistence farmers, some plots are owned by "telephone farmers" who live far away and manage their plots remotely - such plots are excluded from our sample.

[22] Several months later we also followed up with a subset of the farmers by phone, to ask whether they remembered the terms of the insurance and whether they regretted their take-up decision, as discussed below.

treatments U1 and U2 and comparing them to pay-at-harvest treatment H. The table shows that the randomization achieved balance across most observed covariates; only age is significantly different when comparing the bundled upfront group U to H. We confirm below that the experiment results are robust to the inclusion of baseline controls.

3.3 Experimental results

Our main outcome of interest is insurance take-up. Take-up rates have been consistently low across a wide range of geographical settings and insurance designs (Cole et al. 2013a, Elabed et al. 2013, Mobarak and Rosenzweig 2012). Yet gains from insurance could be large, both directly and indirectly - farmers are subject to substantial income risk from which they are unable to shield consumption, and previous studies suggest that when farmers are offered agricultural insurance they increase their investment levels (Karlan et al. 2014, Cole et al. 2013b). The central hypothesis tested in this paper is that low take-up is in part be due to the intertemporal transfer in insurance, which differentiates standard insurance products from their purely intratemporal ideal.

The regression model we use compares the binary indicator for insurance take-up – T_{if}, defined for farmer i in field f – across the three treatment groups, controlling for field fixed effects:

$$T_{if} = \alpha + \beta Discount_i + \gamma Harvest_i + \eta_f + \epsilon_{if} \tag{1}$$

Figure 3 summarizes the take-up rates across the three treatment groups. For groups U2 and H, it also includes 95% confidence intervals for the difference in take-up with U1, obtained from a regression of take-up on treatment dummies.

The first result is that take-up of the full-price, upfront premium is low, at 5%. While low, this finding is consistent with numerous existing crop insurance studies mentioned above. It shows that, in this setting, reducing basis risk (the risk that insurance does not pay out when farmers have bad yields – one of the proposed explanations for low demand for rainfall insurance) by using an area yield double-trigger design is alone not enough to raise adoption.

The second result (the main result of the paper) is that delaying the premium payment until harvest, thus removing the transfer across time, has a large effect on take-up. Take-up of the pay-at-harvest, interlinked insurance contract (H) is 72%, a 67 percentage point increase from the baseline, pay-upfront (U1) level, and one of the highest take-up rates observed for actuarially fair crop insurance. The result shows that, in our setting, farmers *do* want risk reduction, they just do not want to pay for it upfront.

The third result, which allows us to benchmark the importance of the second, is that a 30% price cut in the upfront premium has no statistically significant effect on take-up rates. The effect's

point estimate is 1 percentage point, and even at the upper bound of the confidence interval, take-up only increases by 8 percentage points. While this upper bound is consistent with substantial price elasticity of demand (given the low baseline take-up) it suggests that medium-sized subsidies have limited scope to increase demand in absolute terms.

Table 2 presents regression analysis of these treatment effects, and shows that they remain stable across a variety of specifications. Column (1) is the basic specification, which includes fixed effects at the field level, the stratification unit. As in Figure 3, the pay-at-harvest product (H) has 67 percentage points higher take-up than the "full-price" pay-upfront product (U1), significant at the 1% level, whereas the 30% price cut product (U2) has just 0.4 percentage point higher take-up, far from significant. The difference between the pay-at-harvest (H) and the reduced price pay-upfront (U2) products is also significant at the 1% level. Column (2) pools the upfront treatments U1 and U2, consistent with the specification we use later in the heterogeneity analysis. Columns (3) and (4) further add controls for plot and farmer characteristics, respectively, and column (5) includes both types of controls.

Farmer understanding One key question for the interpretation of the high take-up rate is whether farmers understood what they were signing up for. There are two reasons to believe they did. First, as mentioned above, farmers were asked questions to test their understanding of the product before it was offered to them. Second, several months after the recruitment, we called back 76 farmers who had signed up for the pay-at-harvest insurance, in two waves. In the first wave of 40 farmers, we began by reminding the farmers of the terms of the insurance product (the deductible premium and the double trigger design) and then checked that the terms were what the farmers had understood when originally visited. All farmers said they were. We then asked the farmers if they would sign up again for the product if offered next season. 32 (80%) said they would while 3 (7.5%) said they would not. The remaining 5 (12.5%) stated that their choice would depend on the outcome of the current cycle. In the second wave of 36 farmers, we did not prompt the farmers about the insurance terms, but instead asked farmers to explain them to us. 25 (69%) were able to do so. Of this second wave of farmers, after reminding those who had forgotten the terms, 28 (85%) said they would sign up for the product if offered next season.

To summarize, the results in this section show that pay-at-harvest insurance, enabled by inter-linking product and insurance markets, has high take-up at actuarially fair price levels, while its standard, pay-upfront equivalent has low take-up (even with a substantial price cut), consistent with experience in other settings.

4 An intertemporal model of insurance demand

To understand the forces behind the experiment results, we develop a model which captures both the cross-state and cross-time transfers in insurance. We begin by setting up a background intertemporal model, without insurance, into which we then introduce the insurance products.[23] We first consider the case where contracts are perfectly enforceable, and then allow for imperfect enforcement. The model shows how the channels interact to affect insurance demand (and for whom) and motivates our subsequent experiments and empirical tests to identify them. Proofs and derivations are in the appendix.

4.1 Background

The background model is a buffer-stock savings model, as in Deaton (1991), with the addition of present-biased preferences and cyclical income flows (representing agricultural seasonality).

Time and state We use a stochastic discrete-time, infinite horizon model. The probability distribution over states is assumed to be memoryless and cyclical.

Utility Individuals have time-separable preferences and maximize present-biased expected utility $u(c_t) + \beta\Sigma_{i=1}^{\infty}\delta^i\mathbb{E}[u(c_{t+i})]$ as in Laibson (1997).[24] We assume that $u(.)$ satisfies $u' > 0$, $u'' < 0$, $\lim_{c\to 0} u'(c) = \infty$ and $u''' > 0$, and that $\beta \in (0, 1]$ and $\delta \in (0, 1)$.[25]

Intertemporal transfers Households have access to a risk-free asset with constant rate of return R and are subject to a borrowing constraint. As in Deaton (1991), we assume $R\delta < 1$.

Income and wealth Households have state-dependent income in each period y_t. We assume $y_t > 0\ \forall t \in \mathbb{R}^+$.[26] We denote cash-on-hand once income is received by x_t.

[23] An alternative approach is to use observed investment behavior (in particular the potential returns of risk-free investments which farmers make or forgo) as a sufficient statistic for the cost of the transfer across time. In appendix section A.2 we report basic quantitative bounds for the effect of the transfer across time on insurance demand using this approach.

[24] We note that time-separable preferences equate the elasticity of intertemporal substitution, ψ, and the inverse of the coefficient of relative risk aversion, $\frac{1}{\gamma}$. As such they imply a tight link between preferences over risk and consumption smoothing, both of which are relevant for insurance demand. Recursive preferences allow them to differ (Epstein and Zin 1989), which would provide an additional channel: if $\psi \ll \frac{1}{\gamma}$, then demand for upfront and at-harvest insurance may differ greatly, since the cost of variation in consumption over time would far exceed that of variation across state.

[25] We assume prudence, i.e. $u''' > 0$, as is common in the precautionary savings literature (and as holds for CRRA utility), to ensure that the value of risk reduction is decreasing in wealth, i.e. Lemma 1, part 3. Liquidity constraints strengthen concavity of the value function, and thus the result, but our proof requires prudence.

[26] As a technical assumption we actually assume that y_t is strictly bounded above zero $\forall t$.

Household's problem The household faces the following maximization sequence problem in period t:

$$\begin{aligned} \max_{(c_{t+i})_{i\geq 0}} \quad & u(c_t) + \beta \mathbb{E}[\Sigma_{i=1}^{\infty} \delta^i u(c_{t+i})] \\ \text{s.t. } \forall i \geq 0 \quad & x_{t+i+1} = R(x_{t+i} - c_{t+i}) + y_{t+i+1} \\ & x_{t+i} - c_{t+i} \geq 0 \end{aligned} \tag{2}$$

Denote the value function of the household by V_t, a function of one state variable, cash-on-hand x_t.[27] We assume that households are naive-$\beta\delta$ discounters: they believe that they will be exponential discounters in future periods (and so may have incorrect beliefs about future consumption). There is evidence for such naivete in other settings (DellaVigna and Malmendier 2006) and, with the exception of Proposition 2, all propositions hold with slight modification in the sophisticated-$\beta\delta$ case.[28]

Iterated Euler equation To consider the importance of the timing of premium payment, we will compare the marginal utility of consumption across time periods using the Euler equation:

$$u'(c_t) = \max\{\beta\delta R \mathbb{E}[u'(c_{t+1})], u'(x_t)\} \tag{3}$$

$$= \beta\delta R \mathbb{E}[u'(c_{t+1})] + \mu_t \tag{4}$$

where $\mu_t(x_t)$ is the Lagrange multiplier on the borrowing constraint, and c_{t+1} is period t self's belief about consumption in period $t+1$. Iterating the Euler equation to span more periods gives:

$$u'(c_t) = \beta(R\delta)^H \mathbb{E}[u'(c_{t+H})] + \lambda_t^{t+H} \tag{5}$$

where $\lambda_t^{t+H}(x_t)$ represents distortions in transfers from t to $t+H$ arising from (potential) borrowing constraints:

$$\lambda_t^{t+H} \coloneqq \mu_t + \beta \mathbb{E}[\Sigma_{i=1}^{H-1} (R\delta)^i \mu_{t+i}] \tag{6}$$

The setup provides the following result, which we will use when considering insurance demand.

Lemma 1. *$\forall t \in \mathbb{R}^+$:*

1. *$\frac{d^3 V_t}{dx_t^3}, \frac{d^3 V_t^c}{dx_t^3} > 0$, so the value of risk reduction is decreasing in wealth.*

[27] Since preferences are not time-consistent, V_t is different from the continuation value function, denoted V_t^c, which is the value function at time t, given time $t-1$ self's intertemporal preferences, i.e. without present bias.

[28] The required modification is replacing β by a state-specific discount factor, which is a function of the marginal propensity to consume. Proposition 2 and Lemma 1 may no longer hold, since concavity and uniqueness of the continuation value V_t^c is no longer guaranteed, complicating matters significantly.

2. $\frac{d\lambda_t^{t+H}}{dx_t} < 0$, *i.e. the distortion arising from liquidity constraints is decreasing in wealth.*

The intuition behind part 1 of the lemma is as follows. The value of risk reduction depends on how much the marginal utility of consumption varies across states of the world. Two things dictate this. First, how much marginal utility varies for a given change in consumption; this drives the comparative static through prudence (i.e. $u''' > 0$). Second, how much consumption varies for a given change in wealth (the marginal propensity to consume). Concavity of the consumption function, another consequence of prudence (Carroll and Kimball 1996), but further strengthened by the borrowing constraint (Zeldes 1989; Carroll and Kimball 2005), reinforces the result.[29]

4.2 Insurance with perfect enforcement

We begin with the case where insurance contracts are perfectly enforceable.

Timing The decision to take up insurance is made in period 0. Any insurance payout is made in period H, the harvest period.

Payouts Farmers can buy one unit of the insurance, which gives state-dependent payout I in period H, normalized so that $\mathbb{E}[I] = 1$. We assume that $y_H + I - 1$ second-order stochastically dominates y_H.[30]

Premiums We consider two timings for premium payment: upfront, at time 0, and at harvest, at time H. If paid at harvest the premium is 1, the expected payout (commonly referred to as the actuarially-fair price). If paid upfront, the premium is R^{-H}. Thus, at interest rate R, upfront and at-harvest payment are equivalent in net present value.

Demand for insurance Farmers buy insurance if the expected benefit of the payout is greater than the expected cost of the premium. Thus, to first order,[31] the take-up decisions are:

$$\text{Take up insurance iff } \begin{cases} \beta\delta^H \mathbb{E}[u'(c_H)] \leq \beta\delta^H \mathbb{E}[Iu'(c_H)] & \text{(pay-at-harvest insurance)} \\ R^{-H}u'(c_0) \leq \beta\delta^H \mathbb{E}[Iu'(c_H)] & \text{(pay-upfront insurance)} \end{cases} \tag{7}$$

For pay-at-harvest insurance, the decision is based on a comparison of the marginal utility of consumption across states (when insurance pays out vs. when it does not). For pay-upfront

[29] Mathematically, the value of a marginal transfer from state $x + \Delta$ to state x, assuming both equally likely, is (one-half times) $V'(x + \Delta) - V'(x) = u'(c(x + \Delta)) - u'(c(x)) \simeq u''(c(x))c'(x)\Delta$. Its derivative w.r.t. x is $\Delta(u'''(c(x))c'(x)^2 + u''(c(x))c''(x))$, which shows the role of both u''' and c''.

[30] Historical simulations using administrative data suggest this assumption is reasonable in our setting. While the second, area-yield based trigger, does lead to basis risk in the insurance product, it only prevents payouts in 26% of cases receiving payouts under the single trigger, as shown in Figure A.1.

[31] We use first order approximations at several points. They are reasonable in our setting for several reasons: the premium is small (3% of average revenues) and the insurance provides low coverage (a maximum payout of 20% of expected revenue); we care about differential take-up by premium timing, so second order effects which affect upfront and at-harvest insurance equally do not matter; both the double trigger insurance design, and the provision of inputs by the company, meant insurance was unlikely to affect input provision, in line with results in section 5.4.

insurance, in contrast, the decision in based on a comparison across both states and time (when insurance pays out in the future vs. today). To relate the two decisions, we use the iterated Euler equation, equation 5, which gives the following.

Proposition 1. *If farmers face a positive probability of being liquidity constrained before harvest, they prefer pay-at-harvest insurance to pay-upfront insurance; otherwise they are indifferent.*[32] *To first order, the difference is equivalent to a proportional price cut in the upfront premium of* $\frac{\lambda_0^H}{u'(c_0)}$ (< 1).

Intuitively, paying the premium upfront, rather than at harvest, is akin to holding a unit of illiquid assets. The cost of doing so is given by the (shadow) interest rate, which depends on whether liquidity constraints may bind before harvest - if not, then asset holdings can simply adjust to offset the difference. As a corollary, intertemporal preferences only matter for the timing of premium payment indirectly, through their effect on liquidity constraints, reflecting the fact that preferences are defined over flows of utility rather than over flows of money.

Liquidity constraints are closely tied to wealth (specifically, to deviations from permanent income, rather than permanent income itself) in the model. Combining Proposition 1 and Lemma 1 gives the following corollary, under the assumption that the product provides just a marginal unit of insurance (so that we can ignore second order effects).

Proposition 2. *The net benefit of pay-at-harvest insurance is decreasing in wealth. So too is the cost of paying upfront, rather than at harvest. Among farmers sure to be liquidity constrained before harvest, the latter dominates, so the benefit of pay-upfront insurance is increasing in wealth.*[33]

Thus, while the benefit of risk reduction (pay-at-harvest insurance) is higher among the poor, they may buy less (pay-upfront) insurance than the rich, because the inherent intertemporal transfer is more costly for them. Liquidity constraints drive both results: the poor are more likely to face liquidity constraints *after* harvest, meaning that they are less able to self-insure risks to harvest income (shocks in income lead to larger shocks in consumption), but they are also more likely to face liquidity constraints *before* harvest, making illiquid investments more costly.

[32] To be precise, being "almost" liquidity constrained is sufficient: the exact condition for preferring pay-at-harvest is that, upon purchasing pay-at-harvest insurance, $x_t - c_t \leq R^{-H+t}$ for some time $t < H$ and for some path.

[33] The general point that the gap between pay-upfront and pay-at-harvest insurance is decreasing in wealth follows from the shadow interest rate being decreasing in wealth. In our model that comes from a borrowing constraint (and wealth is the deviation from permanent income), but it could be motivated in other ways, and models sometimes take it as an assumption.

4.2.1 Delaying premium payment by one month

Consider the same insurance product as above, but with the premium payment delayed by just one period (corresponding to our experiment in Section 5.3, where the delay is one month).

Proposition 3. *The gain in the expected net benefit of insurance from delaying premium payment by one month is, to first order, equivalent to a proportional price cut in the upfront premium of* $\frac{\mu_0}{u'(c_0)}$.

Delaying premium payment by one period only increases demand if the farmer is liquidity constrained. The effect on the expected net benefit of doing so is $R^{-H}\mu_0$, compared to $R^{-H}(\mu_0 + \beta\mathbb{E}[\Sigma_{i=1}^{H-1}(R\delta)^i\tilde{\mu}_i])$ from delaying until harvest time. Thus, when H is large, a one month delay will have a small effect relative to a delay until harvest, unless either liquidity constraints are particularly strong at time 0, or there is present bias. Present bias closes the gap in two ways: first, the effect of future liquidity constraints are discounted by β, and second, the individual naively believes that he will be less likely to be liquidity constrained in the future.

4.3 Insurance with imperfect enforcement

If either side breaks the contract before harvest time, then the farmer does not pay the at-harvest premium, while he would have already paid the upfront premium. Accordingly, imperfect enforcement has implications both for farmer demand for insurance and for the willingness of insurance companies to supply it.

Default We assume that both sides may default on the insurance contract. At the beginning of the harvest period, with probability p_I (unrelated to yield) the insurer defaults on the contract, without reimbursing any upfront premiums.[34] The farmer then learns his yield and, if the insurer has not defaulted, can himself strategically default on any at-harvest premium, subject to some (possibly state dependent) utility cost c_D and the loss of any insurance payouts due.[35]Denoting whether the farmer chooses to pay the at-harvest premium by the (state-dependent) indicator function D_P, then to first order:

$$D_P \coloneqq \mathbb{I}[Iu'(c_H) + c_D \geq u'(c_H)] \tag{8}$$

[34]Such default could represent, for example, the insurer going bankrupt or deciding not to honor contracts. The assumption that it is unrelated to yield is reasonable in our setting, as strategic default by the insurer would be highly costly for the farming company, both legally and in terms of reputational costs. We ignore any insurer default after the farmer's decision to pay the harvest time premium, since it would not have a differential effect by the timing of premium payment.

[35]In practice the farmer may face considerable uncertainty about both yields and insurance payouts when deciding to default, which shrinks the difference between pay-upfront and pay-at-harvest. In our setting, for example, the company harvests the crop, at which point its weight is unknown to the farmer, and the area yield trigger further increases uncertainty.

Demand for insurance Given this defaulting behavior, imperfect contract enforcement drives an additional first-order difference between upfront and at-harvest insurance:

$$\begin{aligned}\text{Difference in net benefit of at-harvest \& upfront} &= \underbrace{R^{-H}\lambda_0^H}_{\text{liquidity constraint term}} + \underbrace{\beta\delta^H p_I E[u'(c_H)]}_{\text{premium saved when insurer defaults before harvest}} \\ &+ \underbrace{\beta\delta^H(1-p_I)\mathbb{E}[(1-D_P)(u'(c_H)-c_D-Iu'(c_H))]}_{\text{premium saved, minus cost of default and loss in insurance payouts, when farmer defaults}}\end{aligned} \tag{9}$$

The size of the difference caused by imperfect enforcement is clearly decreasing in the cost of default, c_D. If the cost of default is high enough, $c_D > \max_s u'(c_H(s))$, the farmer never strategically defaults.

Supply of insurance While the farmer is better off with the pay-at-harvest insurance, the possibility for strategic default means that the insurer may be worse off, which is the most likely reason why pay-upfront insurance is the norm. Whether there exists prices at which either of the two insurance products could be traded in a given setting depends on both c_D and p_I, as well as liquidity constraints and preferences as discussed earlier.[36]

Proposition 4. *If the cost of defaulting for the farmer, c_D, is too low, pay-at-harvest insurance will not be traded. If the probability of insurer default, p_I, is too high, pay-upfront insurance will not be traded.*

4.3.1 Interlinked insurance

Interlinking the insurance contract with the production contract has implications for contractual risk, as it means that default on one entails default on the other.

Default Now the farmer has one default decision to make: whether to default on both the insurance and production contracts. To translate this into the above framework, we define the (now endogenous) cost of farmer default, c_D, to be the value of the production relationship to the farmer relative to his outside option of selling to another buyer (side-selling). This will typically be positive, in which case interlinking helps to enforce the pay-at-harvest premium (this is why credit is often interlinked). However, if the farmer wishes to side-sell for some other reason, for example if the company defaults on aspects of the production contract, then c_D will be negative, in which case interlinking encourages default on the premium. Importantly, selective default by the farmer in order to avoid the pay-at-harvest premium is unlikely with under the interlinked

[36]The cost of strategic default is also key in another type of purely cross-state insurance: risk sharing (Ligon et al. 2002; Kocherlakota 1996). Related to the discussion here, Gauthier et al. (1997) show that enlarging the risk-sharing contracting space so as to allow for ex-ante transfers makes the first-best outcome easier to achieve.

contract, since the premium is only marginal if c_D is close to zero, and so expected default can be priced into the premium.

While unlikely, if pay-at-harvest insurance does affect side-selling, then the following (simple) proposition tells us how. Intuitively, for those with low yields, insurance payouts increase income from the contract, and so decrease the incentive to side-sell, whereas for those with high yields, pay-at-harvest premiums decrease income, and so increase the incentive to side-sell.

Proposition 5. *If pay-at-harvest insurance affects side-selling, it makes those with high yields more likely to side-sell, and those with low yields less likely to side-sell.*

As for the effect on imperfect enforcement on insurance demand, we have the following result, which enables us to relate the impact of ex-ante expectations of default to the impact of a price cut in the upfront premium, a point we return to in Section 5.4:

Proposition 6. *The option to side-sell in the interlinked contract drives a wedge between pay-at-harvest and pay-upfront insurance, bound above by a price cut in the upfront premium of:*

$$\mathbb{P}(\textit{side-sell with pay-at-harvest})\frac{\mathbb{E}[u'(c_H)|\textit{side-sell with pay-at-harvest}]}{\mathbb{E}[u'(c_H)]}$$

4.4 Implications and extensions

The transfer across time in insurance has several implications beyond the focus of this paper. It changes the relationship between insurance and self-insurance, and hence how background risk affects insurance demand: more risk *before* harvest may reduce demand for (pay-upfront) insurance, since insurance ties up liquidity which is needed for self-insurance; while more risk *at or after* harvest may increase demand for insurance, by motivating (precautionary) saving and hence reducing the cost of the transfer across time. When background risk is high, this tension between pay-upfront insurance and self-insurance may explain why insurance demand is often decreasing with risk aversion.[37] Finally, the transfer across time also changes the relationship between insurance and credit: for risk reduction they may be complements, not substitutes.

5 Why does the timing of the premium payment matter?

In this section we present evidence on the channels behind our main results, focusing on the same three as in the model: liquidity constraints, intertemporal preferences, and imperfect contract

[37] For a discussion of the evidence for insurance demand decreasing with risk aversion see Clarke 2016, for example, who propose basis risk as an explanation.

enforcement. We explain why we focus on these channels, and then show for each in turn that all three constrain demand for pay-upfront insurance.

Before beginning, we first note that since demand for pay-at-harvest insurance is high, our results cannot be explained by many of the mechanisms shown to constrain insurance demand in other settings. This includes basis risk (the risk that insurance payouts are not received when needed, because the insurance index is imperfectly correlated with individual loses), preferences over harvest risk (Clarke 2016; Mobarak and Rosenzweig 2012; Elabed et al. 2013) (risk preferences may still matter through imperfect contract enforcement), the presence of informal insurance (Mobarak and Rosenzweig 2012), and lack of information and understanding about insurance (Cai et al. 2015, Handel and Kolstad 2015).

Liquidity constraints, the first channel we consider, introduce a cost of holding the savings implicit in upfront insurance if they may bind *at any time* before harvest, as shown in Proposition 1. Several studies have documented liquidity constraints among similar populations in the region of the study (Duflo et al. 2011; Cohen and Dupas 2010). In Section 5.1 we present evidence for them from heterogeneous treatments effects in the main experiment, and in Section 5.2 we present related evidence from a second experiment.

Intertemporal preferences are the second channel we consider, and we are particularly interested in the role of present bias for three reasons.[38] First, recent evidence shows that present bias can distort intertemporal decisions substantially in similar settings (Loewenstein et al. 2003; Duflo et al. 2011; Schilbach 2015). Second, with present bias, the timing of insurance has additional welfare implications, as future selves may regret the decision to forgo pay-upfront insurance. Third, present bias has implications for insurance design: even slight delays in premium payment may increase demand without the enforcement concerns of pay-at-harvest insurance, as argued in Section 4.2.1. We test such a product in Section 5.3.

Imperfect contract enforcement, the final channel we consider, matters because if either party defaults on the contract before harvesting, then under pay-upfront insurance the premium is paid, whereas under pay-at-harvest insurance it is not. In Section 5.4 we report tests for this channel which are motivated by our model.

[38] As shown in Proposition 1, intertemporal preferences only differentially affect the decision to take up insurance when individuals have a non-zero chance of being liquidity constrained before the next harvest. As shown by Duflo et al. (2011) and Cohen and Dupas (2010), this is likely to be the case for some farmers in our setting. Further, liquidity constraints are an endogenous outcome of the intertemporal optimization problem farmers face, for which intertemporal preferences are of key importance.

5.1 Is upfront payment more costly for the poor & the liquidity constrained?

It is often argued that income variation is more costly for the poor, and so they should have higher demand for risk reduction. Yet the poor demand less insurance. Proposition 2 showed that the transfer across time in insurance is a possible explanation – the poor are more likely to be liquidity constrained, and liquidity constraints increase the cost of paying the premium upfront. If so, in our experiment we would expect the gap between pay-upfront and pay-at-harvest insurance to be higher among the poor.

Here we report how demand for pay-upfront and pay-at-harvest insurance varies by proxies for wealth and liquidity constraints, and thus the heterogeneous treatment effect of removing the transfer across time. The proxies include yield levels in the previous harvest, sugarcane plot size, number of acres cultivated, whether the household owns a cow, access to savings and the portion of income from sugarcane.[39] In order to gain power, we bundle together the two pay-upfront groups (full price and 30% discount), as stated when registering the trial, giving the regression model:

$$T_{if} = \alpha + \beta Harvest_i + \gamma x_i + \delta Harvest_i * x_i + \nu_f + \epsilon_{if} \tag{10}$$

Table 3 presents the results, which show that the treatment effect does vary by proxies for wealth and liquidity constraints. While not all of the interaction coefficient estimates are significant, delaying premium payments until harvest does increase take-up more among less wealthy and more liquidity constrained households, as predicted by proposition 2. For example, the treatment effect is 14 percentage points larger for those who do not own a cow, and 17 percentage points larger for those who would do not have savings to cover an emergency expenditure of Sh 1,000 ($10). Further, also in line with proposition 2, the difference comes from demand for pay-at-harvest insurance being higher among the poor.[40] Of course, these are heterogeneous treatment effects and so cannot be interpreted causally, as there could be confounders.[41] From a policy perspective, the results imply that pay-at-harvest insurance is particularly beneficial for poorer farmers, who are typically in greater need of novel risk management options.

[39] Time since the last maize harvesting season would have been another interesting proxy, but we have little variation in it in our experiment, given the short time frame.

[40] There is less margin for take-up heterogeneity in Pay-Upfront insurance, given its low average take-up, but the two predictions of the model hold: both take-up of Pay-At-Harvest and the gap between Pay-At-Harvest and Pay-Upfront are larger among the poor. Further, existing studies on Pay-Upfront insurance typically find lower take-up among the poor and the liquidity constrained (Cole et al. 2013a).

[41] Also, the different proxies are obviously not independent, although pairwise correlations are all less than 0.27 (except for the two access to emergency savings variables).

5.2 Do people buy upfront insurance, given enough cash to do so?

In line with the importance of liquidity constraints, when we surveyed farmers in the pay-upfront group about why they did not purchase insurance, their main reason was lack of cash. In this section we present a second experiment which investigates this, by asking: if farmers did have the cash to buy upfront insurance, would they do so?

5.2.1 Experimental design

In the experiment, which targeted 120 farmers, we cross cut the pay-upfront and pay-at-harvest treatments of the main experiment with a cash drop treatment (with stratification again at the field level). Under the cash drop, during the baseline survey enumerators gave farmers an amount of cash slightly larger than the price of the insurance premium, around an hour before company outreach workers offered farmers the insurance product. The treatment mimics closely one of the arms in Cole et al. (2013a). This cross-cut design allows us to test whether the impact of the cash drop varies across the pay-upfront vs. pay-at-harvest groups, as well as assessing the relative impact of the cash drop compared to the premium deferral. Appendix Table A.3 shows that the treatment groups were balanced.

Before presenting results, we consider how this cash treatment may affect demand. First, in the pay-upfront group, it ensures that farmers have enough cash to pay the premium if they wish to, removing any hard cash constraint and thus addressing the most commonly cited reason for not purchasing upfront insurance.[42] Yet, while the cash drop eases liquidity constraints, it need not remove them entirely - an individual is liquidity constrained if they are not able to borrow any more at the market interest rate; after receiving the cash drop, farmers may still have wanted to borrow more. In the pay-at-harvest group, in contrast, the cash drop may affect demand through a small wealth effect, but not through a liquidity effect.

The cash may also affect demand through a reciprocity effect - a standard concern with cash drop designs – whereby farmers buy the insurance product just to reciprocate the cash gift. We tried to minimize any reciprocity by having the survey enumerator give the cash gift at the beginning of the meeting – in contrast, the insurance product is offered by a company outreach worker, at the end of the meeting. We try to control for reciprocity using our cross-cut design, based on the assumption that reciprocity affects demand of pay-upfront and pay-at-harvest insurance equally. Finally, we note that the cash does not affect contractual risk, so that any resulting treatment

[42] We note that this also addresses another potential channel in the main experiment: that farmers feel somehow pressured to buy insurance (for instance through social desirability bias), and not having the cash is a convenient excuse not to buy insurance when farmers have to pay upfront, which is no longer credible when they can pay at harvest. Giving farmers the cash also renders the argument non-credible.

effect is not driven by imperfect contract enforcement.[43]

The experiment is best interpreted as answering whether pay-upfront insurance is the marginal expenditure (given cash which removes any hard cash constraints). Evidence from other settings suggests that the answer may be no: when interlinking insurance with credit, Gine and Yang (2009) and Banerjee et al. (2014) find that demand for credit actually decreases when bundled with insurance. If pay-upfront insurance was the marginal expenditure, if anything we would expect the opposite.

5.2.2 Experimental results

We estimate the following regression model:

$$T_{if} = \alpha + \beta Harvest_{if} + \gamma Cash_{if} + \nu Harvest_{if} * Cash_{if} + \eta_f + \epsilon_{if} \tag{11}$$

Figure 4 presents the results. First, it is reassuring to note that, in this different sample, the comparison between the pay-upfront and pay-at-harvest groups resembles that of the main experiment. Take-up for the upfront group is slightly larger (13%), but, again, introducing at-harvest payment raises take-up dramatically (up to 76%). Second, the cash drop raises substantively the take-up rate in the upfront group (up to 33%), suggesting farmers may have faced cash constraints. However, the impact of the cash drop is much smaller than that of the harvest time premium, meaning that many who would purchase pay-at-harvest insurance would not purchase pay-upfront insurance even if they do have the cash to do so – they would prefer to use the money for other purposes (e.g. consumption, labor payments, school fees). Third, the cash drop also has an impact on take-up rates in the pay-at-harvest group (from 76% to 88%). Our model predicts, if anything, a (very small) *negative* wealth effect on demand, so that this is likely a reciprocity effect as discussed above (and mentioned in Cole et al. 2013a). The difference in impact of the cash drop between the pay-at-harvest and pay-upfront groups is 8%, which is small. While imprecisely estimated, we take this as evidence that the cash drop had little effect on take-up of pay-upfront insurance beyond the reciprocity effect, especially relative to delaying premium payment until harvest time.

Table 4 confirms the patterns described above. Column (1) presents the basic level impact of the cash drop and pay-at-harvest treatments, from a regression with fixed effects at the field level, the stratification unit. We add additional controls in column (2). In both specifications, we reject the null of equality of the two treatments at the 1% level (p-value .00004). The coefficient on Cash is significant at the 10% level in column (1) and remains similar in size but loses some

[43] Ignoring any second order effects of the cash drop on side-selling, which are likely very small given the size of the cash drop

precision as we add more controls. In columns (3) and (4), we look at the interaction between the two treatments. The coefficient on the interaction is always negative, as we would expect, but it is small and insignificant. It is imprecisely estimated, but even at the upper bound of the (very wide) confidence interval the interaction can only account for around half of the difference between pay-upfront and pay-at-harvest insurance.

To summarize, the results show that cash drops do relatively little to close the gap between pay-upfront and pay-at-harvest insurance – pay-upfront insurance is not the marginal expenditure. This tells us that if liquidity constraints are important for our results, then they cannot be removed with a small cash drop - farmers must be substantially liquidity constrained, with more pressing uses for the cash. The next experiment tests for liquidity constraints in a different way, by asking whether they can be removed with a small delay in premium payment.

5.3 Does delaying the premium payment by one month increase take-up?

If present bias is a major driver of our results, then delaying the premium payment by just a short amount of time, one month for example, may increase take-up. Here we describe a third experiment which did exactly this.

5.3.1 Experimental design

We randomly allocated a sample of 120 farmers to two treatment groups (with stratification again at the field level). Both groups were offered a choice between either a cash payment, equal to the insurance premium, or free enrollment in the insurance. The difference between the treatment groups was *when* farmers would receive whatever they chose: in the first treatment group, *Receive Choice Now*, farmers were told that they would receive it immediately; while in the second group, *Receive Choice in One Month*, farmers were told that they would receive it (plus interest) in one month's time.

Offering the choice between insurance for free or cash, rather than the choice between buying or not buying insurance, allowed us to isolate the role of intertemporal preferences in two ways. First, it ensured that the choice in the *Receive Choice in One Month* group could be enforced (since premium payments did not rely on the farmer paying out of her own pocket). Second, it relaxed any hard cash constraints, ensuring the farmer could take-up the insurance if she wanted to, just like a cash drop.

What does the experiment test? First, it tests liquidity constraints - as shown in Section 4.2.1, for a one month delay in premium payment to make a difference, the farmer must be liquidity constrained at the time of the experiment. But, as clarified by the model, liquidity constraints

are not an exogenous parameter, but are themselves a function of the fundamentals of the model, including (quasi-hyperbolic) discount rates. We claim that a large effect on take-up of a one-month delay would be evidence of present bias, rather than time-consistent exponential discounting. The argument is as follows. If farmers are exponential discounters, then for a one month delay to have a large effect on the net benefit of insurance, they would have to have low δ (especially as a 30% price cut in the main experiment had little effect). However, the same low δ would mean that farmers would not buy insurance even with a one month delay – such a product still transfers income across approximately eleven months. In contrast, under present bias, a low *beta* leads to a large difference between paying now and in one month, without making insurance paid for in one month unattractive. Another potential explanation would have been that credit constraints vary across time periods (Dean and Sautmann 2014), and the experiment just happened to take place at a time of large and very short-run liquidity constraints (for example due to an aggregate shock). However, we ran the experiment across two months (plus a one-month pilot beforehand) and the results, presented below, are stable across these periods.

Appendix Table A.4 reports the balance test across the two groups. We note that, due to the small sample size, there are significant imbalances across the two groups in the share of men, the acres of land cultivated and plot size, and emergency savings for Sh5,000; pairwise correlations of these variables are all positive (except one). As discussed below, results are robust to the inclusion of these variables as controls.

5.3.2 Experimental results

Figure 5 shows that the take-up share in the *Receive Choice in One Month* group is 72%, compared to a baseline of 51% in the *Receive Choice Now* group. This 21 percentage point increase shows that a change of only one month in the timing of the premium payment has a large impact on insurance take-up. While the experimental design does not allow us to directly distinguish between time-consistent and time-inconsistent discounting *directly*,[44] the large effect is inconsistent with exponential discounting, as argued above. In contrast, it is consistent with present bias, as the *Receive Choice in One Month* treatment provides farmers with a commitment device on how to use the cash transfer, potentially overcoming their time inconsistency.

Table 5 confirms these results across different specifications. The gap between the two treatments begins statistically significant at 5% and becomes statistically significant at 1% when adding farmer controls. We note that the point estimate raises from 0.23 in the baseline specification with

[44] We could not test time inconsistency by allowing farmers to revise their commitment one month later because any new information received during the month (for instance on expected yield) would have potentially changed farmers' decisions even under time-consistent discounting.

field fixed effects (Column 1) to 0.29 when adding both set of controls, though the difference in the two estimates is not statistically significant. This suggests that, if anything, accounting for the baseline imbalances reported above increases the estimate of the impact of requiring farmers to sign up in advance.

We note that the design mitigates the traditional trust concerns associated to standard time preferences experiments (Andreoni and Sprenger 2012). In the *Receive Choice in One Month* treatment, both the cash transfer and the insurance sign-up depend on the field officer revisiting the field, so there are no differential trust concerns across the two choices. It is still possible, though implausible, that farmers think field officers are more likely to return if they choose insurance. However, visits are organized at the field level, not the individual level, so officers meet multiple households in a given visit, and more importantly, farmers have the contact info of the relevant company field staff and IPA staff.

While present bias can lead to under subscription in pay-upfront insurance, one might think that it could also lead to over subscription and hence future regret in pay-at-harvest insurance. While we believe that this is a real possibility with the sale of goods on credit, where benefits are borne immediately, in the case of insurance there is no clear immediate benefit to subscription. On the contrary, pay-at-harvest insurance eliminates the time gap between cost and benefit that standard insurance products introduce. In line with this argument, as discussed above, in follow-up calls with 40 farmers who took-up the pay-at-harvest insurance, only 7.5% of farmers said they would not take-up the product again.

We note that the baseline take-up for the *Receive Choice Now* group is larger than the take-up in the group *Pay Upfront + Cash* in the cash-constraints experiment. This difference should be interpreted with caution. The two groups are drawn from different samples and so are not directly comparable: the cash-constraints experiment occurred in late Summer 2014, while this experiment was implemented in Spring 2015, shortly after the end of the dry season (December-March), when the risk of low harvest may have been more salient. With this caveat, the difference could also be explained by a literature dating back to Knetsch and Sinden (1984):[45] *Receive Choice Now* represents the Willingness to Accept, whereas *Pay Upfront + Cash* represents the Willingness to Pay (without the wealth effect) and may include an endowment effect from handing farmers the cash at the start of the visit.[46]

Before moving to the next channel, we note that in the main experiment we elicited measures

[45] See Horowitz and McConnell (2002) for a summary of the literature.

[46] If so, the 21 p.p. effect of a one month delay reported above may be a lower bound for the effect when farmers have to pay from their own wealth.

of preferences over the timing of cash flows, using standard (Becker-DeGroot) Money Earlier or Later questions (Cohen et al. 2016).[47] We did not find heterogeneous treatment effects by these Required Rate of Return variables, as shown in table A.2. It is fairly common to find no such effects, which could be due to measurement issues, limited statistical power, or the fact that standard lab-experiment measures in a given domain (e.g. the timing of cash disbursements) may fail to hold predictive power on other domains, such as how to use that cash.[48] Given resource constraints, and because our main experimental variation was on timing, we did not elicit measures of risk aversion.

5.4 Imperfect enforcement

Anticipation that either party may default before harvest drives a wedge between take-up of pay-upfront and pay-at-harvest insurance, as shown in Section 4.3.1. Here we consider the importance of this channel. While we find evidence that counterparty risk mattered for overall levels of take-up, we find no evidence for a differential effect by the timing of the premium payment, in spite of significant side-selling ex-post.

Before the farmers in our study were due to harvest, financial problems of the company led to the closure of the factory for several months. During the closure the company did not harvest cane, and the resulting backlog caused severe harvesting delays afterwards, leading to uncertainty among farmers as to when harvesting would happen, if at all. As a result, unsurprisingly, only 48% of our farmers harvested with the company. Those that did not either side-sold or uprooted the crop (for brevity we refer to this as side-selling below). Figure 6 plots the harvesting rate by sublocation, and for comparison also plots a lower bound for it historically.[49] It shows that the rate was much lower than usual, and that it varied substantially by sublocation. While those who harvested with the company did receive any insurance payouts due, those who side-sold were ineligible.

The widespread default ex-post underlines the trust required by standard pay-upfront insurance, and raises two important questions: (i) did pay-at-harvest insurance induce side-selling; and (ii) were expectations of default responsible for the difference in take-up, ex-ante?

[47] A recent experimental literature considers what such questions elicit, and suggests difficulties with using them to measure intertemporal preferences (Andreoni and Sprenger 2012, Augenblick et al. 2015, Cohen et al. 2016)

[48] For instance, Kaur et al. (2015) find no correlation between lab experiment measures of time inconsistency and workers' choices on effort and labor contracts.

[49] The historical measure of the harvesting rate is a lower bound on the true harvesting rate because of the data we had to construct it. It is constructed as the proportion of farmers who previously harvested a Plant or Ratoon 1 cycle who appear in the data as harvesting the subsequent cycle. However, some of these farmers would have uprooted the crop after harvesting, and thus never begun the subsequent cycle.

5.4.1 Did insurance affect side-selling?

We can rule out any sizeable effect of insurance on side-selling, in line with the design of the insurance product and the assumptions and results of our model. Given the low take-up of pay-upfront insurance, Figure 7 effectively reports the Intent-To-Treat of offering pay-at-harvest insurance on harvesting with the company, showing no level-effect on side-selling in spite of high take-up. But insurance could still have affected *who* side-sold. If so, Proposition 5 showed that pay-at-harvest insurance makes those with low yields less likely to side-sell and those with high yields more likely to, so yield conditional on selling to the company should be higher among the Pay-Upfront group. Figure 8 shows it was not.[50]

5.4.2 Did anticipation of default affect take-up differentially?

Given the extent of side-selling, it is particularly important for us to consider how important ex-ante expectations of contract risk were in driving our main result. Here we present two sets of results which suggest that the role was limited. Before doing so, we note that our two mechanism experiments and heterogeneous treatment effects showed that liquidity constraints and present bias were important channels, and also that in the *Receive Choice in One Month* treatment, where the insurance product was fully exposed to contract risk, take-up reached 72%.

Our first evidence for a limited role for contract risk considers heterogeneous treatment effects of delaying the premium payment, by plausible proxies for ex-ante priors of default. If anticipation of default did drive a difference in take-up between pay-upfront and pay-at-harvest insurance, and there was heterogeneity in priors for the probability of default, then we would expect a take-up regression to show an interaction between proxies for priors and pay-at-harvest time premiums (similar to positive correlation tests for adverse selection in the insurance literature, Einav and Finkelstein 2011). We consider two such proxies for prior probabilities of default. First, in the baseline survey, we asked respondents about their trust in, and relationship with, the company. Table A.1 shows that while some of these measures do predict overall levels of take-up (consistent with a belief that the company will not make insurance payouts even if the production contract is upheld),[51] they do not predict take-up differentially by premium timing. Second, we consider actual harvesting rates ex-post, both of individual farmers and in the local area (Figure 6 shows it had substantial geographical variation), and both in the current season and in the previous season.

[50] Besides side-selling, one might also worry that insurance induced moral hazard. However, moral hazard, if present, would work in the same direction as selective side-selling, lowering yields in the pay-at-harvest treatment group. In addition, partial side-selling is unlikely, both because of high transportation costs and because of monitoring by company outreach workers.

[51] Indeed, some farmers did mention trust as a reason why they did not buy insurance.

Using harvesting in the current season as a proxy relies on the assumption that actual harvesting ex-post was (negatively) correlated with the ex-ante probability of default, and requires the caveat that we are conditioning on an ex-post variable. Table 6 shows that we do not find heterogeneous treatment effects for any of these proxies for ex-ante expectations of contract default.

Our second evidence for a limited role for contract risk relies on Proposition 6, which allows us to bound the differential effect of expectations of default on take-up, by the effect of a price cut in the upfront premium. Specifically, by a proportional price cut equal to the expected probability of side-selling weighted by the relative marginal utility of consumption when side-selling. Yet, in our main experiment, a 30% price cut had almost no effect on take-up of upfront insurance, suggesting a low price elasticity. Thus, for imperfect enforcement to account for much of our result, ex-ante expectations of either the probability of default, or of the marginal utility of consumption conditional on default, would have had to be extremely high, calling in to question why farmers entered the production contract with the company to begin with.

Finally, we note that counterparty risk from (non-selective) default by the buyer should have little effect on demand for Pay-At-Harvest insurance - if anything, it should reduce demand slightly, through increased precautionary savings. This suggests that, even absent any risk of buyer default, take-up of Pay-At-Harvest insurance would have been high. Yet, while counterparty risk may have reduced take-up for Pay-Upfront insurance, it is the high take-up of Pay-At-Harvest insurance, rather than the low take-up of Pay-Upfront insurance, which is an outlier compared to other insurance studies.

5.5 Other channels

We conclude this section by briefly discussing several additional potential channels, several of which are interesting and warrant future work.

The at-harvest premium is a deduction, while the upfront premium is a payment; this difference suggests several (behavioral) channels which are not directly about timing. First, according to prospect theory (Kahneman and Tversky 1979; Kőszegi and Rabin 2007), farmers may be more sensitive to losses than gains. While a thorough application of the theory is beyond the scope of this paper (and would require detailing how reference points are set), intuitively upfront payments may fall in the loss domain, while at-harvest payments, being deductions, may be perceived as lower gains. Second, according to relative thinking (Tversky and Kahneman 1981, Azar 2007), farmers may make choices based on relative quantities, rather than absolute quantities. Being small relative to harvest revenues, the at-harvest premium could appear smaller than the upfront

premium.[52] Salience Theory offers a similar argument: under a multiple time period interpretation of Bordalo et al. (2012), diminishing sensitivity means that the upfront period may be more salient than harvest period, since income will be higher in the latter. Finally, inputs were already charged as deductions from harvest revenues in our setting, so pay-at-harvest could have seemed like the default (although we note that the high take-up of pay-at-harvest insurance, not the low take-up of pay-upfront insurance, is the outlier in our results compared to other studies).

The large effect of just a one month delay in premium payment, however, does point to the direct importance of timing, which could arise in several ways beyond those captured in our model. First, numerous empirical studies find a jump in demand at zero prices (Cohen and Dupas 2010); a similar, zero-price *today* effect could help explain our results.[53] Second, Andreoni and Sprenger (2012) report expected utility violations when certain and uncertain outcomes are combined – pay-upfront insurance combines a certain payment with an uncertain payout, whereas both are uncertain in pay-at-harvest insurance. Third, at-harvest and upfront payments may have different implications for bargaining in other interactions within the household or within informal risk sharing networks (Jakiela and Ozier 2016; Kinnan 2017). Finally, while unlikely in our setting, allowing farmers to pay at harvest rather than upfront for insurance may provide a positive signal of the quality of the insurance.

6 Policy implications

Almost all insurance products transfer income across time. The resulting mechanisms, shown to affect insurance demand in our experiment, are known to shape financial decisions across many diverse settings. In the final section of the paper, we discuss the policy implications, first for crop insurance and then for insurance more generally.

6.1 Crop insurance

From a policy perspective, boosting crop insurance take-up is an ongoing challenge.[54] This paper shows that the timing of the premium payment matters, and that pay-at-harvest insurance is a promising solution, which warrants replication in other settings. While the enforcement mechanism we used could be used in most contract farming settings (whose presence is growing

[52] We thank Nathan Nunn for pointing out this explanation.

[53] Such an effect would be an alternative explanation for the finding in Tarozzi et al. (2014) that offering anti-malarial bednets through loans has results in a large increase in take-up, and would also explain the prevalence of zero down-payment financing options for many consumer purchases, such as cars and furniture.

[54] An interesting recent literature (Clarke 2016) shows that basis risk is so high in some index insurance products that farmers should not buy them. We assume that the policy maker has a good product.

steadily in developing countries, UNCTAD 2009), the wider applicability of the idea depends on the answer to two questions.

First, are there other ways to enforce pay-at-harvest premium payments? US Federal Crop Insurance (FCI) is one example – historically it is a pay-at-harvest insurance - but it operates with strong legal institutions and government backing. More generally, credit provides a promising comparison, since it faces a stricter enforcement constraint than cross-state insurance (in the latter, net payment is only due in good states of the world), yet often achieves very low default rates. Perhaps methods used for credit, and in particular microfinance, such as relational contracting, group liability, and collateral, could be adopted for cross-state insurance?[55]

Second, do premiums actually need to be paid at the subsequent harvest, or are there other timings which would still boost take-up while being easier to enforce? Our *One Month Experiment* showed that even a slight delay can increase take-up substantially. But seasonality may be important too – as in Duflo et al. (2011), farmers may be less liquidity constrained at the *previous* harvest time than at planting (and potentially also less affected by scarcity Mani et al. 2013) - although in our experiment we met farmers just a few weeks after harvesting, suggesting any such effects would have been very short lived. Relatedly, while we have considered the timing of insurance premiums, the timing of *payouts* may also matter. Times are likely to be hardest for farmers in the hungry season following a bad harvest; farmers may prefer insurance payouts then.

We conclude the policy discussion for crop insurance by noting several other benefits of interlinking it with contract farming. First, since farmers already contract with the company, administrative costs would be lower and trust may be higher. Second, contract farming schemes often collect detailed plot-level data, which could help cost insurance products.[56] Third, insurance renewal is often low, with high dropout among farmers not receiving a payout in the first season (Cole et al. 2014, Cai et al. 2016). With interlinking, farmers could credibly sign up for insurance contracts covering multiple seasons, increasing their chance of receiving a payout before policy renewal.

6.2 Other insurance products

The transfer across time is almost ubiquitous in insurance products; it is most likely to affect insurance demand when the shadow interest rate is high or when the time period involved is long.

[55] An alternative approach would be to offer a loan and pay-upfront insurance at the same time, but unbundled. However, under present bias, doing so may have negative welfare implications. Further, enforcing repayment of the loan would be harder, and limited liability could reduce the incentive to buy insurance through the standard asset substitution problem (Jensen and Meckling 1976).

[56] Data limitations are a fundamental constraint in the design of area yield products (Elabed et al. 2013), which displayed lower basis risk than rainfall index insurance in our setting

This has several policy implications. First, insurance contracts should be designed and marketed with insurees' paths of liquidity in mind. For example, households could be offered to purchase insurance directly from cash transfers or EITC payments (potentially with pre-commitment). Second, the transfer across time may help to explain low take-up of rare-disaster insurance and front-loaded dynamic insurance contracts such as life insurance (Pauly et al. 1995; Finkelstein et al. 2005; Handel et al. 2015), for which the intertemporal transfer is particularly long. Finally, wishing to remove the transfer across time (as is done, for example, in social insurance and in the FCI) may provide another justification for government intervention in insurance markets, if they are better able than private providers to enforce premium payments ex post.

7 Conclusion

By requiring that the premium be paid upfront, standard insurance contracts introduce a fundamental difference between the goal of insurance and what insurance products do in practice: they not only transfer income across states, they also transfers income across time. We have argued that this difference is at the heart of several explanations offered for the low take-up of insurance, such as liquidity constraints, present bias, and trust in the insurer. In addition, once the temporal dimension of insurance contracts is taken into account, we have shown that a standard borrowing constraint can resolve the puzzlingly low demand for insurance among the poor – while the poor have greater demand for risk reduction, they face a higher cost of paying the premium upfront.

In the context of crop insurance, where seasonality makes the transfer across time particularly costly, the difference can be removed by charging the premium at harvest time rather than upfront. Doing so in our experiment, by charging the premium as a deduction from harvest revenues in a contract farming setting, increased take-up by 67 percentage points, with the effect largest among the poorest. We discussed numerous possible channels for this large effect, and presented several pieces of evidence which show that two of the three most natural ones play a role. Heterogeneous treatment effects suggest that liquidity constraints mattered, and a second experiment shows that they ran deeper than simply not having the cash to pay the premium. A third experiment found that even a small delay in premium payment increased demand substantially, showing the role of present bias, and providing further evidence for liquidity constraints. Lastly, while contractual risk *may* have driven a difference between take-up of pay-upfront and pay-at-harvest insurance, in our setting we find no evidence that it did, across multiple tests, in spite of a financial shock which led to high levels of default ex-post.

From a policy perspective, our results may have broad implications. For crop insurance, where

boosting demand has proven difficult, we showed that timing matters and proposed pay-at-harvest insurance as a promising potential solution. Whether it could work outside of contract farming settings remains an important question. More broadly, the transfer across time is almost ubiquitous in insurance products. The effect on the demand for other types of insurance, and on risk management more generally, are interesting questions for future work.

References

Ahmed, Shukri, Craig McIntosh, and Alexandros Sarris. 2017. "The Impact of Commercial Rainfall Index Insurance: Experimental Evidence from Ethiopia." *Unpublished*.

Andreoni, James, and Charles Sprenger. 2012. "Risk preferences are not time preferences." *The American Economic Review*, 102(7): 3357–3376.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2015. "Working Over Time: Dynamic Inconsistency in Real Effort Tasks." *The Quarterly Journal of Economics*.

Azar, Ofer H. 2007. "Relative thinking theory." *Journal of Socio-Economics*, 36(1): 1–14.

Banerjee, Abhijit, Esther Duflo, and Richard Hornbeck. 2014. "Bundling Health Insurance and Microfinance in India: There Cannot Be Adverse Selection If There Is No Demand." *American Economic Review*, 104(5): 291–97.

Bardhan, Pranab K. 1980. "Interlocking Factor Markets and Agrarian Development: A Review of Issues." *Oxford Economic Papers*, 32(1): 82–98.

Bardhan, Pranab K. 1989. *The economic theory of agrarian institutions.*: Clarendon Press Oxford.

Becker, Bo, and Victoria Ivashina. 2015. "Reaching for Yield in the Bond Market." *The Journal of Finance*, 70(5): 1863–1902.

Bell, Clive. 1988. "Credit markets and interlinked transactions." In *Handbook of Development Economics*. eds. by Hollis Chenery, and T.N. Srinivasan, 1 of Handbook of Development Economics: Elsevier, , Chap. 16 763–830.

Blouin, Arthur, and Rocco Macchiavello. 2017. "Strategic Default in the International Coffee Market." *Unpublished*.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice Under Risk." *The Quarterly Journal of Economics*, 127(3): 1243–1285.

Bulow, Jeremy, and Kenneth Rogoff. 1989. "Sovereign Debt: Is to Forgive to Forget?" *The American Economic Review*, 79(1): 43–50.

Cai, Jing, Alain De Janvry, and Elisabeth Sadoulet. 2015. "Social Networks and the Decision to Insure." *American Economic Journal: Applied Economics*, 7(2): 81–108.

Cai, Jing, Alain de Janvry, and Elisabeth Sadoulet. 2016. "Subsidy Policies and Insurance Demand." Working Paper 22702, National Bureau of Economic Research.

Carroll, Christopher, and Miles Kimball. 1996. "On the Concavity of the Consumption Function." *Econometrica*, 64(4): 981–992.

Carroll, Christopher, and Miles Kimball. 2005. "Liquidity Constraints and Precautionary Saving." *Unpublished*.

Carter, Michael R, Lan Cheng, and Alexander Sarris. 2011. "The impact of interlinked index insurance and credit contracts on financial market deepening and small farm productivity." In *Annual Meeting of the American Applied Economics Association, Pittsburgh PA, July*. 24–26.

Casaburi, Lorenzo, and Rocco Macchiavello. 2016. "Firm and Market Response to Saving Constraints: Evidence from the Kenyan Dairy Industry." *Unpublished*.

Casaburi, Lorenzo, and Tristan Reed. 2017. "Competition in Agricultural Markets: An Experimental Approach." *Unpublished*.

Clarke, Daniel J. 2016. "A Theory of Rational Demand for Index Insurance." *American Economic Journal: Microeconomics*, 8(1): 283–306.

Cohen, Jessica, and Pascaline Dupas. 2010. "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment." *The Quarterly Journal of Economics*, 125(1): 1–45.

Cohen, Jonathan D., Keith Marzilli Ericson, David Laibson, and John Myles White. 2016. "Measuring Time Preferences." Working Paper 22455, National Bureau of Economic Research.

Cole, Shawn A., and Wentao Xiong. 2017. "Agricultural Insurance and Economic Development." *Annual Review of Economics*, 9(1): 235–262.

Cole, Shawn, Xavier Gine, Jeremy Tobacman, Petia Topalova, Robert Townsend, and James Vickery. 2013a. "Barriers to Household Risk Management: Evidence from India." *American Economic Journal: Applied Economics*, 5(1): 104–35.

Cole, Shawn, Xavier Gine, and James Vickery. 2013b. "How does risk management influence production decisions? Evidence from a field experiment." Policy Research Working Paper Series 6546, The World Bank.

Cole, Shawn, Daniel Stein, and Jeremy Tobacman. 2014. "Dynamics of Demand for Index Insurance: Evidence from a Long-Run Field Experiment." *American Economic Review*, 104(5): 284–90.

Conning, Jonathan, and Christopher Udry. 2007. "Rural financial markets in developing countries." *Handbook of agricultural economics*, 3 2857–2908.

Dean, Mark, and Anja Sautmann. 2014. "Credit constraints and the measurement of time preferences." *Available at SSRN 2423951*.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica*, 59(5): 1221–1248.

Deaton, Angus, and Guy Laroque. 1992. "On the behaviour of commodity prices." *The Review of Economic Studies*, 59(1): 1–23.

DellaVigna, Stefano, and Ulrike Malmendier. 2006. "Paying Not to Go to the Gym." *American Economic Review*, 96(3): 694–719.

Dercon, Stefan, Jan Willem Gunning, and Andrew Zeitlin. 2011. "The demand for insurance under limited credibility: Evidence from Kenya." In *International Development Conference, DIAL*.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review*, 101(6): 2350–90.

Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives*, 25(1): 115–38.

Elabed, Ghada, Marc F. Bellemare, Michael R. Carter, and Catherine Guirkinger. 2013. "Managing basis risk with multiscale index insurance." *Agricultural Economics*, 44(4-5): 419–431.

Epstein, Larry G., and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica*, 57(4): 937–969.

Fafchamps, Marcel, Christopher Udry, and Katherine Czukas. 1998. "Drought and saving in West Africa: are livestock a buffer stock?" *Journal of Development Economics*, 55(2): 273 – 305.

Finkelstein, Amy, Kathleen McGarry, and Amir Sufi. 2005. "Dynamic Inefficiencies in Insurance Markets : Evidence from Long-Term Care Insurance." *American Economic Review*.

Gauthier, Celine, Michel Poitevin, and Patrick Gonzalez. 1997. "Ex ante payments in self-enforcing risk-sharing contracts." *Journal of Economic Theory*, 76(1): 106–144.

Ghani, T., and T Reed. 2017. "Relationships, Risk and Rents: Evidence from a Market for Ice." *Unpublished*.

Gine, Xavier, and Dean Yang. 2009. "Insurance, credit, and technology adoption: Field experimental evidence from Malawi." *Journal of Development Economics*, 89(1): 1 – 11.

Handel, Benjamin, Igal Hendel, and Michael Whinston. 2015. "Equilibria in Health Exchanges : Adverse Selection vs . Reclassification Risk." *Econometrica*, 83(4): 1261–1313.

Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review*, 105(8): 2449–2500.

Horowitz, John K, and Kenneth E McConnell. 2002. "A review of WTA/WTP studies." *Journal of environmental economics and Management*, 44(3): 426–447.

Jakiela, Pamela, and Owen Ozier. 2016. "Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies." *The Review of Economic Studies*, 83(1): , p. 231.

Jensen, Michael, and William Meckling. 1976. "Theory of the firm: Managerial behavior, agency costs and ownership structure." *Journal of Financial Economics*, 3(4): 305 – 360.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect theory: An analysis of decision under risk." *Econometrica: Journal of the Econometric Society* 263–291.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. "Agricultural Decisions after Relaxing Credit and Risk Constraints." *The Quarterly Journal of Economics*, 129(2): 597–652.

Kaur, Supreet, Michael Kremer, and Sendhil Mullainathan. 2015. "Self-Control at Work." *Journal of Political Economy*, 123(6): 1227–1277.

Kinnan, Cynthia. 2017. "Distinguishing barriers to insurance in Thai villages." *Unpublished*.

Knetsch, Jack L., and J. A. Sinden. 1984. "Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value." *The Quarterly Journal of Economics*, 99(3): 507–521.

Kocherlakota, Narayana R. 1996. "Implications of Efficient Risk Sharing without Commitment." *The Review of Economic Studies*, 63(4): 595–609.

Kőszegi, Botond, and Matthew Rabin. 2007. "Reference-dependent risk attitudes." *The American Economic Review* 1047–1073.

Laibson, David. 1997. "Golden eggs and hyperbolic discounting." *The Quarterly Journal of Economics* 443–477.

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

Liu, Yanyan, Kevin Chen, Ruth Hill, and Chengwei Xiao. 2016. "Delayed Premium Payment, Insurance Adoption, and Household Investment in Rural China." *IFPRI Discussion Paper*(01306): .

Liu, Yanyan, and Robert J. Myers. 2016. "The Dynamics of Microinsurance Demand in Developing Countries Under Liquidity Constraints and Insurer Default Risk." *Journal of Risk and Insurance*, 83(1): 121–138.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection bias in predicting future utility." *The Quarterly Journal of Economics* 1209–1248.

Ludwig, Jens, Jeffrey R. Kling, and Sendhil Mullainathan. 2011. "Mechanism Experiments and Policy Evaluations." *Journal of Economic Perspectives*, 25(3): 17–38.

Macchiavello, Rocco, and Ameet Morjaria. 2014. "Competition and Relational Contracts: Evidence from Rwanda's Mills." *Unpublished*.

Macchiavello, Rocco, and Ameet Morjaria. 2015. "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports." *American Economic Review*, 105(9): 2911–45.

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao. 2013. "Poverty Impedes Cognitive Function." *Science*, 341(6149): 976–980.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al. 1995. *Microeconomic theory*. 1: Oxford university press New York.

Minten, Bart, Anneleen Vandeplas, and Johan FM Swinnen. 2011. "The broken broker system?: Transacting on agricultural wholesale markets in India (Uttarakhand)."Technical report, International Food Policy Research Institute (IFPRI).

Mobarak, Ahmed Mushfiq, and Mark Rosenzweig. 2012. "Selling Formal Insurance to the Informally Insured." Working Papers 1007, Economic Growth Center, Yale University.

Pauly, Mark V, Howard Kunreuther, and Richard Hirth. 1995. "Guaranteed renewability in insurance." *Journal of Risk and Uncertainty*, 10(2): 143–156.

Rampini, Adriano A, Amir Sufi, and S Viswanathan. 2014. "Dynamic risk management." *Journal of Financial Economics*, 111(2): 271–296.

Rampini, Adriano A, and S Viswanathan. 2010. "Collateral, Risk Management, and the Distribution of Debt Capacity." *The Journal of Finance*, 65(6): 2293–2322.

Rampini, Adriano A., and S. Viswanathan. 2016. "Household Risk Management." Working Paper 22293, National Bureau of Economic Research.

Rosenzweig, Mark R, and Kenneth I Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investment in Bullocks in India." *Journal of Political Economy*, 101(2): 223–44.

Schilbach, Frank. 2015. "Alcohol and Self-Control: A Field Experiment in India." *Unpublished*.

Stokey, Nancy L, and Edward C Lucas, Robert L with Prescott. 1989. *Recursive methods in economic dynamics.*: Harvard University Press.

Tarozzi, Alessandro, Aprajit Mahajan, Brian Blackburn, Dan Kopf, Lakshmi Krishnan, and Joanne Yoong. 2014. "Micro-loans, Insecticide-Treated Bednets, and Malaria: Evidence from a Randomized Controlled Trial in Orissa, India." *American Economic Review*, 104(7): 1909–41.

Thaler, Richard H, and Shlomo Benartzi. 2004. "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving." *Journal of Political Economy*, 112(S1): S164–S187.

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

Tversky, Amos, and Daniel Kahneman. 1981. "The framing of decisions and the psychology of choice." *Science*, 211(4481): 453–458.

Udry, Christopher. 1994. "Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria." *The Review of Economic Studies*, 61(3): 495–526.

UNCTAD. 2009. "World Investment Report 2009 - Transnational Corporations, Agricultural Production and Development."

Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, 97(2): 305–46.

Figures

Figure 1: Insurance vs. Risk Reduction

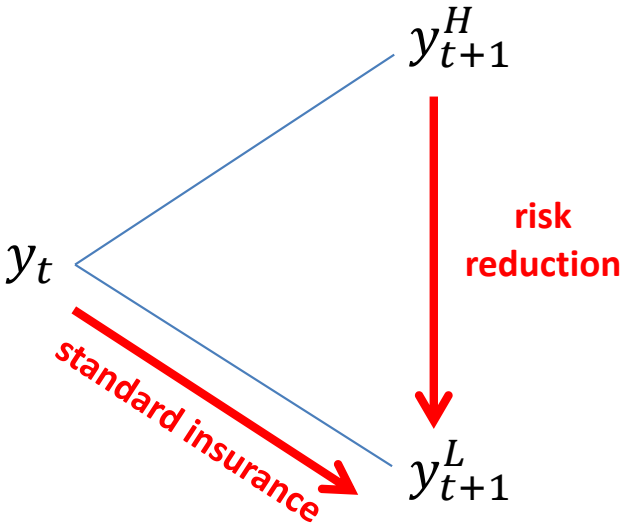
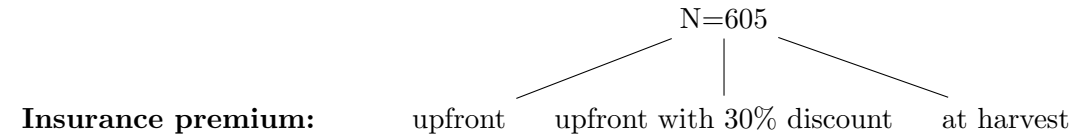


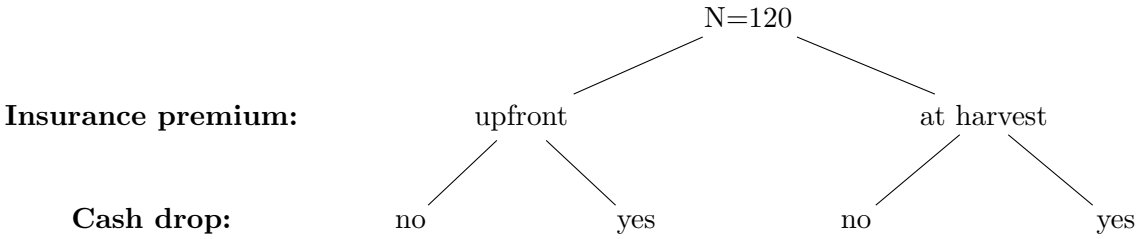
Figure 2: Experimental Design

(a) Design of Main Experiment



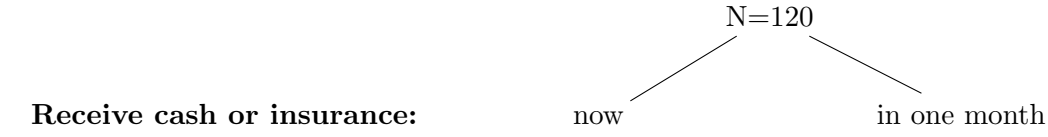
Notes: The experimental design randomized 605 farmers (approximately) equally across three treatment groups. All farmers were offered an insurance product; the only thing varied across treatment groups was the premium. In the first group (U1), farmers were required to pay the ("actuarially-fair") premium upfront, as is standard in insurance contracts. In the second group (U2), premium payment was again required upfront, but farmers received a 30% discount relative to (U1). In the third group (H), the full-priced premium would be deducted from farmers' revenues at (future) harvest time, including interest charged at the same rate used for the inputs the company supplies on credit (1% per month). Randomization across these treatment groups occurred at the farmer level and was stratified by Field, an administrative unit of neighboring farmers.

(b) Design of Cash Constraints Experiment



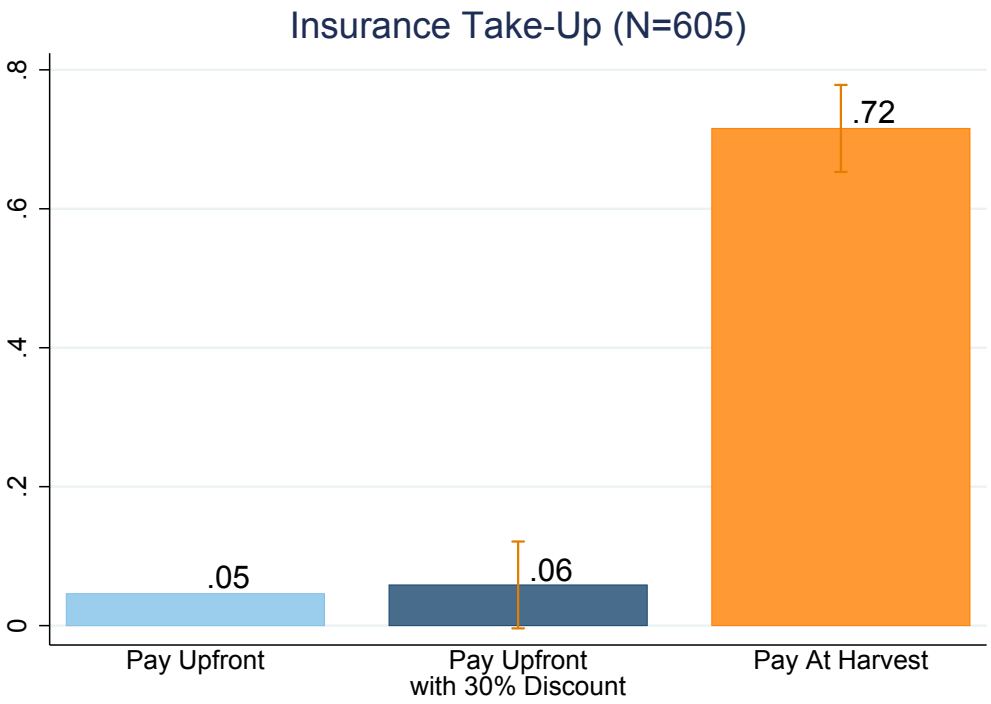
Notes: The experimental design randomized 120 farmers (approximately) equally across four treatment groups. The design cross-cut two treatments: pay-upfront vs. pay-at-harvest insurance, as in the main experiment, and a cash drop. At the beginning of individual meetings with farmers, those selected to receive cash were given an amount which was slightly larger than the insurance premium, and then at the end of the meetings farmers were offered the insurance product. Randomization across these treatment groups occurred at the farmer level and was stratified by Field.

(c) Design of Present Bias Experiment



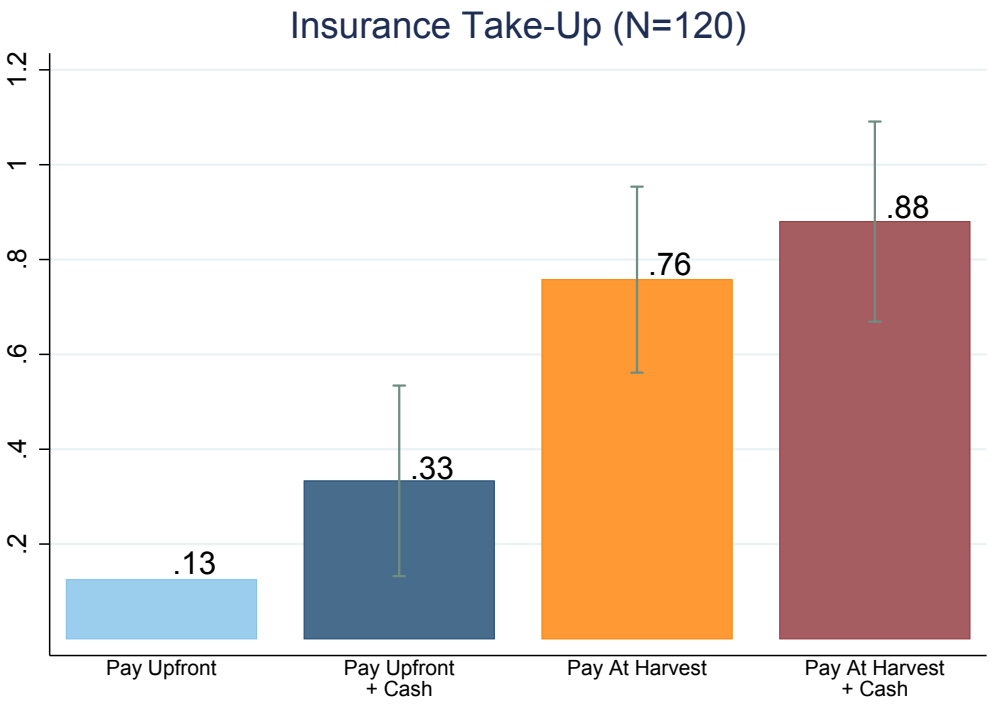
Notes: The experimental design randomized 120 farmers (approximately) equally across two treatment groups. Farmers in both groups were offered a choice between either a cash payment, equal to the "full-priced" insurance premium, or free enrollment in the insurance. Both groups had to make the choice during the meeting, but there was a difference in when it would be delivered. In the first treatment group, the *Receive Choice Now* group, farmers were told that they would receive their choice immediately. In the second group, the *Receive Choice in One Month* group, farmers were told that they would receive their choice in one month's time (the cash payment offered to farmers in this case included an additional month's interest). Randomization across these treatment groups occurred at the farmer level and was stratified by Field.

Figure 3: Main Experiment: Insurance Take-Up by Treatment Group



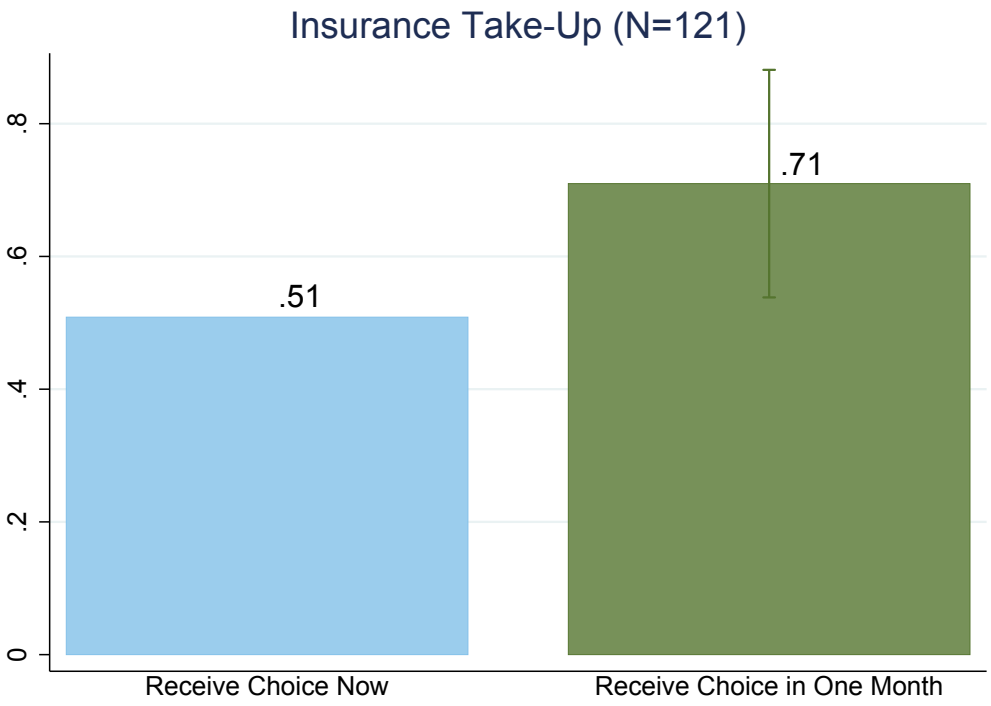
Notes: The figure shows insurance take-up rates across the three treatment groups in the main experiment. In the *Pay Upfront* group, farmers had to pay the full-price premium when signing up to the insurance. In the *Pay Upfront + 30% Discount* group, farmers also had to pay the premium at sign-up, but received a 30% price reduction. In the *Pay At Harvest* group, if farmers signed up to the insurance, then the premium (including accrued interest at 1% per month) would be deducted from their revenues at (future) harvest time. The bars report 95% confidence intervals from a regression of takeup on dummies for the treatment groups.

Figure 4: Cash Constraints Experiment: Insurance Take-Up by Treatment Group



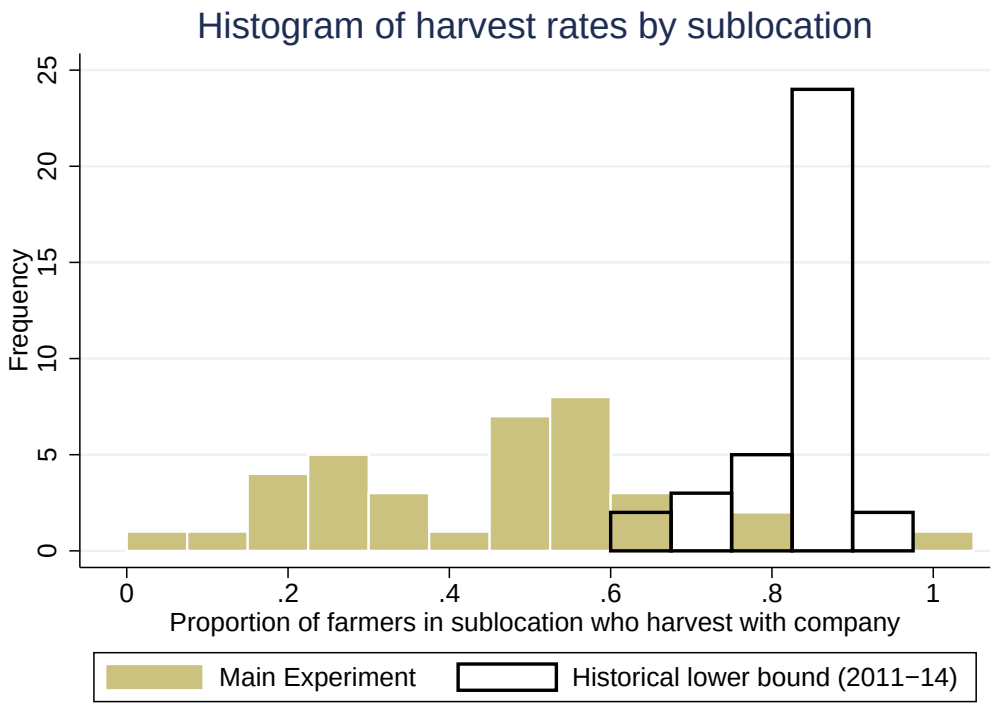
Notes: The figure shows insurance take-up rates across the four treatment groups in the cash constraints experiment. In the *Pay Upfront* group, farmers had to pay the premium when signing up for the insurance. In the *Pay Upfront + Cash* group, farmers were given a cash drop slightly larger than the cost of the premium, and had to pay the premium at sign-up. In the *Pay At Harvest* group, if farmers signed up for the insurance then the premium (including accrued interest at 1% per month) would be deducted from their revenues at (future) harvest time. In the *Pay At Harvest + Cash* group, farmers were given a cash drop equal to the cost of the premium and premium payment was again through deduction from harvest revenues. The bars report 95% confidence intervals from a regression of takeup on dummies for the treatment groups.

Figure 5: Present Bias Experiment: Insurance Take-Up by Treatment Group



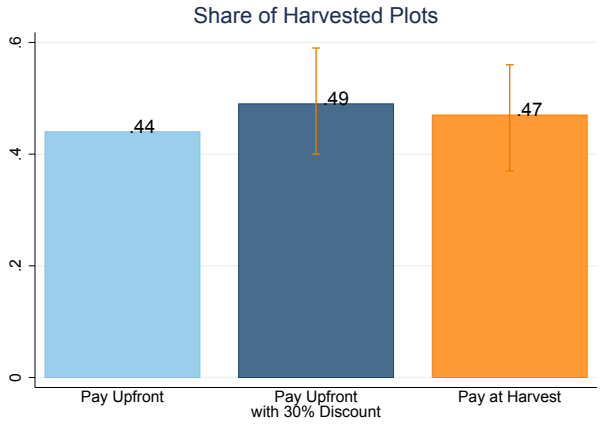
Notes: The figure shows insurance take-up rates across the two treatment groups in the present bias experiment. In the *Receive Now* group, farmers chose between an amount of money equal to the premium and free subscription to the insurance, knowing that they would receive their choice straight away. In the *Receive in One Month* group, farmers made the same choice, but knowing that they would receive whatever they chose *one month later*. The bars report 95% confidence intervals from a regression of takeup on dummies for the treatment groups.

Figure 6: Histogram of Harvesting With Company, by Sublocation



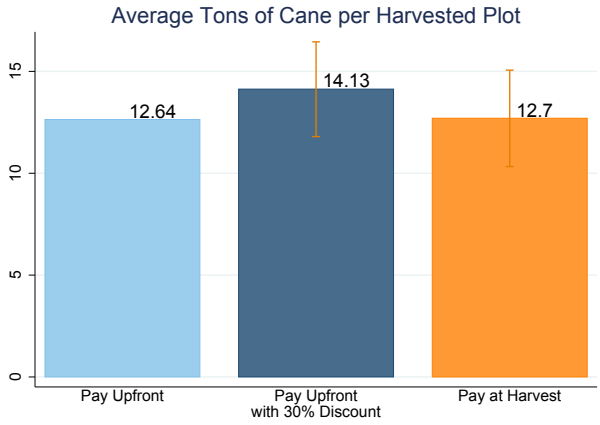
Notes: The histogram shows the proportion of farmers who harvested with the company in the sublocations in which we undertook the experiment. The data is by sublocation and we plot separate histograms for the main experiment (which is just for the farmers in our sample, who were due to harvest approximately twelve months after our experiment) and for the three year period prior to the experiment, from 2011 to 2014 (which is for all farmers in the sublocations). The historical measure is a lower bound on the harvest rate, since it is calculated as the proportion who harvested in the previous cycle who do not harvest this cycle, some of whom will not have grown cane this cycle. We note two things from the histograms. First, harvesting with the company is much lower during the experiment than historically, in line with the financial troubles at the company. Second, there is a large amount of geographic variation in the harvesting rate among farmers in our sample.

Figure 7: Proportion of Farmers Harvesting with the Company in Main Experiment



Notes: The figure shows the proportion of farmers from the main experiment who subsequently harvested with the company, as agreed under the contract. The bars report 95% confidence intervals from a regression of harvesting rates on dummies for the treatment groups.

Figure 8: Harvest Weight Conditional on Harvesting with the Company in Main Experiment



Notes: The figure shows the harvest weight, conditional on harvesting with the company, for farmers in the main experiment. The bars report 95% confidence intervals from a regression of harvest yields on dummies for the treatment groups.

Tables

Table 1: Main Experiment: Balance Table, Baseline Variables

	Upfront [U1]	Upfront-30% [U2]	Harvest [H]	P-value [U1-U2]	P-value [U1-H]	P-value [U2-H]	P-value [U-H]	N
Plot Size (Acres)	.309	.340	.319	.012**	.187	.219	.875	605
	(.128)	(.147)	(.133)					
Previous Yield	53.9	56.8	55.7	.23	.112	.67	.581	605
	(17.2)	(18.0)	(17.7)					
Man	.670	.734	.671	.082*	.732	.387	.762	573
	(.471)	(.442)	(.470)					
Age	48.3	48.3	45.4	.979	.06*	.02**	.025**	570
	(13.5)	(14.3)	(11.7)					
Land Cultivated (Acres)	2.79	3.13	2.81	.576	.421	.254	.852	562
	(3.34)	(4.01)	(2.98)					
Own Cow(s)	.762	.805	.803	.289	.25	.908	.508	569
	(.426)	(.397)	(.398)					
Portion of Income from Cane	3.30	3.38	3.23	.451	.817	.089*	.25	569
	(1.09)	(1.09)	(1.18)					
Savings for Sh1,000	.284	.282	.333	.979	.273	.44	.24	566
	(.452)	(.451)	(.472)					
Savings for Sh5,000	.089	.138	.130	.198	.345	.829	.558	565
	(.286)	(.346)	(.338)					
Expected Yield	72.5	70.4	66.8	.686	.522	.547	.533	573
	(100.)	(31.0)	(49.4)					
Expected Yield in Good Year	83.5	85.4	83.1	.814	.852	.908	.898	572
	(47.8)	(30.8)	(56.0)					
Expected Yield in Bad Year	53.4	54.0	52.3	.935	.889	.986	.937	571
	(40.6)	(21.5)	(47.2)					
Good Relationship with Company	.362	.343	.300	.919	.357	.933	.554	570
	(.482)	(.476)	(.459)					
Trust Company Field Assistants	3.02	2.84	2.80	.245	.11	.786	.312	569
	(1.01)	(1.01)	(1.09)					
Trust Company Managers	2.46	2.35	2.44	.449	.999	.598	.743	567
	(1.11)	(1.06)	(1.12)					

Notes: The table presents the baseline balance for the Main Experiment. *Plot Size* and *Previous Yield* are from the administrative data of the partner company and are available for each of the 605 farmers in our sample. The rest of the variables are from the baseline survey. These are missing for 32 farmers who denied consent to the survey. In addition, a handful of other values for specific variables is missing because of enumerator mistakes or because the respondent did not know the answer or refused to provide an answer. *Previous Yield* is measured as tons of cane per hectare harvested in the cycle before the intervention. *Man* is a binary indicator equal to one if the person in charge of the sugarcane plot is male. *Own Cow(s)* is a binary indicator equal to one if the household owns any cows. *Portion of Income from Cane* takes value between 1 ("None") to 6 ("All"). *Savings for Sh 1,000 (Sh 5,000)* is a binary indicator that equals one if the respondent says she would be able to use household savings to deal with an emergency requiring an expense of Sh 1,000 (Sh 5,000). 1 USD= 95 Sh. *Good Relationship with the Company* is a binary indicator that equals one if the respondent says she has a "good" or "very good" relationship with the company (as opposed to "bad" or "very bad"). *Trust Company Field Assistants* and *Trust Company Managers* are defined on a scale 1 ("Not at all") to 4 ("Completely"). P-values are based on specifications which include field fixed effects (since randomization was stratified at the field level). *p<0.1, **p<0.05, ***p<0.01.

Table 2: Main Experiment: Treatment Effects on Take-Up

	(1)	(2)	(3)	(4)	(5)
Pay Upfront with 30% Discount	0.004		0.013	0.003	0.015
	[0.033]		[0.033]	[0.032]	[0.033]
Pay At Harvest	0.675***	0.673***	0.680***	0.686***	0.694***
	[0.033]	[0.028]	[0.033]	[0.032]	[0.032]
Plot Controls	N	N	Y	N	Y
Farmer Controls	N	N	N	Y	Y
Mean Y Control	0.046	0.052	0.046	0.046	0.046
Observations	605	605	605	605	605

Notes: The table presents the results of the Main Experiment. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. Specification (2) bundles together treatment groups U1 (Pay Upfront) and U2 (Pay Upfront with 30% discount) as baseline group. *Plot Controls* are *Plot Size* and *Previous Yield*. *Farmer Controls* are all of the other controls reported in the balance table, Table 1. For each of the plot controls, we also include a dummy equal to one if there is a missing value (and recode missing values to an arbitrary value), so to keep the number of observations unchanged. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

Table 3: Main Experiment: Heterogeneous Treatment Effects by Wealth and Liquidity Constraints Proxies

	(1) Land Cultivated (Acres)	(2) Own Cow(s)	(3) Previous Yield	(4) Plot Size (Acres)	(5) Portion of Income from Cane	(6) Savings for Sh1,000	(7) Savings for Sh5,000
X*Pay At Harvest	-0.065**	-0.139*	-0.079**	-0.001	0.053*	-0.174**	-0.131
	[0.033]	[0.078]	[0.031]	[0.031]	[0.028]	[0.069]	[0.097]
X	-0.000	0.066	0.015	-0.022	-0.004	0.006	-0.016
	[0.017]	[0.044]	[0.020]	[0.019]	[0.016]	[0.043]	[0.059]
Pay At Harvest	0.706***	0.822***	0.673***	0.672***	0.540***	0.764***	0.725***
	[0.029]	[0.068]	[0.028]	[0.028]	[0.096]	[0.035]	[0.031]
Mean Y Control	0.052	0.052	0.052	0.052	0.052	0.052	0.052
Mean X	0.000	0.791	0.000	-0.000	3.311	0.300	0.120
S.D. X	1.000	0.407	1.000	1.000	1.126	0.459	0.326
Observations	562	569	605	605	569	566	565

Notes: The table shows heterogenous treatment effects on take-up from the Main Experiment, by different proxies for liquidity constraints and wealth. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance, and in each column the relevant heterogeneity variable (X) is reported in the column title. Treatments U1 (Pay Upfront) and U2 (Pay Upfront with 30% discount) are bundled together as baseline group, as specified in the pre-analysis plan. *Plot Size* and *Previous Yield* are from the administrative data of the partner company and are available for each of the 605 farmers in our sample. The rest of the variables are from the baseline survey. These are missing for 32 farmers who denied consent to the survey. In addition, a handful of other values for specific variables are missing because of enumerator mistakes or because the respondent did not know the answer or refused to provide an answer. *Land cultivated* is the standardized total area of land cultivated by the household. *Own Cow(s)* is a binary indicator for whether the household owns any cows. *Previous Yield* is the standardized tons of cane per hectare harvested in the cycle before the intervention. *Plot size* is the standardized area of the sugarcane plot. *Portion of Income from Cane* takes value between 1 ("None") to 6 ("All"). *Savings for Sh 1,000 (Sh 5,000)* is a binary indicator that equals one if the respondent says she would be able to use household savings to deal with an emergency requiring an expense of Sh 1,000 (Sh 5,000). 1 USD = 95 Sh. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

Table 4: Cash Constraints Experiment: Treatment Effects on Take-Up

	(1)	(2)	(3)	(4)
Pay At Harvest	0.603***	0.589***	0.635***	0.635***
	[0.077]	[0.078]	[0.105]	[0.107]
Cash	0.132*	0.128	0.167	0.177
	[0.079]	[0.079]	[0.110]	[0.111]
Pay At Harvest * Cash			-0.071	-0.100
			[0.156]	[0.159]
Plot Controls	N	Y	N	Y
Mean Y Control	0.125	0.125	0.125	0.125
P-value: Pay at Harvest = Cash	0.000	0.000	0.000	0.000
Observations	120	120	120	120

Notes: The table presents the results of the Cash Constraints Experiment. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. The baseline (omitted) group is the *Pay Upfront* group, where farmers had to pay the premium upfront and did not receive a cash drop. *Plot Controls* are *Plot Size* and *Previous Yield*. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

Table 5: Intertemporal Preferences Experiment: Treatment Effect on Take-Up

	(1)	(2)	(3)	(4)
Receive in One Month	0.233**	0.237**	0.286***	0.293***
	[0.089]	[0.092]	[0.107]	[0.109]
Plot Controls	N	Y	N	Y
Farmer Controls	N	N	Y	Y
Mean Y Control	0.508	0.508	0.508	0.508
Observations	121	121	121	121

Notes: The table presents the results of the Present Bias Experiment. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. The baseline (omitted) group is the *Receive Now* group, where farmers chose between an amount of money equal to the premium and free subscription to the insurance. In the *Receive Choice in One Month* group, farmers made the same choice, but were told that what chose would be delivered one month later (plus one month's interest if they chose cash). *Plot Controls* are *Plot Size* and *Previous Yield*. *Farmer Controls* are all the other controls reported in the main balance table, Table 1. For each of the plot controls, we also include a dummy equal to one if there is a missing value (and recode missing values to an arbitrary value), so to keep the number of observations unchanged. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

Table 6: Take-Up by (Subsequent) Harvest Status

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Pay at Harvest	0.675***	0.684***	0.592***	-0.054	0.707***	0.673***	0.680***	0.594***	0.096	0.695***
	[0.033]	[0.054]	[0.107]	[0.407]	[0.046]	[0.028]	[0.046]	[0.091]	[0.358]	[0.040]
Pay Upfront with 30% discount	0.004	0.008	-0.004	-0.316	0.025					
	[0.033]	[0.054]	[0.107]	[0.408]	[0.047]					
Pay at Harvest*Share Harvested in Field		-0.023					-0.018			
		[0.101]					[0.088]			
Pay at Harvest*Share Harvested in Subloc			0.189					0.180		
			[0.233]					[0.198]		
Pay at Harvest*Past Share Harvested in Subloc				0.869*					0.687	
				[0.484]					[0.425]	
Pay at Harvest*Plot Harvested					-0.071					-0.047
					[0.070]					[0.060]
Pay Upfront 30% Discount*Share Harvested in Field		-0.009								
		[0.100]								
Pay Upfront 30% Discount*Share Harvested in Subloc			0.017							
			[0.234]							
Pay Upfront 30% Discount*Past Share Harvested in Subloc				0.382						
				[0.484]						
Pay Upfront 30% Discount*Plot Harvested					-0.047					
					[0.071]					
Plot Harvested					0.041					0.018
					[0.056]					[0.044]
Mean Y Control	0.046	0.046	0.046	0.046	0.046	0.046	0.046	0.046	0.046	0.046
Observations	605	605	605	605	605	605	605	605	605	605

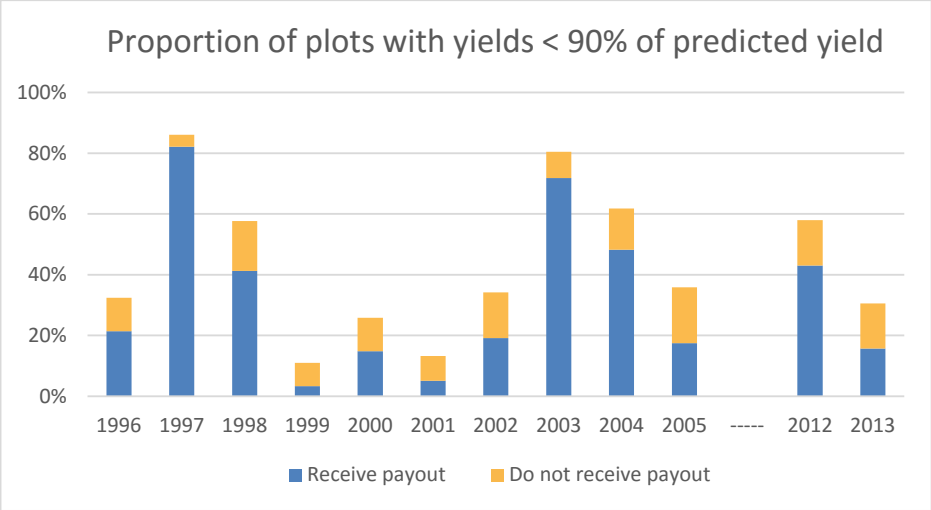
Notes: This table presents how take-up in the Main Experiment varies with the interaction of treatment group and subsequent harvesting behavior approximately twelve months later. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. *Share harvested in Field* is the proportion of farmers in the Field (an administrative, geographic unit) who harvest with the company. *Share harvested in Subloc* is the proportion of farmers in the Sublocation (a geographic identifier which is coarser than Field) who harvest with the company. *Past share harvested in Subloc* is the same variable, but instead covering the time period 2011-14, before the experiment, when side-selling was lower. *Plot harvested* is a binary indicator for whether the farmer harvests his plot with the company. Specifications (6)-(10) bundle the two pay-upfront treatment groups as baseline group. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

.

A Appendix figures and tables

A.1 Risk profile of the insurance

Figure A.1: Simulation of Insurance Payouts Based on Historical Data



Notes: The diagram shows what proportion of farmers would have received a positive payout from the insurance in previous years, and gives a sense of the basis risk of the insurance product. The numbers are based on simulations using historical administrative data on yields. The total bar height is the proportion of people who would have received an insurance payout under a single trigger design. It is broken down into those who still receive a payout when the second, area yield based trigger is added, and those who do not. We do not have historical data for the years 2006-2011.

Table A.1: Main Experiment: Heterogeneous Treatment Effects by Trust

	(1) Good Relationship with Company	(2) Trust Company Field Assistants	(3) Trust Company Managers
X*Pay At Harvest	-0.062	0.022	0.029
	[0.070]	[0.029]	[0.028]
X	0.087**	0.034*	0.027
	[0.040]	[0.018]	[0.017]
Pay At Harvest	0.726***	0.654***	0.640***
	[0.035]	[0.087]	[0.073]
Mean Y Control	0.052	0.052	0.052
Mean X	0.335	2.889	2.423
S.D. X	0.472	1.045	1.101
Observations	570	569	567

Notes: The table shows heterogeneities of the treatment effects of the pay-at-harvest premium on insurance take-up in the main experiment, by different proxies for trust toward the company. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. *Upfront Payment* and *Upfront Payment with 30% discount* treatment groups are bundled together as baseline group, as outlined in the pre-analysis plan. The relevant heterogeneity variable is reported in the column title. The notes of Table 1 provide a definition of the variables used in the heterogeneity analysis. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

Table A.2: Main Experiment: Heterogeneous Treatment Effects by Required Rates of Return

	(1)	(2)	(3)
	RRR on inputs	RRR 0 to 1 week	RRR 0 to 1 week minus RRR 1 to 2 weeks
X*Pay At Harvest	-0.124	0.099	0.001
	[0.141]	[0.114]	[0.152]
X	0.073	0.035	0.121
	[0.081]	[0.065]	[0.091]
Pay At Harvest	0.761***	0.685***	0.716***
	[0.054]	[0.042]	[0.029]
Mean Y Control	0.052	0.052	0.052
Mean X	0.324	0.269	-0.043
S.D. X	0.228	0.278	0.211
Observations	561	563	561

Notes: The table shows heterogeneities of the treatment effect of the pay-at-harvest premium on insurance take-up in the main experiment, by preferences in Money Earlier or Later experiments. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. *Upfront Payment* and *Upfront Payment with 30% discount* treatment groups are bundled together as baseline group, as outlined in the pre-analysis plan. The relevant heterogeneity variable is reported in the column title. These variables come from responses to hypothetical (Becker-DeGroot) choices over earlier or later cash transfers, from which we deduce three Required Rates of Returns. 'RRR for inputs' is the required rate of return which would (hypothetically) make farmers indifferent between paying for inputs upfront and having them deducted from harvest revenues. 'RRR 0 to 1 week' is the required rate of return to delay receipt of a cash transfer by one week. 'RRR 0 to 1 week - RRR 1 to 2 weeks' is the difference between the rates of return required to delay receipt of a cash transfer from today to one week from now, and from one week from now to two weeks from now. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

Table A.3: Cash Constraints Experiment: Balance Table

	Upfront [U]	Upfront + Cash [U + Cash]	Pay at Harvest [H]	Pay at Harvest + Cash [H + Cash]	P-value [H - U]	P-value [Cash - No cash]	N
Plot Size	.301	.290	.283	.282	.18	.967	120
	(.107)	(.092)	(.121)	(.088)			
Yield	54.3	57.8	61.4	54.1	.758	.745	120
	(18.4)	(17.9)	(14.8)	(17.0)			

Notes: The table presents baseline balancing for the Cash Constraints Experiment. *Previous Yield* is measured as tons of cane per hectare harvested in the cycle before the intervention. There are fewer covariates for this experiment as it did not have an accompanying survey, so we only have covariates from administrative data. P-values are based on specifications which include field fixed effects (since randomization was stratified at the field level). *p<0.1, **p<0.05, ***p<0.01.

Table A.4: Present Bias Experiment: Balance Table

	Receive Now	Receive in One Month	p-value	N
Plot Size	.328	.290	.085*	121
	(.109)	(.106)		
Yield	58.0	57.8	.571	121
	(20.1)	(21.3)		
Man	.793	.590	.009***	119
	(.408)	(.495)		
Age	48.3	47.7	.573	119
	(12.8)	(11.9)		
Land Cultivated (Acres)	3.81	2.67	.02**	118
	(3.87)	(1.72)		
Own Cow(s)	.844	.852	.987	119
	(.365)	(.357)		
Portion of Income from Cane	3.62	3.32	.193	119
	(1.12)	(.943)		
Savings for Sh1,000	.327	.295	.526	119
	(.473)	(.459)		
Savings for Sh5,000	.155	.065	.056*	119
	(.365)	(.249)		
Expected Yield	77.7	87.5	.47	119
	(65.3)	(38.4)		
Expected Yield in Good Year	95.1	109.	.322	119
	(70.7)	(48.4)		
Expected Yield in Bad Year	63.0	69.4	.682	119
	(61.7)	(32.0)		
Good Relationship with Company	.310	.316	.622	118
	(.466)	(.469)		
Trust Company Field Assistants	3.10	2.83	.315	119
	(1.02)	(1.01)		
Trust Company Managers	2.15	2.11	.32	119
	(1.13)	(1.03)		

Notes: The table presents baseline balancing for the Present Bias Experiment. *Plot Size* and *Previous Yield* are from the administrative data of the partner company and are available for each of the 605 farmers in our sample. The rest of the variables are from the baseline survey. These are missing for 2 farmers who denied consent to the survey. In addition, a handful of other values for specific variables is missing because of enumerator mistakes or because the respondent did not know the answer or refused to provide an answer. *Previous Yield* is measured as tons of cane per hectare harvested in the cycle before the intervention. *Man* is a binary indicator equal to one if the person in charge of the sugarcane plot is male. *Own Cow(s)* is a binary indicator equal to one if the household owns any cows. *Portion of Income from Cane* takes value between 1 ("None") to 6 ("All"). *Savings for Sh 1,000 (Sh 5,000)* is a binary indicator that equals one if the respondent says she would be able to use household savings to deal with an emergency requiring an expense of Sh 1,000 (Sh 5,000). 1 USD= 95 Sh. *Good Relationship with the Company* is a binary indicator that equals one if the respondent says she has a "good" or "very good" relationship with the company (as opposed to "bad" or "very bad"). *Trust Company Field Assistants* and *Trust Company Managers* are defined on a scale 1 ("Not at all") to 4 ("Completely"). P-values are based on specifications which include field fixed effects (since randomization was stratified at the field level). *p<0.1, **p<0.05, ***p<0.01.

A.2 Bounding the effect of the transfer across time

Households are both consumers and producers. The implications of this dual role have long been considered in development economics. In particular, in the presence of market frictions, separation may no longer hold, so that production and consumption decisions can no longer be considered separately (Rosenzweig and Wolpin 1993; Fafchamps et al. 1998). Above we considered the household's full dynamic problem, which incorporates discount factors and stochastic consumption paths. Often, however, we can apply a sufficient-statistic style approach, where we rely on observed behavior to tell us what we need to know, without having to estimate all of the parameters of the full optimization problem. In the case of intertemporal decisions, an individual's investment behavior, and in particular the interest rates of investments they do and do not make, can serve this role.

In this section we consider what observed investment behavior can tell us about hypothetical insurance take-up decisions, given the intertemporal transfer in insurance. Empirically, investment decisions may be easier to observe than discount factors and beliefs about consumption distributions (which are needed if we consider the full dynamic problem), and other studies provide evidence on interest rates in similar settings - both for investments made and for investments forgone. Using a simplified version of the model developed above, we consider under which conditions farmers would and would not take up insurance, given information on their other investment behavior.

To simplify, we now assume that at harvest time there are just two states of the world, the standard state h and the low state l, with the low state happening with probability p.[57] We assume that insurance is perfect - it only pays out in the low state (at time H), and that it is again actuarially fair. To simplify notation, in this section we denote by R the interest rate on the insurance *covering the whole period from the purchase decision until harvest time.* We also assume CRRA utility, so that $u(c) = c^{1-\gamma}/(1-\gamma)$.

Under this setup, the expected net benefit of a marginal unit of standard, upfront insurance is:

$$\beta\delta^H R\mathbb{E}[c_H(y_l)^{-\gamma}] - c_0^{-\gamma}$$

Consider first the case that the farmer forgoes a risk-free investment over the same time period which has rate of return R'. Then, first we know that paying upfront is at least as costly as a price increase in pay-at-harvest insurance of $\frac{R'}{R}$, and second we know that:

$$\beta\delta^H R'(p\mathbb{E}[c_H(y_l)^{-\gamma}] + (1-p)\mathbb{E}[c_H(y_h)^{-\gamma}]) - c_0^{-\gamma} < 0$$

Substituting this into the expected benefit of upfront insurance, we can deduce that farmers will not purchase standard insurance if:

$$\begin{aligned} & R\mathbb{E}[c_H(y_l)^{-\gamma}] < R'(p\mathbb{E}[c_H(y_l)^{-\gamma}] + (1-p)\mathbb{E}[c_H(y_h)^{-\gamma}]) \\ \Leftrightarrow & \frac{\mathbb{E}[c_H(y_l)^{-\gamma}]}{\mathbb{E}[c_H(y_h)^{-\gamma}]} < \frac{1-p}{\frac{R}{R'}-p} \end{aligned}$$

So, the farmer will not purchase insurance if under all consumption paths:

$$c_H(y_h) < Ac_H(y_l)$$

with A given by:

$$A = \left(\frac{1-p}{\frac{R}{R'}-p}\right)^{\frac{1}{\gamma}}$$

[57] Note that the following can be easily generalized so that these two states represent average outcomes when insurance does not and does pay out respectively.

Unsurprisingly, A is increasing in the (relative) forgone interest rate R/R', and decreasing in the CRRA γ. Also, A is increasing in the probability of the low state, p, suggesting that the intertemporal transfer is less of a constraint on insuring rarer events.

Similarly, we can consider the case where the farmer *makes* an investment over the period with risk-free interest rate R'. Under the same logic, we first know that a price raise of pay-at-harvest insurance of $\frac{R'}{R}$ is at least as costly as paying upfront, and second we also know the farmer will purchase insurance if, for all consumption paths:

$$c_H(y_h) > Ac_H(y_l)$$

The following tables report A for various values of R'/R, p, and γ. The tables thus reports how much consumption must vary between good and bad harvests in order to be sure about farmers' decisions to buy perfect insurance, given their investment decisions. In the case of forgone investments, it tells us the largest variation in consumption for which we can be sure that the farmer will still not buy perfect insurance; in the case of made investments, it tells us the smallest variation in consumption for which we can be sure that the farmer will buy perfect insurance. We note that A represents variation in *consumption* between states at harvest time - not variation in income, which is likely to be significantly larger. The effect can be sizeable. For example, for a risk which has a 20% chance of occurring, if the forgone investment has risk-free rate of return 50% higher than the interest rate charged on the insurance, then farmers with CRRA of 1 will forgo a perfect insurance product even when the consumption in the good state is 71.4% higher than consumption in the bad state.

$\gamma = 1$

		p				
		0.01	0.05	0.1	0.2	0.4
$\frac{R'}{R}$	1	1	1	1	1	1
	1.1	1.101	1.106	1.112	1.128	1.179
	1.2	1.202	1.213	1.227	1.263	1.385
	1.5	1.508	1.541	1.588	1.714	2.250
	2	2.020	2.111	2.250	2.667	6.000
	3	3.062	3.353	3.857	6.000	∞

$\gamma = 2$

		p				
		0.01	0.05	0.1	0.2	0.4
$\frac{R'}{R}$	1	1	1	1	1	1
	1.1	1.049	1.052	1.055	1.062	1.086
	1.2	1.097	1.101	1.108	1.124	1.177
	1.5	1.228	1.241	1.260	1.309	1.500
	2	1.421	1.453	1.500	1.633	2.449
	3	1.750	1.831	1.964	2.449	∞

$\gamma = 5$

		p				
		0.01	0.05	0.1	0.2	0.4
$\frac{R'}{R}$	1	1	1	1	1	1
	1.1	1.019	1.020	1.022	1.024	1.033
	1.2	1.038	1.039	1.042	1.048	1.067
	1.5	1.086	1.090	1.097	1.114	1.176
	2	1.151	1.161	1.176	1.217	1.431
	3	1.251	1.274	1.310	1.431	∞

A.3 Proofs and derivations

A.3.1 Background

States Each period t, which we will typically think of as one month, has a set of states S_t, corresponding to different income realizations. The probability distribution over states is assumed to be memoryless, so that $P(s_t = s)$ may depend on t, but is independent of the history at time t, $(s_i)_{i<t}$. We assume that the probability distribution of outcomes is cyclical, of period N, so that $S_t = S_{t+N}$ and $P(s_t = s) = P(s_{t+N} = s)\ \forall t, s$.

Income and wealth We denote wealth at the beginning of each period by w_t, so that $x_t = w_t + y_t$.

Dynamic programming problem

$V_t(x_t)$, the time t self's value function, is the solution to the following recursive dynamic programming problem:

$$V_t(x_t) = \max_{c_t} u(c_t) + \beta\delta\mathbb{E}_s[V^c_{t+1}(x_{t+1})] \tag{A.1}$$

subject to, for all $i \geq 0$, $\quad x_{t+i+1} = R(x_{t+i} - c_{t+i}) + y_{t+i+1}$

$$x_{t+i} - c_{t+i} \geq 0$$

where $V^c_t(x_t)$, the continuation value function, is the solution to equation A.1, but with $\beta = 1$, i.e.

$$V^c_t(x_t) = \max_{c_t} u(c_t) + \delta\mathbb{E}_s[V^c_{t+1}(x_{t+1})] \tag{A.2}$$

Because of the cyclicality of the setup, the functions $V_t(.) = V_{t+N}(.)$ and $V^c_t(.) = V^c_{t+N}(.)\ \forall t$.

Lemma A.1. *$\forall t \in \mathbb{R}^+$:*

1. *V_t, V^c_t exist, are unique, and are concave.*
2. *$\frac{dc_t}{dx_t} < 1$, so investments (and wealth in the next period) are increasing in wealth.*

Proof of Lemma A.1

Part (1) Since V^c is the solution to a recursive dynamic programming problem with convex flow payoffs, concave intertemporal technology, and convex choice space, theorem 9.6 and 9.8 in Stokey and Lucas (1989) tell us that V^c exists and is strictly concave. To expand further, the proofs, which are similar in method to subsequent proofs below, are as follows.

Existence & Uniqueness. Blackwell's sufficient conditions hold for the Bellman operator mapping V^c_{t+1} to V^c_t: monotinicity is clear; discounting follows by the assumption that $\delta R < 1$ - taking $a \in \mathbb{R}$, $V^c_{t+1} + a$ is mapped to $V^c_{t+1} + \delta R a$; the flow payoff $(u(c_t))$ is bounded and continous by assumption; compactness of the state-space is problematic, but given $\delta R < 1$ the stock of cash-on-hand will not amass indefinitely, so we can bound the state space with little concern (Stokey and Lucas (1989) provide more formal, technical methods to deal with the problem. Since it is not the focus of the paper, we do not go into more details). Thus, the Bellman operator is a contraction mapping, and iterating this operator implies the mapping from V^c_{t+N} to V^c_t is a contraction mapping also. V^c_t is a fixed point of this mapping, and thus exists and is unique by the contraction mapping theorem.

Concavity. Assume V^c_{t+N} is concave. Then, V^c_{t+N-1} is strictly concave, since the utility function is concave and the state space correspondence in convex, by standard argument (take $x_\theta = \theta x_a + (1-\theta)x_b$, expand out the definition of $V^c_{t+N-1}(x_\theta)$ and use the concavity of V^c_{t+N-1} and

the strict concavity of $u(.)$). Iterating this argument, we thus have that V_t^c is concave. Therefore, since there is a unique fixed point of the contraction mapping from V_{t+N}^c to V_t^c, that fixed point must be concave (since we will converge to the fixed point by iterating from any starting function; start from a concave function).

Part (2)

$$V_t(x_t) = max_c\, u(c) + \beta\delta\mathbb{E}[V_{t+1}^c(R(x_t - c) + y_{t+1})]$$

Since V_{t+1}^c is concave, this is a convex problem, and the solution satisfies:

$$u'(c_t) = \max\{\beta\delta R\mathbb{E}[V_{t+1}^{c'}(R(x_t - c_t) + y_{t+1})], u'(x_t)\}$$

Define $a(x_t) = x_t - c(x_t)$. Take $x_t' > x_t$, and suppose $a_t'(x_t') < a_t(x_t)$. Since $a_t' \geq 0$, we must have $a_t > 0$. Now, $a_t' < a_t$ implies $c_t' > c_t$, so $u'(c_t') < u'(c_t) = \beta\delta R\mathbb{E}[V^{c'}(Ra_t + y)] \leq \beta\delta R\mathbb{E}[V^{c'}(Ra_t' + y)] \leq u'(c_t')$. Contradiction. Thus $a_t'(x_t) \geq 0$. Since $V^{c'}(Ra_t + y_{t+1}) = u'(c_{t+1})$, the concavity of V^c also implies that c_{t+1} is increasing in x_t in the sense of first order stochastic dominance.

Proof of Lemma 1

Part (1) The intuition for the result is that $V_t^{c'} = u'(c_t(x_t))$ (combining the first order condition with the envelope condition), and u' and c are convex by prudence (with the convexity of c strenghtened by the borrowing constraint). The proof relies on showing that the mapping from $V_{t+1}^{c'}$ to $V_t^{c'}$ conserves convexity, $\forall t \in \mathbb{R}^+$. Then the proof follows as in 1 above: V_t^c is the fixed point of a contraction mapping which conserves convexity of the first derivative, hence $V_t^{c'}$ must be convex. We show that the mapping preserves convexity as follows, which is based on Deaton and Laroque (1992):

Suppose $V_{t+1}^{c'}$ is convex.

$$\begin{aligned} V_t^{c'}(x_t) &= u'(c_t) \\ &= \max\{\delta R\mathbb{E}[V_{t+1}^{c'}(R(x_t - c_t) + y_{t+1})], u'(x_t)\} \end{aligned}$$

Define G by $G(q, x) = \delta R\mathbb{E}[V_{t+1}^{c'}(R(x_t - u'^{-1}(q) + y_{t+1})]$.

G is convex in q and x: u' is convex and strictly decreasing, so u'^{-1} is convex (and so $-u'^{-1}$ is concave); $V_{t+1}^{c'}$ is convex and decreasing, so $V_{t+1}^{c'}(R(x_t - u'^{-1}(q)) + y_{t+1})$ convex in q and x (since f convex decreasing and g concave $\Rightarrow f \circ g$ convex); expectation is a linear operator (and hence preserves convexity).

Now $V_t^{c'} = \max[G(V_t^{c'}(x_t), x_t), u'(x_t)]$, or, defining $H(q, x) = \max\{G(q, x) - q, u'(x) - q\}$, then $V_t^{c'}$ is the solution in q of $H(q, x) = 0$.

H is convex in q and x, since it is the max of two functions, each of which are convex in q and x. Take any two x and x' and $\lambda \in (0, 1)$. Then $H(V_t^{c'}(x), x) = H(V_t^{c'}(x'), x') = 0$. Thus, by the convexity of H, $H(\lambda V_t^{c'}(x) + (1 - \lambda)V_t^{c'}(x'), \lambda x + (1 - \lambda)x') \leq 0$. Now, since H is decreasing in q, that means that $V_t^{c'}(\lambda x + (1 - \lambda)x') < \lambda V_t^{c'}(x) + (1 - \lambda)V_t^{c'}(x')$, i.e. $V_t^{c'}$ is convex.

Part (2) Clearly $\frac{d\mu_t}{dx_t} \leq 0$. Also, the distribution of x_{t+1} is increasing in the distribution of x_t, is the sense of first order stochastic dominance, by iterating Lemma A.1 part (2). Hence the result holds by the law of iterated expectations.

A.3.2 Insurance with perfect enforcement

Proof of Propostion 1

In the following, denote by a_t the assets held at the end of period t, so that $a_t = x_t - c_t$.

Suppose farmers have zero probability of being liquidity constrained before the next harvest when they buy pay-upfront insurance. Denote their (state-dependent) path of assets until harvest by $(a_t^U)_{t<H}$, given that they have purchased pay-upfront insurance. By the assumption that the farmers will not be liquidity constrained before harvest, $a_t^U > 0 \ \forall t < H$ and for all histories $(s_i)_{i\leq t}$. Now, suppose instead of pay-upfront insurance, they had been offered pay-at-harvest insurance. If they invest the money they would have spent on pay-upfront insurance in assets instead, so $a_t^H(s) = a_t^U(s) + R^{-H-t}$, then they can pay the pay-at-harvest premium at harvest time and have the same consumption path as in the case of pay-upfront, so they must be at least as well off. Similarly, suppose they optimally hold $(a_t^D)_{t<H}$ in the pay-at-harvest case. If instead offered upfront insurance, they can use some of these assets to instead buy insurance, so that $a_t^U(s) = a_t^D(s) - R^{-H-t}$. Since, by assumption $a_t^U(s) > 0$, doing so they can again follow the same consumption path as in the case of pay-at-harvest insurance, so pay-upfront insurance is at least as good as at-harvest insurance. Thus the farmer is indifferent between pay-upfront and pay-at-harvest insurance. As an aside, we note that this holds true even in the sophisticated $\beta\delta$ case, since so long as the farmer is not liquidity constrained he is passing forward wealth, meaning that paying the insurance at harvest time doesn't give him any extra ability to constrain his choices at harvest time than what he already has.

To first order, at time 0 the net benefit of pay-at-harvest insurance is $\beta\delta^H\mathbb{E}(Iu'(c_H)) - \beta\delta^H\mathbb{E}(u'(c_H))$, and of pay-upfront is $\beta\delta^H\mathbb{E}(Iu'(c_H)) - \beta\delta^H\mathbb{E}(u'(c_H)) - R^{-H}\lambda_0^H$ (note that the envelope theorem applies because, in the sequence problem, the insurance payout I does not enter any constraints before time H. This would no longer be the case if borrowing constraints were endogenous to next period's income). Thus the difference between the two is $R^{-H}\lambda_0^H$. Consider a pay-upfront insurance product which had premium $(1 - \frac{\lambda_0^H}{u'(c_0)})R^{-H}$. The net benefit would be

$$\begin{aligned}
&\beta\delta^H\mathbb{E}(Iu'(c_H)) - (1 - \frac{\lambda_0^H}{u'(c_0)})R^{-H}u'(c_0) \\
=&\beta\delta^H\mathbb{E}(Iu'(c_H)) - (u'(c_0) - \lambda_0^H)R^{-H} \\
=&\beta\delta^H\mathbb{E}(Iu'(c_H)) - \beta\delta^H\mathbb{E}(u'(c_H))
\end{aligned}$$

This is the net benefit of pay-at-harvest insurance.

Proof of Propostion 2

The net benefit of the pay-at-harvest insurance is $\beta\delta^H\mathbb{E}(V_H^c(w_H+y_H+I-1)) - \beta\delta^H\mathbb{E}(V_H^c(w_H+y_H))$. How this changes wrt x_0 is given by:

$$\begin{aligned}
&\frac{d}{dx_0}[\beta\delta^H\mathbb{E}(V_H^c(w_H + y_H + I - 1)) - \beta\delta^H\mathbb{E}(V_H^c(w_H + y_H))] \\
=&\frac{dw_H}{dx_0}\beta\delta^H[\mathbb{E}(V_H^{c'}(w_H + y_H + I - 1)) - \mathbb{E}(V_H^{c'}(w_H + y_H))]
\end{aligned}$$

Now, $\frac{dw_H}{dx_0} \geq 0$, by iterating lemma 1 back from period H to period 0. Also, $y_H + I - 1$ strictly second order stochastic dominates y_H by assumption, and $V_H^{c'}$ is strictly convex ($V^{c'''} > 0$ by lemma 1), so $\mathbb{E}(V_H^{c'}(w_H + y_H + I - 1)) - \mathbb{E}(V_H^{c'}(w_H + y_H)) < 0$. Thus, the value of pay-at-harvest insurance is decreasing with wealth.

The reduction in net utility from insurance arising from upfront premium payment is $R^{-H}\lambda_0^H$, by proposition 1. By lemma 1, this is also decreasing in wealth.

If the farmer is certain to be liquidity constrained before the next harvest, when starting with x_0, then his wealth at the start of the next harvest w_H will be the same as if he started with x_0', for any $x_0' < x_0$. This is because wealth in the next period is decreasing in wealth this period, so by the time the farmer has exhausted his wealth starting at x_0, he will also have exhausted his wealth starting at x_0'. Now, since the income process is memoryless, once the agent has exhausted his wealth, his distribution of wealth at the next harvest is the same, irrespective of his history. Thus the farmer has the same value of deductible insurance, regardless of whether he starts with x_0 or x_0', but the extra cost of the intertemporal transfer in the upfront insurance starting from x_0' means that the farmer has a lower value of upfront insurance.

Proof of Proposition 3

The proof is essentially the same as that of the second half of proposition 1.

A.3.3 Insurance with imperfect enforcement

Outside option $o(s_H, w_H)$

If the farmer chooses to sell to the company he receives profits $y(s)$ (comprising revenues minus a deduction for inputs provided on credit) plus any insurance payout $I(s)$, minus the insurance premium in the case of pay-at-harvest insurance. He also receives continuation value $r_C(s)$ from the relationship with the company, which is possibly state dependent. If he chooses to side sell, he receives outside option $o(s)$ [58], and saves the deductions for inputs provided on credit and for the deductible insurance premium, but loses the continuation value and any insurance payout. We abstract from any impact of insurance on the choice of input supply, since, as argued before, the choice set is limited, the double trigger design of the insurance was chosen to minimize moral hazard, and, as reported below, we see no evidence of moral hazard in the experimental data.

Default We will solve the farmer's problem backwards, starting with the decision of whether to side-sell conditional on the company not having defaulted on the farming contract. All decisions are as anticipated at time 0. We define the (endogenous) cost of side-selling in when the farmer does not have insurance as c_D, where we purposely use the same notation as above:

$$c_D = \mathbb{E}[V_H^c(w_H + o(w_H))] - \mathbb{E}[V_H^c(w_H + y_H)] \tag{A.3}$$

Proof of Proposition 5

Consider the decisions to sell to the company (i.e. not to side-sell). Denote the indicator functions for these decisions by D, with a subscript representing whether or not the insurer has already defaulted on the insurance contract, and a supercript denoting whether the farmer holds insurance, and if so the type of the insurace.

If the insurer has not already defaulted, they are:

$$\begin{aligned} D_I &= \mathbb{I}[c_D \geq 0] && \text{without insurance} \\ D_I^U &= \mathbb{I}[Iu'(c_H) + c_D \geq 0] && \text{with pay-upfront insurance} \\ D_I^D &= \mathbb{I}[Iu'(c_H) + c_D \geq u'(c_H)] && \text{with pay-at-harvest insurance} \end{aligned}$$

[58] We don't have detailed information on payments under side selling, but anecdotal evidence suggests that side sellers pay significantly less than the contract company, so a natural assumption would be that $o(s) = \alpha y(s)$, where $\alpha < 1$

If the insurer has already defaulted, they are:

$$\begin{aligned} D_D =& \mathbb{I}[c_D \geq 0] && \text{without insurance} \\ D_D^U =& \mathbb{I}[c_D \geq 0] && \text{with pay-upfront insurance} \\ D_D^D =& \mathbb{I}[c_D \geq u'(c_H)] && \text{with pay-at-harvest insurance} \end{aligned}$$

Since $I(s)u'(c_H(s))$ and $u'(c_H(s))$ are non-negative, and $Iu'(c_H)$ and $(I-1)u'(c_H)$ are larger when yields are low, the results follow.

Proof of Proposition 6

The basic intuition is that the extra loss from paying upfront is at most the premium when the farmer side-sells - if insurance did not change the decision to side-sell, then it is exactly the premium, if it did change the decision to side-sell, then by revealed preference the farmer loses at most the premium.

Formally, consider the net benefit of insurance, which is the benefit of the payout minus the cost of the premium payment. With perfect enforcement, we know that pay-at-harvest insurance is equivalent to upfront insurance with a percentage price cut of $\frac{\lambda_0^H}{u'(c_0)}$. With imperfect enforcement, denote the net benefit of pay-upfront insurance product by S_U, and the net benefit of pay-at-harvest insurance by S_D. Then:

$$\begin{aligned} \mathbb{E}[S_D - S_U] =& (1-p_I)(\Sigma_{d^D,d^U\in\{0,1\}}\mathbb{P}[D_I^U = d^U, D_I^D = d^D]\mathbb{E}[S_D - S_U | D_I^U = d^U, D_I^D = d^D]) \\ &+ p_I(\Sigma_{d^D,d^U\in\{0,1\}}\mathbb{P}[D_D^U = d^U, D_D^D = d^D]\mathbb{E}[S_D - S_U | D_D^U = d^U, D_D^D = d^D]) \end{aligned}$$

Now, $D_D^U \geq D_D^D$ and $D_I^U \geq D_I^D$. Also

$$\mathbb{E}[S_D - S_U | D_I^U = 1, D_I^D = 1] = \mathbb{E}[S_D - S_U | D_D^U = 1, D_D^D = 1] = 0$$

This leaves the cases where both default, or where pay-at-harvest defaults and pay-upfront doesn't. Conditional on $D_I^U = 0,\ D_I^D = 0$, or $D_D^U = 0,\ D_D^D = 0$, we have

$$S_D - S_U = \beta\delta^H u'(c_H)$$

When $D_I^U = 1, D_I^D = 0$, then

$$S_D - S_U = \beta\delta^H(u'(c_H) - (1-p_I)Iu'(c_H) - c_D) \leq \beta\delta^H u'(c_H)$$

Thus:

$$\begin{aligned} \mathbb{E}[S_D - S_U] \leq& (1-p_I)(\mathbb{P}[D_I^U = D_I^D = 0] + P[D_I^U = 1, D_I^D = 0])\beta\delta^H\mathbb{E}[u'(c_H)|D_I^D = 0] \\ &+ p_I(\mathbb{P}[D_D^U = D_D^D = 0] + \mathbb{P}[D_D^U = 1, D_D^D = 0])\beta\delta^H E[u'(c_H)|D_D^D = 0] \end{aligned}$$

with strict inequality iff $\mathbb{P}[D_I^U = 1, D_I^D = 0] > 0$. The right hand side can be rewritten to give:

$$\begin{aligned} \Leftrightarrow \mathbb{E}[S_D - S_U] \leq& (1-p_I)\mathbb{P}[D_I^D = 0]\beta\delta^H\mathbb{E}[u'(c_H)|D_I^D = 0] \\ &+ p_I\mathbb{P}[D_D^D = 0]\beta\delta^H\mathbb{E}[u'(c_H)|D_D^D = 0] \\ \Leftrightarrow \mathbb{E}[S_D - S_U] \leq& \mathbb{P}(\text{side-sell with at-harvest})\beta\delta^H\mathbb{E}(u'(c_H)|\text{side-sell with at-harvest}) \end{aligned}$$

We compare this to the surplus effect on the net benefit of upfront insurance of a further proportional price reduction of $\mathbb{P}(\text{side-sell with at-harvest})\frac{\mathbb{E}(u'(c_H)|\text{side-sell with at-harvest})}{\mathbb{E}(u'(c_H))}$, which is:

$$\begin{aligned} &\mathbb{P}(\text{side-sell with at-harvest})\frac{\mathbb{E}(u'(c_H)|\text{side-sell with at-harvest})}{\mathbb{E}(u'(c_H))}\mathbb{E}(u'(c_H)) \\ =&\mathbb{P}(\text{side-sell with at-harvest})\mathbb{E}(u'(c_H)|\text{side-sell with at-harvest}) \end{aligned}$$