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

Lorenzo Casaburi
University of Zurich

Jack Willis
Harvard University

JOB MARKET PAPER

Please click here for the latest version

December 17, 2016

Abstract

The gains from insurance arise from the transfer of income across states. Yet, by requiring that the premium be paid upfront, standard insurance products also transfer income across time. We show that this intertemporal transfer can help explain low insurance demand, especially among the poor, and in a randomized control trial in Kenya we test a crop insurance product which removes it. The product is interlinked with a contract farming scheme: as with other inputs, the buyer of the crop offers the insurance and deducts the premium from farmer revenues at harvest time. The take-up rate is 72%, compared to 5% for the standard upfront contract, and take-up is highest among poorer farmers. Additional experiments and outcomes indicate that liquidity constraints, present bias, and counterparty risk are all important constraints on the demand for standard insurance. Finally, evidence from a natural experiment in the United States, exploiting a change in the timing of the premium payment for Federal Crop Insurance, shows that the transfer across time also affects insurance adoption in developed countries.

*Lorenzo Casaburi, lorenzo.casaburi@econ.uzh.ch. Jack Willis, jackwillis@fas.harvard.edu. We wish to thank Emily Breza, John Campbell, Michael Carter, Shawn Cole, Pascaline Dupas, Siddharth George, Ed Glaeser, Oliver Hart, Nathan Hendren, Michael Kremer, David Laibson, Guilherme Lichand, Craig McIntosh, Mushfiq Mubarak, Sendhil Mullainathan, Nathan Nunn, Matthew Rabin, Gautam Rao, Martin Rotemberg, Frank Schilbach, Andrei Shleifer, Tavneet Suri, Argyris Tsiaras, Chris Udry, Andrew Weiss and seminar audiences at the Basis Technical Meeting, Barcelona Summer Forum, BREAD pre-conference, Columbia, CSAE Oxford, Harvard, IGC International Growth Week, NEUDC, NOVAFRICA, PACDEV, and Stanford for useful comments. We are grateful to Noah Mambo, Mayara Silva and Rachel Steinacher for excellent research assistance and to the management of the contract farming company for their continued partnership. The project was funded by SIEPR, USAID BASIS and an anonymous donor, to whom we are very grateful. All errors are our own.

1 Introduction

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

This paper provides experimental evidence on the consequences of the transfer across time in insurance, by evaluating a crop insurance product which eliminates it. Crop insurance offers large potential welfare gains in developing countries, as farmers face risky incomes and have little savings to self-insure. Yet demand for crop insurance has remained persistently low, in spite of heavy subsidies, product innovation, and marketing campaigns (Cole et al. 2013a). 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 yields.[2]

The insurance product we study eliminates the transfer across time by charging the premium at harvest, rather than upfront. We work in partnership with a Kenyan contract farming company, one of the largest agri-businesses in East Africa, which contracts around 80,000 small-holder farmers to grow sugarcane. As is standard in contract farming, the company provides inputs to the farmers on credit, deducting the costs from farmers' revenues at harvest time. We tie an insurance contract to the production contract and use the same mechanism to enforce premium payment: the company offers the insurance product,[3] and deducts the premium (plus interest) at harvest.

Our first experiment shows that delaying the premium payment until harvest time results in a large increase in insurance demand. In the experiment we offered insurance to 605 of the farmers contracting with the company and randomized the timing of the premium payment.[4] Take-up of the standard, upfront insurance was 5%: low, but not out of line with results for other "actuarially-

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

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

[3] The insurance product has a double-trigger design (Elabed et al. 2013), providing partial yield insurance, conditional on both plot yield and local area yield being below given thresholds, and on harvesting with the company.

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

fair" insurance products in similar settings.[9] In contrast, when the premium was due at harvest (including interest at 1% per month, the rate which the company charges on loans for inputs), take-up was 72%, among the highest take-up rates seen for agricultural insurance in similar settings. To benchmark this difference, in a third treatment arm we offered a 30% price discount on the upfront insurance premium. Take-up among this third group was 6%, not significantly different from take-up under the full-price upfront premium. Taken together, these results show that farmers *do* have high demand for insurance, but they have a low willingness to pay for it upfront.

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

Two additional experiments provide further evidence on channels. In the first, we dig deeper into the role of liquidity constraints. Not having cash was the most common reason given by farmers who did not buy pay-upfront insurance in the main experiment. To test this, in this experiment we gave a subset of farmers cash, before offering them insurance later in the same meeting (similar to Cole et al. 2013a). The cash gift, being slightly larger than the cost of the premium, ensured that farmers did have liquidity to purchase the insurance if they wished to. However, as acknowledged by Cole et al. (2013a), it may also have induced reciprocity. To address this, we cross-cut the cash treatment with a pay-upfront vs pay-at-harvest treatment. The difference-in-differences effect of the cash was 8%, small and not significant, showing that pay-upfront insurance was not the marginal expenditure, and ruling out the explanation that farmers simply did not have cash.[5] However, because the cash gift may not have completely removed liquidity constraints (farmers may still have wanted to borrow), two possible explanations remain: farmers were not liquidity constrained to begin with, or farmers were very liquidity constrained and hence had other uses for

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

cash. The next experiment helps separate the two: it should only find an effect in the latter case.

In the second experiment on channels, we consider the effect of a small delay in the premium payment, such that payment is not due immediately at sign-up.[6] The design targets present bias, which has important implications for optimal contract design and for welfare interpretations. In the experiment, we compare insurance take-up across two groups. In both groups, during the visit, farmers had to choose between a cash payment, equal to the insurance premium, and free enrollment in the insurance. Farmers in the first group were told they would receive their choice immediately, whereas farmers in the second group were told they would receive their choice in one month's time.[7] Delaying delivery this way, by just one month, increased insurance take-up by 21 percentage points. We show that the size of this effect is inconsistent with standard exponential discounting - if the discount rate was high, then why buy insurance in one month? In contrast, it is consistent with present bias (combined with liquidity constraints).

The final channel we consider is imperfect contract enforcement. If either party defaults before harvest time, then the farmer does not pay the premium at harvest, whereas the upfront premium is sunk. Tying the contracts together means that, for the farmer, defaulting on the insurance requires side-selling (selling to other buyers), and vice versa.[8] This has two implications. First, it reduces strategic default on the harvest-time premium, the natural concern with removing the transfer across time, since farmers typically value the production relationship. In keeping with this, there was no significant difference in side-selling, or in yield conditional on not side-selling, across pay-upfront and pay-at-harvest treatment groups. Second, however, it can *induce* default: if the farmer side-sells for other reasons, he automatically defaults on the insurance contract. In our setting, before harvest, the company faced severe financial difficulties and temporarily shut-down their factory, resulting in long delays and uncertainty in harvesting. Because of this, twelve months after our experiment began, there was widespread side-selling: 52% of farmers side-sold or uprooted, compared to a historical rate of less than 10%.

In spite of the large default rates ex-post, three arguments suggest that, ex-ante, any differential effect on take-up by the timing of the premium was limited. First, using our model, we bound the differential effect on take-up by the effect of a price cut on take-up for upfront insurance. In particular, a percentage price cut, of size given by the expected probability of side-selling, times the relative (expected) marginal utility of consumption conditional on side-selling, has a larger

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

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

[8] If the farmer side-sells neither at-harvest premium payments, nor payouts, occur.

effect. But the main experiment showed that demand for upfront insurance is inelastic, so for this channel to be important one of these two terms would have had to be very large. Second, while survey measures of trust in the company *are* correlated with overall insurance take-up, their interactions with the timing of the premium are not, suggesting that concerns of the company defaulting on insurance payouts *after* harvesting were more prevalent than expectations of side-selling. Third, assuming ex-ante expectations of side-selling are predictive of actual side-selling, then the correlation between take-up and actual side-selling should vary by premium timing. For both individual and local average side-selling, in the data it does not.

In the final contribution of the paper, we present evidence of external validity from a natural experiment in the United States. In developed countries, better legal institutions may make the enforcement of cross-state insurance easier, but better functioning financial markets may also make the transfer across time matter less. We provide evidence that the transfer across time does still matter, using a policy reform in the Federal Crop Insurance program, one of the largest in the world. Historically, premiums for the FCI typically *were* due around harvest time. But, starting in 2012, the premium due date was moved earlier for certain crops, requiring agricultural producers to pay premiums from operating funds. Identifying off of variation across time, crops, states, and county characteristics, we show that the change reduced the amount of insurance purchased, and that, in line with the role of liquidity constraints, the effect was concentrated in counties with smaller farms.

This paper adds to several strands of literature. First, many papers have investigated the demand for agricultural insurance and the factors which constrain it (Cole et al. 2013a; Karlan et al. 2014). Demand is generally found to be low, and interventions to increase it typically have small effects in percentage-point terms.[9] Many of the proposed explanations, such as risk preferences and basis risk (Mobarak and Rosenzweig 2012; Elabed et al. 2013; Clarke 2016), relate to the transfer across states in insurance; we focus on the transfer across time.[10] Several studies have bundled insurance with credit (Gine and Yang 2009; Carter et al. 2011; Karlan et al. 2014; Banerjee et al. 2014), finding that take-up of credit increases little, and in some cases decreases. We effectively do the reverse, bundling credit with insurance. The closest paper to ours, Liu et al. (2016), finds that, for a livestock mortality insurance provided by the government in China,

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

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

delaying premium payment increases take-up from 5% to 16%; Liu and Myers (2016) considers the theoretical implications.[11] 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.[12]

Second, 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 Ivashina 2015), rather than on the cost for consumers, our focus. A recent exception is a largely theoretical literature (Rampini and Viswanathan 2010, 2013; Rampini et al. 2014) which argues that firms face a trade-off between financing and insurance. Rampini and Viswanathan (2016) apply similar reasoning to households.[13] These papers are part of a wide literature on how imperfect enforcement affects the set of financial contracts which exists (Bulow and Rogoff 1989; Ligon et al. 2002), to which we add by considering the implications of imperfect enforcement for the timing of insurance premiums.

Finally, our paper adds to a literature on the importance of interlinked contracts, i.e. contracts covering multiple markets, in developing country settings.[14] In particular, our work relates to a body of research that documents the presence of informal insurance agreements in output and credit market contracts (Udry 1994; Minten et al. 2011), and to a recent line of empirical research on the emergence and impact of interlinked transactions (Casaburi and Macchiavello 2016; Casaburi and Reed 2014; Ghani and Reed 2014; Macchiavello and Morjaria 2014, 2015).[15]

The remainder of the paper is organized as follows. Section 2 describes the setting in which the experiment took place and how tying an insurance contract to a production contract affects enforcement. Section 3 presents the main experimental design and results. Section 4 develops an intertemporal model of insurance demand. Section 5 presents evidence on channels, from the main experiment and from two additional experiments. Section 6 presents evidence of external validity from a natural experiment in the U.S. Finally, section 7 discusses the policy relevance of the results, presents ideas for future work, and concludes.

[11] They argue that charging the premium later may help address trust and liquidity constraints.

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

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

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

[15] In particular, Casaburi and Macchiavello (2016) discuss how reputation allows large firms to credibly provide interlinked services to their suppliers. While they focus on the interlinked provision of saving services, some of the arguments apply to insurance products too.

2 Setting, contract farming, and interlinked insurance

We work in partnership with a contract farming company in Western Kenya, one of the largest agri-businesses in East Africa. The company, founded in 1971, operates a large sugar factory which needs a steady supply of sugarcane. To supply the factory, the company contracts around 80,000 small-holder farmers to grow sugarcane. We offer insurance to some of these farmers. The setting is ideal for studying the intertemporal transfer in insurance because, as we explain below, contract farming offers a way to charge the premium at harvest time. Several other aspects also make it a good setting for developing and evaluating an insurance product: a large number of farmers, a long panel of data on production and plot characteristics, and important production risks.

Sugarcane, the crop which we insure in the study, is the main cash crop in the region. It accounts for more than a quarter of the total income for 80% of farmers in our sample, and more than one half of total income for 38% of farmers. It is not seasonal (in the region) and, once planted, lasts upwards of three growing cycles. Growing cycles are long - each one lasts around sixteen months - which means that the difference between paying the insurance premium upfront or at harvest time is particularly stark. The first cycle, called the plant cycle, involves higher input costs and hence lower profits than the subsequent cycles, called the ratoon cycles, and yields decline over cycles. Crop failure is rare, but yields are subject to significant risks from rainfall, climate, pests and cane fire.

Farmers who grow sugarcane are typically poor, but not the poorest in the region. In our sample, 80% of farmers own at least one cow, the average total area of cultivated land is 2.9 acres, and the average sugarcane plot size is 0.8 acres. Very few farmers in the study area have had experience with formal insurance.[16]

2.1 Contract farming

Contract farming involves the signing of a production contract between farmers and buyers, at planting, which both guarantees the farmer a market and binds them to sell to the company, at harvest. It is a production form of increasing prevalence (UNCTAD 2009), and our setting is typical. Farmers, termed outgrowers, enter in to a contract with the firm at planting time, with the contract covering the life of the cane seed, meaning multiple harvests over at least four years. The harvesting is done by company contractors, who subsequently transport the cane to the company factory where it is weighed. Farmers are paid at each harvest, by weight, at a price set by the Kenyan Sugar Board, the regulatory body of the national sugar industry.[17]

[16] Those farmers that do have experience of insurance have mostly had health or funeral insurance.
[17] Refined sugar is subject to import tariffs and quotas in Kenya.

Interlinked credit A major benefit of contract farming is that the company 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,[18] is widespread and our setting is no exception: the company provides inputs such as land preparation, seedcane, fertilizer, and harvesting services, with the costs (plus interest, at 1% per month)[19] deducted from harvest revenues.[20]

Contract enforcement The company must rely on self-enforcement of the contract, and thus the repayment of the input loan. This is because, as is common in developing countries, the cost of direct enforcement with a small farmer is prohibitively high. So, while illegal, farmers may side-sell (i.e. break the contract by selling to another sugarcane company) with little risk of prosecution. By side-selling, farmers avoid repaying the input loan,[21] and are paid immediately upon harvesting, but they are typically paid a significantly lower price for their cane. While the company cannot directly penalize farmers for side-selling, the company will collect any dues owed to it (plus interest) if possible, either from the same plot if it is re-contracted in the future, or from other plots of the farmer if he contracts multiple plots.[22] Since sugarcane is a bulky crop, the cost of transporting the cane to other sugar factories is often prohibitively high, and so side-selling is generally less of a concern than for more easily transported crops. Both the high transportation cost and monitoring by company outreach workers (discussed below) make partial side-selling unlikely.

We do not have direct information on the historical levels of side-selling in our context, but the administrative data we do have allows us to bound it. Prior to our experiment, in 2010-12, an average of 12% of plots which harvested in ratoon 1 did not harvest in ratoon 2, and figure 5 shows how this varied geographically. This is an upper bound on side-selling during this time, since some of these cases will be where farmers uproot the crop immediately after the ratoon 1 harvest (for example because of crop disease or poor yields). Since side-selling is illegal, obtaining detailed information from farmers about its causes and consequence is difficult.

The company is unlikely to strategically default on the contract. The main responsibility for

[18] Interlinkages, where a single contract or transaction between agents spans different markets, are an important feature of many agricultural markets in developing countries (Bardhan 1980; Bell 1988; IFAD 2003)

[19] Inflation in Kenya was around 6% per annum during the study, so the real interest rate on input loans from the company was 7% per annum.

[20] In addition to the risks mentioned above, company delays in these input provisions are a further source of risk for farmers.

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

[22] Any debt owed on a plot remains on the plot even if it is sold, and the debt is collected from future revenues even if the plot is subsequently farmed by another farmer. At the time we ran our experiment, debt collection remained at the plot level: if a farmer defaulted on a loan on one plot, the company would not recover that loan from revenues from other plots farmed by the farmer. However the company changed this policy before harvest time for our farmers, so that defaulted loans on one plot could be recovered from other plots of the same farmer.

the company under the contract is to purchase the farmer's cane at the price set by the Kenyan Sugar Board. Failure to do so would likely be reported to the Board, with consequences for the company, as farmers are well represented politically in the region. Farmers are, however, vulnerable to non-strategic default. If the company were to become insolvent, it would be unable to honor its agreement to purchase the cane, in which case farmers would be forced sell to another buyer. This happened, temporarily, 12 months after the start of our experiment, affecting some of the farmers in our sample. In Section 5.4 we discuss in detail the implications for the interpretation of our results - to summarize that discussion, multiple tests find no evidence that ex-ante anticipation of the shock affected our main results and we bound the size of the role it could have played.

Administration The company has to coordinate with its 80,000 farmers; how it does so 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 must group neighboring plots into administrative units called fields, which can be provided inputs and harvested concurrently. As detailed in Section 3, in our experiments we stratify treatment assignment at the field level. In our study sample, fields contain on average 8.7 plots.

2.2 Interlinked insurance

In standard insurance contracts the farmer pays upfront, so all of the contract risk is placed on the farmer.[23] Charging the premium at harvest time reduces the risk faced by the farmer: if the insurance company defaults before harvest time, at least the farmer will not have to pay the premium. However, it places significant counterparty risk on the insurer, the risk that the farmer does not pay the premium when harvests are good. In contract farming settings, future premium payments can be better enforced using the same mechanism which is already used to enforce repayment of input loans. Namely, the buyer can provide the insurance, and charge the premium as a deduction from harvest revenues.

Tying together the insurance and production contracts in this way, which we refer to as interlinking, increases the cost to the farmer of defaulting on the premium payment. In an interlinked contract, the only way default on the premium payment is to also default on the production contract, by selling to another buyer. As such, default compromises not only access to insurance in future periods, but also to all the other gains from the interaction with the buying company,

[23] Consistent with this, trust has been shown to be an important issue in shaping insurance take-up in other settings (Dercon et al. 2011, Cole et al. 2013a, Liu et al. 2016).

including the current and future purchase guarantees as well as future input provision.[24,25]

However, interlinking the insurance also has a potential cost in terms of enforcement. Namely, counterparty risk in the farming contract spills over into the insurance contract.[26] If a farmer decides to side-sell for reasons unrelated to the insurance contract, doing so mechanically triggers default on the insurance contract. While, under certain conditions, this need not affect the functioning of the insurance market (for example, if default on the product contract is orthogonal to potential insurance payouts), it does limit the states of the world than can be covered by an interlinked insurance product.

Another concern with interlinking insurance with contract farming is it increases side-selling relative to that which would occur without insurance. However, there are two reasons to believe that this effect will be minimal in our setting. First, the value of the insurance premium is likely to be much smaller than pre-existing input loans, such as seeds and fertilizer. As argued above, the choice of strategic default depends on the comparison between the static benefits of the default and the continuation value of the relationship; the insurance premium is unlikely to be marginal in this decision.[27] Second, given the insurance design (detailed in the next section), the farmer has limited information about whether he will receive the insurance payout at the time when he has to decide whether to sell to the company. In keeping with this, Section 5.4 shows indeed that interlinked insurance did not increase side-selling.

Finally, we note that in contract farming settings, since many of the inputs are provided by the company, there is less scope for insurance to impact productivity than in many other settings. The only inputs required from the farmer are the use of their land, plus labor for planting, weeding, and protecting the crop. This means that insurance is less likely to induce moral hazard, which would lower productivity, but also that insurance may be less likely to enable risky investments, which would increase productivity.

3 Does the transfer across time affect insurance demand?

This section describes the main experiment of the paper, in which we compare take-up for insurance when the premium is paid upfront to take-up when the premium is paid at harvest time, thus removing the intertemporal transfer. Doing so has a large effect: in our setting, charging

[24]The same argument is often cited as an important advantage of interlinking credit markets and product markets.

[25]Interlinking also has the additional benefit of the insurance being offered by a firm that the farmer is familiar with, which may help with trust.

[26]The same cost exists arises from interlinking credit and product markets.

[27]Further, if the farmers value access to insurance in future years, offering it would increase the continuation value of the relationship with the buyer, and hence insurance provision could actually reduce the occurrence of side-selling.

the premium as a deduction from harvest revenues, rather than upfront, increases take-up by 67 percentage points.

3.1 Experimental design

Treatment groups The experimental design randomized 605 farmers across three treatment groups. In all three treatments the farmers were offered the same insurance product, described below. The only thing that we varied across groups was the premium payment. In the first group (U1), farmers were offered the insurance product and had to pay the premium upfront. The premium was charged at "full price", which across the study spanned between 85% and 100% of the actuarially fair price. In the second group (U2), payment was again required upfront, but farmers received a 30% discount relative to the full price. Finally, in the third group (H), farmers could subscribe for the insurance and have the premium (full-price) deducted from their revenues at harvest time, including interest charged at the same rate used for the inputs the company supplies on credit (1% per month).[28] Randomization across these treatment groups occurred at the farmer level and was stratified by field.

Insurance design The insurance was offered by the company and the terms of the payout of the insurance were the same across the different experimental treatments. There was no intensive margin of insurance and farmers could only subscribe for the entire plot, not for parts of it. The insurance product offered was a double-trigger area yield insurance, preferred to a standard rainfall insurance product because risk factors other than rainfall affect yields.[29] Under the double-trigger design, a farmer receives a payout if two conditions are met: first, their plot yield has to be below 90% of its predicted level; second, average yield in their field must be below 90% of its predicted level. The design borrows from other studies that have used similar double trigger products in different settings (Elabed et al. 2013), and relied on rich administrative panel data at the plot level which helped with the design and pricing of the insurance in numerous ways.[30] The product is very much a partial insurance product: in the states where payouts are triggered, it covers half of the plot losses below the 90% trigger, up to a cap of 20% of the predicted production revenues.

[28] For this group, the expected interest was added to the initial premium when marketing the insurance product.

[29] The company collects rainfall information through stations scattered across the company catchment area. However, data quality issue is a concern and the predictive power of this rainfall data is low.

[30] The data included information on production levels, plot size, plot location, and crop cycle, and was available for a subsample of plots for the period 1985-2006 and for the entire company catchment area from 2008 onwards. The data was used to compute the predicted yields at the plot and area level which were needed for the double trigger insurance design. The historical data also made it possible to run simulations of alternative prediction models, and to compare the predictive power and costing of alternative insurance products. The simulations verify that the double trigger product performs well in terms of basis risk (A.1) - the proportion of farmers who receive a payout when the second area-level trigger is added is about 74% of those who would receive it with a single-trigger insurance. Additional calculations suggest the product does much better than an alternative product based on rainfall indexes.

Finally, farmers only receive insurance payouts if they sell their cane to 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. 938 farmers were targeted, 639 (68%) of whom attended. The primary reason (75%) for non-attendance was that the farmers were busy somewhere far away 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 the target farmers mastered very basic related concepts (e.g. the concept of tonnage and of acre). A small number of farmers (5%), typically elderly people, were deemed non-eligible at this stage. The final sample for the randomization was 605 farmers. Comparing to the 333 who did not enter the sample, these farmers had slightly larger plots (0.81 vs. 0.75 acres; p-value=0.015) and similar yields (55.5 vs. 54.5 tons per hectare; p-value=0.41).

After the initial group meeting, the outreach officers described the product in detail in one-to-one meetings with each farmer. They provided plot-specific visual aids concerning insurance triggers and payout scenarios. In order to ensure that the target farmers correctly understood the insurance product, outreach officers verified that farmers could correctly answer basic questions about the product before it was offered to them, e.g. the scenarios under which it would pay out. If farmers could not answer these questions, outreach officers re-explained. Farmers then had three to five business days to subscribe, with premiums collected either immediately or during revisits at the end of this period.[31]

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.[32] The experiment targeted plots in the early stages of the ratoon cycles (in particular the first and second ratoons), i.e. plots which have already harvested at least once. This choice was made because the yield prediction model

[31] In practice, for a substantial share of these farmers, revisits occurred between one and two weeks after the first visit.

[32] The criteria used to select the sample were: plot size - large plots were removed from the sample, to minimize financial exposure from the insurance pilot; plot yields - outliers were excluded, to improve the fit of the predicted yield in the insurance contract; the number of plots in the field - the study targeted only fields with at least five plots, to improve power given the stratified design. the number of plots per farmer - within each field, the few farmers with multiple plots were eligible to receive insurance in only one of those, the smallest one, the number of farmers per plot - plots owned by groups of farmers were excluded from the sample; finally, while contracted farmers are usually subsistence farmers, some plots are owned by "telephone farmers" who live far from the plots and manage them remotely - such plots are excluded from our sample.

performs better for ratoon than for plant cycles.[33] In addition, input costs are lower in the ratoon cycles, leading to fewer deductions from harvest revenues, and thus less incentives for side-selling.

Data collection We combine two sources of data for the analysis: survey data and administrative data. Our survey data comes from a short baseline survey, carried out by our survey team during the outreach-worker visits described.[34,35] As mentioned in section 2, the company keeps administrative data on all of the farmers in the scheme. From the administrative data we get information on previous yields, plot size, and time since last harvest. The data also enables us to track whether or not the farmer sells cane to the company at the end of the cycle, and their yield conditional on doing so.

3.2 Balance

Table 1 provides descriptive statistics for the three treatment groups using both administrative data and survey data, and shows that the experiment was balanced across most covariates. Since stratification occurred at the field level, we report p-values capturing the differences across the groups that are obtained from regressions that include field fixed effects.[36] Consistent with the specification we use for some of our analysis, and our pre-analysis plan, we also report p-values when bundling pay-upfront treatments U1 and U2 and comparing them to pay-at-harvest treatment H. The table suggests that the randomization mostly achieved balance across the observed covariates. In particular, only age is significantly different when comparing the bundled upfront group U to H. We confirm below that the experiment results are robust to the inclusion of baseline controls.

3.3 Experimental results

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

[33] The model performs better for ratoon cycles because previous yields within the same contract, which are available only for ratoon cycles, are a much better predictor of current yields than yields of harvests in a previous contract

[34] The survey was undertaken before randomization and the offering of the insurance product, with the exception of a question on why the farmer chose to take up / not take up insurance, which was asked after.

[35] We also use a second source of survey data. Several months later we followed up with a subset of the farmers by phone, to check whether they remembered the terms of the insurance and whether they regretted their take-up decision, as discussed below.

[36] This also implies that characteristics that do not vary within field, such as location, specific ratoon cycle and average field yield are essentially perfectly balanced across treatment groups.

tested in this paper is that low take-up is in part be due to the intertemporal transfer in insurance, which differentiates standard insurance products from their purely intratemporal ideal.

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

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

Figure 2 summarizes the take-up rates across the three treatment groups. For groups U2 and H, it also includes 95% confidence intervals of the difference in take-up from U1, obtained from a regression of take-up on treatment dummies.

The first result is that take-up of the full-price, upfront premium is low, at 5%. While low, this finding is consistent with several of the existing crop insurance studies mentioned above. The low take-up shows that reducing basis risk (the risk that the insurance will not pay out when farmers have a bad yield – one of the proposed explanations for low demand for rainfall insurance) by using an area yield double-trigger design is not enough to raise adoption, suggesting that the availability of rich plot-level data, one of the main advantages a large firm may bring to insurance, may not be sufficient to generate high demand for insurance.

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

The third result, which allows us to benchmark the importance of the second, is that offering a 30% price discount to the upfront premium has no statistically significant impact on take-up rates. The point estimate for the effect is 1 percentage point, and even when considering the upper bound of the confidence interval, take-up only increases by 8 percentage points. While this upper bound is consistent with substantial demand price elasticity (given the low baseline take-up) it suggests that medium-sized subsidies have limited scope to prompt large increases in demand in absolute terms in this setting. In subsection 5.4 we use this result to quantifying the possible importance of imperfect contract enforcement for our results.

Table 2 presents regression analysis of these treatment effects, and shows that they remain stable across a variety of specifications. Column (1) reports the coefficient from the simple regression used to generate the histogram in Figure 2. As mentioned above, the pay-at-harvest product

(H) has 67 percentage points higher take-up relative to the "full-price" pay-upfront product (U1), which is significant at the 1% level, whereas the 30% price cut product (U2) has just 1 percentage point higher take-up, which is far from significant. The difference between the pay-at-harvest (H) and the reduced price pay-upfront (U2) products is also significant at the 1% level. Column (2) adds fixed effects at the field level, the stratification unit. The results are virtually unchanged. Column (3) pools the upfront treatments U1 and U2, consistent with the specification we use later in the heterogeneity analysis. Columns (4) and (5) further add controls for plot and farmer characteristics, respectively. Finally, Column (6) 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. We believe that they did for two reasons. The first reason is, as mentioned above, farmers were asked questions to test their understanding of the product before it was offered to them. The second reason is, several months after the recruitment, we called back 76 farmers who had signed up for the pay-at-harvest insurance. We did so in two waves. In the first wave of 40 farmers, we began by reminding the farmers of the terms of the insurance product (the deductible premium and the double trigger design) and then checked that the terms were what the farmers had understood when originally visited. All farmers said they were. We then asked the farmers if they would sign up again for the product if offered next season. 32 (80%) said they would while 3 (7.5%) said they would not. The remaining 5 (12.5%) stated that their choice would depend on the outcome of the current cycle. In the second wave of 36 farmers, we did not prompt the farmers about the insurance terms, but instead asked farmers to explain them to us. 25 (69%) were able to do so. Of this second wave of farmers, after reminding those who had forgotten the terms, 28 (85%) said they would sign up for the product if offered next season.

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

4 An intertemporal model of insurance demand

In this section we develop a model which captures both the cross-state and cross-time transfers in insurance. To do so the model must be dynamic, so we begin by setting up a background intertemporal model, without insurance, into which we then introduce an insurance product. We first consider the case when the insurance contract is perfectly enforceable, and then we allow for

imperfect enforcement. The model shows how the different channels interact to affect insurance demand (and for whom) and motivates subsequent experiments and empirical tests to identify them. Proofs and derivations are in the appendix.

4.1 Background

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

Time and state We use a discrete-time, infinite horizon model. Each period t, which we will typically think of as one month, has a set of states corresponding to different income realizations. The probability distribution over states is assumed to be memoryless and cyclical, representing, for example, cyclical agricultural incomes.

Utility Individuals have time-separable preferences[37] and maximize present-biased expected utility $u(c_t) + \beta\Sigma_{i=1}^{\infty}\delta^i\mathbb{E}[u(c_{t+i})]$ as in Laibson (1997). We assume that $u(.)$ satisfies $u' > 0$, $u'' < 0$, $lim_{c\to 0}u'(c) = \infty$ and $u''' > 0$.[38] We also assume that $\beta \in (0, 1]$ and $\delta \in (0, 1)$.

Intertemporal transfers Households have access to a risk-free asset with constant rate of return R and are subject to a borrowing constraint.[39] As in Deaton (1991), we assume $R\delta < 1$.

Income and wealth Households have state-dependent income in each period y_t. We assume $y_t > 0\ \forall t \in \mathbb{R}^+$.[40] We denote wealth at the beginning of each period by w_t, and cash-on-hand once income is received by $x_t = w_t + y_t$.

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

$$\max_{(c_{t+i})_{i\geq 0}} u(c_t) + \beta\mathbb{E}[\Sigma_{i=1}^{\infty}\delta^i u(c_{t+i})] \tag{2}$$

subject to, for all $i \geq 0$,

[37] We note that time-separable preferences equate the elasticity of intertemporal substitution, ψ, and the inverse of the coefficient of relative risk aversion, $\frac{1}{\gamma}$. As such they imply a tight link between preferences towards risk and towards consumption smoothing, both of which are relevant for the demand for upfront insurance. Recursive preferences would enable us to untangle the two (Epstein and Zin 1989), which would provide an additional interpretation for our results. Namely, if $\psi \ll \frac{1}{\gamma}$, then we may expect a large gap between the demand for upfront and deductible insurance, since the cost of variation in consumption over time would greatly exceed that of variation in consumption across state.

[38] We assume prudence, i.e. $u''' > 0$, as is common in the precautionary savings literature (and as holds for CRRA utility), to ensure that the value of risk reduction is decreasing in wealth, i.e. Lemma 1, part 3. Prudence and liquidity constraints interact, strengthening the result, but our proof requires prudence to ensure the result holds with strict inequality.

[39] The results are identical if instead there is a borrowing limit $\underline{a}$, so that $x_t - c_t \geq \underline{a}\ \forall t$. Also, the model may be extended to assume a private investment technology $F(k)$, or similarly a wealth-dependent interest rate. Then the individual is liquidity constrained at time t iff $F'(k_t) > R$.

[40] As a technical assumption (to avoid complications arising from zero consumption), we actually assume that y_t is strictly bounded above zero $\forall t$.

$$x_{t+i+1} = R(x_{t+i} - c_{t+i}) + y_{t+i+1}$$
$$x_{t+i} - c_{t+i} \geq 0$$

Denote the value function of the household by V_t. Since income is memoryless, V_t is a function of just one state variable, cash-on-hand x_t. Since preferences are not time-consistent, V_t is different from the continuation value function, denoted V_t^c, which is the value function at time t, given time $t-1$ self's intertemporal preferences, i.e. without present bias. We assume that households are naive-$\beta\delta$ discounters, in that they believe that they will be exponential discounters in future periods (and thus may have incorrect beliefs about future consumption functions). There is evidence for such naivete in other settings (DellaVigna and Malmendier 2006), and with the exception of Proposition 2, the comparative statics with respect to wealth, all of the propositions below follow through with slight modification in the sophisticated-$\beta\delta$ case.[41]

Iterated Euler equation We are interested in the effect of changing the timing of the premium payment, so we need to compare the marginal utility of consumption across time periods. We can do so using the Euler equation, given in this model by:

$$u'(c_t) = \max\{\beta\delta R\mathbb{E}[u'(\tilde{c}_{t+1})], u'(x_t)\} \tag{3}$$
$$= \beta\delta R\mathbb{E}[u'(\tilde{c}_{t+1})] + \mu_t \tag{4}$$

where $\mu_t(x_t)$ is the Lagrange multiplier on the borrowing constraint, and the tilde in $\tilde{c}_{t+1}$ represents that the decision is with respect to period-t self's beliefs about consumption in period $t+1$. We can iterate the Euler equation to span more periods:

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

where $\lambda_t^{t+H}(x_t)$ is a measure of the extent to which transfers between time t and time $t+H$ are distorted by (potential) borrowing constraints,[42] defined as follows:

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

For the rest of this section we omit the tildes for notational simplicity. Since all decisions are being made in period 0, all variables are with respect to time 0 beliefs. The following lemma provides the main results from the model which will be useful when considering the demand for insurance.

[41] The required modification is replacing β by a state-specific discount factor, which is a function of the marginal propensity to consume. Proposition 2 and Lemma 1 parts 1 and 3 may no longer hold. Concavity and uniqueness of the continuation value, V_t^c, is no longer guaranteed, and we have not yet been able to prove the comparative statics of insurance demand with respect to wealth, because of the additional complication of the state-dependent discount factor.

[42] $R^H + \frac{\lambda_t^{t+H}}{\beta\delta^H E[u'(\tilde{c}_{t+H})]}$ is the shadow interest rate that the household faces for such transfers.

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

1. V_t, V_t^c *exist, are unique, and are concave.*
2. $\frac{dc_t}{dx_t} < 1$*, so investments (and wealth in the next period) are increasing in wealth.*
3. $\frac{d^3V_t}{dx_t^3}$, $\frac{d^3V_t^c}{dx_t^3} > 0$*, so the value of risk reduction is decreasing in wealth.*
4. $\frac{d\lambda_t^{t+H}}{dx_t} < 0$*, i.e. the distortion arising from liquidity constraints is decreasing in wealth.*

The intuition behind part 3 of the lemma is as follows. The value of risk reduction depends on how much the marginal utility of consumption varies across states of the world. Two things drive this. First, how much the marginal utility varies for a given change in consumption; this gives rise to the comparative static through the standard force of prudence (i.e. $u''' > 0$). Second, how much consumption varies for a given change in wealth, i.e. the marginal propensity to consume. Concavity of the consumption function, a consequence of prudence (Carroll and Kimball 1996) and strengthened by the borrowing constraint (Zeldes 1989; Carroll and Kimball 2005), further drives the comparative static.[43]

4.2 Insurance with perfect enforcement

We begin with the case where insurance contracts are perfectly enforceable. For tractability, at several points we make first order approximations for the costs and benefits of the insurance contracts. Such approximations are likely to be reasonable in our setting, as the insurance product offered only provides partial coverage.[44]

Timing We assume that the decision of whether to take up insurance is made in period 0, and any insurance payout is made in period H, the harvest period. Since the take-up decision is our main focus, and is made in period 0, all expectation operators in this section are with respect to information at time 0 and we drop the time subscript below.

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

[43] Mathematically, consider the value of a marginal transfer of a state with $x + \Delta$ to a state with x, assuming both are equally likely. The value of this transfer is (one-half times) $V'(x+\Delta) - V'(x) = u'(c(x+\Delta)) - u'(c(x)) \simeq u''(c(x))c'(x)\Delta$. The derivative of this with respect to wealth, x, is $\Delta(u'''(c(x))c'(x)^2 + u''(c(x))c''(x))$, which shows the role of both u''' and c''.

[44] The premium is small (3% of average revenues) and correspondingly the insurance product has relatively low coverage (maximum payout of 20% of expected revenue). In addition, there are two other reasons to believe that a first order approximation is reasonable in our setting. First, we are concerned with differential take-up according to the timing of the premium payment, so that second order effects which affect upfront and deductible insurance equally play no role. Second, the double trigger of the insurance is designed to minimize moral hazard, and many of the inputs are provided by the company. Thus there is less scope for the insurance to affect input provision. In line with this, as we report in section 5.4, we see no impact of insurance on yield.

stochastically dominates y_H.[45]

Premium We consider two possibilities for the timing of premium payment: upfront, at time 0; and at harvest, at time H. If paid in period H, the insurance premium is 1, equal to the expected payout of the insurance (commonly refered to as the actuarially-fair price).[46] If paid upfront, in period 0, the insurance premium is R^{-H}. Thus, at interest rate R, upfront and at-harvest payment are equivalent in net present value.

Demand for insurance The farmer buys insurance if the expected benefit of the payout is greater than the expected cost of the premium. Thus, using first-order expansions of the value functions, the envelope theorem, and first order conditions, to first order the take-up decisions are:[47]

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

This clarifies the difference between the take-up decisions for pay-upfront and pay-at-harvest insurance: for pay-at-harvest insurance, the decision is a comparsion of the marginal utility of consumption across states (future states when insurance does pay out vs. when it does not); whereas for pay-upfront insurance, the decision is a comparison of marginal utility across both states and time (future states when insurance does pay out vs. today). To relate the two decisions, we use the iterated Euler equation, equation 5, which gives the following proposition.

Proposition 1. *If farmers are certain they will not be liquidity constrained before harvest (or almost liquidity constrained - the exact condition is that, upon purchasing pay-at-harvest insurance, $x_t - c_t > R^{-H+t}$ for all times $t < H$ and for all paths) then they are indifferent between the two insurance contracts. If, however, farmers face a positive probability of being (almost) liquidity constrained, the expected net benefit of pay-at-harvest insurance is higher than that of upfront insurance. To first order:*

$$\begin{matrix} \textit{Difference between expected net benefit} \\ \textit{of pay-at-harvest and pay-upfront} \end{matrix} = R^{-H}\lambda_0^H \tag{8}$$

This difference is equivalent to a proportional price cut in the upfront premium of $\frac{\lambda_0^H}{u'(c_0)}$ (< 1).

[45] Historical simulations using administrative data suggest this assumption is reasonable in our setting - while the insurance product does involve basis risk because of the second, area-yield based trigger, this trigger only prevents payouts to 26% of those who would receive the payout under the single trigger, as shown in Figure A.1.

[46] We follow the convention of the literature and denote by the actuarially fair price the price at which a risk-neutral insurance company would break even, assuming the risk is uncorrelated with market risk and assuming zero administrative costs.

[47] The first order expansion used is $\beta\delta^H \mathbb{E}[V_H^c(w_H + y_H + I)] - \beta\delta^H \mathbb{E}[V_H^c(w_H + y_H)] \approx \beta\delta^H \mathbb{E}[Iu'(c_H)]$ and the first order conditions are $V_t^{c'}(x_t) = u'(c_t)$. We can use the envelope theorem because insurance payouts I do not enter into constraints before time H.

The intuition for the result is that, since paying the premium at harvest rather than upfront is akin to holding an extra unit of assets until harvest, any reduction in cost is purely down to differences in the (shadow) interest rate faced by the individual and the interest rate offered by the insurer. If, under pay-at-harvest insurance, liquidity constraints will not be close to binding before harvest, then asset holdings can simply adjust to offset the difference. In particular, switching from paying at harvest time to paying upfront, the individual can follow the same consumption paths so long as they do not now hit the borrowing constraint before harvest.

One corollary is that to drive a difference between pay-upfront and pay-at-harvest insurance, intertemporal preferences must act through liquidity constraints. If discount rates are high then liquidity constraints are likely to be (close to) binding in some states of the world, even with a precautionary motive for holding assets. A second corollary follows from the relationship between liquidity constraints and wealth in the model. Combining Proposition 1 and Lemma 1 gives the following result, under the assumption that the insurance product provides just a marginal unit of insurance (so that we can ignore second order and higher effects).

Proposition 2. *Demand for pay-at-harvest insurance is decreasing in wealth. However, the wedge between pay-upfront and pay-at-harvest insurance is also decreasing in wealth. Among farmers who are sure to be liquidity constrained before the next harvest, the latter dominates and hence their demand for pay-upfront insurance is increasing in wealth.*[48]

The proposition tell us that, although demand for risk reduction (pay-at-harvest insurance) is higher among the poor, they may buy less standard (pay-upfront) insurance than the rich, because the inherent intertemporal transfer is more costly for them. The intuition is as follows. The value of risk reduction is higher among the poor, because the poor are more likely to face liquidity constraints after harvest, meaning that they are less able to self-insure - liquidity constraints make consumption smoothing in response to shocks harder (so that shocks in income lead to larger shocks in consumption).[49] At the same time, the poor are also more likely to face liquidity constraints before harvest, which makes paying the premium upfront more costly than paying at harvest.

More generally, we can consider how the cost of the transfer across time varies with the parameters of the model. Higher discount rates (i.e. lower β or δ) make the transfer across time more costly, as they make the borrowing constraint more likely to bind, and to bind more tightly when it does bind. Higher harvest risk increases the importance of precautionary savings, meaning

[48] The general point that the gap between pay-upfront and pay-at-harvest insurance is decreasing in wealth follows from the shadow interest rate being decreasing in wealth. In our model that comes from a borrowing constraint, but it could be motivated in other ways. Indeed, models sometimes take it as an assumption (e.g. Dean and Sautmann 2014).

[49] The value of risk reduction is also higher among the poor for the standard reason of prudence.

farmers save more, thus reducing the likelihood that liquidity constraints bind before harvest, and thus making the transfer across time less costly. The effect of risk *before* harvest however is less clear, as the inherent savings in upfront insurance are illiquid, and thus less helpful in smoothing wealth shocks before harvest.

In the above, we have developed the full dynamic problem, and considered the cost of the intertemporal transfer in terms of the model's fundamental parameters. An alternative approach is to use observed investment behavior, in particular the potential returns of (risk-free) investments which the farmers makes or forgoes, as a sufficient statistic for the cost of the transfer across time. In appendix section A.1 we develop this approach further to bound the effect of the transfer across time on insurance demand.

4.2.1 Delaying premium payment by one month

Here we show that, with present bias, delaying the premium payment by just a short time may increase demand significantly, but only if the farmer is liquidity constrained. We consider the same insurance product as above, but with the premium payment instead delayed by one month (corresponding to our experiments in Section 5.3). Thus the premium is due at period 1, and costs $R^{-(H-1)}$. The following holds:

Proposition 3. *The gain in the expected net benefit of insurance from delaying upfront premium payment by one month is, to first order,* $R^{-H}\lambda_0^1 = R^{-H}\mu_0$. *It is the same as that from a proportional price cut in the upfront premium of* $\frac{\mu_0}{u'(c_0)}$.

The difference between a one month delay in premium payment and a delay until harvest time is the difference between $\lambda_0^1 = \mu_0$ and $\lambda_0^t = \mu_0 + \beta\mathbb{E}[\Sigma_{i=1}^{H-1}(R\delta)^i\tilde{\mu}_i]$. Thus, since harvest is many months away (i.e. H is large), the effect of a one month delay is likely to be very small relative to the effect of a delay until harvest, unless either liquidity constraints are particularly strong this month, or there is significant present bias. Present bias shrinks the difference between the two delays (equal to $\beta\mathbb{E}[\Sigma_{i=1}^{H-1}(R\delta)^i\tilde{\mu}_i]$) in two ways. First, the effect of future liquidity constraints are discounted by β, and second, the individual naively believes that he will be an exponential discounter in the future, and hence is less likely to be liquidity constrained.

4.3 Insurance with imperfect enforcement

In small-holder farming settings, contracts are unlikely to be perfectly enforceable, with both sides facing counterparty risk. This introduces another difference between paying upfront and paying at harvest: if the contract breaks before harvest time, then the farmer does not pay the

premium due at harvest, while he would have already paid the upfront premium. Accordingly, imperfect enforcement has implications both for farmer demand for insurance and for the willingness of insurance companies to supply it, both of which we discuss below. We use the same setup as above, and introduce the possibility of default on both sides.

Default *Farmer.* The farmer may strategically default on the harvest time premium, subject to some (possibly state dependent) utility cost c_D of breaking the contract.

Insurer. Before harvest, with probability p_I unrelated to yield, there is a stochastic shock which causes the insurer to default on the contract.[50] In the case of default, upfront premium payments are not reimbursed. We ignore the possibility of insurer default which occurs after the farmer's decision to pay the harvest time premium, since it would not have a differential effect by the timing of the premium payment.[51]

Timing In period 0, the farmer decides whether to sign up for the insurance contract, and in the case of pay-upfront insurance pays the premium. At the beginning of the harvest period (period H), with probability p_I there is a shock and the insurer defaults on the insurance contract. If the insurer has not defaulted, the farmer then learns what the insurance payout would be (and his yield) before, in the case of pay-at-harvest insurance, deciding whether to default on the premium at utility cost c_D.[52] Finally, if the contract still stands, the insurer pays out any insurance payments due.

The farmer faces two choices: whether to buy the insurance and, in the case of the pay-at-harvest insurance, whether to pay the premium at harvest time. Consider first the choice of whether to pay the premium at harvest time. Denoting the choice to pay by the (state-dependent) indicator function D_P, to first order:

$$D_P \coloneqq \mathbb{I}[Iu'(c_H) + c_D \geq u'(c_H)] \tag{9}$$

[50] We motivate the assumption that insurer default is exogenous, and in particular unrelated to the payout due, by the assumption that if the insurer defaults, he defaults on payments due to all farmers, not selectively. Such default could represent, for example, the insurer going bankrupt or deciding not to honor contracts. The assumption is reasonable in our setting, since the strategic default by the insurer would likely be highly costly for the farming company, both legally and in terms of reputational costs.

[51] Such default could either be assumed to be accounted for within I, in which case $\mathbb{E}I < 1$, or could be given by a second stochastic shock. Results below are robust to inclusion of such default.

[52] In practice he may have considerable uncertainty about both: the company harvests the crop, increasing uncertainty about the yield, and the area yield trigger in the insurance design increases uncertainty about the insurance payout. Such uncertainty shrinks the difference between paying upfront and at harvest compared to that which is derived below.

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

$$\begin{aligned}\text{Difference in net benefit of at-harvest \& upfront} &= \underbrace{R^{-H}\lambda_0^H}_{\text{liquidity constraint term}} + \underbrace{\beta\delta^H p_I E[u'(c_H)]}_{\text{premium saved when insurer defaults before harvest}} \\ &+ \underbrace{\beta\delta^H(1-p_I)\mathbb{E}[(1-D_P)(u'(c_H)-c_D-Iu'(c_H))]}_{\text{premium saved, minus cost of default and loss in insurance payouts, when farmer defaults}}\end{aligned} \tag{10}$$

It is easy to show that the size of the difference caused by imperfect enforcement is decreasing in the cost of default, c_D. If the cost of default is high enough, $c_D > \max_s u'(c_H(s))$, the farmer never strategically defaults and the additional difference is just that driven by insurer default.

Supply of insurance While the farmer is better off with the pay-at-harvest insurance, the possibility for strategic default by the farmer means that the insurer may be worse off. This is the most likely reason why premiums are typically charged upfront - if $c_D = 0$, in the case of pay-at-harvest insurance premiums received are never greater than payouts made, and so the insurer cannot make a profit at any premium price. But, it is not always clear cut - which of the two insurance products could actually be traded depends on both c_D and p_I, as well as liquidity constraints and preferences as discussed earlier. We discuss the supply of insurance further in the Appendix, Section A.2. Intuitively, a high risk of insurer default, p_I, makes pay-upfront insurance harder to trade (as does liquidity constrained farmers), and a small cost of farmer default, c_D, makes pay-at-harvest insurance harder to trade.

4.3.1 Interlinked insurance

Interlinking the insurance contract with the production contract has implications for contractual risk, since default on one entails default on the other. Typically, farmers face a cost of defaulting on the production contract, and so such interlinking raises the cost of defaulting on the harvest time premium, c_D. In addition, interlinking may also reduce the risk of the insurer defaulting in the eyes of the farmer, p_I, since the farmer is familiar with the company. However, interlinking may encourage default on the insurance contract, if the farmer wants to default on the production contract for other reasons. Given these forces, in order to understand enforcement of the insurance, we need to understand enforcement of the production contract, and how the two interact. To do so we modify the model developed above, to tie together the insurance and production contracts, and see how doing so affects default, and ultimately the decision to purchase insurance.

As in the imperfect enforcement case, the farmer has two decisions to make: whether to sign up for insurance, and whether to default on the contract at harvest time. The main difference is in what drives the decision to default, and what defaulting involves: default now implies more than just not paying the premium, it also involves selling to another buyer. We refer to doing so as side-selling, and assume that it results in a state-dependent outside option of $o(w_H)$, net of any loss caused by the loss of the relationship with the buyer. We defer further discussion of this outside option to the appendix, but note that in our setting this outside option is generally sufficiently low to avoid farmers defaulting on the input loans provided by the company, which are substatially larger than the cost of the insurance premium.

Default The timing is as above, and again we solve the farmer's problem backwards, starting with the decision of whether to side-sell conditional on the company not having defaulted on the farming contract. All decisions are as anticipated at time 0. First consider the case without insurance. We can define the (endogenous) cost of side-selling in this case as c_D, where we purposely use the same notation as above:

$$c_D = \mathbb{E}[V_H^c(w_H + o(w_H))] - \mathbb{E}[V_H^c(w_H + y_H)] \tag{11}$$

Unlike in the imperfect enforcement example above, c_D may now be positive or negative. If the farmer values the relationship with the company, and will be paid more by the company, it will be positive. However, if the farmer doesn't value the relationship, and would be paid more by selling to another company, it will be negative.

Whereas above the farmer may only default in the case of pay-at-harvest insurance, now he may have an incentive to default in all three cases: without insurance, with pay-upfront insurance, and with pay-at-harvest insurance. Without insurance, he defaults if $c_D \geq 0$. With insurance, the net payout of the insurance contract enters the decision (I in the case of pay-upfront, $I - 1$ in the case of pay-at-harvest), but seeing that the absolute value of c_D is likely to be large relative to both the insurance payout and the premium, insurance is unlikely to affect the decision to default. This is important: defaulting because $c_D \ll 0$ does not matter for the functioning of the insurance market, whereas selective defaulting to avoid the premium does. While insurance is unlikely to affect side-selling, if it does, then the following (simple) proposition tells us how.

Proposition 4. ***Pay-at-harvest insurance*** *makes those with high yields more likely to side-sell, and those with low yields less likely to side-sell.* ***Pay-upfront insurance*** *does not affect side-selling of those with high yields, and makes those with low yields less likely to side-sell.*

The intuition is that, compared to the case with no insurance, those with pay-upfront insurance may receive the additional insurance payout if they do not side-sell. In the case of pay-at-harvest insurance, there is also a force in the other direction: farmers do not have to pay the premium if they side-sell.

Demand for insurance The main difference from the previous case is now it is the expected *marginal* cost of defaulting induced by the insurance which matters. For the sake of brevity we omit expressions for the general case here, but note that if insurance does not affect the decision to side-sell, (in which case there is only one indicator function D for harvesting with the company), then the take-up decisions simplify considerably:

$$\text{Take up insurance iff } \begin{cases} \beta\delta^H \mathbb{E}[Du'(c_H)] \leq \beta\delta^H(1-p_I)\mathbb{E}[DIu'(c_H)] & \text{(pay-at-harvest)} \\ R^{-H}u'(c_0) \leq \beta\delta^H(1-p_I)\mathbb{E}[DIu'(c_H)] & \text{(pay-upfront)} \end{cases} \tag{12}$$

Regardless of whether insurance does affect the decision to side-sell, we have the following result, which allows us to bound the effect that imperfect enforcement has on take-up by that of a price cut in the upfront premium.

Proposition 5. *The risk of contract default in the interlinked contract drives a wedge between demand for pay-at-harvest and pay-upfront insurance. The size of the wedge is bound above by that of a price cut in the upfront insurance of*

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

This result enables us to relate the impact of ex-ante expectations of default to the impact of a price cut in the upfront premium, on which we have experimental evidence from the main experiment. We return to this point in Section 5.4.

5 Why does the timing of the premium payment matter?

The model in the previous section showed that liquidity constraints, intertemporal preferences, and imperfect contract enforcement could all drive the experimental results presented in Section 3. In this section we provide evidence which suggests that all three of these channels play an important role in constraining demand for upfront insurance.

Before further discussing the mechanisms which can explain our results, we first note several which cannot. Since demand for pay-at-harvest insurance is high, and the only difference with upfront insurance is the timing of the premium payment, we can rule out many of the mechanisms shown to constrain insurance demand in other settings. This includes basis risk (the risk that

insurance does not pay out when needed) and the risk preferences of farmers (Clarke 2016; Mobarak and Rosenzweig 2012; Elabed et al. 2013), the presence of informal insurance (Mobarak and Rosenzweig 2012), and farmer understanding of insurance (Cai et al. 2015).

Liquidity constraints are a natural candidate for the gap in demand between pay-upfront and pay-at-harvest insurance because, as shown in Proposition 1, they introduce a cost of holding the savings which are implicit in upfront insurance. Several studies have documented liquidity constraints among similar populations in the region of the study (Duflo et al. 2011; Cohen and Dupas 2010) and, also shown in Proposition 1, liquidity constraints may matter even if they do not bind at the time farmers are offered insurance – they drive a difference between paying upfront and paying at harvest if there is some chance they bind before harvest. Liquidity constraints are also tightly tied to wealth, and Proposition 2 showed that, as a result, they are a larger constraint on demand for upfront insurance among the poor.

Intertemporal preferences are another important candidate for explaining the large difference in demand for upfront vs. deduction insurance.[53] Among intertemporal preferences, we are particularly interested in present bias. This is for three reasons: first, a recent literature shows evidence that present bias is responsible for significant distortions in intertemporal decisions in similar settings (Loewenstein et al. 2003; Duflo et al. 2011); second, with present bias, the type of insurance offered has welfare implications, since the decision to forgo upfront insurance for present consumption is not time consistent, and future selves may regret it; third, if present bias is driving our results, the implications for insurance design are very different – as argued in Section 4.2.1, a product under which the premium is paid after take-up, but before harvest, may still have high demand without the contract enforcement concerns of a pay-at-harvest insurance.

Finally, imperfect contract enforcement could be responsible for the observed difference between pay-upfront and pay-at-harvest insurance. If either party defaults on the contract before harvesting, the upfront premium is sunk, whereas the farmer does not pay the premium at harvest. As shown in Section 4, anticipation of this possibility can mechanically drive a wedge between the pay-upfront and pay-at-harvest insurance products.

Before presenting evidence on individual channels, 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).[54] However, we did not detect heterogeneous

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

[54] A recent experimental literature considers what such questions elicit, and suggests difficulties with using them

treatment effects by these Required Rate of Return variables, as shown in table A.2. This is possibly driven by measurement issues, limitations associated with the hypothetical nature of the questions, and limited statistical power. In addition, standard lab-experiment measures in a given domain (e.g. the timing of cash disbursement) may fail to hold predictive power on other domains, such as how to use that money.[55]

In the rest of this section we do the following. We begin with evidence on liquidity constraints, first from treatment heterogeneities in the main experiment, and then from a second experiment. Then we present a third experiment which targets present bias, before presenting the evidence we have for the importance of contract risk. Finally, we briefly discuss other, behavioral channels, which could play a role and which we would like to investigate in future work.

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

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

We report how demand for the two types of insurance varies in our experiment with proxies for wealth and liquidity constraints, and thus the heterogeneous treatment effect of the transfer across time. The proxies come from both administrative data and from the baseline survey, and include yield levels in the previous harvest, sugarcane plot size, number of acres cultivated, whether the household owns a cow, access to savings and the portion of income from sugarcane. In order to gain power, we bundle together the two upfront groups (full price and 30% discount).[56] Thus, the regression model is:

$$T_{if} = \alpha + \beta Deduction_i + \gamma x_i + \delta Deduction_i * x_i + \nu_f + \epsilon_{if} \tag{13}$$

Table 3 presents the results, which suggests that the treatment effect does indeed vary by wealth and liquidity constraints. While not all of the interaction coefficient estimates are significant, the results in the table show that less wealthy and more liquidity constrained households are

to measure intertemporal preferences (Andreoni and Sprenger 2012a, Augenblick et al. 2015, Cohen et al. 2016)

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

[56] We mentioned the option of bundling the two upfront treatment groups when registering the trial.

differentially more likely to take-up the insurance when the premium is to be paid through harvest deduction, as predicted by proposition 2. For example, delaying premium payment until harvest time has an effect which is 14 percentage points larger for those who do not own a cow, and 18 percentage points larger for those who would do not have savings to cover an emergency requiring a Sh 1,000 ($10) outlay. Further, also in line with proposition 2, the effect comes from demand for pay-at-harvest insurance being higher among the poor. From a policy perspective, this result implies that pay-at-harvest insurance may be particularly beneficial for poorer farmers, who are typically in stronger need of novel risk management options.

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

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

5.2.1 Experimental design

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

Before presenting results, we first consider what this cash treatment does, and how we might expect farmers to respond. The first thing we note is that the cash ensures that farmers do have enough cash to pay for the insurance if they wish to, removing any hard cash constraint, and addressing the most commonly cited reason for not purchasing upfront insurance. But, while the cash drop eases liquidity constraints, it need not remove them entirely - an individual is liquidity constrained if they are not able to borrow any more at the market interest rate; after receiving the cash drop, farmers may still have wanted to borrow more.

The experiment is thus best interpreted as answering whether upfront insurance is the marginal expenditure, given an increase in cash which removes any hard cash constraints. Evidence from other settings suggests that the answer may be no. When interlinking insurance with credit, Gine

and Yang (2009) and Banerjee et al. (2014) find that demand for credit actually decreases when bundled with insurance. If insurance was the marginal expenditure, if anything we would expect the opposite.

One concern with cash drop designs is that they may induce a reciprocity effect – to reciprocate the cash gift, farmers may be more likely to buy the insurance product. We tried to minimize the risk of reciprocity by having the survey enumerator give the cash gift at the beginning of the meeting – in contrast, the insurance product is offered by a company outreach worker, at the end of the meeting. We discuss further the role of reciprocity below, but note for now that it can affect demand regardless of the timing of the premium payment, hence the cross-cut design of our experiment to help to control for it.

Finally, we note that contractual risk is held constant across cash and non-cash treatment groups, so that any differential take-up is not driven by contractual risk.[57]

5.2.2 Experimental results

We use the following regression model:

$$T_{if} = \alpha + \beta Discount_{if} + \gamma Cash_{if} + \nu Discount_{if} * Cash_{if} + \eta_f + \epsilon_{if} \tag{14}$$

Figure 3 presents the results. First, it is reassuring to note that, in this different sample, the comparison between the pay-upfront and pay-at-harvest groups resembles that of the main experiment. Take-up for the upfront group is slightly larger (13%), but, again, introducing at-harvest payment raises take-up dramatically (up to 76%). Second, the cash drop raises substantively the take-up rate in the upfront group (up to 33%), suggesting it does reduce liquidity constraints. However, the impact of the cash drop is much smaller than that of the harvest time premium - among those who do purchase pay-at-harvest insurance, many do not purchase pay-upfront insurance even if they do have the cash to do so. For these individuals upfront insurance is not the marginal expenditure - farmers prefer to use the additional money for other purposes (e.g. consumption, labor payments, school fees). Third, the cash drop also has an impact on take-up rates in the pay-at-harvest group (from 76% to 88%). While in a more complex model this could potentially be a wealth effect, in our simple model the wealth effect should if anything reduce insurance demand.[58] As mentioned above, this could also be reciprocity, whereby some farmers may feel obliged to purchase the insurance after receiving the transfer, as discussed in Cole et al. (2013a). The difference in impact of the cash drop between farmers offered the at-harvest premium and those offered the upfront

[57] Ignoring any second order effects that the cash drop may have on side-selling, which are likely to be very small given the size of the cash drop

[58] This wealth effect is likely to be small.

premium is 8%, which is small. While it is imprecisely estimated, we take this as evidence that the cash drop had relatively little effect on take-up beyond that caused by reciprocity.[59]

Table 4 confirms the patterns described above. Column (1) presents the basic level impact of the cash drop and pay-at-harvest treatments. We add field fixed effects in column (2) and additional controls in column (3). Across the specifications, we can always reject the null on the equality of the two treatments at the 1% level. The coefficient on Cash is large and significant at the 5% level in column (1) and remains similar in size but loses some precision as we add more controls. In columns (4) to (6), 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 it cannot account for half of the difference between pay-upfront and pay-at-harvest insurance in the main experiment.

To summarize, the results show that cash drops do relatively little to close the gap between pay-upfront and pay-at-harvest insurance. There are two potential explanations: farmers are not liquidity constrained, or farmers are very liquidity constrained and hence insurance is not the marginal expenditure. The next experiment will help to disentangle the two, since it should only find an effect if farmers are liquidity constrained.

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

If present bias is a major driver of our results, then delaying the premium payment by just a short amount of time may affect take-up. Here we describe a third experiment in which we did exactly this. We begin by introducing the experimental design, referencing the earlier theoretical discussion of Section 4.2.1 to explain exactly what the experiment tests. Then we present the results, which suggest that present bias does play an important role.

5.3.1 Experimental design

The aim of the experiment was to compare insurance take-up when the premium had to be paid at sign-up, to insurance take-up when the premium payment was delayed until one month later. However, for reasons explained below, the design was slightly different. A sample of 120 farmers was randomly allocated to two treatment groups.[60] Both groups were offered a choice between either a cash payment, equal to the insurance premium, or free enrollment in the insurance. The

[59] Interpreting the difference-in-difference is also made more difficult because we have to make an assumption on how reciprocity affects the probability of take-up. We use a linear probability model, and so implicitly assume that reciprocity affects the probability of take-up linearly.

[60] Randomization was again stratified at the field level.

difference between the treatment groups was when they would receive what they chose. In the first treatment group, the *Receive Choice Now* group, farmers were told that they would receive what they chose immediately. In the second group, the *Receive Choice in One Month* group, farmers were told that they would receive what they chose in one month's time (the cash payment offered to farmers in this case also included one month of interest).

Offering the choice between cash or the insurance for free, rather than the choice to buy insurance, allowed us to isolate the role of intertemporal preferences in two ways. First, it ensured that the choice in the *Receive Choice in One Month* group could be enforced (since the premium payment the following month did not rely on the farmer paying out of her own pocket). Second, like a cash drop, it relaxed any hard cash constraints, ensuring the farmer could take-up the insurance if he wanted to. In standard models, take-up under this choice should be the same as take-up when the farmer is given the choice to purchase insurance following a cash drop, as in the cash constraints experiment above. But a literature dating back to Knetsch and Sinden (1984)[61] suggests the two may be different: the former is akin to the Willingness to Accept, whereas the latter is the Willingness to Pay and may include an endowment effect.

As shown in Section 4.2.1, for a one month delay in premium payment to have made a difference, the farmer must have been liquidity constrained at the time of the experiment. If the effect is large, it suggests one of three things, two of which we argue were unlikely in our setting. First, credit constraints may have varied across time periods (Dean and Sautmann 2014), and the experiment could have 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 experiments across two months (plus a one-month pilot beforehand) and the results, presented below, are stable across these periods, suggesting this explanation is unlikely. Second, if farmers were heavy exponential discounters, i.e. they had low δ, then they would have preferred to pay one month later. However, even with a one month delay, insurance still involves a transfer across a large time period, and so in this case the farmer would still have been unlikely to buy it. Thus, if we find a large effect, we are left with just the third possibility: farmers were present biased.

Appendix Table A.5 reports the balance test across the two groups. We note that, due to the small sample size, there are significant imbalances across the two groups in the share of men and the acres of land cultivated. As discussed below, results are robust to the inclusion of these variables as controls.

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

5.3.2 Experimental results

Figure 4 shows that the take-up share in the *Receive Choice in One Month* group is 72%, compared to a baseline of 51% in the *Receive Choice Now* group.[62] This 21 percentage point increase suggests that a shift of only one month in the timing of the cash transfer (while keeping the net present value constant) has a large impact on insurance take-up. While the experimental design does not allow us to distinguish between time-consistent and time-inconsistent discounting *directly*,[63] the large effect is inconsistent with exponential discounting, as argued above. In contrast, it is consistent with present bias, as the *Receive Choice in One Month* treatment provides farmers with a commitment device on how to use the cash transfer, potentially overcoming their time inconsistency.

Table 5 confirms these results across different specifications. The gap between the two treatments remains statistically significant at 1% when adding field fixed effects, plot controls, farmer controls and both set of controls together. We note that the point estimate raises from 0.25 in the baseline specification with field fixed effects (Column 2) to 0.33 when adding both set of controls, though the difference in the two estimates is not statistically significant. This suggests that, if anything, accounting for the baseline imbalances reported above increases the estimate of the impact of requiring farmers to sign up in advance.

We note that the design mitigates the traditional trust concerns associated to standard time preferences experiments (Andreoni and Sprenger 2012b). In the *Receive Choice in One Month* treatment, both the cash transfer and the insurance sign-up depend on the field officer revisiting the field, so there are no differential trust concerns across the two choices. It is still possible, though implausible, that a farmer may think the field officers are more likely to return if she chooses the insurance. However, visits are organized at the field level and thus revisits cannot depend on individual choices. Further, and probably most importantly, the respondents have the contact info of the relevant company field staff (and, in most cases, of the IPA staff).

While present bias can lead to under subscription in pay-upfront insurance, one might think

[62] We note that the baseline take-up for the "Receive Choice Now" group is larger than the take-up in the group *Upfront Premium+Cash* in the Liquidity Constraints experiment, which could be due to several factors. First, there could be variation across farmer characteristics as the two experiments targeted different samples. Second, as mentioned above, this experiment estimates the willingness to accept for upfront insurance, whereas the liquidity constraint experiment estimates the willingness to pay, given liquidity. In particular, there may be an endowment effect in the latter associated with handing the cash to farmers at the start of the visit. Third, while the Liquidity Constraints experiment occurred early in late Summer 2014, the Intertemporal Preferences experiment was implemented in Spring 2015, shortly after the end of the dry season (December-March). It is possible this could make the risk of low harvest more salient for the farmers.

[63] In particular, sample size limitations prevented us from running an additional treatment where the cash transfer is postponed by two months, rather than one, and concerns about new information being a confounder prevented us from asking farmers to revisit their decision one month later.

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

5.4 Imperfect enforcement

In the case of the harvest time premium, if either party defaults on the contract before harvesting, the farmer does not pay the premium. Anticipation of this possibility mechanically drives a wedge between take-up of the pay-upfront and pay-at-harvest insurance products, as shown in Section 4.3.1. While the two experiments described above show that counterparty risk does not fully explain our results, since they hold it almost constant across treatments, it is likely that farmers had a non-zero prior for the probability of contract default. In this section we attempt to understand the importance of this channel.[64] We first discuss what happened under the contracts, and then detail the evidence that we have for the effect of imperfect enforcement on take-up decisions. While we find some evidence that counterparty risk did matter for overall levels of take-up, we find no evidence for a differential effect by the timing of the premium payment.

Before the farmers in our study were due to harvest, the contract-farming company ran into serious financial difficulties. These financial difficulties led to the closure of the factory for several months, during which time the company did not harvest any cane, and to subsequent severe delays in harvesting due to the backlog. Consequently, the farmers whose cane was mature during this time faced uncertainty around harvesting: when it would happen, and whether it would happen at all. As a result, unsurprisingly, a significant proportion of our farmers side-sold (i.e. did not sell to the company). Figure 6 shows that the rate of side-selling was 52% across the three experiments. Figure 5 is a histogram of the harvest rate (one minus the side-selling rate) by sublocation, and also for comparison shows a histogram of a lower bound[65] for the historical harvest rate in the same

[64] In the main experiment we collected several variables related to trust and farmers' relationship with the company. However, the heterogeneity analysis does not deliver any clear conclusion. In addition to limited power, it is possible that the variables do not properly capture the specific expectations and trust concerning the insurance product. We report the results in Table A.1. While the failure to detect heterogeneity in the treatment effect along these variables poses an obvious caveat, qualitative evidence from the survey provides suggestive evidence that the interlinked insurance contract may help to address trust issues and we hope it will motivate further work on the topic.

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

locations. The figure shows that default was much higher than historical rates, and underlines the fact that farmers face counterparty risk under upfront insurance contracts.

The widespread default ex-post raises two important questions: did the insurance product induce side-selling; and were expectations of default responsible for the difference in take-up, ex-ante?

5.4.1 Insurance did not affect side-selling

There is no evidence that insurance affected side-selling, in line with the design of the insurance product and the assumptions and arguments of our model. Figure 6 shows that the share of plots which side-sold is similar across treatment groups, in spite of very different levels of insurance take-up. However, it is not the *level* of side-selling, but whether or not it is selective which is important for whether pay-at-harvest insurance could be offered in equilibrium, as shown in Subsection 4.3. Proposition 4 showed that, with pay-at-harvest insurance, those with low yields are less likely to side-sell and those with high yields are more likely to, and these two effects could cancel each other out. However, if that were the case, we would expect yield conditional on selling to the company to be higher among the upfront insurance group, and we see no evidence of this in Figure 7 (we can rule out a 15% standard deviation effect on yield). Further, yield (or more precisely, insurance payout) conditional on not side-selling is what matters for the functioning of the market. The fact that pay-at-harvest insurance has little effect on yield is again consistent with some of the features of the insurance product that may inhibit selection (i.e., the double trigger and the limited scope for changing inputs).[66]

5.4.2 Did anticipation of default affect take-up differentially?

Given the extent of side-selling, it is particularly important for us to consider how important ex-ante expectations of contract risk were in driving our main result. We use two main pieces of evidence to argue that the role was limited. Before doing so, we first reiterate that the additional experiments explained above hold constant contractual risk, and thus show that both liquidity constraints and present bias are part of the story. Further, in the *Receive Choice in One Month* treatment, which is a completely upfront product and thus fully exposed to contract risk, take-up reached 72%, suggesting that the expected probability of default was unlikely to be high.

Our first argument for why the role of contract risk was limited relies on Proposition 5, which tells us that we can bound the differential effect on take-up of ex-ante expectations on contract

[66]Besides selection concerns, One might also worry that insurance induced moral hazard. However, moral hazard, if present, would work in the same direction as conditional side selling, in that the pay-at-harvest treatment group would show higher levels of moral hazard and thus lower yields.

default by that of a price cut in the upfront premium. In particular, the differential effect is less than that of a proportional price cut in the premium, of the expected probability of side-selling weighted by the relative marginal utility of consumption when side-selling. However, in our main experiment, a 30% price cut made almost no difference to take-up of upfront insurance, suggesting a low price elasticity. Thus, unless ex-ante expectations of either the probability of default or of the marginal utility of consumption in the case of default are extremely high, imperfect enforcement is unlikely to account for a large part of the difference in take-up between pay-upfront and pay-at-harvest treatments. We cannot rule out high marginal utility in the case of default - it is possible that default is associated with significant financial stress for the farmer – but it would have to be very high to explain a large fraction of our results.[67]

Our second argument for the limited role of contract risk considers heterogeneous treatment effects of delaying the premium payment, by plausible proxies for ex-ante expectations of the probability of default. If the ex-ante probability of default did drives a difference in take-up between pay-upfront and pay-at-harvest insurance, and there was heterogeneity in the probability, then, conditional on other covariates, we would expect a take-up regression to show a positive interaction between proxies for the probability of default and the harvest time premium. We consider two such proxies for the probability of default. First, in the baseline survey, we asked respondents about their trust in, and relationship with, the company. Table A.1 shows that while some of these measures do predict overall levels of take-up (consistent with a belief that the company will not make insurance payouts even if the production contract is upheld), they do not predict take-up differentially by premium timing. Second, we consider actual side-selling, both of individual farmers and of local averages (Figure 5 shows that there was substantial geographical variation in side-selling), and both in the season in question and in the previous season. For measures of side-selling in the season in question, this is based on the assumption that the actual realization of side-selling was correlated with the ex-ante probability of side-selling, and comes with the caveat that we are conditioning on an ex-post variable. Table A.3 shows that also we find no evidence for heterogeneous treatment effects for any of these proxies for ex-ante expectations of contract default.

[67]The difference with the effect of liquidity constraints is that the liquidity constraints channel is subject to the effect of present bias, as well as any other risks in the period between the take-up decision and harvest time which may lead the household to be liquidity constrained, and the one month experiment showed that just a one month delay has a significant effect.

5.5 Other channels

This section briefly discusses several additional channels which could also drive a difference in take-up between upfront and deduction premium treatment groups. We have no evidence for these channels and leave an experimental analysis of them to future work.

First, according to relative thinking (Tversky and Kahneman 1981, Azar 2007), people may make choices based on relative differences in costs, even when the rational model dictates that they should only consider absolute differences. In our setting, the premium could be perceived as a large amount when farmers consider it as an isolated expense, but as a small amount once farmers consider it as a share of their harvest revenues.[68] As such, the large difference in take-up between pay-upfront and pay-at-harvest insurance could be more to do with the latter being charged as a deduction, rather than the timing per se. A similar explanation is also offered by Salience Theory. If we interpret the model of Bordalo et al. (2012) as one of multiple time periods, diminishing sensitivity suggests that the upfront payment period may be more salient than the harvest time payment period; since income will be higher at harvest time. Salience also provides a similar explanation to quasi-hyperbolic discounting for why charging the premium in the future may result in higher demand than charging it at take-up, even without the future premium being charged as a deduction.

A second mechanism which could be responsible is prospect theory (Kahneman and Tversky 1979; Kőszegi and Rabin 2007). While a thorough application of the theory is beyond the scope of this paper,[69] intuitively upfront payments may fall in the *loss* domain, while deduction payments, at least when yield are high, may be perceived as *lower gains*. The fact that farmers may be more sensitive to losses than gains may then partially explain the large response to the timing of premium payment. This channel is again more related to the harvest time premium being charged as a deduction, rather than being paid later.

Since the pay-at-harvest insurance requires no payment at sign up, there is a third possible mechanism: a zero-price *today* effect. Empirical studies find a jump in demand at zero price across a wide range of settings (Cohen and Dupas 2010). A possibility, about which we are unaware of any papers, is that there is a similar effect when there is a zero *upfront* price, i.e. no payment to be paid at purchase. Such an effect could be an alternative explanation for the finding in Tarozzi et al. (2014) that offering anti-malarial bednets through loans has a large effect on take-up. It would also help to explain the prevalence of zero down-payment financing options for many consumer

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

[69] Such an application would require defining how the reference point is set, both for the time at which the upfront premium is paid and for harvest time.

purchases, such as cars as furniture.

6 External validity

The previous results show that charging the premium upfront, rather than at the same time as any payout would be made, reduced demand significantly for an agricultural insurance product in Kenya, and that liquidity constraints and intertemporal preferences play an important role. A wide body of work documents that these mechanisms also shape financial decisions in the developed world. For instance, a literature starting with Deaton (1991) and Zeldes (1989) points at the relationship between liquidity constraints and saving decisions, and Laibson (1997) sparked a body of empirical investigations on the impact of time-inconsistency on savings. Further, a recent literature on payday lending shows that some pay extremely high interest rates.

In the final section of the paper, we consider the implications of the intertemporal transfer in insurance contracts in developed countries, where better legal institutions may make the enforcement of cross-state insurance easier, but better functioning financial markets may also make the transfer across time matter less. We present empirical evidence from a natural experiment that, in spite of this, the timing of the premium payment also affects adoption of crop insurance among farmers in the U.S..

6.1 Evidence from U.S. Federal Crop Insurance

We exploit a natural experiment concerning the timing of the U.S. Federal Crop Insurance (FCI) premium payment. The FCI is the largest insurance scheme in the world, with a total of around 2 million policies sold in 2010. The program, which offers both crop-yield and crop-revenue insurance, is heavily subsidized by the federal government (Wright 2014). Historically, under the FCI farmers pay the insurance premium around harvest time, similar to the pay-at-harvest premium in our experiment. But this changed for some crops, in some states, in 2012.

6.1.1 Empirical design

The 2008 Farm Act[70] (the relevant part of which was implemented in 2012) moved the premium payment earlier in the season for certain crops, typically to around two to three months before harvesting. Crucially for our identification strategy, the change in timing varied across crops and states. For corn, the most common grain grown in the United States, the billing date shifted earlier throughout the country (from October to August). For wheat, the second most common grain,

[70] *Public Law 110-234.*

the billing date changed only for states growing mostly spring wheat (North West and Midwest), but not for those states growing winter wheat (Central and South).

To test the effect of this change on crop insurance adoption, we exploit the variation over time and across states and crops to implement a triple difference approach at the county level. For this purpose, we obtained data from the *Risk Management Agency* on the number of crop insurance policies sold between 2009 and 2014 by crop and county. We focus on states that harvest at least 500,000 acres of both corn and wheat in 2010, and we include only counties that grew both grains in the same year. The final sample includes 11 states and 899 counties.

Equation 15 shows the estimating equation of our triple-difference approach:

$$IHS(Policies)_{cskt} = \beta_1 Post_t * Treat_{sk} + \eta_{ck} + \eta_{st} + \eta_{kt} + \eta_{sk}t + \epsilon_{sckt}, \tag{15}$$

where the dependent variable is the inverse hyperbolic sine transformation of the number of insurance policies sold in county c in state s in crop k in year t. In years after 2012, the dummy variable $Treat_{sk}$ equals one if the state-crop pair is exposed to the reform. The model includes county-crop fixed effects (η_{ck}), state-year fixed effects (η_{st}), crop-year fixed effects (η_{kt}), and a time trend estimated separately for each state-crop ($\eta_{sk} * t$). We cluster standard errors by state-crop, as that is the level of our treatment.

β_1 identifies the effect of the earlier premium payment on take-up under the identifying assumption that, upon including all of the fixed effects listed above, county-level take-up of crop insurance had common trends across both corn and wheat, absent the policy change.

To assess the importance of wealth and liquidity constraints in this new setting, we also test whether the reform had differential impact by the average plot size in the county-crop, measured in acres. To do so we augment the model to estimate a heterogeneous treatment effect term ($Post_t * Treat_{sk} * AvgPolicySize_{ck}$):

$$\begin{aligned} IHS(Policies)_{sckt} = {} & \beta_1 Post_t * Treat_{sk} + \beta_2 Post_t * Treat_{sk} * IHS(AvgPolicySize)_{ck} \\ & + \eta_{ck} + \eta_{st} + \eta_{kt} + \eta_k * Post_t * IHS(AvgPolicySize)_{ck} \\ & + \eta_s * Post_t * IHS(AvgPolicySize)_{ck} + \eta_{sk} * IHS(AvgPolicySize)_{ck} * t + \epsilon_{sckt} \end{aligned} \tag{16}$$

This model allows average plot size to impact the dependent variable differentially in each state and crop after 2012 (terms $\eta_k * Post_t * IHS(AvgPolicySize)_{ck}$ and $\eta_s * Post_t * IHS(AvgPolicySize)_{ck}$). It also includes the interaction between the state-crop specific trends and the average policy size by county-crop ($\eta_{sk} * IHS(AvgPolicySize)_{ck} * t$).

6.1.2 Empirical results

Table 7, Column (1) estimates equation 15. It shows that the change in timing reduced the number of insurance policies sold in the county-crop by approximately 4.3%. Column (2) estimates equation 16. An increase of 10% in the average county plot size reduces the negative impact of the reform on insurance adoption by 0.52%. This results is confirmed in Column (3) where we add state-crop-year fixed effects.[71] Column (4) replaces the average county-crop plot size variable with a dummy for whether the county-crop average plot size is above the median for that state-crop. The regression confirms the findings of column (2) and suggests the negative impact of Column (1) is entirely driven by counties with farms below median size, in line with Proposition 2, which says that the difference in timing should only matter for those who risk facing liquidity constraints. Column (5) shows that this result is again robust to the inclusion of state-crop-year fixed effects.

The negative and statistically significant impact of the reform on insurance adoption suggests that the mechanisms driving the results of our experiment in Kenya may also affect risk management choices for farmers in developed countries, albeit with a lower intensity. This could be why the FCI premium has historically been due around harvest time. As in our experiment, the effect of the 2012 reform was largest among smaller farmers, in line with Proposition 2 and with the role of liquidity constraints. The effect is unsurprisingly much smaller than in our experiment. There are several natural explanations: the change in timing is much smaller (and does not interact with any present bias, since in both cases premiums are due in the future); farmers in the US may be less likely to be liquidity constrained; and finally, very importantly, late premium payments for FCI are only penalized by a penalty interest rate of 1.25% per month (within reason – beyond a certain date insurance access is revoked for subsequent seasons), which bounds the size of the possible effect of earlier payment.

7 Conclusion

By requiring that the premium be paid upfront, standard insurance contracts introduce a fundamental difference between their goal and what they do in practice: they not only transfer income across states, they also transfer income across time. We have argued that this transfer across time 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. 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

[71] In this model, we drop state-year fixed effects, crop year fixed effects and the state-crop trends.

demand for risk reduction, they face a higher cost of paying the premium upfront.

In the context of crop insurance, the transfer across time can be removed by charging the premium at harvest time, rather than upfront. Doing so in our experiment, by charging the premium as a deduction from harvest revenues in a contract farming setting, resulted in a high level of insurance take-up at actuarially fair prices. In contrast, when the premium was charged upfront, take-up for the same insurance product was 67 percentage points lower, and the effect of the timing was largest among the poorest.

We discussed the numerous channels which could be responsible for our main experimental result, and presented evidence which showed that two of the three most natural ones play a role. Heterogeneous treatment effects suggested that liquidity constraints mattered, and a first experiment on channels showed that they ran deeper than simply not having the cash to pay the premium. A second experiment on channels provided further evidence for the importance of liquidity constraints, and showed that present bias was also a significant constraint on the demand for upfront insurance, since take-up rose substantially when the premium payment was delayed by one month. Finally, we considered the role of contractual risk. A lack of trust in the insurance provider is a common reason given for not buying not buying insurance in similar settings, and ours was no different: heterogeneous treatment effects showed that insurance take-up was higher overall among those who trusted the company. But, while contractual risk *may* drive a difference between take-up of pay-upfront and pay-at-harvest insurance, in our setting, we find no evidence across multiple tests for such a differential effect, in spite of a financial shock which led to high levels of default among our farmers before harvest.

In the final contribution of the paper, we showed that the intertemporal transfer has implications for crop insurance in other settings, using a natural experiment from the U.S. It may well also have implications for other types of insurance. The most immediate comparison is with other production insurance products, but implications may be broader, as most insurance products have an intertemporal transfer. For example, the transfer across time may help to explain the low take-up for rare-disaster insurance – by definition payouts of such insurance products are rare, meaning long expected delays between payment and payout, and thus potentially large effects of small differences in interest rates faced by insurance companies and their clients. Two exceptions, i.e. insurance products which do not have a transfer across time, are the Federal Crop Insurance program detailed above and social insurance. The government is involved in the provision of both, which could in part be because cross-state insurance is hard to enforce. An alternative policy would be for the government to simply provide a loan, and allow individuals to buy standard

insurance if they want to. However, under present bias, this alternative may have negative welfare implications. Further, whereas cross-state insurance only has to be paid for in the good state of the world, the loan would have to be paid for in bad states of the world too, and limited liability could reduce the incentive to buy insurance through the standard asset substitution problem (Jensen and Meckling 1976).

From a policy perspective, boosting crop insurance take-up is an ongoing challenge. The results in this paper show that changing the timing of the premium payment represents a promising idea for doing so, which warrants replication in other settings. While the enforcement mechanism used in the paper could be used in most contract farming settings, whose presence is growing steadily in developing countries (UNCTAD 2009), two questions which we hope to investigate in future work are particularly important to understand whether the idea could be applied more widely.

First, in order to boost take-up, does the premium need to be paid at harvest time, or is there another timing which would have a similar effect? This question is important because of enforcement: the earlier the premium is paid, the less there is room for moral hazard on its payment. Our one month experiment suggests that just delaying the premium payment slightly may have a sizeable effect on take-up. But seasonality may be important too - if the premium was paid at the previous harvest time, rather than early on in the agricultural season, farmers may be less liquidity constrained, and thus more willing to purchase insurance. A related question is whether there are other timings of the insurance *payout* which farmers would prefer. Even in a bad season farmers may well have liquidity at harvest time, only running out later in the season - especially under present bias. If farmers are sophisticated, then they may prefer an insurance product which pays out later in the season, when times are hardest.

Second, if the premium is charged at (subsequent) harvest time, are there other ways to enforce payment by the farmer, beyond interlinking the product and insurance markets? A comparison to credit suggests there could be: the enforcement constraint for cross-state insurance is, if anything, easier than that for credit, since in the case of insurance a net payment is only required in good states of the world. Whether methods used to ensure repayment in credit, such as relational contracting, group liability, and credit scores could be used for cross-state insurance is an important question.

In addition to improving the enforcement of a premium paid at harvest, we note finally that interlinking insurance and production contracts can benefit insurance design and administration in several other ways. The first concerns data availability. For administrative purposes, large buyers in contract farming schemes often collect detailed plot-level data, including output and farm sizes.

These records, which can span for decades, can address data limitations for costing insurance products, a fundamental constraint in the design of area yield products (Elabed et al. 2013). This could be particularly relevant as area yield insurance may display lower basis risk than rainfall index insurance, as was the case in our setting. Second, administrative costs are likely to be lower. For example, company field assistants already visit farmers' plots at contract inception, and in an interlinked contract that visit could include insurance recruitment, at minimal additional cost. This would reduce the gap between actuarially fair premiums and market premiums, which can be large.[72] Third, getting people to renew their insurance contracts has been a challenge in other settings, with high dropout rates among farmers who do not receive an insurance payout in the first season (Cole et al. 2014, Cai et al. 2016). Under an interlinked contract, farmers could credibly sign up for insurance contracts which cover multiple seasons, increasing the chance that they receive an insurance payout before policy renewal is due. Finally, evidence that farmers increase their productive investments and output when insured (Karlan et al. 2014) provides another rationale for interlinked contracts. If product buyers are partial residual claimants on farmers' production, they make a profit on the additional quantities produced. Thus, unlike a third party insuring agent, they may not need to break even on the insurance sales.

[72] In Karlan et al. (2014), market prices are 50% higher than actuarially fair prices.

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Casaburi, Lorenzo, and Tristan Reed. 2014. "Incomplete Pass-Through and Interlinked Transactions." *Unpublished*.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Ghani, T., and T Reed. 2014. "Competing for Relationships: Markets and Informal Institutions in Sierra Leone." *Unpublished*.

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

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

IFAD. 2003. *Agricultural Marketing Companies as Sources of Smallholder Credit in Eastern and Southern Africa: Experience, Insights and Potential Donor Role.*: Eastern and Southern Africa Division, IFAD.

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

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

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

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

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

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

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

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

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

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

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

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

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.

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

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

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

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

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

Rampini, Adriano A, and S Viswanathan. 2013. "Collateral and capital structure." *Journal of Financial Economics*, 109(2): 466–492.

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

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

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

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

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

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

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

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

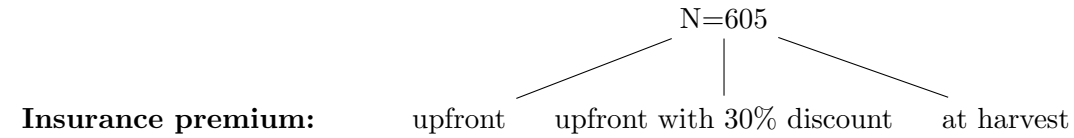
UNCTAD. 2009. "World Investment Report 2009 - Transnational Corporations, Agricultural Production and Development."

Wright, Brian Davern. 2014. "Multiple Peril Crop Insurance." *Choices*, 29(3): .

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

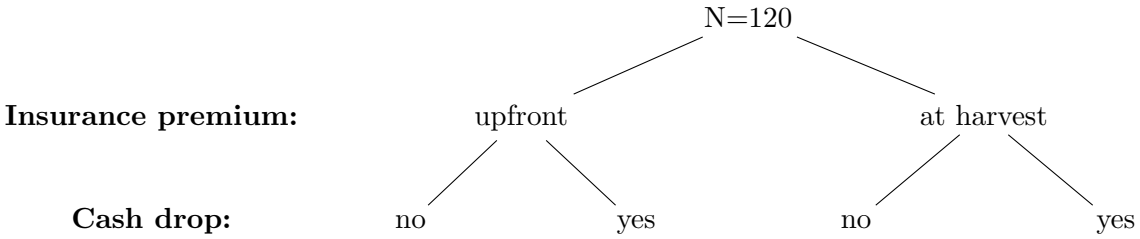
Figure 1: Experimental design

(a) Design of Main Experiment



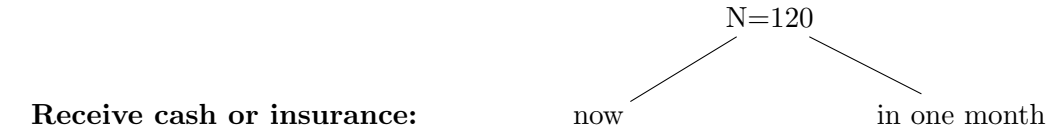
Notes: The experimental design randomized 605 farmers (approximately) equally across three treatment groups. All farmers were offered an insurance product; the only thing varied across treatment groups was the premium. In the first group (U1), farmers were required to pay the ("actuarially-fair") premium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

(b) Design of Cash Constraints Experiment



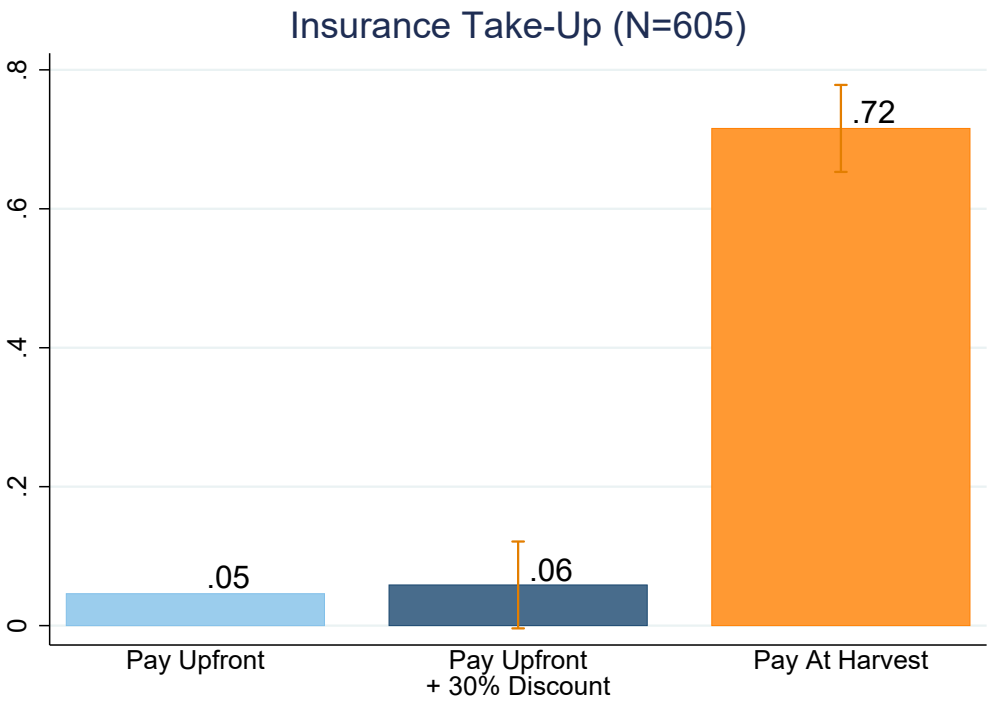
Notes: The experimental design randomized 120 farmers (approximately) equally across four treatment groups. The design cross-cut two treatments: pay-upfront vs. pay-at-harvest insurance, as in the main experiment, and a cash drop. At the beginning of individual meetings with farmers, those selected to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

(c) Design of Present Bias Experiment



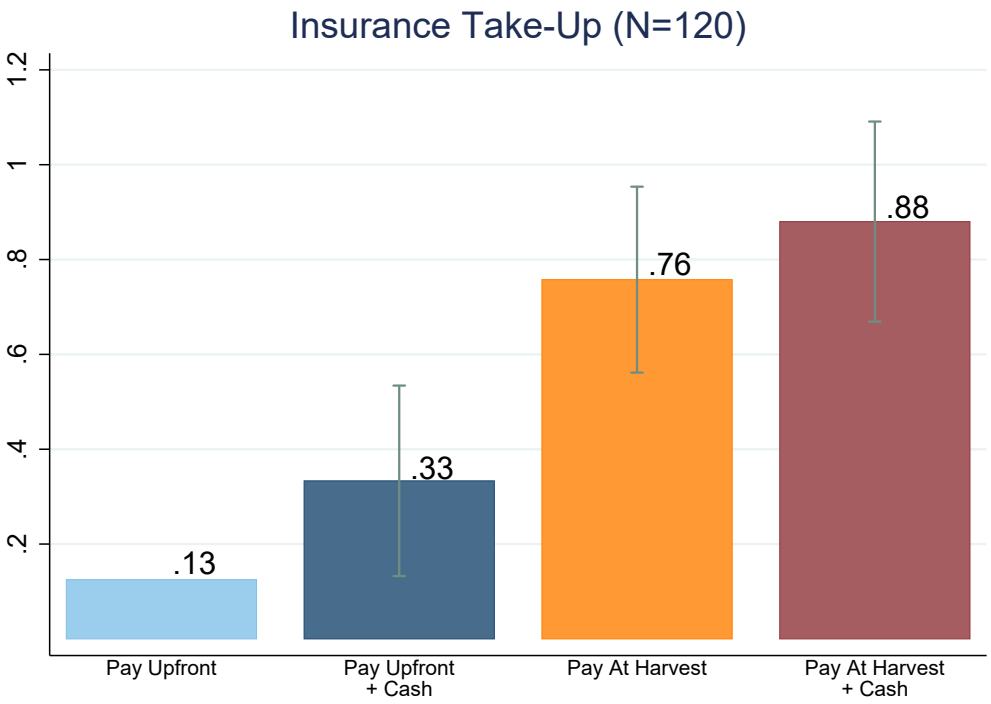
Notes: The experimental design randomized 120 farmers (approximately) equally across two treatment groups. Farmers in both groups were offered a choice between either a cash payment, equal to the "full-priced" insurance premium, or free enrollment in the insurance. Both groups had to make the choice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

Figure 2: Main Experiment: Insurance Take-Up by Treatment Group



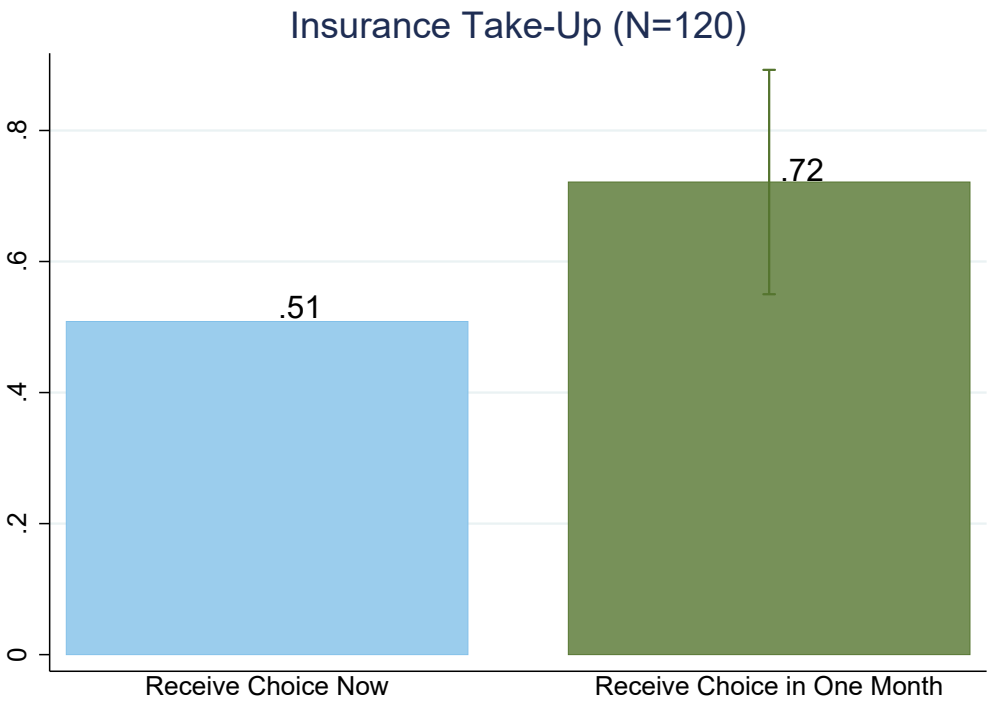
Notes: The figure shows insurance take-up rates across the three treatment groups in the main experiment. In the *Pay Upfront* group, farmers had to pay the full-price premium when signing up to the insurance. In the *Pay Upfront + 30% Discount* group, farmers also had to pay the premium at sign-up,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

Figure 3: Cash Constraints Experiment: Insurance Take-Up by Treatment Group



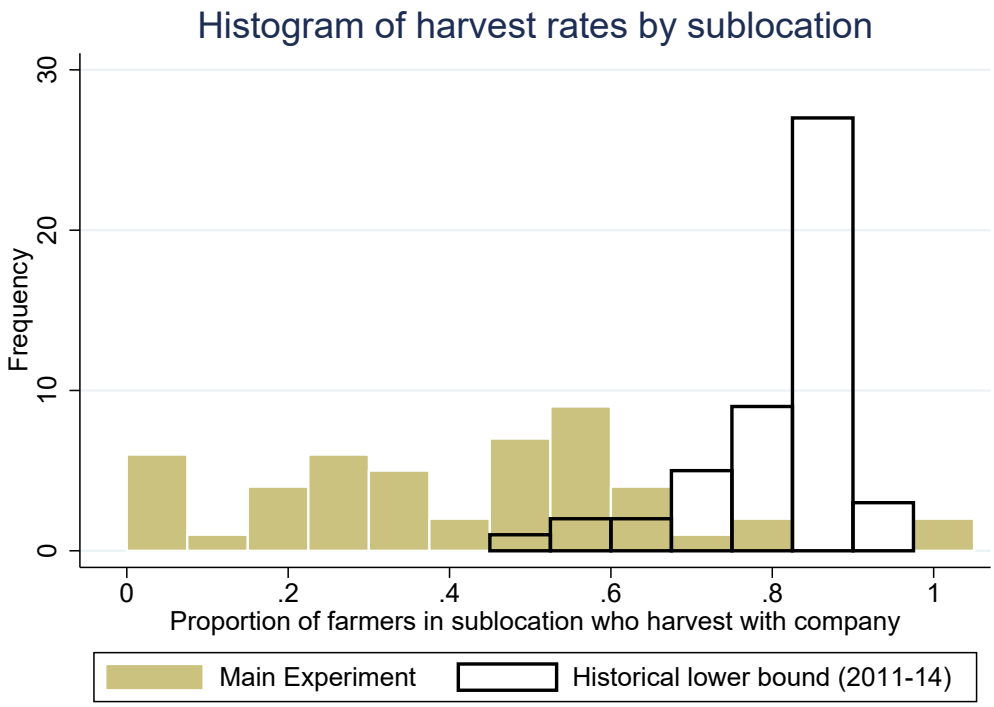
Notes: The figure shows insurance take-up rates across the four treatment groups in the cash constraints experiment. In the *Pay Upfront* group, farmers had to pay the premium when signing up for the insurance. In the *Pay Upfront + Cash* group, farmers were given a cash drop slightly larger than the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

Figure 4: Present Bias Experiment: Insurance Take-Up by Treatment Group



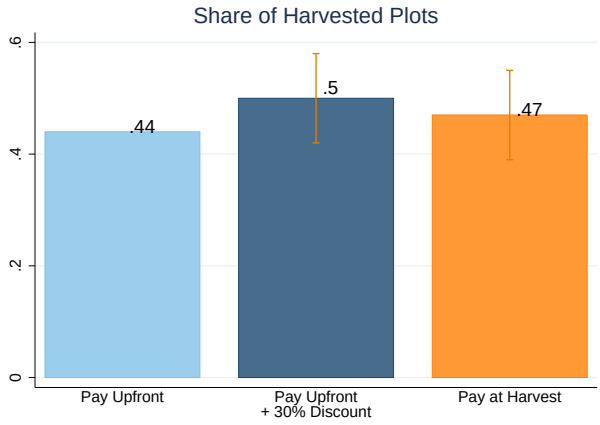
Notes: The figure shows insurance take-up rates across the two treatment groups in the present bias experiment. In the *Receive Now* group, farmers chose between an amount of money equal to the premium and free subscription to the insurance, knowing that they would receive their choice straight away.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

Figure 5: Histogram of Harvesting With Company, by Sublocation



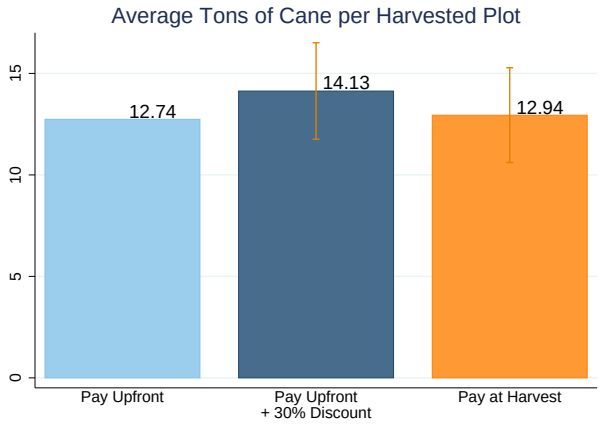
Notes: The histogram shows the proportion of farmers who harvested with the company in the sublocations in which we undertook the experiment. The data is by sublocation and we plot separate histograms for the main experiment (which is just for the farmers in our sample, who were due to harvest approximately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

Figure 6: Contractual Default: Proportion Who Harvest with Company in Main Experiment



Notes: The figure shows the proportion of farmers from the main experiment who subsequently harvested with the company, as agreed under the contract. In the *Pay Upfront* group, farmers had to pay the full-price premium when signing up for the insurance. In the *Pay Upfront + Discount* group, farmers had to pay the premium at sign-up but received a 30% price reduction. In the *Pay At Harvest* group, premium payment was through deduction from (future) harvest revenues, and included the accrued interest. The bars capture 95% confidence intervals.

Figure 7: Contractual Default: Harvest Weight Conditional on Harvesting with Company in Main Experiment



Notes: The figure shows the harvest weight, conditional on harvesting with the company, for farmers in the main experiment. In the *Pay Upfront* group, farmers had to pay the full-price premium when signing up for the insurance. In the *Pay Upfront + Discount* group, farmers had to pay the premium at sign-up but received a 30% price reduction. In the *Pay At Harvest* group, premium payment was through deduction from (future) harvest revenues, and included the accrued interest. The bars capture 95% confidence intervals.

Table 1: Main Experiment: Balance Table, Baseline Variables

	Upfront [U1]	Upfront-30% [U2]	Harvest [H]	P-value [U1-U2]	P-value [U1-H]	P-value [U2-H]	P-value [U-H]	N
Plot Size (Acres)	.309	.340	.319	.012**	.187	.219	.875	605
	(.128)	(.147)	(.133)					
Previous Yield	53.9	56.8	55.7	.23	.112	.67	.581	605
	(17.2)	(18.0)	(17.7)					
Man	.670	.731	.671	.114	.732	.387	.779	571
	(.471)	(.444)	(.470)					
Age	48.3	48.3	45.4	.983	.06*	.019**	.023**	568
	(13.5)	(14.2)	(11.7)					
Land Cultivated (Acres)	2.79	3.15	2.81	.563	.421	.259	.858	560
	(3.34)	(4.02)	(2.98)					
Own Cow(s)	.761	.808	.803	.22	.209	.952	.452	566
	(.427)	(.394)	(.398)					
Portion of Income from Cane	3.30	3.38	3.23	.503	.817	.103	.274	567
	(1.09)	(1.09)	(1.18)					
Savings for Sh1,000	.284	.279	.333	.936	.273	.362	.207	564
	(.452)	(.450)	(.472)					
Savings for Sh5,000	.089	.139	.130	.198	.345	.829	.559	563
	(.286)	(.347)	(.338)					
Expected Yield	72.5	70.5	66.8	.674	.522	.577	.543	571
	(100.)	(31.1)	(49.4)					
Expected Yield in Good Year	83.5	85.5	83.1	.844	.852	.944	.88	570
	(47.8)	(30.8)	(56.0)					
Expected Yield in Bad Year	53.4	54.0	52.3	.958	.889	.972	.916	569
	(40.6)	(21.6)	(47.2)					
Good Relationship with Company	.362	.341	.300	.821	.357	.933	.553	568
	(.482)	(.475)	(.459)					
Trust Company Field Assistants	3.02	2.83	2.80	.163	.11	.877	.347	567
	(1.01)	(1.01)	(1.09)					
Trust Company Managers	2.46	2.34	2.44	.319	.999	.513	.678	565
	(1.11)	(1.05)	(1.12)					

Notes: The table presents the baseline balance for the Main Experiment. *Previous Yield* is measured as tons of cane per hectare harvested in the cycle before the intervention. *Man* is a binary indicator equal to one if the person in charge of the sugarcane plot is male. *Own Cow(s)* is a binary indicator equal to one if the household owns any cows. *Portion of Income from Cane* takes value between 1 ("None") to 6 ("All"). *Savings for Sh 1,000 (Sh 5,000)* is a binary indicator that equals one if the respondent says she would be able to use household savings to deal with an emergency requiring an expense of Sh 1,000 (Sh 5,000). 1 USD= 95 Sh. *Good Relationship with the Company* is a binary indicator that equals one if the respondent says she has a "good" or "very good" relationship with the company (as opposed to "bad" or "very bad"). *Trust Company Field Assistants* and *Trust Company Managers* are defined on a scale 1 ("Not at all") to 4 ("Completely"). P-values are based on specifications which include field fixed effects (since randomization was stratified at the field level). *p<0.1, **p<0.05, ***p<0.01.

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

	(1)	(2)	(3)	(4)	(5)	(6)
30% Discount	0.013	0.004		0.013	0.008	0.020
	[0.032]	[0.033]		[0.033]	[0.032]	[0.033]
Pay At Harvest	0.670***	0.675***	0.673***	0.680***	0.687***	0.697***
	[0.032]	[0.033]	[0.028]	[0.033]	[0.032]	[0.032]
Field FE	N	Y	Y	Y	Y	Y
Plot Controls	N	N	N	Y	N	Y
Farmer Controls	N	N	N	N	Y	Y
Mean Y Control	0.046	0.046	0.052	0.046	0.046	0.046
Observations	605	605	605	605	605	605

Notes: The table presents the results of the Main Experiment. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. Specification (3) bundles together treatment groups U1 (Pay Upfront) and U2 (Pay Upfront with 30% discount) as baseline group. *Plot Controls* are *Plot Size* and *Previous Yield*. *Farmer Controls* are all of the other controls reported in the balance table, Table 1. Stratification occurred at the field level. *p<0.1, **p<0.05, ***p<0.01.

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

X =	Land Cultivated (standardized)	Own Cow(s)	Previous Yield (standardized)	Plot Size (standardized)	Portion of Income from Cane	Savings for Sh1,000	Savings for Sh5,000
X * Pay At Harvest	-0.065**	-0.141*	-0.079**	-0.001	0.052*	-0.175**	-0.131
	[0.033]	[0.079]	[0.031]	[0.031]	[0.028]	[0.069]	[0.097]
X	-0.000	0.070	0.015	-0.022	-0.004	0.008	-0.016
	[0.017]	[0.045]	[0.020]	[0.019]	[0.017]	[0.043]	[0.059]
Pay At Harvest	0.705***	0.823***	0.673***	0.672***	0.542***	0.764***	0.725***
	[0.029]	[0.069]	[0.028]	[0.028]	[0.096]	[0.035]	[0.031]
Mean Y Control	0.052	0.052	0.052	0.052	0.052	0.052	0.052
Mean X	0	0.792	0	0	3.309	0.300	0.121
S.D. X	1	0.407	1	1	1.126	0.459	0.326
Observations	560	566	605	605	567	564	563

Notes: The table shows heterogenous treatment effects on take-up from the Main Experiment, by different proxies for liquidity constraints. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance, and in each column the relevant heterogeneity variable (X) is reported in the column title. Treatments U1 (Pay Upfront) and U2 (Pay Upfront with 30% discount) are bundled together as baseline group, as specified in the pre-analysis plan. *Land cultivated* is the standardized total area of land cultivated by the household. *Own Cow(s)* is a binary indicator for whether the household owns any cows. *Previous Yield* is the standardized tons of cane per hectare harvested in the cycle before the intervention. *Plot size* is the standardized area of the sugarcane plot. *Portion of Income from Cane* takes value between 1 ("None") to 6 ("All"). *Savings for Sh 1,000 (Sh 5,000)* is a binary indicator that equals one if the respondent says she would be able to use household savings to deal with an emergency requiring an expense of Sh 1,000 (Sh 5,000). 1 USD = 95 Sh. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

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

	(1)	(2)	(3)	(4)	(5)	(6)
Pay At Harvest	0.593***	0.603***	0.589***	0.633***	0.635***	0.635***
	[0.074]	[0.077]	[0.078]	[0.100]	[0.105]	[0.107]
Cash	0.167**	0.132*	0.128	0.208**	0.167	0.177
	[0.074]	[0.079]	[0.079]	[0.102]	[0.110]	[0.111]
Pay At Harvest * Cash				-0.086	-0.071	-0.100
				[0.148]	[0.156]	[0.159]
Field FE	N	Y	Y	N	Y	Y
Plot Controls	N	N	Y	N	N	Y
Mean Y Control	0.125	0.125	0.125	0.125	0.125	0.125
P-value: Deductible = Cash	0.000	0.000	0.000	0.000	0.000	0.000
Observations	120	120	120	120	120	120

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

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

	(1)	(2)	(3)	(4)	(5)
Receive Choice in One Month	0.213**	0.252***	0.255***	0.294***	0.301***
	[0.087]	[0.090]	[0.092]	[0.109]	[0.111]
Field FE	N	Y	Y	Y	Y
Plot Controls	N	N	Y	N	Y
Farmer Controls	N	N	N	Y	Y
Mean Y Control	0.508	0.508	0.508	0.508	0.508
Observations	120	120	120	120	120

Notes: The table presents the results of the Present Bias Experiment. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. The baseline (omitted) group is the *Receive Now* group, where farmers chose between an amount of money equal to the premium and free subscription to the insurance. In the *Receive Choice in One Month* group, farmers made the same choice, but were told that what chose would be delivered one month later (plus one month's interest if they chose cash). *Plot Controls* are *Plot Size* and *Previous Yield*. *Farmer Controls* are all the other controls reported in the main balance table, Table 1. Stratification occurred at the field level. *p<0.1, **p<0.05, ***p<0.01.

Table 6: Main Experiment: Take-Up by Harvest Status

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
H	0.675***	0.694***	0.585***	-0.054	0.719***	0.673***	0.688***	0.585***	0.096	0.707***
	[0.033]	[0.054]	[0.109]	[0.407]	[0.046]	[0.028]	[0.047]	[0.092]	[0.358]	[0.040]
U2	0.004	0.012	-0.001	-0.316	0.026					
	[0.033]	[0.055]	[0.109]	[0.408]	[0.047]					
H*Share Harvested in Field		-0.046					-0.037			
		[0.103]					[0.089]			
H*Share Harvested in Subloc			0.207					0.202		
			[0.238]					[0.202]		
H*PAST Share Harvested in Subloc				0.869*					0.687	
				[0.484]					[0.425]	
H*Plot Harvested					-0.096					-0.072
					[0.070]					[0.060]
U2*Share Harvested in Field		-0.018								
		[0.101]								
U2*Share Harvested in Subloc			0.010							
			[0.240]							
U2*PAST Share Harvested in Subloc				0.382						
				[0.484]						
U2*Plot Harvested					-0.048					
					[0.071]					
Plot Harvested					0.039					0.015
					[0.057]					[0.044]
Field FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Mean Y Control	0.046	0.046	0.046	0.046	0.046	0.046	0.046	0.046	0.046	0.046
Observations	603	603	603	603	603	603	603	603	603	603

Notes: This table presents take-up during the experiment by subsequent harvesting behavior approximately twelve months later. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. *H* is a binary indicator for the *Pay At Harvest* treatment group. *U2* is a binary indicator for the *Pay Upfront with 30% discount* treatment group. *Share harvested in Field* is the proportion of farmers in the Field (an administrative, geographic unit) who harvest with the company. *Share harvested in Subloc* is the proportion of farmers in the Sublocation (a geographic identifier which is coarser than Field) who harvest with the company. *PAST Share harvested in Subloc* is the same variable, but instead covering the time period 2011-14, before the experiment, when side-selling was lower. *Plot harvested* is a binary indicator for whether the farmer harvests his plot with the company. Specifications (6)-(10) bundle groups treatments U1 (Upfront Premium at full price) and U2 (Upfront Premium at 30% discount) as baseline group. Regressions contain field fixed effects and are clustered at the field level. *p<0.1, **p<0.05, ***p<0.01.

.

Table 7: U.S. Federal Crop Insurance Regression Table

	(1)	(2)	(3)	(4)	(5)
Post 2012*Treatment	-0.039***	-0.281***		-0.060***	
	(0.015)	(0.094)		(0.021)	
Post 2012*Treatment*IHS(AvgPlotSize)		0.047***	0.043***		
		(0.016)	(0.016)		
Post 2012*Treatment*(AvgPlotSize > Median)				0.043**	0.043**
				(0.017)	(0.017)
R^2	0.265	0.312	0.317	0.297	0.302
Dependent Variable Mean	5.771	5.771	5.771	5.771	5.771
County*Crop FE	Y	Y	Y	Y	Y
State*Year FE	Y	Y	N	Y	N
Crop*Year FE	Y	Y	N	Y	N
State*Crop Trend	Y	Y	N	Y	N
State*Crop Trend * Size	N	Y	Y	Y	Y
State*Crop*Year FE	N	N	Y	N	Y
Observations	10753	10753	10753	10753	10753

Notes: The table presents the affect of the change in the premium billing date for U.S. Federal Crop Insurance on insurance adoption. Data are at the county-crop-year level. The dependent variable is the inverse hyperbolic sine ($\approx log(2) + log(x)$) of the number of policies sold in the county-crop-year.[73] *Treatment* is a binary indicator equal to one if the billing date for the county-crop is earlier than the harvesting period from 2012 onward. *AvgPlotSize* is the average plot size in the county. Column (1) shows the basic difference-in-differences by state-crop (Equation (15) in the paper). Columns (2)-(5) show the heterogeneity by average plot size in the county-crop (Equation (16) in the paper). Columns (2)-(3) present heterogeneity by the inverse hyperbolic sign of the average plot size in the county-crop. Columns (3)-(4) present heterogeneity by a binary indicator that is equal to one if the average plot size in the county-crop is above the median in the state-crop. Standard errors clustered by crop-state. *p<0.1, **p<0.05, ***p<0.01.

[73]The inverse hyperbolic sign is defined as $log(x + (x^2 + 1)^{1/2})$, which is approximately equal to $log(2) + log(x)$, but is defined when the number of policies sold is zero.

A Appendix

A.1 Bounding the effect of the transfer across time

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

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

To simplify, we now assume that at harvest time there are just two states of the world, the standard state h and the low state l, with the low state happening with probability p.[74] We assume that insurance is perfect - it only pays out in the low state (at time H), and that it is again actuarially fair. To simplify notation, in this section we denote by R the interest rate on the insurance *covering the whole period from the purchase decision until harvest time.* We also assume CRRA utility, so that $u(c) = c^{1-\gamma}/(1-\gamma)$.

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

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

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

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

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

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

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

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

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

with A given by:

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

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

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

$$c_H(y_h) > A c_H(y_l)$$

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

$\gamma = 1$

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

$\gamma = 2$

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

$\gamma = 5$

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

A.2 Supply of insurance

While the farmer is better off with the pay-at-harvest insurance, strategic default means that the insurer may be worse off. Suppose the insurer could set the price of the insurance, p. Under pay-at-harvest insurance at price p, their expected profit conditional on them not defaulting themselves is:

$$\begin{aligned}\text{Expected profit for insurer} &= \mathbb{E}[(p-I)D_p(p)] \\ &= \mathbb{P}[D_p(p)=1]\mathbb{E}[p-I|p-I \leq \frac{c_D}{u'(c_H)}]\end{aligned}$$

This makes it clear that imperfect enforcement may prevent pay-at-harvest insurance from being offered when c_D is small, since strategic default means the insurer never makes a profit, even as he increases price. This is likely the reason why pay-at-harvest insurance is not offered more often. However, charging the premium upfront places all of the risk of default on the farmer, and as we have seen may also activate liquidity constraints.

When $c_D > 0$, which timing of premium payment would prevail in equilibrium is not clear-cut. Intuitively, the risk of insurer default pushes towards pay-at-harvest, whereas the risk of farmer default pushes towards upfront. To the extent that insurer risk is not subject to moral hazard or adverse selection, for example if it is caused by stochastic shocks which occur after take-up, then it may be priced in to the insurance, suggesting upfront may be optimal. But liquidity constraints, as well as the second order effects (not modelled here) of exposing the farmer to risk over whether the insurer will default, push the other way, as does moral hazard on the insurer and adverse selection over insurance companies which are likely to default.

To demonstrate simply how the different forces push towards upfront or at-harvest contracts, we consider the above setup, and assume that insurers know ex-ante whether they will default or not, so are of two types: proportion $1-p_I$ who never default, and proportion p_I who always default before harvest. Further assume that the farmer cannot tell ex-ante which type he faces. Under this setup, we allow the price p of the insurance to vary, and consider under which conditions a price p could exist under which trade would occur.

For pay-upfront insurance, this requires

$$pR^{-H}u'(c_0) \leq \beta\delta^H(1-p_I)\mathbb{E}[Iu'(c_H)]$$

and[75]

$$p \geq 1$$

For pay-at-harvest insurance, it requires

$$\mathbb{E}[D_p(p)pu'(c_H) + (1-D_P(p))c_D] \leq \mathbb{E}[D_p(p)Iu'(c_H)]$$

and

$$\mathbb{E}[p-I|p-I \leq \frac{c_D}{u'(c_H)}] \geq 0$$

where $D_p(p) = \mathbb{I}[p-I \leq \frac{c_D}{u'(c_H)}]$. Given these conditions, the following follows easily

Proposition 6.
Pay-at-harvest insurance *$\exists \underline{c}_D > 0$ for which there is no price at which pay-at-harvest insurance would be traded when the cost of default $c_D(s) < \underline{c}_D \; \forall s$. However, if $c_D > max_s\, u'(c_H(s)) \; \forall s$, then deductible insurance could be traded at price 1.*
Pay-upfront insurance *$\exists \underline{p}_I < 1$ such that if $p_I > \underline{p}_I$, there is no price at which pay-upfront insurance would be traded.*

[75] If $p < 1$ the defaulting insurers would still offer insurance, but the farmers would know that only defaulting insurers were offering insurance, and hence would not take it up.

This proposition shows that the cost of strategic default for the insured, c_D, and the probability of the insurer defaulting, p_I, are key considerations for which insurance products could be exist in equilibrium (and if so how much insurance they could provide). The cost of strategic default is similarly found to be important in a literature discussing another type of purely cross-state insurance: risk sharing (Ligon et al. 2002; Kocherlakota 1996). The literature considers both how c_D is made large enough for informal risk sharing to exist, typically through relational contracting and punishment mechanisms, and how the size of c_D dictates the extent of risk sharing which can be achieved. Related to the discussion here, Gauthier et al. (1997) consider how enlarging the contracting space in risk sharing, so as to allow for ex-ante transfers, enlarges the set of parameters for which the first-best outcome is achievable.

A.3 Proofs and derivations

A.3.1 Background

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

Dynamic programming problem

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

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

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

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

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

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

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

Proof of Lemma 1

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

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

Concavity. Assume V^c_{t+N} is concave. Then, V^c_{t+N-1} is strictly concave, since the utility function is concave and the state space correspondence in convex, by standard argument (take $x_\theta = \theta x_a + (1-\theta)x_b$, expand out the definition of $V^c_{t+N-1}(x_\theta)$ and use the concavity of V^c_{t+N-1} and the strict concavity of $u(.)$). Iterating this argument, we thus have that V^c_t is concave. Therefore, since there is a unique fixed point of the contraction mapping from V^c_{t+N} to V^c_t, that fixed point must be concave (since we will converge to the fixed point by iterating from any starting function; start from a concave function).

Part (2)

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

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

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

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

Part (3) The proof relies on showing that the mapping from $V^{c'}_{t+1}$ to $V^{c'}_t$ conserves convexity, $\forall t \in \mathbb{R}^+$. Then the proof follows as in 1 above: V^c_t is the fixed point of a contraction mapping which conserves convexity of the first derivative, hence $V^{c'}_t$ must be convex. We provide two methods to show that the mapping preserves convexity, the first of which is based on Deaton and Laroque (1992), and the second of which assumes $V^{c'''}_t$ exists and gives more intuition for where the convexity is coming from.

Version 1 Suppose $V^{c'}_{t+1}$ is convex.

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

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

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

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

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

Version 2 Suppose $V^{c'}_{t+1}$ is convex.

$$\begin{aligned} V^c_t(x_t) &= max_c\ u(c) + \delta\mathbb{E}[V^c_{t+1}(R(x_t - c) + y_{t+1})] \\ \Rightarrow V^{c'}_t(x_t) &= \delta R\mathbb{E}V^{c'}_{t+1}(x_{t+1}) + c'_t(x_t)\mu_t \\ \Rightarrow V^{c''}_t(x_t) &= \delta R^2\mathbb{E}V^{c''}_{t+1}(1 - c'_t(x_t)) + c''_t(x_t)\mu_t + c'_t(x_t)\mu'_t \\ \Rightarrow V^{c'''}_t(x_t) &= \delta R^3\mathbb{E}V^{c'''}_{t+1}(1 - c'_t(x_t))^2 - \delta R^2\mathbb{E}V^{c''}_{t+1}c''_t(x_t) + c'''_t(x_t)\mu_t + 2c''_t(x_t)\mu'_t + c'_t(x_t)\mu''_t \end{aligned}$$

Now, the first order condition is

$$u'(c_t) = \beta\delta R\mathbb{E}V^{c'}_{t+1}(x_{t+1}) + \mu_t$$

Consider first the case where $\mu_t > 0$. Then $c(x_t) = x_t$ and $v^{c'}_t(x_t) = u'(c(x_t))$. So $v^{c'}_t(x_t) = u'(x_t)$ thus $v^{c'''}_t(x_t) = u'''(x_t) > 0$.

So, assume $\mu = 0$. Differentiating the FOC wrt x gives

$$u''(c(x_t))c'(x_t) = \beta\delta R^2 \mathbb{E}V_{t+1}^{c''}(1 - c'(x_t))$$
$$\Rightarrow c'(x_t) = \frac{\beta\delta R^2 \mathbb{E}V_{t+1}^{c''}(x_t)}{\beta\delta R^2 \mathbb{E}V_{t+1}^{c''}(x_t) + u''(c(x_t))}$$

Differentiating again with respect to w gives

$$c''(x_t) = \frac{(\beta\delta R^3 \mathbb{E}V_{t+1}^{c'''}(1-c'))(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'') - (\beta\delta R^2 \mathbb{E}V_{t+1}^{c''})(\beta\delta R^3 \mathbb{E}V_{t+1}^{c'''}(1-c') + u'''c')}{(\beta\delta R^2 \mathbb{E}V_{+}^{c''} + u'')^2}$$
$$= \frac{u''\beta\delta R^3 \mathbb{E}V_{t+1}^{c'''}(1-c') - u'''c'\beta\delta R^2 \mathbb{E}V_{t+1}^{c''}}{(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'')^2}$$

Substituting in to the equation for $V^{c'''}$ gives:

$$V_t^{c'''}(x_t) = \delta R^3 \mathbb{E}V_{+}^{c'''}(\frac{u''}{\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u''})^2 - \delta R^2 \mathbb{E}V_{t+1}^{c''} \frac{u''\beta\delta R^3 \mathbb{E}V_{t+1}^{c'''}(1-c') - u'''c'\beta\delta R^2 \mathbb{E}V_{t+1}^{c''}}{(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'')^2}$$
$$= \frac{u''^2\delta R^3 \mathbb{E}V_{t+1}^{c'''}(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'') - \beta\delta^2 R^5 u''^2 \mathbb{E}V_{t+1}^{c''}\mathbb{E}V_{t+1}^{c'''} + \beta\delta^2 R^4 u'''c'(\mathbb{E}V_{t+1}^{c''})^2(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'')}{(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'')^3}$$
$$= \frac{u''^3\delta R^3 \mathbb{E}V_{t+1}^{c'''}}{(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'')^3} + \frac{\beta\delta^2 R^4 u'''c'(\mathbb{E}V_{t+1}^{c''})^2}{(\beta\delta R^2 \mathbb{E}V_{t+1}^{c''} + u'')^2}$$
$$\geq 0$$

A.3.2 Insurance with perfect enforcement

Proof of Propostion 1

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

Suppose farmers have zero probability of being liquidity constrained before the next harvest, when they buy pay-upfront insurance. Denote their (state-dependent) path of assets until harvest by $(a_t^u)_{t<H}$, given that they have purchased pay-upfront insurance. By the assumption that the farmers will not be liquidity constrained before harvest, $a_t^U > 0 \ \forall t < H$ and for all histories $(s_i)_{i\leq t}$. Now, suppose instead of pay-upfront insurance, they had been offered pay-at-harvest insurance. If they invest the money they would have spent on pay-upfront insurance in assets instead, so $a_t^H(s) = a_t^U(s) + R^{-H-t}$, then they can pay the pay-at-harvest premium at harvest time and have the same consumption path as in the case of pay-upfront, so they must be at least as well off. Similarly, suppose they optimally hold $(a_t^D)_{t<H}$ in the pay-at-harvest case. If instead offered upfront insurance, they can use some of these assets to instead buy insurance, so that $a_t^U(s) = a_t^D(s) - R^{-H-t}$. Since, by assumption $a_t^U(s) > 0$, doing so they can again follow the same consumption path as in the case of pay-at-harvest insurance, so pay-upfront insurance is at least as good as at-harvest insurance. Thus the farmer is indifferent between pay-upfront and pay-at-harvest insurance.

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

pay-upfront insurance product which had premium $(1 - \frac{\lambda_0^H}{u'(c_0)})R^{-H}$. The net benefit would be

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

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

Proof of Propostion 2

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

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

Now, $\frac{dw_H}{dx_0} \geq 0$, by iterating lemma 1 back from period H to period 0. Also, $y_H + I - 1$ strictly second order stochastic dominates y_H by assumption, and $V_H^{c'}$ is strictly convex ($V^{c'''} > 0$ by lemma 1), so $\mathbb{E}(V_H^{c'}(w_H + y_H + I - 1)) - \mathbb{E}(V_H^{c'}(w_H + y_H)) < 0$. Thus, the value of pay-at-harvest insurance is decreasing with wealth. Now, the reduction in net utility from insurance arising from upfront premium payment is $R^{-H}\lambda_0^H$, by proposition 1. By lemma 1, this is also decreasing in wealth.

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

Proof of Proposition 3

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

A.3.3 Insurance with imperfect enforcement

Outside option $o(s_H, w_H)$

If the farmer chooses to sell to the company he receives profits $y(s)$ (comprising revenues minus a deduction for inputs provided on credit) plus any insurance payout $I(s)$, minus the insurance premium in the case of pay-at-harvest insurance. He also receives continuation value $r_C(s)$ from the relationship with the company, which is possibly state dependent. If he chooses to side sell, he receives outside option $o(s)$ [76], and saves the deductions for inputs provided on credit and for

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

the deductible insurance premium, but loses the continuation value and any insurance payout. We abstract from any impact of insurance on the choice of input supply, since, as argued before, the choice set is limited, the double trigger design of the insurance was chosen to minimize moral hazard, and, as reported below, we see no evidence of moral hazard in the experimental data.

Proof of Proposition 4

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

If the insurer has not already defaulted, they are:

$$\begin{aligned} D_I =& \mathbb{I}[c_D \geq 0] && \text{without insurance} \\ D_I^U =& \mathbb{I}[(1-p_{s_2})Iu'(c_H) + c_D \geq 0] && \text{with pay-upfront insurance} \\ D_I^D =& \mathbb{I}[(1-p_{s_2})Iu'(c_H) + c_D \geq u'(c_H)] && \text{with pay-at-harvest insurance} \end{aligned}$$

If the insurer has already defaulted, they are:

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

Since $(1-p_{s_2})I(s)u'(c_H(s))$ and $u'(c_H(s))$ are non-negative, and $Iu'(c_H)$ and $(I-1)u'(c_H)$ are larger when yields are low, the results follow.

Proof of Proposition 5

Consider the net benefit of insurance, which is the benefit of the payout minus the cost of the premium payout. With perfect enforcement, we know that at-harvest insurance is equivalent to upfront insurance with a percentage price cut of $\frac{\lambda_0^H}{u'(c_0)}$. Denote the net benefit of such an upfront insurance product by S_U, and by S_D the net benefit of at-harvest insurance. Then[77]:

$$\begin{aligned} \mathbb{E}[S_D - S_U] =& (1-p_{s_1})(\Sigma_{d^D,d^U \in \{0,1\}} \mathbb{P}[D_I^U = d^U, D_I^D = d^D]\mathbb{E}[S_D - S_U | D_I^U = d^U, D_I^D = d^D]) \\ &+ p_{s_1}(\Sigma_{d^D,d^U \in \{0,1\}} \mathbb{P}[D_D^U = d^U, D_D^D = d^D]\mathbb{E}[S_D - S_U | D_D^U = d^U, D_D^D = d^D]) \end{aligned}$$

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

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

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

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

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

$$S_D - S_U = \beta\delta^H(u'(c_H) - (1-p_{s_2})Iu'(c_H) - c_D) \leq \beta\delta^H u'(c_H)$$

[77] See proof of Proposition 4 for the definitions of the D.

Thus:

$$\begin{aligned}\mathbb{E}[S_D - S_U] \leq &(1-p_{s_1})(\mathbb{P}[D_I^U = D_I^D = 0] + P[D_I^U = 1, D_I^D = 0])\beta\delta^H \mathbb{E}[u'(c_H)|D_I^D = 0] \\ &+ p_{s_1}(\mathbb{P}[D_D^U = D_D^D = 0] + \mathbb{P}[D_D^U = 1, D_D^D = 0])\beta\delta^H E[u'(c_H)|D_D^D = 0]\end{aligned}$$

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

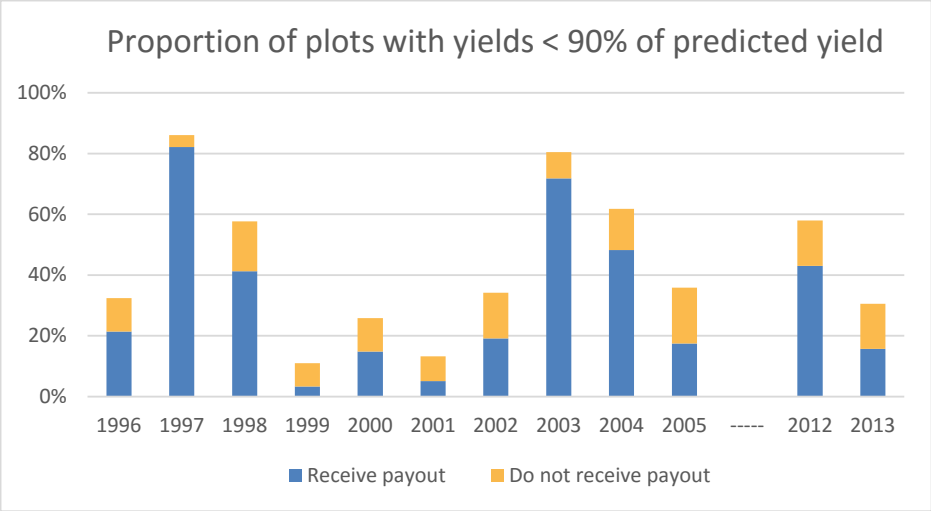
$$\begin{aligned}\Leftrightarrow \mathbb{E}[S_D - S_U] \leq &(1-p_{s_1})\mathbb{P}[D_I^D = 0]\beta\delta^H \mathbb{E}[u'(c_H)|D_I^D = 0] \\ &+ p_{s_1}\mathbb{P}[D_D^D = 0]\beta\delta^H \mathbb{E}[u'(c_H)|D_D^D = 0] \\ \Leftrightarrow \mathbb{E}[S_D - S_U] \leq &\mathbb{P}(\text{side-sell with at-harvest})\beta\delta^H \mathbb{E}(u'(c_H)|\text{side-sell with at-harvest})\end{aligned}$$

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

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

A.4 Risk profile of the insurance

Figure A.1: Simulation of insurance payout based on historical data



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

A.5 Additional experimental results

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

	(1) Good Relationship with Company	(2) Trust Company Field Assistants	(3) Trust Company Managers
X * Pay At Harvest	-0.063	0.021	0.027
	[0.070]	[0.029]	[0.028]
X	0.088**	0.035*	0.028
	[0.041]	[0.019]	[0.017]
Pay At Harvest	0.726***	0.656***	0.642***
	[0.035]	[0.087]	[0.073]
Mean Y Control	0.052	0.052	0.052
Mean X	0.335	2.885	2.418
S.D. X	0.472	1.045	1.099
Observations	568	567	565

Notes: The table shows heterogeneity of the treatment effect of the pay-at-harvest premium on insurance take-up in the main experiment by different proxies for trust toward the company. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. Treatments U1 (Upfront Payment) and U2 (Upfront Payment with 30% discount) are bundled together as baseline group, as outlined in the pre-analysis plan. The relevant heterogeneity variable is reported in the column title. The notes of Table 1 provide a definition of the variables used in the heterogeneity analysis. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

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

	RRR on inputs	RRR 0 to 1 week	RRR 0 to 1 week − RRR 1 to 2 weeks
X * Pay At Harvest	-0.123	0.099	0.003
	[0.142]	[0.114]	[0.153]
X	0.071	0.035	0.119
	[0.082]	[0.065]	[0.091]
Pay At Harvest	0.760***	0.684***	0.716***
	[0.054]	[0.042]	[0.029]
Mean Y Control	0.052	0.052	0.052
Mean X	0.326	0.270	-0.043
S.D. X	0.228	0.278	0.211
Observations	559	561	559

Notes: The table shows heterogeneity of the treatment effect of the pay-at-harvest premium on insurance take-up in the main experiment, by preferences in Money Earlier or Later experiments. The dependent variable is a binary indicator equal to one if the farmer took-up the insurance. Treatments U1 (Upfront Payment) and U2 (Upfront Payment with 30% discount) are bundled together as baseline group. The relevant heterogeneity variable is reported in the column title. These variables come from responses to hypothetical (Becker-DeGroot) choices over earlier or later cash transfers, which give various Required Rates of Returns. 'RRR for inputs' is the required rate of return which would (hypothetically) make farmers indifferent between paying for inputs upfront and having them deducted from harvest revenues. 'RRR 0 to 1 week' is the required rate of return to delay receipt of a cash transfer by one week. 'RRR 0 to 1 week - RRR 1 to 2 weeks' is the difference between the rates of return required to delay receipt of a cash transfer from today to one week from now, and from one week from now to two weeks from now. All columns include field fixed effects. *p<0.1, **p<0.05, ***p<0.01.

Table A.3: Main Experiment: Harvest Status vs. Wealth & Liquidity Constraint Proxies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Yield	0.005***							0.005***
	[0.001]							[0.001]
Plot Size		0.329**						0.429***
		[0.151]						[0.159]
Land Cultivated (Acres)			0.002					0.001
			[0.005]					[0.006]
Any Cow				-0.018				-0.035
				[0.052]				[0.052]
Portion of Income from Cane					0.030*			0.012
					[0.018]			[0.018]
Savings for Sh1,000						0.000		-0.072
						[0.048]		[0.056]
Savings for Sh5,000							0.064	0.112
							[0.062]	[0.076]
Constant	0.190***	0.361***	0.473***	0.490***	0.377***	0.479***	0.470***	0.042
	[0.072]	[0.050]	[0.023]	[0.045]	[0.060]	[0.022]	[0.017]	[0.109]
Field FE	Y	Y	Y	Y	Y	Y	Y	Y
Observations	603	603	559	564	565	562	561	552

Notes: This table presents ex-post harvesting behavior by proxies for wealth and liquidity constraints at baseline. The dependent variable is whether the farmer harvested with the company, a binary indicator. The independent variables are as described in Table 3. *p<0.1, **p<0.05, ***p<0.01.

Table A.4: Liquidity Constraints Experiment: Balance Table

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

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

Table A.5: Present Bias Experiment: Balance Table

	Receive Now	Receive in One Month	p-value	N
Plot Size	.328	.292	.093*	120
	(.109)	(.105)		
Yield	58.0	57.7	.512	120
	(20.1)	(21.5)		
Man	.793	.590	.009***	119
	(.408)	(.495)		
Age	48.3	47.7	.573	119
	(12.8)	(11.9)		
Land Cultivated (Acres)	3.47	2.67	.067*	117
	(2.92)	(1.72)		
Any Cow	.844	.852	.987	119
	(.365)	(.357)		
Portion of Income from Cane	3.62	3.32	.193	119
	(1.12)	(.943)		
Savings for Sh1,000	.327	.295	.526	119
	(.473)	(.459)		
Savings for Sh5,000	.155	.065	.056*	119
	(.365)	(.249)		
Expected Yield	77.7	87.5	.47	119
	(65.3)	(38.4)		
Expected Yield in Good Year	95.1	109.	.322	119
	(70.7)	(48.4)		
Expected Yield in Bad Year	63.0	69.4	.682	119
	(61.7)	(32.0)		
Good Relationship with Company	.310	.316	.622	118
	(.466)	(.469)		
Trust Company Field Assistants	3.10	2.83	.315	119
	(1.02)	(1.01)		
Trust Company Managers	2.15	2.11	.32	119
	(1.13)	(1.03)		

Notes: The table presents baseline balancing for the Present Bias Experiment. *Previous Yield* is measured as tons of cane per hectare harvested in the cycle before the intervention. *Man* is a binary indicator equal to one if the person in charge of the sugarcane plot is male. *Own Cow(s)* is a binary indicator equal to one if the household owns any cows. *Portion of Income from Cane* takes value between 1 ("None") to 6 ("All"). *Savings for Sh 1,000 (Sh 5,000)* is a binary indicator that equals one if the respondent says she would be able to use household savings to deal with an emergency requiring an expense of Sh 1,000 (Sh 5,000). 1 USD= 95 Sh. *Good Relationship with the Company* is a binary indicator that equals one if the respondent says she has a "good" or "very good" relationship with the company (as opposed to "bad" or "very bad"). *Trust Company Field Assistants* and *Trust Company Managers* are defined on a scale 1 ("Not at all") to 4 ("Completely"). P-values are based on specifications which include field fixed effects (since randomization was stratified at the field level). *p<0.1, **p<0.05, ***p<0.01.

Table A.6: U.S. Crop Insurance Robustness Regression Table

	DD across states (spring vs. winter wheat)		DD within states (corn vs. winter wheat)		DDD			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post 2012*Treatment	-0.081***	-0.081***	0.044	0.044	0.023	-0.081***	-0.089***	-0.039***
	(0.025)	(0.025)	(0.034)	(0.034)	(0.031)	(0.025)	(0.022)	(0.015)
R^2	0.075	0.095	0.147	0.182	0.132	0.190	0.254	0.265
Dependent Variable Mean	5.480	5.480	5.653	5.653	5.771	5.771	5.771	5.771
County*Crop FE	Y	Y	Y	Y	Y	Y	Y	Y
Year FE	N	Y	N	Y				
Crop*Year FE					N	Y	Y	Y
State*Year FE					N	N	Y	Y
State*Crop Trend					N	N	N	Y
Observations	5378	5378	8337	8337	10753	10753	10753	10753

Notes: The table provides further details for the triple difference estimates provided in Table 7 for the effect of the change in the premium billing date for U.S. Federal Crop Insurance on insurance adoption. Columns (1)-(4) give the difference-in-differences effects which go into the triple difference estimate and columns (5)-(8) present robustness checks for the estimate. Data are at the county-crop-year level. The dependent variable is the inverse hyperbolic sine ($\approx log(2) + log(x)$) of the number of policies sold in the county-crop-year. *Treatment* is a binary indicator equal to one if the billing date for the county-crop is earlier than the harvesting period from 2012 onward. *AvgPlotSize* is the average plot size in the county. Columns (1) and (2) report a difference-in-differences estimate, which compares insurance take-up for wheat across states: spring wheat states (for which the premium moved before harvest) vs. winter wheat states (for which the premium did not move before harvest time). Columns (3) and (4) report an alternative difference-in-differences estimate, which compares insurance take-up across crops within (winter wheat) states: corn (treated) vs winter wheat (not treated). Columns (5)-(8) report robustness of the DDD estimate to the inclusion of fixed effects. Standard errors clustered by crop-state. *p<0.1, **p<0.05, ***p<0.01.