American Economic Review 2018, 108(12): 3778–3813
https://doi.org/10.1257/aer.20171526

Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya†

By Lorenzo Casaburi and Jack Willis*

The gains from insurance arise from the transfer of income across states. Yet, by requiring that the premium be paid up front, 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 percent, compared to 5 percent for the standard pay-up-front 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 find that enabling farmers to commit to pay the premium just 1 month later increases demand by 21 percentage points. (*JEL* G22, I32, O13, O16, Q12, Q14)

In the textbook model of insurance, income is transferred across states of the world, from good states to bad. In practice, however, many insurance products also transfer income across time: the premium is paid up front 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

* Casaburi: University of Zurich, Schönberggasse 1, 8001 Zurich, Switzerland (email: lorenzo.casaburi@econ.uzh.ch); Willis: Columbia University, 1022 IAB, 420 W 118th Street, New York, NY 10027 (email: jw3634@columbia.edu). This paper was accepted to the *AER* under the guidance of Esther Duflo, Coeditor. We wish to thank the anonymous referees, as well as Vojta Bartos, Emily Breza, John Campbell, Michael Carter, Shawn Cole, Brian Dillon, David Dorn, 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, Eric Verhoogen, Andrew Weiss, and seminar audiences at ABCA conference, BASIS Technical Meeting, Barcelona Summer Forum, BREAD pre-conference, Cambridge, Columbia, Cornell (Dyson), CREST, CSAE Oxford conference, DIAL conference, ECBE conference, Harvard, IGC International Growth Week, IIES, INSEAD, Lindau meetings, NBER/BREAD development meeting, 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 partnership. The project was funded by SIEPR, USAID BASIS, and an anonymous donor, to whom we are very grateful. The main experiment was registered at the AEA RCT registry, ID AEARCTR-0000486, and the paper has completed the J-PAL pre-publication replication exercise. All errors are our own. We declare that we have no relevant or material financial interests that relate to the research described in this paper.

† Go to https://doi.org/10.1257/aer.20171526 to visit the article page for additional materials and author disclosure statement(s).

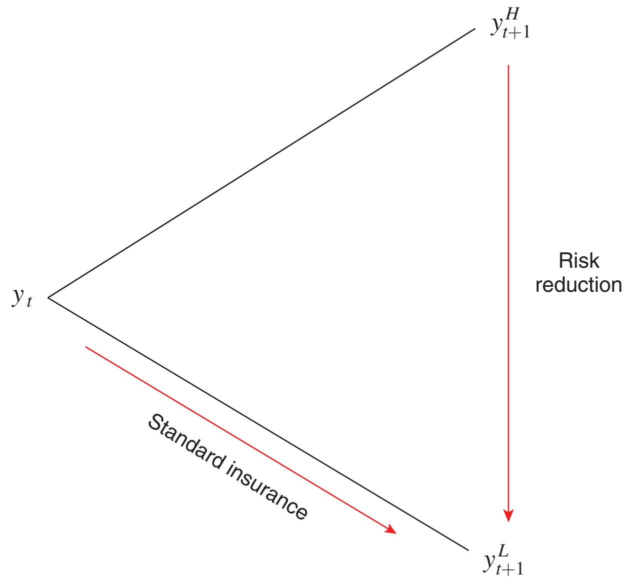

FIGURE 1. INSURANCE VERSUS RISK REDUCTION

Notes: In the textbook model, insurance offers *risk reduction*: income is transferred across states of the world, from good states to bad. In practice, however, standard insurance products also feature *a transfer of income across time*: the premium is paid up front with certainty, and any payouts are made in the future, if a bad state occurs.

additional factors, including liquidity constraints, intertemporal preferences, and trust. Moreover, the ability to *self*-insure depends on these same factors, so that charging the premium up front may reduce demand for insurance precisely when the potential gains are largest, for example among the poor.

This paper provides experimental evidence on the consequences of the transfer across time common 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 many attempts to increase adoption through 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.[1]

The insurance product we study eliminates the transfer across time by charging the premium at harvest, rather than up front. 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 directly from farmers' revenues at harvest time. In the experiment, the company

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

uses the same mechanism to collect insurance premiums: it offers the insurance product early in the growing cycle and deducts the premium (plus interest) at harvest, 14–16 months later.

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. Take-up of the standard, pay-up-front insurance was 5 percent: low, but not out of line with results for other "actuarially-fair" insurance products in similar settings. In contrast, when the premium was due at harvest time (including interest at 1 percent per month, the rate which the company charges on input loans), take-up was 72 percent, substantially higher than for other insurance products in similar settings. To benchmark this difference, in a third treatment arm we offered a 30 percent price discount on the pay-up-front premium. Take-up among this third group was 6 percent, not significantly different from take-up under the full-price up-front premium.

While the results show that delaying the premium payment increases insurance demand dramatically, charging premiums up front ensures that they are paid. Our contract farming setting facilitates at-harvest premium collection, as the company does the harvesting and can deduct premiums directly from farmers' payments. Yet, even in this setting, contract enforcement was an important issue. Before the plots in our sample were to be harvested, the company announced severe financial difficulties and temporarily shut down their factory. This resulted in long delays and uncertainty in harvesting, and in turn to unprecedented side-selling (52 percent versus a loose upper bound of 12 percent historically), highlighting the challenge of delaying premium payments in settings with imperfect contract enforcement.

Given this episode, it is especially important to consider the channels driving the high demand for pay-at-harvest insurance vis-à-vis the low demand for pay-up-front insurance. To do so, and to lay out the broader implications of the transfer across time in insurance, we first develop an intertemporal model of insurance demand. Then, guided by the model, we present additional empirical evidence, including from two small mechanism experiments (Ludwig, Kling, and Mullainathan 2011), on the role of liquidity constraints, time preferences, and imperfect contract enforcement.

The model shows that the transfer across time in insurance can help explain why the poor demand so little of it. The model is based on a buffer-stock saving model (Deaton 1991). Liquidity constraints are central and play a dual role. First, they make paying the premium up front 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 up front. Heterogeneous treatment effects show that both predictions hold, for the poor and for the liquidity constrained, in the main experiment. Finally, the model also shows how imperfect contract enforcement affects the demand for pay-up-front and pay-at-harvest insurance, and allows us to bound the relative effect, which we return to below.

In the first mechanism experiment, we test for a primitive form of cash constraints: did farmers just not have the cash to pay up front? To do so, we gave cash to a subset of farmers (enough to cover the premium), before offering them insurance later in the same meeting (similar to Cole et al. 2013). The cash drop did increase take-up of pay-up-front insurance, by 20 percentage points, but take-up remained much lower than take-up of pay-at-harvest insurance (in spite of any reciprocity effects that the cash gift may have induced).[2] This shows that our main result was not simply driven by farmers not having cash; rather, for most, *pay-up-front* insurance was not the marginal expenditure.

In the second mechanism experiment, we test whether allowing farmers to *pre-commit* to buying insurance in one month's time increases take-up. In the experiment, farmers had to choose between a cash payment, equal to the insurance premium, and free enrollment in the insurance. One-half of farmers were told they would receive their choice immediately, whereas the other half of farmers were told they would receive their choice in one month's time.[3] When farmers made choices for the subsequent month, insurance take-up increased by 21 percentage points. The increase suggests that farmers are present-biased as it is difficult to explain its magnitude with standard exponential discounting.[4] Moreover, the result provides additional (implicit) evidence of liquidity constraints, as time preferences should only matter when farmers cannot borrow at the market rate, given the experimental design.

Lastly, we consider the role of imperfect contract enforcement. The company's default on the *sales* contract implies an automatic default on the *insurance* contract. What role did any potential anticipation of such a default play in our results? First, the model clarifies that expectation of company default would reduce demand for pay-up-front insurance but would not increase demand for pay-at-harvest insurance (since neither premium nor payout would be paid under default). Yet, compared with other studies, the high take-up of pay-at-harvest insurance is the outlier, rather than the low take-up of pay-up-front. Second, the model also shows that we can bound the *differential* effect of imperfect enforcement on take-up by the *direct* effect of a price cut in the pay-up-front premium; in particular, that of a proportional price cut of the expected probability of side-selling times the relative (expected) marginal utility of consumption conditional on side-selling. Yet, the main experiment showed that demand for pay-up-front insurance was not very elastic. Thus, for default to be an important channel, one of these two terms would have had to be large; other results suggest they were not. Third, both survey measures of trust in the company and historical side-selling rates at the local level are uncorrelated with the effect of delaying premium payment on take-up, suggesting that any anticipation of company default played little role (insofar as it varied among farmers). To summarize, while it is certainly possible that expectation of default played some role in the differential

[2] To try to control for reciprocity, we cross-cut the cash treatment with a pay-up-front versus pay-at-harvest treatment, and the difference-in-differences effect of the cash was only 8 percent.

[3] Giving farmers the choice between the premium in cash 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.

[4] Under exponential discounting, the one-month delay in cash payment would substantially raise take-up only if the discount factor was very low. However, if this were the case, the insurance would still be unattractive, given that it only benefits the farmer more than one year after the cash payment.

take-up we observe in the main experiment, the evidence (including from the mechanism experiments, where take-up in the 1-month treatment was 72 percent, in spite of farmers facing the full default risk) suggests it was not the main driver.

This paper adds to several strands of literature. First, numerous papers have investigated the factors that constrain demand for agricultural insurance (Cole et al. 2013; Karlan et al. 2014). Yet, in spite of many innovative efforts, demand has remained stubbornly low, even at highly subsidized prices.[5] Most 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 instead focus on the transfer across time. From a policy point of view, the large demand effect we find, especially relative to the small effects of previous interventions, suggest that the challenge of enforcing delayed premium payments warrants serious consideration and future work. We discuss this challenge, along with ideas for addressing it, in Section V.

The timing of payment has been found to be an important determinant of demand for other products in similar settings, such as fertilizer, bednets, cookstoves, and water connections (Duflo, Kremer, and Robinson 2011; Devoto et al. 2012; Fink and Masiye 2012; Tarozzi et al. 2014; Levine et al. 2018). These products are investments, where the main benefits are the high expected returns. For insurance, in contrast, the main benefit is reducing consumption variation: making payments when times are good, in order to receive payouts when times are bad. If the timing is wrong, there are no benefits.

While we study the impact of offering loans to buy insurance on credit, several studies have tried insuring loans (Giné and Yang 2009; Carter, Cheng, and Sarris 2011; Karlan et al. 2011), finding that demand for loans increases little, and in some cases decreases, when doing so. 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 percent to 16 percent; 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.[6]

The transfer across time in insurance is studied implicitly in finance, but the focus is on how insurance companies benefit by investing the premiums (Becker and

[5] Karlan et al. (2014) find the highest take-up rates at actuarially fair prices among these studies, at around 40 percent; but most find significantly lower rates, for example Ahmed, McIntosh, and Sarris (2017); Giné, Townsend, and Vickery (2008), and Dercon et al. (2014) (in the untrained group) all find take-up rates of less than 5 percent. Moreover, the extensive margin alone does not give the whole story. Those studies which do find higher take-up rates typically allowed farmers to choose how many units of insurance to buy (in small increments) and few buy a meaningful amount of insurance on this intensive margin. For example, in Cole et al. (2013) those which did buy insurance almost universally bought just $1 worth, the minimal amount (Mobarak and Rosenzweig 2012 find a similar result).

[6] 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 up front will be costly. Second, we work with crop insurance, where seasonality increases the importance of the transfer across time.

Ivashina 2015), rather than on the cost for consumers, our focus. A recent exception is a largely theoretical literature (Rampini and Viswanathan 2010; Rampini, Sufi, and Viswanathan 2014) which argues that firms face a trade-off between financing and insurance. Rampini and Viswanathan (2016) apply similar reasoning to households.[7] These papers are part of a wide literature on how imperfect enforcement affects the set of financial contracts that exists (Bulow and Rogoff 1989; Ligon, Thomas, and Worrall 2002), to which we add by considering the implications of imperfect enforcement for the timing of insurance premiums.

Finally, delayed premium payments should certainly not be considered a panacea, even when they are enforceable. If the insurance product itself is not attractive, people will not buy it, regardless of the timing. For example, in Banerjee, Duflo, and Hornbeck (2014), take-up for health insurance was low even when subscribers could pay for it in installments, possibly because potential clients anticipated subsequent mismanagement in the claim and reimbursement process. Further, for several important forms of insurance, the premium is already not charged up front. Examples include Federal Crop Insurance in the United States, which is pay-at-harvest, futures contracts, and social security. Assessing the quantitative importance of the premium timing in other settings is an important open area for future research and policy consideration.

The remainder of the paper is organized as follows. Section I describes the setting in which the experiment took place. Section II presents the main experimental design and results. Section III develops an intertemporal model of insurance demand, which provides comparative statics and directs subsequent experiments. Section IV presents evidence on channels, from the main experiment and from two additional mechanism experiments. Section V discusses the policy implications of our results, both for crop insurance and for insurance markets more generally, and presents ideas for future work. Section VI concludes.

I. 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 one-quarter of total income for 80 percent of farmers in our sample. It has a long growing cycle (around 16 months), leading to a long transfer across time in pay-up-front insurance, and it is not seasonal. Once planted, crops last upward 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 percent own at least 1 cow, the average total 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.

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

A. *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 agribusinesses in East Africa and contracts around 80,000 small-holder farmers. As is standard in contract farming (a production form of increasing prevalence: see 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 percent per month interest on the loans.

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, 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.[8] Our administrative data do not tell us historical levels of side-selling, but do allow us to bound them. In the 3 years before the experiment, an average of 12 percent 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 to 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 IVD we

[8] Debts remain on plots even if plots are sold and are collected from future revenues regardless of who farms the plot.

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 II, we stratify treatment assignment at the field level in our experiments. Fields typically contain three to ten plots.

B. *Interlinked Insurance*

In standard insurance contracts, farmers pay the premium up front and so bear all of the contract risk.[9] Pay-at-harvest insurance reduces the contract risk farmers 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 defaulting on them.[10] 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 IIIC 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 preexisting 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).[11] Second, given the insurance design

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

[10] Our paper adds to the literature on interlinked contracts in developing country settings (Bardhan 1991). In particular, our work relates to research documenting informal insurance agreements in output and credit market contracts (Udry 1994; Minten, Vandeplas, and Swinnen 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 forthcoming, Casaburi and Reed 2017, Ghani and Reed 2017).

[11] 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.

(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 IVD 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.

II. 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 up front 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.

A. *Experimental Design*

Treatment Groups.—The experimental design randomized 605 farmers across 3 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 up front, at "full price" (which across the study spanned between 85 percent and 100 percent of the actuarially fair price). In the second group (U2), premium payment was again up front, but farmers received a 30 percent 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 percent per month).[12] Randomization occurred at the plot level and was stratified by field (i.e., groups of nearby plots). Farmers were unlikely to be sampled more than once, but in the few (3.5 percent) cases where they were, only 1 of their plots was retained in the sample.

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. Under the design, a farmer received a payout if 2 conditions were met: first, if their plot yield was below 90 percent of its predicted level; and second, if the average yield in their field was below 90 percent of

[12] The interest was added to the initial premium when marketing the insurance product. This interest rate is somewhat lower than that often charged in similar settings: 20–30 percent per year is common for small-scale agricultural loans, reflecting how farmers' crops act as (imperfect) collateral. However, comparing take-up under the 30 percent price-cut treatment versus under the pay-at-harvest treatment is equivalent to the considering the effect of the premium delay under a substantially higher interest rate. Inflation in Kenya was around 6 percent per annum during the study, so the real interest rate was 6 percent per annum.

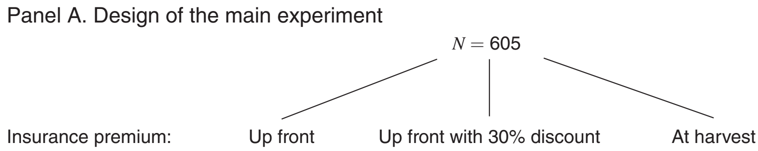

Notes: The experimental design randomized 605 farmers (approximately) equally across 3 treatment groups. All farmers were offered an insurance product; the only thing varied across treatment groups was the premium. In the first group (U1), farmers were required to pay the (actuarially-fair) premium up front, as is standard in insurance contracts. In the second group (U2), premium payment was again required up front, but farmers received a 30 percent 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 percent per month). Randomization across these treatment groups occurred at the farmer level and was stratified by field, an administrative unit of neighboring farmers.

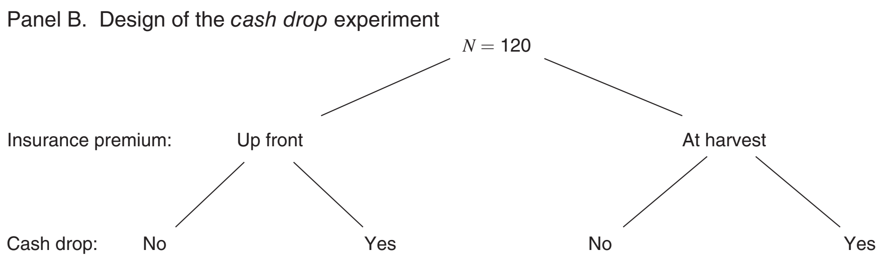

Notes: The experimental design randomized 120 farmers (approximately) equally across 4 treatment groups. The design cross-cut two treatments: pay-up-front versus pay-at-harvest insurance, as in the main experiment, and a cash drop. At the beginning of individual meetings with farmers, those selected 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.

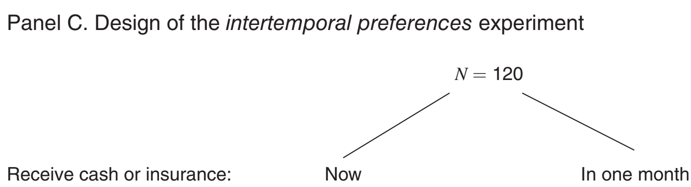

Notes: The experimental design randomized 120 farmers (approximately) equally across 2 treatment groups. Farmers in both groups were offered a choice between either a cash payment, equal to the "full-priced" insurance premium, or free enrollment in the insurance. Both groups had to make the choice during 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 2. Experimental Design

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.[13] The product was very much a partial insurance product: in the states where payouts were triggered, it covered one-half of plot losses below the 90 percent trigger, up to a cap of 20 percent 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 percent) of whom attended; the primary reason (75 percent) 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 percent), typically elderly, were deemed non-eligible at this stage. The resulting sample for randomization comprised 605 farmers. Compared to the 333 who did not enter the sample, they had slightly larger plots (0.81 versus 0.75 acres; p-value $= 0.015$) and similar yields (22.2 versus 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 one to two weeks to subscribe, with premiums collected either immediately or during revisits at the end of this period.

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.[14] 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

[13]The data included production, plot size, plot location, and crop cycle, and were available for a subsample of contracted plots for 1985–2006 and for all contracted plots from 2008 onward. The data were used to compute predicted yields at the plot and area level, which were needed for the double trigger insurance design. The historical data were 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 (online Appendix Figure A.1): 74 percent 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.

[14]The criteria used to select the sample were plot size: large plots were removed from the sample, to minimize the insurers' 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.

at least once. This choice was made because the yield prediction model performed 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 come from a short baseline survey, carried out by our survey team (before farmers were offered insurance) during the outreach worker visits described; 32 of the 605 farmers declined to be surveyed. As mentioned in Section I, the company keeps administrative data on all farmers in the scheme. They give us previous yields, harvest dates, plot size, and growing cycle, and enable us to track whether the farmer sells cane to the company at the end of the cycle, and their yield conditional on doing so.

B. *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 registered plan) we also report *p*-values when bundling pay-up-front 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 up front group U to H. We confirm below that the experiment results are robust to the inclusion of baseline controls.

C. *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. 2013; 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, Giné, and Vickery 2017). The central hypothesis tested in this paper is that low take-up is in part 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:

$$T_{if} = \alpha + \beta Discount_i + \gamma Harvest_i + \eta_f + \epsilon_{if}. \quad (1)$$

Figure 3 summarizes the take-up rates across the 3 treatment groups. For groups U2 and H, it also includes 95 percent 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, up-front premium is low, at 5 percent. While low, this finding is consistent with numerous existing crop insurance

Table 1—Main Experiment: Balance Table, Baseline Variables

	Up-front [U1]	Up-front-30% [U2]	Harvest [H]	*p*-value [U1–U2]	*p*-value [U1–H]	*p*-value [U2–H]	*p*-value [U–H]	*N*
Plot size (acres)	0.309 (0.128)	0.340 (0.147)	0.319 (0.133)	0.012	0.187	0.219	0.875	605
Previous yield	53.9 (17.2)	56.8 (18.0)	55.7 (17.7)	0.23	0.112	0.67	0.581	605
Man	0.670 (0.471)	0.734 (0.442)	0.671 (0.470)	0.082	0.732	0.387	0.762	573
Age	48.3 (13.5)	48.3 (14.3)	45.4 (11.7)	0.979	0.06	0.02	0.025	570
Land cultivated (acres)	2.79 (3.34)	3.13 (4.01)	2.81 (2.98)	0.576	0.421	0.254	0.852	562
Own cow(s)	0.762 (0.426)	0.805 (0.397)	0.803 (0.398)	0.289	0.25	0.908	0.508	569
Portion of income from cane	3.30 (1.09)	3.38 (1.09)	3.23 (1.18)	0.451	0.817	0.089	0.25	569
Savings for Sh 1,000	0.284 (0.452)	0.282 (0.451)	0.333 (0.472)	0.979	0.273	0.44	0.24	566
Savings for Sh 5,000	0.089 (0.286)	0.138 (0.346)	0.130 (0.338)	0.198	0.345	0.829	0.558	565
Expected yield	72.5 (100)	70.4 (31.0)	66.8 (49.4)	0.686	0.522	0.547	0.533	573
Expected yield in good year	83.5 (47.8)	85.4 (30.8)	83.1 (56.0)	0.814	0.852	0.908	0.898	572
Expected yield in bad year	53.4 (40.6)	54.0 (21.5)	52.3 (47.2)	0.935	0.889	0.986	0.937	571
Good relationship with company	0.362 (0.482)	0.343 (0.476)	0.300 (0.459)	0.919	0.357	0.933	0.554	570
Trust company field assistants	3.02 (1.01)	2.84 (1.01)	2.80 (1.09)	0.245	0.11	0.786	0.312	569
Trust company managers	2.46 (1.11)	2.35 (1.06)	2.44 (1.12)	0.449	0.999	0.598	0.743	567

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 1 if the person in charge of the sugarcane plot is male. *Own cow*(*s*) is a binary indicator equal to 1 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 1 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 US$ = 95 Sh. *Good relationship with the company* is a binary indicator that equals 1 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 of 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).

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

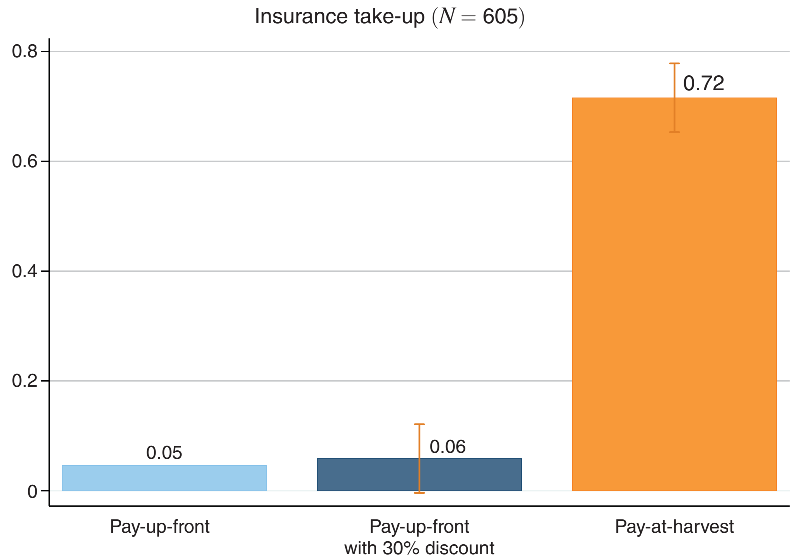

Figure 3. Main Experiment: Insurance Take-Up by Treatment Group

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

on take-up.[15] Take-up of the pay-at-harvest, interlinked insurance contract (H) is 72 percent, a 67 percentage point increase from the baseline, pay-up-front (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 up front.

The third result, which allows us to benchmark the importance of the second, is that a 30 percent price cut in the up-front premium does not have a statistically significant effect on take-up rates. The point estimate of the effect is 1 percentage point; a 20 percent increase in take-up which implies an elasticity of 0.39 $((1/6)/(30/70))$. Estimates of this elasticity in other settings vary substantially (between 0.4 and 2), both across studies and at different points along the demand curve within studies. Our point estimate is at the lower end of this range, though our confidence interval spans it. One possible explanation for our low estimate is that insurance purchase was a binary choice in our study: farmers could only insure their entire plot, not a fraction of it (at an average cost of \$18). In many other studies, including Karlan et al. (2014) and Cole et al. (2013) (where the estimated elasticities were 2 and 1 respectively), farmers could choose how many units of insurance

[15] Ex post, of course, only one state of the world is realized. Therefore, across multiple seasons, in a sense pay-at-harvest insurance does transfer income over time: on net farmers are paying in good seasons and getting paid in bad seasons. However, ex ante, when a farmer is making their take-up decision, pay-at-harvest insurance concerns purely a transfer across states of the world.

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

	(1)	(2)	(3)	(4)	(5)
Pay-up-front 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	No	No	Yes	No	Yes
Farmer controls	No	No	No	Yes	Yes
Mean dependent variable (pay-up-front group)	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 1 if the farmer took up the insurance. Specification 2 bundles together treatment groups U1 (pay-up-front) and U2 (pay-up-front 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 1 if there is a missing value (and recode missing values to an arbitrary value), so to keep the number of observations unchanged. *Mean dependent variable* (*pay-up-front group*) reports the mean of the dependent variable in the pay-up-front group. All columns include field fixed effects.

to buy, in small increments, and those who took-up insurance typically bought just a single unit at low cost (e.g., $1 in Cole et al. 2013). These small purchase decisions may well be more elastic than the larger purchase decision faced by farmers in our experiment.

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-up-front product (U1), significant at the 1 percent level, whereas the 30 percent 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-up-front (U2) products is also significant at the 1 percent level. Column 2 pools the up-front 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 2 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 percent) said they would while 3 (7.5 percent) said they would not. The remaining 5 (12.5 percent) 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 percent) were able to do so. Of this second wave of farmers, after reminding those who had forgotten the terms, 28 (85 percent) 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 interlinking product and insurance markets, has high take-up at actuarially fair price levels, while its standard, pay-up-front equivalent has low take-up (even with a substantial price cut), consistent with experience in other settings.

III. An Intertemporal Model of Insurance Demand

To understand the forces which could be driving the experimental results, and to motivate our subsequent attempts to identify them, we develop a model which captures both the cross-state and cross-time transfers in insurance. We introduce pay-up-front and pay-at-harvest insurance products into a standard dynamic consumption model and allow for three main channels which could affect their relative demand: liquidity constraints, time preferences, and imperfect enforcement.[16] A full exposition of the model, with derivations and proofs (as well as a brief treatment of the supply-side implications) can be found in online Appendix B.

A. *Buffer-Stock Savings Model*

The background model is a buffer-stock savings model, as in Deaton (1991), with the addition of naïve present-biased preferences and cyclical income flows (representing agricultural seasonality).[17] In the model, households receive income in each period, denoted y_t, and have access to a risk-free asset with constant rate of return R, but are subject to a borrowing constraint. Denoting cash-on-hand once income is received by x_t, the household faces the following maximization sequence problem in period t:

$$V_t(x_t) = \max_{(c_{t+i})_{i\geq 0}} u(c_t) + \beta E\left[\sum_{i=1}^{\infty} \delta^i u(c_{t+i})\right], \tag{2}$$

subject to

$$x_{t+i+1} = R(x_{t+i} - c_{t+i}) + y_{t+i+1},$$

$$x_{t+i} - c_{t+i} \geq 0, \quad \forall i \geq 0.$$

[16] 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 online Appendix C we report basic quantitative bounds for the effect of the transfer across time on insurance demand using this approach.

[17] Naïve-$\beta\delta$ discounters believe that they will be exponential discounters in future periods (and so may have incorrect beliefs about future consumption). There is evidence for such naïveté in other settings (DellaVigna and Malmendier 2006). In contrast, sophisticated-$\beta\delta$ discounters realize they will still be present-biased in the future. With the exception of Proposition 2 and online Appendix Lemmata B.1 and B.2 (which may no longer hold, since concavity and uniqueness of the continuation value function is no longer guaranteed), the results below hold in the sophisticated case, but with β replaced by a state-specific discount factor which is a function of the marginal propensity to consume.

The model gives the following iterated Euler equation, which is useful for considering the importance of the timing of premium payment:

$$(3)\qquad u'(c_t) \;=\; \max\{\beta\delta RE[u'(c_{t+1})], u'(x_t)\} \\ \;=\; \beta(R\delta)^H E[u'(c_{t+H})] + \lambda_t^{t+H},$$

where $\lambda_t^{t+H}(x_t)$ represents distortions in transfers from t to $t+H$ arising from potential borrowing constraints (it is a discounted sum of the Lagrange multipliers on the constraints).

The model also implies that both the benefit of risk reduction and the cost of the intertemporal distortion arising from liquidity constraints are decreasing in wealth (i.e., $d^3V_{t+H}/dx_t^3 > 0$ and $d\lambda_t^{t+H}/dx_t < 0$, online Appendix Lemma B.2), which we will rely on when deriving how the demand for insurance varies with wealth. Intuitively, the comparative statics follow from the poor being less able to self-insure: the existence of the borrowing constraint induces precautionary saving and hence a concave consumption function (Zeldes 1989; Carroll and Kimball 2005), so that among the poor, changes in wealth translate into larger changes in consumption.

B. *Insurance with Perfect Enforcement*

We start with the case where insurance contracts are perfectly enforceable. Farmers can buy one unit of insurance, which gives state-dependent payout I at harvest, normalized so that $E[I] = 1$. The purchase decision is made at time 0, while premium payment is either up front, at time 0, or at harvest, at time H. When paid at harvest, the premium is 1—the expected payout (commonly referred to as the actuarially-fair price)—whilst when paid up front the premium is R^{-H}. Under this setup, the farmer's decision to take up insurance is, to first-order:[18]

(4) Take up insurance if and only if

$$\begin{cases} \beta\delta^H E[u'(c_H)] \;\leq\; \beta\delta^H E[Iu'(c_H)] & \text{(pay-at-harvest insurance)} \\ R^{-H}u'(c_0) \;\leq\; \beta\delta^H E[Iu'(c_H)] & \text{(pay-up-front insurance)} \end{cases}.$$

This shows that, for pay-at-harvest insurance, it is variation in marginal utility of consumption across states which matters, whereas for pay-up-front insurance, variation across both states *and* time matters. To compare the two products we use the iterated Euler equation, equation (3).

PROPOSITION 1: *If farmers face a positive probability of being liquidity constrained before harvest, they prefer pay-at-harvest insurance to pay-up-front*

[18] We use first-order approximations at several points. They are reasonable in our setting for a number of reasons: the premium is small (3 percent of average revenues) and the insurance provides low coverage (a maximum payout of 20 percent of expected revenue); we care about differential take-up by premium timing, so second-order effects which affect up-front 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 IVD.

insurance; otherwise they are indifferent. To first order, the difference is equivalent to a proportional price cut in the up-front premium of $\lambda_0^H/u'(c_0)$.

Intuitively, paying the premium up-front, 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 they will not, then asset holdings can simply adjust to offset the difference. Liquidity constraints are therefore key, and intertemporal preferences (being defined over flows of utility rather than over flows of money) only matter through them. But liquidity constraints are also closely tied to wealth in the model (specifically, to deviations from permanent income, rather than permanent income itself). Combining Proposition 1 and online Appendix Lemma B.2 gives the following, for a marginal unit of insurance.

PROPOSITION 2: *The net benefit of pay-at-harvest insurance is decreasing in wealth. So too is the cost of paying up-front, rather than at harvest. Among farmers sure to be liquidity constrained before harvest, the latter dominates, so the net benefit of pay-up-front insurance is increasing in wealth.*[19]

Thus, while the benefit of risk reduction (pay-at-harvest insurance) is higher among the poor, they may buy less (pay-up-front) 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, but they are also more likely to face liquidity constraints *before* harvest, making illiquid investments more costly.

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 IVC, where the delay is one month).

PROPOSITION 3: *The benefit of delaying premium payment by 1 period is, to first order, equivalent to a proportional price cut in the up-front premium of* $\lambda_0^1/u'(c_0)$.

Thus, the one-period delay only increases demand if the farmer is liquidity constrained at time 0. Moreover, comparing Propositions 1 and 3 shows that a one-month delay will have a small effect relative to a delay until harvest (when H is large), unless either liquidity constraints are particularly strong at time 0, or there is present bias. Present bias closes the gap in two ways: first, future liquidity constraints are discounted by β in the discounted sum of Lagrange multipliers, λ_0^H, and second, the individual naïvely believes that he will be less likely to be liquidity

[19] The general point that the gap between pay-up-front 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 such as collateral requirements, and models sometimes take it as an assumption.

constrained in the future. Equivalently, the one-month delay allows the farmer to commit at time 0 to a purchase at time 1, and hence to overcome their present bias.

C. *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 up-front premium. Accordingly, imperfect enforcement has implications for both the demand and supply of pay-at-harvest insurance.

We assume that both sides may default on the insurance contract. At the beginning of the harvest period, with probability p (unrelated to yield) the insurer defaults on the contract, without reimbursing any up-front premiums. The farmer then learns his yield and, if the insurer has not defaulted, can himself default on any at-harvest premium, subject to some utility cost d and the loss of any insurance payouts due.[20] Denoting whether the farmer chooses to pay the at-harvest premium by the (state-dependent) indicator function D_P, then to first-order:

$$D_P \;:=\; \mathbf{1}[Iu'(c_H) + d \;\geq\; u'(c_H)]. \tag{5}$$

Ex ante, anticipation of potential default drives an additional difference between up-front and at-harvest insurance:

(6) Difference in net benefit of at-harvest and up-front

$$= \underbrace{R^{-H}\lambda_0^H}_{\text{liquidity constraint term}} + \underbrace{\beta\delta^H pE[u'(c_H)]}_{\text{premium saved when insurer defaults before harvest}}$$

$$+ \underbrace{\beta\delta^H(1-p)E[(1-D_P)(u'(c_H) - d - Iu'(c_H))]}_{\text{premium saved, minus cost of default and loss in insurance payouts, when farmer defaults}}.$$

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

Interlinked Insurance.—Interlinking the insurance contract with the sales contract affects contractual risk: default on one contract entails default on the other. In the above framework, we can define the (now endogenous) cost to the farmer of defaulting on insurance, d, to be the cost of defaulting on the sales contract by selling to another buyer (side-selling), i.e., the value of the production relationship relative to their outside option. This will typically be positive, in which case interlinking helps to enforce the pay-at-harvest premium (this is why credit is often

[20] In practice the farmer may face considerable uncertainty about both yields and insurance payouts when deciding to default, which shrinks the difference between pay-up-front 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.

interlinked). However, if the farmer wishes to side-sell for some other reason, then d will be negative, in which case interlinking encourages default on the premium.

Importantly, strategic side-selling by the farmer to avoid the pay-at-harvest premium is unlikely, since the premium is only marginal if d is close to zero (thus any expected farmer 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 decrease the incentive to side-sell, whereas for those with high yields, pay-at-harvest premiums increase the incentive to side-sell.

PROPOSITION 4: *If pay-at-harvest insurance affects side-selling, it makes those anticipating high yields more likely to side-sell, and those anticipating low yields less likely to side-sell.*

Strategic or not, ex ante expectations of default drive a wedge between pay-up-front and pay-at-harvest insurance. The following proposition allows us to bound the effect, in Section IVD, and shows how expectations of *non-selective* default (i.e., company default, or farmer default which is independent of yield) work by reducing demand for pay-up-front insurance, rather than increasing demand for pay-at-harvest insurance: intuitively, the net benefit of pay-at-harvest insurance is zero in such cases.

PROPOSITION 5: *The possibility of default in the interlinked contract drives a wedge between pay-at-harvest and pay-up-front insurance, bound above by a price cut in the up-front premium of*

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

Further, in so far as default is non-selective (i.e., independent of yield), it does not affect demand for pay-at-harvest insurance, to first order.

D. *Implications and Extensions*

The transfer across time in insurance has several implications beyond the focus of this paper. For instance, 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-up-front 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 insurance and self-insurance could explain why insurance demand is often decreasing with risk aversion (Clarke 2016). Relatedly, the transfer across time also changes the relationship between insurance and credit: for risk reduction purposes, they may be complements, not substitutes.

IV. 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 enforcement. The evidence comes from two small mechanism experiments, as well as additional results from the main experiment described in Section II.[21] We focus on liquidity constraints and present bias because both have been widely documented among similar populations (Loewenstein, O'Donoghue, and Rabin 2003; Cohen and Dupas 2010; Schilbach forthcoming; and Duflo, Kremer, and Robinson 2011, who estimate that 86 percent of their sample are present-biased at least some of the time). Additionally, present bias has implications for the interpretation of our results, both because future selves may regret the decision to forgo pay-up-front insurance, and because it introduces a role for commitment (Ashraft et al. 2006; Casaburi and Macchiavello forthcoming; Schilbach forthcoming). We focus on imperfect contract enforcement (especially in light of the temporary closure of the factory) since, to the extent that it is anticipated, it mechanically drives a difference between pay-up-front and pay-at-harvest insurance. In online Appendix D, we discuss several additional potential (behavioral) channels, which warrant future work.

A. *Is Up-Front Payment More Costly for the Poor and the Liquidity Constrained?*

Income variation is more costly for the poor, so they should have higher demand for risk reduction. So goes the common argument. Yet the poor typically show *lower* demand for 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 up front. If so, in our experiment we would expect the gap between pay-up-front and pay-at-harvest insurance to be higher among the poor.

Here we report how demand for pay-up-front 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 yields 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. In order to gain power, we bundle together the two pay-up-front groups (full price and 30 percent discount), as stated when registering the trial, giving the regression model:

$$T_{if} = \alpha + \beta Harvest_i + \gamma x_i + \delta Harvest_i \times x_i + \nu_f + \epsilon_{if}. \tag{7}$$

Table 3 presents the results, which show that the treatment effect does indeed vary by proxies for wealth and liquidity constraints. While not all the interaction coefficient estimates are significant, delaying premium payments until harvest

[21] We also 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 (Clarke 2016, Mobarak and Rosenzweig 2012, Elabed et al. 2013), the presence of informal insurance (Mobarak and Rosenzweig 2012), lack of information and understanding about insurance (Cai, De Janvry, and Sadoulet 2015; Handel and Kolstad 2015), and unfavorable claims ratios (Cole et al. 2013).

Table 3—Main Experiment: Heterogenous Treatment Effect by Wealth and Liquidity Constraints Proxies

Heterogeneity variable (X)	Land cultivated (1)	Own cow(s) (2)	Previous yield (3)	Plot size (4)	Portion of income from cane (5)	Savings for Sh 1,000 (6)	Savings for Sh 5,000 (7)
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 dependent variable (pay-up-front group)	0.052	0.052	0.052	0.052	0.052	0.052	0.052
Mean heterogeneity variable (X)	0.000	0.791	0.000	0.000	3.311	0.300	0.120
SD heterogeneity variable (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 heterogeneous 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 1 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-up-front) and U2 (pay-up-front with 30 percent discount) are bundled together as baseline group, as specified in the registered plan. The relevant heterogeneity variable is reported in the column title. *Mean dependent variable* (*pay-up-front group*) reports the mean of the dependent variable in the pay-up-front group. For each of the heterogeneity variables (X), we report their mean (*Mean heterogeneity variable*) and standard deviation (*SD heterogeneity variable*). *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 1 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 US$ = 95 Sh. All columns include field fixed effects.

increases 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.[22] Of course, these are heterogeneous treatment effects and so cannot be interpreted causally, as there could be confounders. Moreover, 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).

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.

[22] There is less margin for take-up heterogeneity in pay-up-front 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-up-front are larger among the poor. Further, existing studies on pay-up-front insurance typically find lower take-up among the poor and the liquidity constrained (Cole et al. 2013).

B. *Do People Buy Pay-Up-Front Insurance, Given Enough Cash to Do So?*

In line with the importance of liquidity constraints, when we surveyed farmers in the pay-up-front group about why they did not purchase insurance, their most common answer was not having the cash. In this section we present a second experiment, which investigates this very basic notion of liquidity constraints by asking: if we gave farmers the cash, would they use it to buy pay-up-front insurance?

In the experiment, half of the 120 targeted farmers were allocated to a cash drop treatment. Under this treatment, during the baseline survey, enumerators gave farmers a cash gift, which was slightly more than the price of the insurance premium (farmers were offered the insurance product around an hour later). The treatment, which mimics closely one of the arms in Cole et al. (2013), ensured that farmers did have enough cash to pay the up-front premium if they wished to. It therefore tests whether pay-up-front insurance was the marginal expenditure, or whether farmers had preferred uses for the money, such as consumption or higher return investments.[23] We also cross cut this cash-drop treatment with a premium-timing treatment: half of the farmers were offered pay-up-front insurance and half pay-at-harvest insurance. Stratification occurred at the field level and online Appendix Table A.2 shows that the treatment groups were balanced.

Figure 4 presents the results. The main finding is that the cash drop did raise take-up of pay-up-front insurance somewhat (from 13 percent to 33 percent), but that the impact of the cash drop was much smaller than that of switching to a harvest-time premium (which raised take-up to 76 percent). Thus, while basic cash constraints may account for a part of our main results, other channels mattered substantially more, consistent with the cash drop being small relative to the plausible cross-period variation in wealth which motivates the model.

The cash gift may also have increased demand through a reciprocity effect, whereby farmers bought insurance just to reciprocate the gift (a standard concern with cash drop designs). The cross-cutting design helps us shed light on this. In the pay-at-harvest group, the cash drop should not have increased demand through a liquidity channel (if anything a small wealth effect should have decreased demand), yet it did lead to higher take-up (from 76 percent to 88 percent), suggesting reciprocity did play a role. Taking the difference-in-differences between the pay-up-front and pay-at-harvest groups goes some way to controlling for reciprocity. This difference-in-differences estimate, for the effect of the cash drop on the demand for pay-up-front insurance, is 8 percent. While imprecisely estimated, we take this as further evidence that basic cash constraints had little effect on take-up of pay-up-front insurance, especially relative to delaying premium payment until harvest time.

Table 4 confirms these results. 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 2 treatments at the 1 percent level

[23] We note that the design 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 up front, which is no longer credible when either they can pay at harvest, or when they are given the cash.

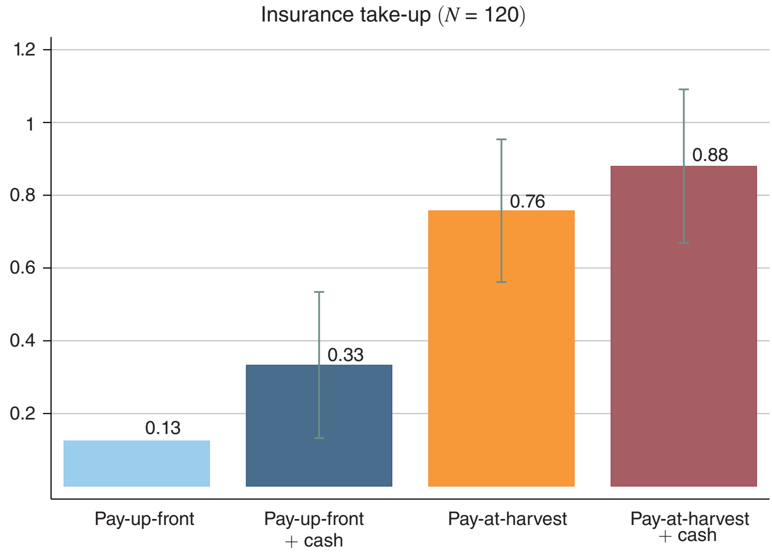

Figure 4. Cash Drop Experiment: Insurance Take-Up by Treatment Group

Notes: The figure shows insurance take-up rates across the four treatment groups in the cash drop experiment. In the *pay-up-front* group, farmers had to pay the premium when signing up for the insurance. In the *pay-up-front* + *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 percent 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 percent confidence intervals from a regression of take-up on dummies for the treatment groups.

Table 4—Cash Drop 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	No	Yes	No	Yes
Mean dependent variable (pay-up-front group)	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 drop experiment. The dependent variable is a binary indicator equal to 1 if the farmer took up the insurance. The baseline (omitted) group is the pay-up-front group, where farmers had to pay the premium upfront and did not receive a cash drop. *Mean dependent variable* (*pay-up-front group*) reports the mean of the dependent variable in the pay-up-front group. *Plot controls* are plot size and previous yield. All columns include field fixed effects.

(p-value 0.00004). The coefficient on Cash is significant at the 10 percent level in column 1 and remains similar in size but loses some 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 one-half of the difference between pay-up-front and pay-at-harvest insurance.

C. *Does Take-Up Increase If Farmers Can Commit Today to Buy Insurance in the Near Future?*

The third experiment tests whether take-up increases if farmers can *commit* up front to buying insurance using a windfall paid in the near future, i.e., in one month (versus buying insurance up front using a windfall paid up front). Similar small differences in timing have been shown to increase savings in other settings, such as save-more-tomorrow programs (Thaler and Benartzi 2004).

We randomly allocated a sample of 120 farmers to 2 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 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. Online Appendix Table A.3 reports the balance test across the 2 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 Sh 5,000; all but one of the pairwise correlations between these variables are positive. As discussed below, results are robust to the inclusion of these variables as controls.

Three aspects of the design allow us to isolate the role of intertemporal preferences. First, offering the choice between insurance for free or cash, rather than the option to buy insurance, allowed farmers to pre-commit their future income. In particular, 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 (as in the cash drop experiment) ensuring the farmer could take-up the insurance if she wanted to. Third, delaying both the cash gift *and* the insurance sign-up in the receive-choice-in-one-month treatment mitigated the traditional trust concerns associated with standard time preference experiments (Andreoni and Sprenger 2012). In particular, it meant that receipt of either choice depended on the field officer revisiting the household, removing differential trust concerns.[24]

Figure 5 shows that take-up in the receive-choice-in-one-month group was 72 percent, compared to a baseline of 51 percent in the receive-choice-now group. This 21 percentage point increase shows that a change of only 1 month in the timing of the premium payment had a large impact on insurance take-up.[25] Table 5 confirms

[24] It is still possible, though implausible, that farmers thought field officers were more likely to return if they chose 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.

[25] We note that the baseline take-up for the receive-choice-now group is larger than the take-up in the group pay-up-front + cash in the cash-constraints experiment. The two groups are drawn from different samples and so

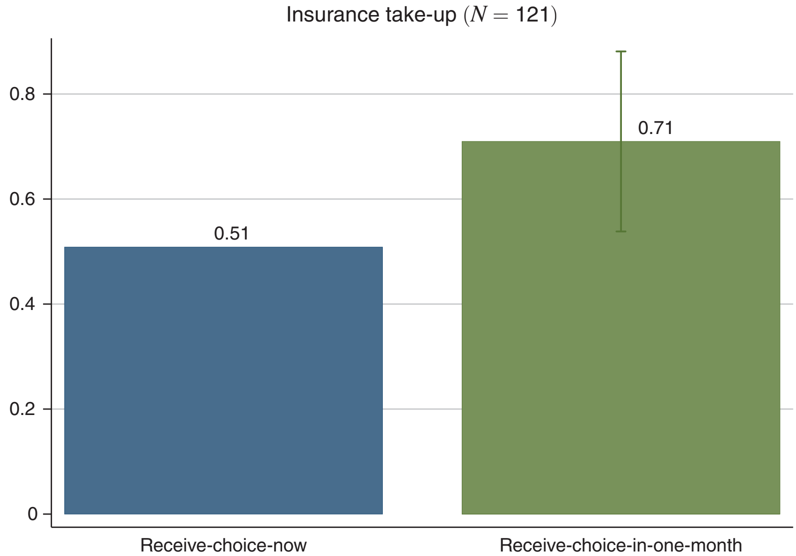

FIGURE 5. INTERTEMPORAL PREFERENCES EXPERIMENT: INSURANCE TAKE-UP BY TREATMENT GROUP

Notes: The figure shows insurance take-up rates across the two treatment groups in the intertemporal preferences experiment. In the *receive-choice-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-choice-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 percent confidence intervals from a regression of take-up on dummies for the treatment groups.

TABLE 5—INTERTEMPORAL PREFERENCES EXPERIMENT: TREATMENT EFFECT ON TAKE-UP

	(1)	(2)	(3)	(4)
Receive-choice-in-one-month	0.233	0.237	0.286	0.293
	[0.089]	[0.092]	[0.107]	[0.109]
Plot controls	No	Yes	No	Yes
Farmer controls	No	No	Yes	Yes
Mean dependent variable (receive now group)	0.508	0.508	0.508	0.508
Observations	121	121	121	121

Notes: The table presents the results of the intertemporal preferences experiment. The dependent variable is a binary indicator equal to 1 if the farmer took up the insurance. The baseline (omitted) group is the receive-choice-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). *Mean dependent variable* (*receive-choice-now* group) reports the mean of the dependent variable in the receive-choice-now group. *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 1 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.

are not directly comparable. However, the difference is larger than we might have expected, and ultimately we do not know why. One hypothesis is that 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. Another hypothesis relates to a literature dating back to Knetsch and Sinden (1984): receive-choice-now may have captured the willingness-to-accept, while pay-up-front + cash may have captured the willingness-to-pay, including an endowment effect from handing farmers the cash at the start of the visit.

these results across different specifications. The gap between the 2 treatments begins statistically significant at 5 percent and becomes statistically significant at 1 percent when adding farmer controls. We note that the point estimate raises from 0.23 in the baseline specification with field fixed effects (column 1) to 0.29 when adding both set of controls, though the difference in the 2 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.

The large effect is hard to explain with exponential discounting, but it is consistent with present bias. 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 very low δ (especially as a 30 percent 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 1-month delay, since such a product still transfers income over more than 1 year (as the harvest cycle is 16–18 months). In contrast, under present bias, a low β leads to a large difference between paying now and in one month, without making insurance paid for in one month unattractive.[26] This relies on the receive-choice-in-one-month treatment providing farmers with a commitment device on how to use their cash windfall, and hence allowing them to overcome any time inconsistency.[27] Importantly, the large effect also shows that farmers were liquidity constrained at the time of the experiment. If farmers could freely borrow money at the prevailing interest rate, then their time preferences would be irrelevant and take-up should not have differed between the two treatment groups.

While present bias can lead to under-subscription in pay-up-front insurance, might it lead to over-subscription and hence future regret in pay-at-harvest insurance? While 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 percent of farmers said they would not take up the product again.

Before moving to the next channel, we note that in the main experiment we elicited measures of preferences over the timing of cash flows, using standard (Becker-DeGroot) Money Earlier or Later questions (Cohen et al. 2016). We did not find heterogeneous treatment effects by these Required Rate of Return variables, as shown in online Appendix Table A.1. It is not uncommon 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

[26] 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.

[27] A competing explanation could 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.

disbursements) may fail to hold predictive power on other domains, such as how to use that cash.[28]

D. *Imperfect Enforcement*

Anticipation that either party may default before harvest drives a wedge between take-up of pay-up-front and pay-at-harvest insurance, as shown in Section IIIC. Here we consider the importance of this channel. While we find some 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.

As discussed in the introduction, 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 afterward, leading to uncertainty among farmers as to when harvesting would happen (if at all). As a result, only 48 percent of farmers in our sample harvested with the company. Those that did not either side-sold or uprooted the crop (for brevity, below we refer to both cases as side-selling). Those who harvested with the company received any insurance payouts due, while those who side-sold were ineligible. Figure 6 plots the harvesting rate by sublocation, and for comparison also plots a loose lower bound for it historically.[29] The figure shows that the harvesting rate was much lower than historically (and that it varied substantially by sublocation). Moreover, we see no evidence of similar firm crashes historically in the limited harvesting data that we have (a small subsample of plots, spanning substantial parts of the last 20 years), in line with our understanding from discussions with company management.

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

Did Insurance Affect Side-Selling.—We can rule out any sizable effect of insurance on side-selling, in line with the partially index-based design of the insurance product (and the results of our model). Given the low take-up of pay-up-front insurance, panel A of Figure 7 effectively reports the intent-to-treat of offering pay-at-harvest insurance on harvesting with the company. It shows that there was no level-effect on side-selling, in spite of high take-up. However, insurance could still have affected *who* side-sold. If so, Proposition 4 showed that, were side-selling strategic, pay-at-harvest insurance would make those with low yields less likely to side-sell and those with high yields more likely to. Hence, yield conditional on harvesting with the

[28] For instance, Kaur, Kremer, and Mullainathan (2015) find no correlation between lab experiment measures of time inconsistency and workers choices on effort and labor contracts. A recent experimental literature considers what such questions elicit, and suggests limitations with using them to measure intertemporal preferences (Andreoni and Sprenger 2012; Augenblick, Niederle, and Sprenger 2015; Cohen et al. 2016).

[29] The historical measure of the harvesting rate is a loose 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 do not 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.

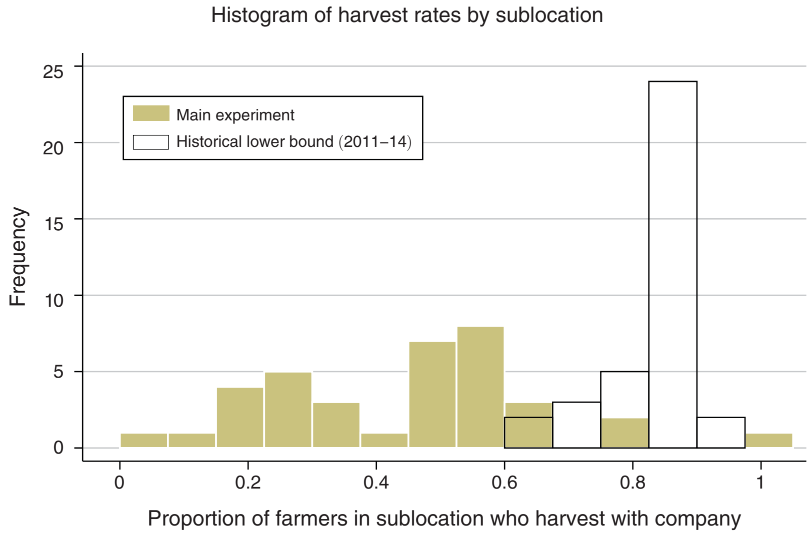

Figure 6. Main Experiment: Histogram of Harvesting with Company, by Sublocation

Notes: The histogram shows the proportion of farmers who harvested with the company in the sublocations in which we undertook the main experiment. The data are 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 12 months after our experiment) and for the 3-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.

company would be higher among the pay-up-front group. Panel B of Figure 7 shows there was no difference across groups, suggesting that the decision to side-sell was not strategic. While we do not observe yields among farmers who side-sold, and hence ultimately cannot say whether side-selling was selective, the lack of a strategic response to insurance, as well as anecdotal evidence that side-selling was largely driven by whether illegal harvesters happened to come to the farmers location, both provide suggestive evidence that side-selling was unrelated to yields.

The null result for the intent-to-treat effect of pay-at-harvest insurance on yields also shows that insurance did not induce moral hazard. If present, moral hazard would also have led to lower yields in the pay-at-harvest treatment group (the same direction as any effect of selective side-selling, so we can rule out the two effects canceling each other out). The lack of moral hazard is in line with many of the inputs being provided by the company, as well as the partially index-based design of the insurance product.

Did Anticipation of Default Affect Take-Up Differentially?—Given the extent of side-selling ex post, it is particularly important to consider the role of ex ante expectations of contract risk in driving the difference in take-up between pay-up-front and pay-at-harvest insurance. Here we present three sets of evidence which suggest that the role was limited. Before doing so, we note that we have already shown that other channels matter, in our two mechanism experiments and heterogeneous treatment

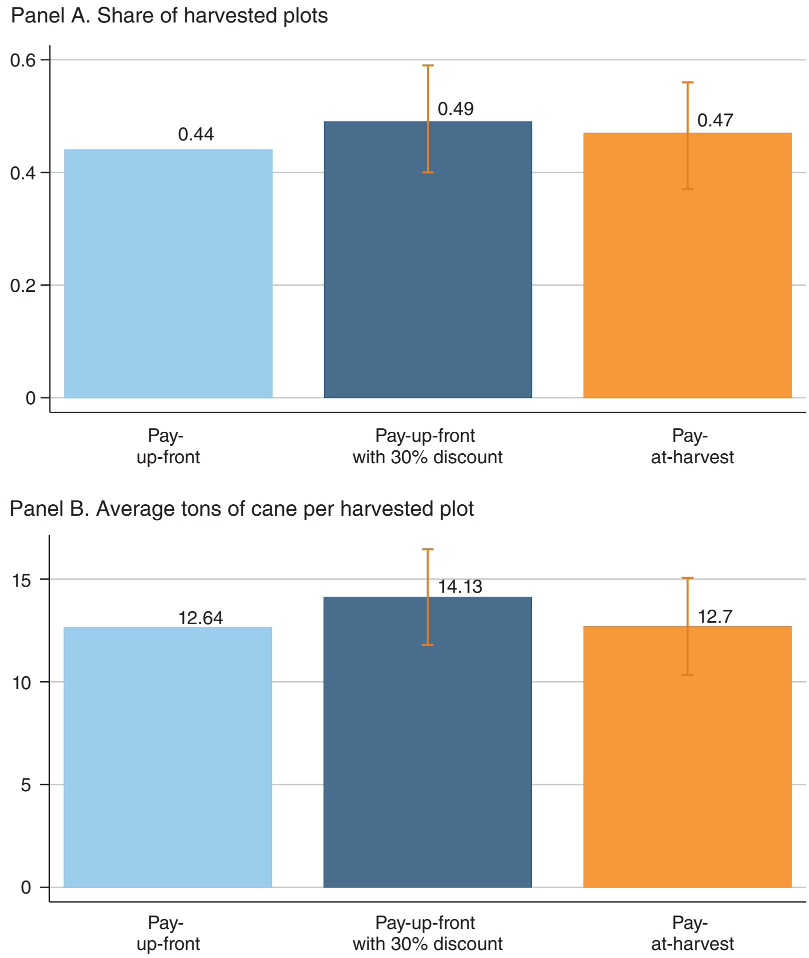

Figure 7. Main Experiment: Harvesting with the Company

Notes: The figure shows harvesting outcomes in the main experiment, by treatment group. Panel A reports the proportion of farmers from the main experiment who subsequently harvested with the company, as agreed under the contract. Panel B reports harvest weight (in tons), conditional on harvesting with the company. The bars report 95 percent confidence intervals from a regression of harvesting rates on dummies for the treatment groups.

effects. For instance, in the *Receive Choice in One Month* treatment of the second mechanism experiment, take-up reached 72 percent, even though farmers still lose the premium if the company defaults.

Our first evidence that the role played by the anticipation of default was limited is based on the intuition that expectations of default are like a price cut in the pay-at-harvest premium relative to the pay-up-front premium. Proposition 5 develops this idea and allows us to bound the differential effect of expectations of default on take-up by the effect of a price cut in the up-front 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 percent price cut had almost no effect on take-up of up-front insurance, ruling out a high price elasticity. Thus, for imperfect enforcement to account for much of our main 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.

Our second piece of evidence is that anticipation of *non-selective* default cannot explain the high demand for pay-at-harvest insurance. As shown in the model, while anticipation of non-selective default would have decreased demand for pay-up-front insurance, it would have had no first-order effect on the demand for pay-at-harvest insurance: intuitively, pay-at-harvest insurance has neither cost nor benefit under default.[30] As discussed above, our evidence suggests that default was indeed non-selective, in which case the high take-up of pay-at-harvest insurance—much higher than take-up for pay-up-front crop insurance in other settings (as well as in ours)—would still have been observed even absent risk of buyer default.

Our final evidence for a limited role of contract risk considers heterogeneous treatment effects of delaying the premium, by plausible proxies for ex ante priors of default. If anticipation of default did drive a difference in take-up between pay-up-front and pay-at-harvest insurance (and there was heterogeneity in prior probabilities of default) then a take-up regression may 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 6, columns 1–3 shows that while some of these measures do predict overall levels of take-up, they do not predict take-up differentially by premium timing (consistent with a belief that the company would not make insurance payouts *even if* the production contract is upheld, and also with some farmers mentioning trust as a reason why they did not buy insurance). Second, we consider harvesting rates in the previous season in the local area (Figure 6 shows it had substantial geographical variation). Table 6, columns 4 and 5 show that we also do not find heterogeneous treatment effects by historical rates of contract default at the field or sublocation level, respectively.

V. Policy Implications

Most insurance products transfer income across time. We have shown, in our setting, that this transfer across time reduces demand for crop insurance and that it does so through several mechanisms. Further, these same mechanism are known to shape financial decisions across a diverse range of other settings, suggesting that our results may have broad policy implications. In this final section of the paper we turn to these policy implications: first for crop insurance, and then for insurance more generally.

A. *Crop Insurance*

From a policy perspective, boosting the take-up of crop insurance is an ongoing challenge. This paper shows that demand is very sensitive to the timing of the

[30] If anything, default could reduce demand slightly through increased precautionary savings, a second-order effect.

Table 6—Main Experiment: Heterogeneous Treatment Effect by Proxies for Expectations of Default

Heterogeneity variable (*X*)	Good relationship with company (1)	Trust company field assistants (2)	Trust company managers (3)	Past share of plots harvested in field (4)	Past share of plots harvested in sublocation (5)
X × pay-at-harvest	−0.062	0.022	0.029	−0.228	0.681
	[0.070]	[0.029]	[0.028]	[0.297]	[0.425]
X	0.087	0.034	0.027		
	[0.040]	[0.018]	[0.017]		
Pay-at-harvest	0.726	0.654	0.640	0.852	0.101
	[0.035]	[0.087]	[0.073]	[0.261]	[0.358]
Mean dependent variable (pay-up-front group)	0.052	0.052	0.052	0.052	0.052
Mean heterogeneity variable (*X*)	0.335	2.889	2.423	0.873	0.839
SD heterogeneity variable (*X*)	0.472	1.045	1.101	0.099	0.068
Observations	570	569	567	556	605

Notes: The table shows heterogeneities of the treatment effects of the pay-at-harvest premium on insurance take-up in the main experiment, by five baseline variables (*X*): three different proxies for trust toward the company (columns 1–3) and the historical harvest rate in the field of the plot (column 4) and the sublocation of the plot (column 5). The name of the heterogeneity variable (*X*) is reported in the column title. The dependent variable is a binary indicator equal to 1 if the farmer took up the insurance. Pay-up-front payment and pay-up-front payment with 30% discount treatment groups are bundled together as baseline group, as outlined in the registered plan. *Mean dependent variable* (*pay-up-front group*) reports the mean of the dependent variable in the pay-up-front group. The relevant heterogeneity variable is reported in the column title. For each of the heterogeneity variables (*X*), we report their mean (*Mean heterogeneity variable*) and standard deviation (*SD heterogeneity variable*). The notes of Table 1 provide a definition of the trust variables used in the heterogeneity analysis. The two variables *past share of plots harvested in field* and *past share of plots harvested in sublocation* capture the share of plots that completed the harvest with the company in the field and sublocation, respectively, in the 2011–2014 period. The coefficients on the level of *past share of plots harvested in field* and *past share of plots harvested in sublocation* are missing because field fixed effects absorb them. All columns include field fixed effects.

premium payment and suggests that pay-at-harvest insurance may be a promising solution, *if premium payment can be enforced*. For our experiment, we used a collection method which relied on our contract farming setting: tying the insurance contract to a sales contract. Yet, the (unrelated) financial problems at the company still led to widespread default. Even if our mechanism can work in contract farming settings (whose presence is growing steadily in developing countries, UNCTAD 2009), the wider applicability of the idea depends on the answer to the following 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?[31]

[31] An alternative approach would be to offer a loan and pay-up-front 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).

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, Kremer, and Robinson (2011), and 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.

B. *Other Insurance Products*

The transfer across time is a feature of many 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. 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, Kunreuther, and Hirth 1995; Finkelstein, McGarry, and Sufi 2005; Handel, Hendel, and Whinston 2015), for which the intertemporal transfer is particularly long. Finally, the desire 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.

VI. Conclusion

By requiring that the premium be paid up front, 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 transfer 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 up front.

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 up front. 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 channels play a role. Namely, 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. Last, while contractual risk *may* have driven a difference between take-up of pay-up-front 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 have potentially broad implications. For crop insurance, where boosting demand has proven difficult, we showed that timing matters and proposed pay-at-harvest insurance as a potential solution *if* it can be enforced, which remains an important question. More broadly, the transfer across time is present in most 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." *American Economic Review* 102 (7): 3357–76.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2015. "Working over Time: Dynamic Inconsistency in Real Effort Tasks." *Quarterly Journal of Economics* 130 (3): 1067–115.

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. 1991. *The Economic Theory of Agrarian Institutions.* Oxford: Clarendon Press.

Becker, Bo, and Victoria Ivashina. 2015. "Reaching for Yield in the Bond Market." *Journal of Finance* 70 (5): 1863–901.

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

Bulow, Jeremy, and Kenneth Rogoff. 1989. "Sovereign Debt: Is to Forgive to Forget?" *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.

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

Carter, Michael, Lan Cheng, and Alexander Sarris. 2011. "The Impact of Interlinked Index Insurance and Credit Contracts of Financial Market Deepening and Small Farm Productivity." Unpublished.

Casaburi, Lorenzo, and Rocco Macchiavello. Forthcoming. "Demand and Supply of Infrequent Payments as a Commitment Device: Evidence from Kenya." *American Economic Review*.

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

Casaburi, Lorenzo, and Jack Willis. 2018. "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya: Dataset." *American Economic Review.* https://doi.org/10.1257/aer.20171526.

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." *Quarterly Journal of Economics* 125 (1): 1–45.

Cohen, Jonathan D., Keith Marzilli Ericson, David Laibson, and John Myles Whites. 2016. "Measuring Time Preferences." NBER Working Paper 22455.

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

Cole, Shawn, Xavier Giné, and James Vickery. 2017. "How Does Risk Management Influence Production Decisions? Evidence from a Field Experiment." *Review of Financial Studies* 30 (6): 1935–70.

Cole, Shawn A., and Wentao Xiong. 2017. "Agricultural Insurance and Economic Development." *Annual Review of Economics* 9: 235–62.

Dean, Mark, and Anja Sautmann. 2014. "Credit Constraints and the Measurement of Time Preferences." Brown University Department Economics Working Papers 2014–1.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59 (5): 1221–48.

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." Unpublished.

Dercon, Stefan, Ruth Vargas Hill, Daniel Clarke, Ingo Outes-Leon, and Alemayehu Seyoum Taffesse. 2014. "Offering Rainfall Insurance to Informal Insurance Groups: Evidence from a Field Experiment in Ethiopia." *Journal of Development Economics* 106: 132–43.

Devoto, Florencia, Esther Duflo, Pascaline Dupas, William Pariente, and Vincent Pons. 2012. "Happiness on Tap: Piped Water Adoption in Urban Morocco." *American Economic Journal: Economic Policy* 4 (4): 68–99.

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): 419–31.

Fink, Günther, and Feli Masiye. 2012. "Assessing the Impact of Scaling-Up Bednet Coverage through Agricultural Loan Programmes: Evidence from a Cluster Randomised Controlled Trial in Katete, Zambia." *Journal of Urology* 106 (6): 660–67.

Finkelstein, Amy, Kathleen McGarry, and Amir Sufi. 2005. "Dynamic Inefficiencies in Insurance Markets: Evidence from Long-Term Care Insurance." *American Economic Review* 95 (2): 224–28.

Ghani, Tarek, and Tristan Reed. 2018. "Relationships on the Rocks: Dynamic Contracts in a Market for Ice." Unpublished.

Giné, Xavier, Robert Townsend, and James Vickery. 2008. "Symposium on Access to Finance: Patterns of Rainfall Insurance Participation in Rural India." *World Bank Economic Review* 22 (3): 539–66.

Giné, 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 R., Igal Hendel, and Michael D. Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection Versus Reclassification Risk." *Econometrica* 83 (4): 1261–313.

Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for 'Humans': Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review* 105 (8): 2449–500.

Jensen, Michael C., and William H. Meckling. 1976. "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure." *Journal of Financial Economics* 3 (4): 305–60.

Karlan, Dean, Ed Kutsoati, Margaret McMillan, and Chris Udry. 2011. "Crop Price Indemnified Loans for Farmers: A Pilot Experiment in Rural Ghana." *Journal of Risk and Insurance* 78 (1): 37–55.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. "Agricultural Decisions after Relaxing Credit and Risk Constraints." *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): 1227–77.

Knetsch, Jack L., and J. A. Sinden. 1984. "Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value." *Quarterly Journal of Economics* 99 (3): 507–21.

Levine, David, Theresa Beltramo, Garrick Blalock, Carolyn Cotterman, and Andrew M. Simmons. 2018. "What Impedes Efficient Adoption of Products? Evidence from Randomized Sales Offers for Fuel-Efficient Cookstoves in Uganda." *Journal of the European Economic Association.* https://doi.org/10.1093/jeea/jvx051.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies." *Review of Economic Studies* 69 (1): 209–44.

Liu, Yanyan, Kevin Chen, Ruth Vargas Hill, and Chengwei Xiao. 2016. "Delayed Premium Payment, Insurance Adoption, and Household Investment in Rural China." International Food Policy Research Institute Discussion Paper Series 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): 121–38.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics* 118 (4): 1209–48.

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): 676–80.

Minten, Bart, Anneleen Vandeplas, and Johan Swinnen. 2011. "The Broken Broker System?: Transcending on Agricultural Wholesale Markets in India." International Food Policy Research Institute Discussion Paper 1143.

Mobarak, A. Mushfiq, and Mark Rosenzweig. 2012. "Selling Formal Insurance to the Informally Insured." Yale University Department of Economics Working Paper 1007.

Pauly, Mark V., Howard Kunreuther, and Richard Hirth. 1995. "Guaranteed Renewability in Insurance." *Journal of Risk and Uncertainty* 10 (2): 143–56.

Rampini, Adriano A., Amir Sufi, and S. Viswanathan. 2014. "Dynamic Risk Management." *Journal of Financial Economics* 111 (2): 271–96.

Rampini, Adriano A., and S. Viswanathan. 2010. "Collateral, Risk Management, and the Distribution of Debt Capacity." *Journal of Finance* 65 (6): 2293–322.

Rampini, Adriano A., and S. Viswanathan. 2016. "Household Risk Management." NBER Working Paper 22293.

Schilbach, Frank. Forthcoming. "Alcohol and Self-Control: A Field Experiment in India." *American Economic Review*.

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 (1): S164–87.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* 62 (3): 539–91.

Udry, Christopher. 1994. "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria." *Review of Economic Studies* 61 (3): 495–526.

United Nations Conference on Trade and Development. 2009. *World Investment Report: Transnational Corporations, Agricultural Production and Development.* New York: United Nations Conference on Trade and Development.

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