The Feasibility of Picture-Based Crop Insurance (PBI): Smartphone Pictures for Affordable Crop Insurance

F. Ceballos[1]; B. Kramer[1]; M. Robles[2]

1: International Food Policy Research Institute, Markets, Trade, and Institutions, United States of America, 2: The World Bank, , United States of America

Corresponding author email: f.ceballos@cgiar.org

Abstract:

This paper describes and tests the feasibility of Picture-Based Crop Insurance (PBI), a new way to deliver affordable and easy-to-understand insurance. Under PBI, loss assessments are based on damage visible from a time-series of pictures taken by the farmer using regular smartphones. PBI aims at boosting uptake, trust, and understanding of insurance by reducing basis risk as well as costs of—and delays in—loss assessment, and by engaging farmers to participate directly, with one's own pictures being more tangible than other indices. Results from a pilot implementation in the rice-wheat belt of India speak to PBI being a feasible and valuable alternative to existing insurance products. Damage is visible from smartphone pictures, farmers can take pictures of sufficient quality for loss assessment, and PBI helps reduce severe downside basis risk at minimal cost.

Acknowledegment: We gratefully acknowledge Braulio Britos and Matt Krupoff for excellent research assistance; and Azad Mishra, Siddhesh Karekar, Dr. Mann S. Toor, Koen Hufkens, Michael Mann, and Eli Melaas. We received valuable suggestions from participants at Göttingen University, PARM K-Sharing and Learning Workshop; IPA's 3rd Annual Researcher Gathering on Financial Inclusion and Social Protection; CEAR Academic Pre-Conference in Microinsurance; and the Regional Dialogue on The Role of Agricultural Research in the Design, Implementation, and Evaluation of Agricultural Insurance. Funding support was provided by the CGIAR Research Program on Policies, Institutions, and Markets (PIM) and the International Initiative for Impact Evaluation (3ie).

JEL Codes: O13, O16

#1094

IAAE

1. Introduction

Around the developing world, small farmers in rural areas generally lack formal protection from natural hazards such as drought, extreme heat, excess rainfall, hail, or pests and diseases. Because their farms are too small and too remote, insurers cannot provide reliable indemnity coverage at affordable rates (Hazell, Pomareda, and Valdes, 1986). The absence of formal risk management mechanisms makes farmers vulnerable to extreme weather shocks and leads to underinvestment in productivity-enhancing technologies (Barrett and McPeak, 2006; Cai, 2013; Cai et al., 2009; Cole, Giné, and Vickery, 2017; Dercon and Hoddinott, 2004; Karlan et al., 2014; Mobarak and Rosenzweig, 2012).

In the past few decades, index-based insurance arose as a potential solution to these issues. Index insurance pays out benefits based on a predetermined index or proxy for losses resulting from weather and other catastrophic events (e.g. excess rainfall or extreme temperatures). Because the index determines payouts, there is no need for insurance claims adjusters to assess damage on individual fields, making claims settlement processes cheaper, faster, and more objective. Yet, demand for these products is typically low, a finding which has been explained in part by high levels of basis risk (meaning that the index and plot-level damage are not sufficiently correlated), limited trust in insurance providers, and a lack of understanding of the insurance product (Cole et al., 2013; Hill, Robles, and Ceballos, 2016; Matul et al. 2013; Mobarak and Rosenzweig, 2012).

In response to the challenges, the index insurance community has been exploring alternatives in two important fronts: (i) reducing basis risk to increase the value proposition of index-based insurance, and (ii) participatory approaches that can help draw farmers into the insurance process to increase understanding and trust in the product. For instance, one strand of work has started using high-resolution satellite imagery whilst validating the methodology and improving take-up through participatory community meetings (Carter et al., 2008; Chantarat et al., 2013). However, in many settings, basis risk will remain an issue even when using high-resolution imagery due to intercropping on small plots. A second strand has proposed adding an extra layer of protection on top of a traditional index product, such as gap insurance (Berhane et al., 2015), which allows for audit-based payouts if the index does not trigger, while a sufficient proportion of farmers in an area claim to have suffered losses. A key challenge in delivering gap insurance is that there is no documentation of pre-damage and pre-audit crop conditions.

This paper describes a new way to deliver affordable and easy-to-understand crop insurance that bundles and aims to improve upon both strands of work: Picture-Based Insurance (PBI). PBI provides insurance coverage for damage that is visible in a time-series of both pre-damage and post-damage pictures taken by the farmer using regular, low-cost smartphones. By taking regular georeferenced pictures using their own smartphones, farmers can reliably document damage after a natural calamity and provide evidence that the crop was managed appropriately until that point. This can help reduce overcome information asymmetries and bring down the high costs of plot-level loss verification that have challenged traditional indemnity insurance. PBI is participatory and tangible, and it can deliver plot-level assessments of damage, removing key barriers in the demand for existing index insurance products, including basis risk, low trust, and poor understanding. As such, PBI can help deliver gap insurance, with pictures of insured crops serving as input for audits when farmers report losses despite the index not having triggered.

The aim of PBI is to combine key advantages of both index-based insurance—timely compensation without expensive loss assessments—and indemnity insurance—minimum basis risk and an easy to understand product. This innovation comes in a timely manner. PBI takes advantage of the general trend of increasing smartphone ownership in developing countries, improved penetration of low-cost mobile internet services among smallholder farmers, and recent advances in image processing for near-surface remote sensing through digital repeat photography.

PBI is a novel concept that has not been tested before. Hence, this paper describes the implementation and results of a formative evaluation addressing key knowledge gaps around the feasibility of this approach. This feasibility study targeted 750 smallholder wheat farmers in Haryana and Punjab, two states in northwest India, and was designed to answer three main questions: (i) to what extent are farmers willing to participate in the insurance process by regularly uploading geo-referenced pictures of their plots; (ii) to what extent is damage visible in the smartphone camera data, that is, do images taken by smallholder farmers using their own phones contain visible characteristics that are predictive of crop damage; and (iii) does PBI reduce basis risk, or, in other words, does PBI offer improved protection against crop damage compared with conventional weather-index insurance products?

Overall, the results speak to PBI being a feasible and valuable alternative to existing insurance products. We find that farmers are able and willing to use the smartphone application and upload enough pictures of sufficient quality for loss assessment. Damage was visible from smartphone pictures and could be quantified by local expert agronomists; picture-based damage estimates are strongly correlated with yields and improve upon weather-based indices. We also perform simulations showing that PBI is a cost-effective add-on to provide gap insurance in the context of an area yield-based index, which is the main index used in India's national crop insurance scheme. Based on these findings, we conclude that PBI offers a promising alternative to existing insurance products for poor farmers.

The next section discusses the study context and procedures. Section 3 describes the different data sources on which the subsequent analyses rely on. Section 4 reviews the available evidence for each of the questions presented above. Section 5 brings the evidence together and provides concluding remarks and some avenues for future research.

2. Context and Procedures

In this section, we provide more information on the context of the feasibility study and the procedures that were followed. The next section will first describe the study region in which we conducted the formative evaluation, as well as the sampling procedures used to identify study participants. We then describe the insurance products that were tested as part of the study, including the PBI product and the weather index-based product that was used for comparison purposes. The final part of this section provides a detailed description of the study procedures.

2.1 Study Context and Sampling

The study was conducted in six districts in the states of Haryana and Punjab. 50 villages were randomly selected within a radius of five kilometers from available weather stations (to limit spatial basis risk), subject to the condition that the village had at least 40 households, 40 main cultivators, or a total population of over 140 individuals during the 2011 Indian Agricultural Census (to capture enough farming households within each village).

In each village, we conducted a listing exercise of all farming households and randomly selected—among those owning a smartphone and planning to grow at least two acres of wheat during the upcoming Rabi season—15 farmers per village for a baseline survey. These farmers were equally distributed across three categories: operating less than five acres, operating five to ten acres, and operating ten to fifteen acres of farmland. Our focus on relatively smaller farmers was motivated by external validity considerations in terms of the representative farmer population in other states. In addition, the PBI approach appears more valuable and relevant for smallholder farmers than for larger farmers, for whom plot sizes are large enough, and land cover sufficiently homogenous, for high-resolution satellite imagery to capture plot conditions.

2.2 Insurance products

All farmers in the baseline survey were offered insurance for one acre of their wheat crop during the upcoming Rabi 2016/17 growing season. In all 50 study villages, the product included a weather index-based component (WBI), which triggered payouts in case of unseasonal rains or above-normal temperatures during February-April (around harvest).[1] For each index, payouts were triggered once the index exceeded a strike value, which was set to the 70th percentile based on historical weather data, and payouts were linearly increasing in the index until reaching an exit value, set to the 99th percentile. For index levels at or above the exit value, farmers would receive the total sum insured. The WBI product would make payments for one of the two indices, whichever triggered the highest payout. The sum insured for this product was 13,000 Indian Rupees (Rs.) or, at the current exchange rate of Rs. 65 per USD, 200 US dollars per acre.

For every weather station, we also randomly selected one of the two villages—or 25 villages in total—to receive *in addition* to the WBI component a picture-based insurance component (PBI), which provided coverage against visible damage in pictures taken throughout the Rabi season.[2] Farmers were informed that to determine payouts, independent experts would inspect their pictures for visible damage due to risks beyond their control such as lodging, hail storms, unseasonal rains, pests and disease, or wild animals. In the absence of existing loss assessment algorithms, this procedure was transparent and acceptable to participating farmers. Farmers were told that damage

[1] These risks the main weather-related risks reported by farmers during focus group discussions. The trigger values were set to the 70th percentile based on historical weather data, and exit values were set to the 99th percentile.

[2] We randomized the type of insurance product offered to farmers in order to test whether PBI affects farmer behavior. The findings from that behavioral experiment are discussed in Ceballos and Kramer (2018).

below 20 percent would not trigger payouts; damage between 20 to 50 percent would trigger a payout of Rs. 3,900; damage between 50 and 75 percent would trigger a payout of Rs. 7,800; and damage above 75 percent would trigger a payout of Rs. 13,000, equal to the insured sum. Farmers would receive the maximum payout of the WBI and PBI components, whichever was higher.

2.3 Procedures

During July and August 2016, we conducted a baseline survey among the 15 selected farmers in each village. In October 2016, we invited these farmers to village sessions in which they were informed that they would receive—free-of-charge—agricultural insurance for one acre of their wheat crop during the upcoming Rabi 2016/17 growing season. The type of product—WBI or the combined WBI and PBI product—was randomized at the village level. All in all, 592 farmers (approximately 12 per village) agreed to provide crop pictures, of which 296 farmers received the WBI product, with the remaining 296 farmers receiving the combined WBI and PBI product. While only this latter group of farmers was insured under PBI, *all* farmers were told that their WBI coverage would be conditional on following the picture-taking protocol during the entire season. Farmers also received a data plan to upload the pictures, conditional on following this protocol.

For each farmer, the protocol for taking pictures entailed capturing repeat photographs of the same portion of a randomly-selected field (i.e. a *site*) three times a week.[3] In addition, pictures were to be taken between 10am and 2pm in order to maintain appropriate and comparable lighting levels across all images. Due to app compatibility issues with older Android versions in the farmer's phones at the launch of the project, all farmers who had agreed to take pictures were provided with a low-cost Android smartphone. They were also provided with a set of two inexpensive poles: an *auxiliary* pole, which served as a tripod to help maintain a fixed position from where to place the phone and take the repeat pictures, and a *reference* pole, which served as a fixed reference in the plot to aid with the framing of the picture (see panel A of Figure 1).

In an initial visit, project staff would download the smartphone app to farmers' phones, enter their unique IDs, take an initial picture in the randomly selected sites, and train farmers on how to use the app and take repeat pictures. Farmers were told that pictures had to be taken from the same

[3] Due to technical problems in the initial roll-out of the WheatCam app, it was decided to relax this criteria, and to consider farmers for insurance pay-outs if they had taken an initial picture in 2016 and at least 2 pictures in 2017.

spot, pointing at the same direction every time. The app facilitated this task through geotags and visual aids. Geotags were used to issue warnings if the repeat picture was being taken at a location different from the initial picture. Further, when taking a repeat picture, the app displayed the initial picture as a "ghost" image (a mildly transparent image), allowing the farmer to align static features in the landscape (such as distant trees or structures as well as the reference pole in the field) with those same elements in the initial picture, thus ensuring an almost identical view frame throughout the season (see panel B of Figure 1). Valid pictures were uploaded to a server and processed by the research team. Farmers could reach out throughout the season to project staff for troubleshooting in case they encountered any problems with the app or protocol.

Figure 1. Visual aids for maintaining a fixed view frame through the growing season

Panel A. Reference and Auxiliary Poles

Panel B. Ghost image

At the end of the season, an independent panel of wheat experts evaluated the time-series of pictures and estimated a percentage of crop damage for each of the available plots. At least three different experts reviewed each time-series. Assessments were first done individually and the median assessment was used to determine insurance payouts. If, however, large disagreement existed between the experts' assessments, a final damage estimate was agreed upon through consensus, and in those cases, the amount reached through consensus was used to determine insurance payouts. Assessments were anonymous; with no access to the farmer's personal details or type of insurance coverage. Expert loss assessments for farmers with damage assessments of more than 20 percent were passed on to the insurance company, which reviewed the existing evidence and issued payments directly into the farmers' bank accounts.

3. Data

In this section, we describe the primary sources of information that we will be using in the analyses below. Most importantly, we conducted baseline and endline surveys with all available farmers during, respectively, August 2016 and April 2017. The baseline survey inquired about an array of farm and household characteristics, such as a roster of plots and specific plot-level characteristics, cultivation practices, input use, adoption of several agricultural technologies, household composition, income, risk attitudes and perceptions, among many others. The endline survey gathered further data on cultivation practices over the 2016/17 rabi season, including input use and self-reported wheat output, perceptions about the insurance product received, and experience with taking the pictures and using the smartphone app.

In order to obtain an objective measure of wheat yields for our season of interest, we carried out crop cutting exercises (CCEs) at all sites in which a farmer had taken at least two pictures over the season. Although this type of measure can still suffer from a number of shortfalls, it is nevertheless considered the gold standard for estimating plot-level yields. At each site, the process consisted of sampling two different square meters (visible in the picture): one to the left and one to the right of the reference pole. The heads of the wheat plants falling inside these sampled square meters were then weighted and recorded. Our final yield estimate for a given field is then the average between the left and right yields. The CCEs were carried out right before harvest, during the first half of April 2017. CCEs were not used to determine insurance payouts, and farmers were only informed about our efforts to conduct CCEs on the day itself.

Finally, we use data stemming from the expert assessments of the time-series of wheat pictures at each site. Once the Rabi season was over and the time-series of pictures had been processed and cleaned, each crop site was individually reviewed by three randomly-assigned wheat experts. For each site, the experts would assess whether the crop was damaged. They would also indicate the loss percentage, the cause of the damage, and in which picture the damage could be first observed. In addition, each expert indicated what percentage of the visible damage was due to unavoidable hazards or due to mismanagement by the farmer. Finally, for those sites in which experts disagreed substantially on the amount of damage, the experts jointly discussed the assessments and reached a consensus about the loss percentage.

4. Analysis and Results

This section reports findings regarding three key knowledge gaps concerning the feasibility of the picture-based insurance approach. First, we will discuss whether farmers are willing and able to send in a sufficiently large number of crop pictures through a smartphone app for loss assessment purposes. Second, we assess whether the smartphone pictures contain visible characteristics that capture damage events, that is, whether damage can be quantified accurately from smartphone camera data. Third, we will analyze to what extent PBI reduces basis risk compared with alternative index insurance approaches.

4.1 Farmer's Ability and Willingness to Take Repeat Pictures

A first prerequisite for the feasibility of PBI is that there is sufficient camera data available at the time of loss assessment: whether automated through image processing algorithms, or through visual inspection by wheat experts, for picture-based loss assessment we require a sufficient number of images for a given site from which to determine the overall damage (if any) suffered by the crop. In the present study, we requested farmers to not only take post-damage pictures but instead to take pictures continuously throughout the season, so that they would document pre-damage conditions, which we believe is important in making the system tamper-proof and reducing scope for moral hazard (see Ceballos and Kramer, 2018). For this, farmers need to be willing and able to take pictures of their fields regularly and with a sufficient level of quality.

Out of the full sample of 592 farmers who agreed to send in pictures on a regular basis and were trained on using the smartphone app, 475 farmers (80.2 percent) uploaded at least one valid picture during the season.[4] Panel A of Figure 2 shows the distribution of the number of pictures taken for this sample of 475 farmers. Of them, the large majority (more than 83 percent) took at least six pictures throughout the season—or roughly one picture per month. A comparable number of farmers (80.5 percent) uploaded at least two pictures in 2017, making them eligible for loss

[4] Valid pictures are considered pictures of sufficient quality showing an unobstructed view of the same portion of the selected farmer's field (including the reference pole).

assessments (not shown in Figure 2). More than 59 percent of them took pictures twice a month or more, resulting in a high-quality time series that can be used to develop image processing algorithms for automated loss assessment.

Figure 2. Picture-taking activity

Panel A. Total number of picture taken by farmer

Panel B. Time of the day at which pictures were taken

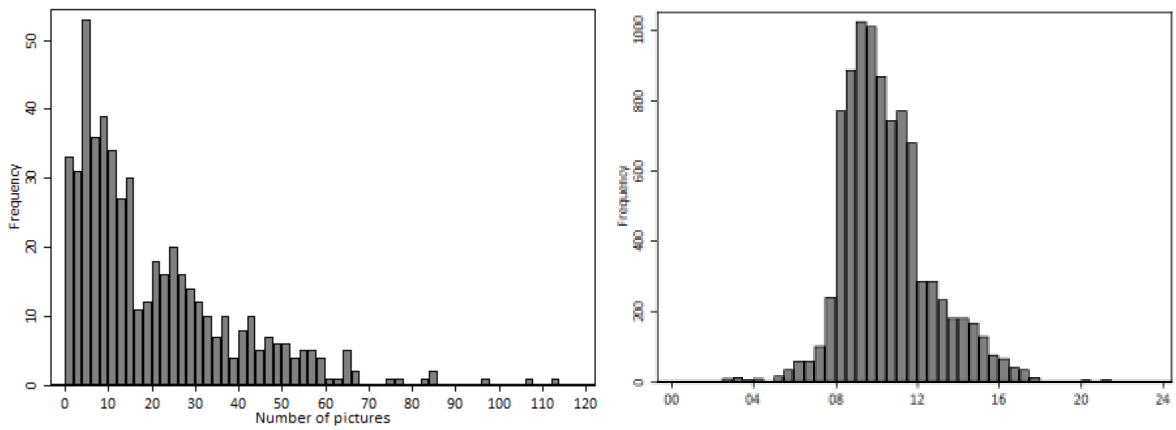

Panel C. Number of farmers taking at least one picture per week

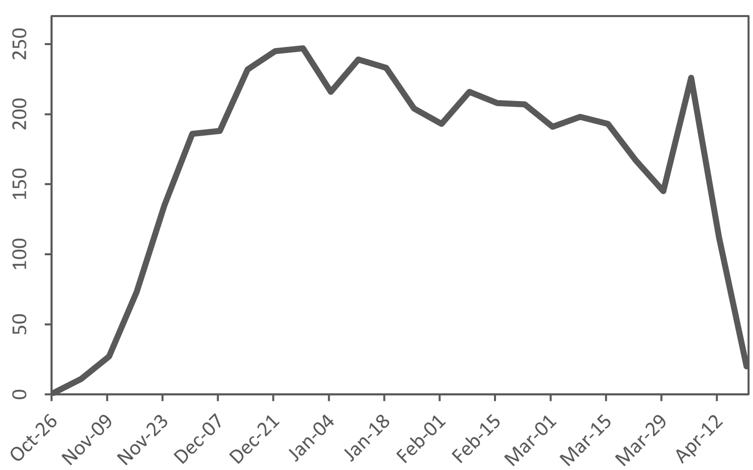

With respect to the time of the day at which pictures were taken, Panel B shows that farmers took pictures across a broader range of times than initially requested. This can be explained in part by the fact that most plots are not located close to the farmer's home, due to which farmers only visit their plots at certain times of the day. Moreover, during the colder months of December and January, fog could strongly reduce the visibility in Haryana and Punjab, especially in the morning time. As a result, farmers often needed to wait until the fog had cleared before taking a picture.

In order to get an idea of the pattern of picture-taking activity over time, Panel C shows the number of farmers who took at least one picture in a given calendar week throughout the season. The pattern is encouraging, with sustained submissions from an average of 200 farmers weekly, except for the beginning of the season and the post-harvest period.

Finally, we analyze whether farmers' ability and willingness to take pictures for insurance purposes depends on observable farmer characteristics. To that end, we study two types of participation. First, we analyze attendance of village sessions, during which farmers were given more information about the insurance products, to explore which characteristics may determine farmers' interest in insurance ex-ante. Second, we analyze the number of pictures uploaded conditional on taking at least one (initial) picture (indicating that project staff successfully installed the technology on a farmer's smartphone), as a proxy for continued participation. Differences in participation across demographic or socioeconomic dimensions could indicate that PBI is more appropriate or inclusive for specific segments of the population.

The first column in Table 1 shows the results from an ordinary least squares regression (OLS) using as dependent variable a dummy indicator for whether a farmer attended the village session. Farmers belonging to a scheduled or other backward caste and farmers with larger households were less likely to participate. Wheat yields from the previous season are also negatively correlated with attending the village session, perhaps related to farming ability (where higher-ability farmers may value insurance less) or to a recency bias (where those farmers who did not recently experience problems with their wheat crops tend to dismiss the probability of future hazards). Interestingly, farmers more dependent on income from crops and those perceiving wheat yields to be more variable were more interested in attending an insurance-related session. Other variables, including farmer size, age, and education level, experience with smartphones, and a measure of farmer's progressiveness (as captured by having adopted laser-land levelling in any of his plots in the past) are not significantly related to the probability of attending a village session.

The other columns in Table 1 show alternative specifications assessing the relationship between a farmer's characteristics and his likelihood of taking at least one repeat picture and his propensity to keep taking pictures, conditional on having an initial picture —that is, having the app installed in their phone and having been shown how to use it by field staff. Overall, the results are consistent across specifications. Belonging to a lower caste is a strong determinant for nonparticipation and reduced picture-taking, both in the intensive and extensive margins. In terms of age, we expected a priori higher technology acceptance among younger farmers and, as such, higher engagement among the youth. The results are however not consistent with this perspective: it is the oldest age tercile in our sample (those above 50 years) that seem to have participated differentially more compared with the middle tercile (those between 30 and 50 years). Interestingly, farmers who did not own their insured plot tended to take more pictures, perhaps related to the fact that conventional insurance products available in the market are linked to land ownership, but that the present study made an exception for these farmers. Finally, a puzzling finding is that farmers whose plots are located farther from their homes tended to take more pictures. This could perhaps be related to more established routines for visiting their plots, though we do not have data to test this hypothesis.

Table 1. Factors behind attending a village session, taking pictures, and number of pictures taken

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Attended vill. sess.	Took pictures		Number of pictures		No. of pictures (cond. on taking)		Took at least 10 pictures	
	OLS	OLS		Tobit		OLS		OLS	
PBI village	0.031	-0.022	-0.022	-0.488	-0.804	0.070	-0.337	-0.064	-0.063
	(0.044)	(0.017)	(0.017)	(1.753)	(1.613)	(1.561)	(1.526)	(0.050)	(0.049)
Burning conditionality	-0.059	0.055	0.046	0.399	-0.265	-1.222	-1.389	0.122	0.106
	(0.062)	(0.037)	(0.039)	(3.063)	(3.100)	(3.055)	(2.992)	(0.093)	(0.094)
Landholdings (HAs)	-0.004	-0.005	-0.005	-0.162	-0.291	0.040	-0.148	0.009	0.007
	(0.004)	(0.003)	(0.004)	(0.283)	(0.285)	(0.295)	(0.293)	(0.008)	(0.008)
Age is under 30 years	-0.027	0.047	0.055*	1.563	1.571	-0.256	-0.458	-0.020	-0.009
	(0.049)	(0.029)	(0.029)	(2.376)	(2.307)	(2.590)	(2.506)	(0.057)	(0.056)
Age is over 50 years	0.074	0.088***	0.089***	2.400	2.253	-0.116	-0.135	0.055	0.057
	(0.048)	(0.032)	(0.033)	(2.213)	(2.158)	(2.691)	(2.631)	(0.054)	(0.053)
Highest level of education	0.015	-0.006	-0.006	-0.453	-0.424	-0.294	-0.246	-0.004	-0.004
	(0.009)	(0.008)	(0.009)	(0.504)	(0.509)	(0.585)	(0.572)	(0.012)	(0.013)
Belongs to sched./OB caste	-0.162**	-0.098*	-0.107*	-4.550**	-7.269***	-2.254	-4.692**	-0.179*	-0.226**
	(0.063)	(0.050)	(0.054)	(2.307)	(2.031)	(2.036)	(2.169)	(0.100)	(0.103)
Perception of yield variability	0.015*	-0.006	-0.003	-0.881	-0.627	-0.750	-0.534	-0.009	-0.005
	(0.008)	(0.007)	(0.007)	(0.536)	(0.544)	(0.570)	(0.564)	(0.013)	(0.014)
Household size	0.014**	-0.001	-0.001	-0.051	0.091	0.060	0.206	0.005	0.006
	(0.006)	(0.005)	(0.005)	(0.360)	(0.352)	(0.415)	(0.405)	(0.010)	(0.010)
Takes pictures on phone often/very often	-0.034	0.030	0.034	0.610	0.710	0.143	0.230	0.098	0.092
	(0.049)	(0.028)	(0.029)	(2.631)	(2.528)	(2.818)	(2.740)	(0.065)	(0.065)
Has network signal often/very often	-0.016	-0.039	-0.036	-1.273	-0.830	0.129	0.471	-0.045	-0.049
	(0.050)	(0.024)	(0.026)	(3.411)	(3.486)	(3.489)	(3.607)	(0.074)	(0.077)
Ever used LLL	0.009	0.021	0.008	0.738	1.243	-0.312	0.635	-0.064	-0.052
	(0.038)	(0.035)	(0.031)	(2.611)	(2.529)	(2.898)	(2.849)	(0.059)	(0.056)
Wheat yield Rabi 2015/16	-0.019*	-0.005	-0.007	0.559	0.367	0.802	0.678	0.016	0.012
	(0.010)	(0.006)	(0.006)	(0.549)	(0.555)	(0.509)	(0.518)	(0.013)	(0.013)
Share of income from crops	0.234**	0.075	0.086	1.869	3.088	-0.625	0.389	-0.191	-0.150
	(0.103)	(0.093)	(0.098)	(7.274)	(7.560)	(8.276)	(8.714)	(0.137)	(0.140)
Share of crop income from wheat	0.075	0.014	0.016	-4.169	-0.952	-8.411	-5.073	-0.050	0.014
	(0.153)	(0.106)	(0.109)	(7.468)	(7.320)	(7.749)	(7.734)	(0.254)	(0.248)
Fraction of land planned to be sowed with wheat	0.234	-0.006	0.002	-22.150	-20.844	-21.045	-20.110	-0.200	-0.205
	(0.183)	(0.131)	(0.136)	(16.571)	(17.427)	(15.379)	(16.324)	(0.276)	(0.293)
Distance from plot to home (minutes)			0.000		0.114***		0.104**		0.003***
			(0.001)		(0.039)		(0.047)		(0.001)
Owns insured plot			-0.018		-11.671**		-10.664*		-0.215**
			(0.070)		(5.038)		(5.648)		(0.085)
Constant	0.462*	1.037***	1.052***	28.684*	37.966**	26.346*	34.568*	0.364	0.555
	(0.242)	(0.163)	(0.178)	(14.935)	(16.521)	(15.697)	(17.327)	(0.357)	(0.361)
Observations	715	461	450	461	450	403	394	461	450
R-squared	0.107	0.124	0.127			0.153	0.174	0.144	0.170

Robust standard errors, clustered at the village level, in parentheses. Weather station fixed effects are included as controls but not reported. Significance levels: *** p<0.01, ** p<0.05, * p<0.1.

All in all, the results indicate that certain common characteristics such as caste and exposure to shocks found to affect participation and risk-management behavior in other contexts are important for PBI adoption. However, other characteristics expected to be related to the frequency or propensity of picture-taking, such as farmer's education or level of experience with smartphones, are statistically insignificant in these specifications. In other words, the concerns that we had a priori about farmers' unwillingness to engage with an innovative product through a relative unfamiliar technology seem to be unfounded in this context.

4.2 Do Pictures Capture Damage Events? Can Damages Be Quantified Accurately?

A second prerequisite for PBI to be feasible is that damage arising from different types of hazards is indeed visible at plain sight in a smartphone picture. We thereby focus on overview pictures of insured plots, taken with enough distance such that a large fraction of the plot as well as structures in the background are visible. Close-up pictures could become subject to tampering too easily.

Initial conversations with local wheat agronomists indicated that pictures would be able to capture most —though not all— hazards. Certain events such as lodging (the bending of the wheat plant due to winds and wet, loose soil), hail, or certain common wheat diseases such as yellow rust would indeed be visible. Other events, such as blight or high temperatures late in the growing season, which can affect grain filling without showing up in the external aspect of the plant, would be much more difficult to identify.

Importantly, farmer perceptions seem to agree with experts' knowledge. Panel A of Figure 3 shows farmer answers from the endline survey to the question of whether pictures can capture damages to Rabi wheat from different hazards. The majority of farmers (more than 80%) believe that damage caused by lodging, hail, and excess rainfall can be "very well" or "fairly well" captured from direct visual inspection of a time-series of pictures. Farmers recognize, however, that other events such as high temperatures or pests and diseases may be harder to recognize in this way. To have a sense of the proportion of total risk that is caused by these hazards leading to non-visible damage, Panel B shows the average degree of farmer concern about these hazards (as elicited during the baseline from asking farmers to place tokens on a board according to the extent to which each hazard worried them) and the average rate of occurrence of the hazard during the Rabi 2016/17 season (self-reported by farmers during the endline survey). In both cases, hazards that are regarded as the most visible through smartphone pictures are also the ones that worry farmers the most and those that occur—at least during the study season—more often. These factors indicate that PBI is very well suited for minimizing basis risk, at least in the case of wheat.

Figure 3. Visibility and concern about different hazards

Panel A. Farmers' perceptions about visibility of hazards in pictures

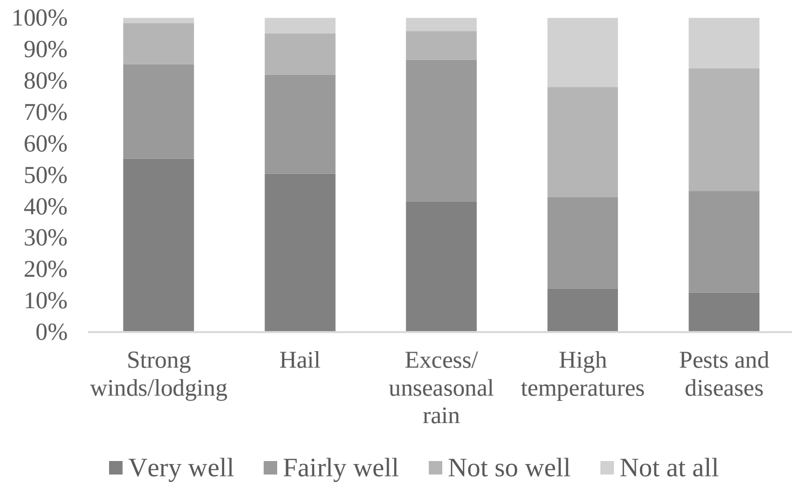

Panel B. Farmers' concern about and last season occurrence rate of hazards

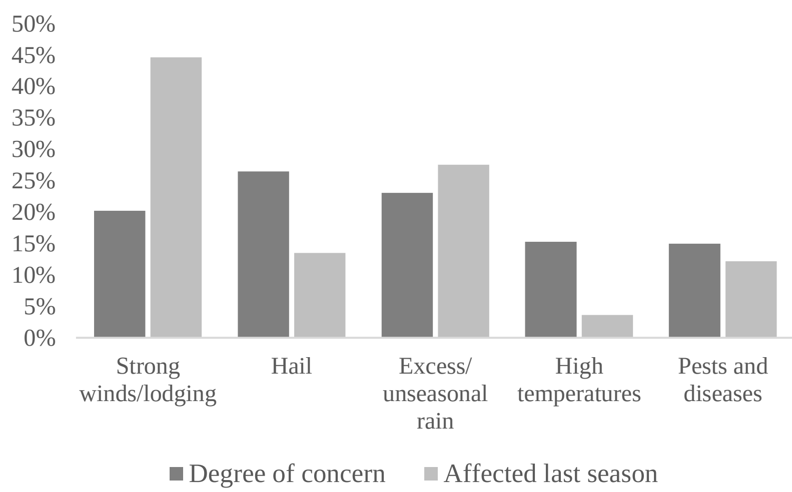

Figure 4 shows a box plot of the expert loss assessments for total damage (i.e. due and not due to mismanagement), ordered by the median assessment within a site (i.e. the assessment that was used for insurance payouts in case of PBI coverage). A few interesting patterns can be seen from the figure. The level of agreement between experts is quite high for low levels damage (under 20 percent), except for sites where the median damage is zero, which shows a few outliers (although most outliers fall below the insurance trigger value of 20 percent, meaning that experts did not reach different conclusions in terms of insurance payouts). For sites with higher visible damage, the degree of disagreement about the exact level of damage is naturally higher. Nonetheless, most experts agree over the approximate region in which the damage falls, and stark outliers are rare. We interpret this consistency across loss assessments as an indication that the wheat experts were able to identify crop losses from direct visual inspection of pictures.

Figure 4. Individual expert loss assessments

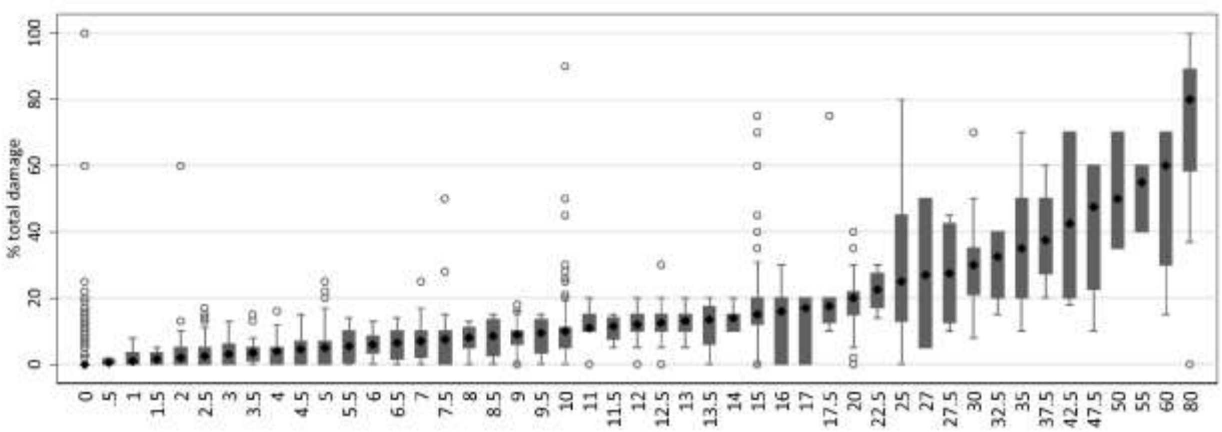

An important question, of course, is whether the quantified damage corresponds with the actual damage present in the crop. For this, we rely on the crop-cutting exercises (CCEs) carried out at the end of the 2016/17 rabi season. Figure 5 shows a scatterplot between the yields from the CCEs and the final expert loss assessment for that field (i.e. joint expert consensus for sites with a lot of disagreement and median assessment for the rest). There is a clear negative relationship between the damage estimated by the experts and CCE yields; that is, experts are generally able to identify damage when damage exists. This negative relationship is mainly driven by the extreme values of yields below 10 quintals per acre, which are the plots where farmers incurred the level of substantial damage that an insurance scheme would want to cover. For higher levels of yield, where damage perhaps was not so extreme that insurance payouts should have been made, this relationship is, however, far from perfect, with some cases where the experts assessed substantial losses that were not reflected in yields and other cases where, despite having found significantly lower yields during the CCEs, the experts were not able to pick it up from the pictures. This is particularly marked at sites with a final assessed damage of zero, which, as discussed above, include some sites with a large degree of disagreement. Further analysis and a closer look to these issues in future seasons is warranted in this regard. Overall, though, the experts were successful at identifying the farmers with especially severe damage, who will have needed insurance payouts the most.

Figure 5. Yields from crop-cutting experiments (CCEs) and expert loss assessments

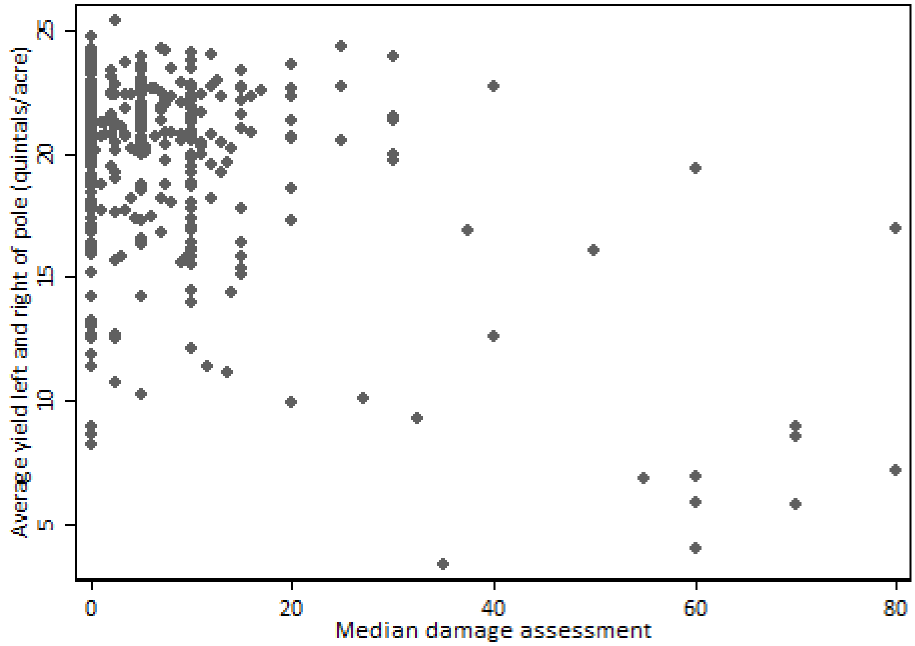

4.3 Does PBI Offer Better Protection than Other Insurance Products?

The results so far are indicative of the potential of basing insurance payouts on a time-series of pictures taken through inexpensive smartphones. Implementing such a system for taking pictures is, however, more costly than relying solely on weather indices to proxy crop damage, both in terms of time and effort spent by the farmers and in terms of resources spent by the project in sustaining the necessary data management system, monitoring, and loss assessment necessary to provide payouts after relevant loss events. Still, such a system is without a doubt less costly than insuring these small rural farmers through more traditional indemnity-based products, which require lengthy and expensive loss verification procedures before submitting payouts.

In this subsection, we analyze whether the coverage provided by PBI is sufficiently more comprehensive than the ones provided by alternative index products to assess whether the additional costs of such a system are justified. Doing so, three types of index insurance are worth noting: (i) weather-index, which bases its payouts on weather variables recorded at nearby weather stations; (ii) satellite-based, which relies on satellite imagery of the area under consideration; and (iii) area-yield, which pays out according to the average yield estimated from a limited number of crop-cutting experiments in a given geographic area (such as a village or block in the case of India).

Recall that, as part of this project, a weather-index based product was designed (WBI) that covered the insured farmer from higher-than-normal night temperatures and (unseasonal) excess rainfall at the end of the wheat season, from February to April. Figure 6 presents a comparison between the average yields from CCEs across farmers that would and would have not received payouts from, alternatively, the PBI and WBI products. The contrast is striking. Panel A shows that, in the case of PBI, average yields amongst farmers for whom the experts' damage assessments indicated losses of 20 to 50% and 50 to 100% were, respectively, 18.2 and 10 quintals per acre, compared to an average of 20 quintals per acre for farmers for whom PBI would not have triggered. In other words, PBI would have on average correctly provided payouts to farmers who suffered a considerable loss in yields of 20% or more. Moreover, the loss estimate from visual inspection by the experts was considerably accurate of the actual average losses as determined through objective crop-cut experiments.

Figure 6. Crop-cutting yields and WBI and PBI payout categories

Panel A. Average yields at different PBI payout categories

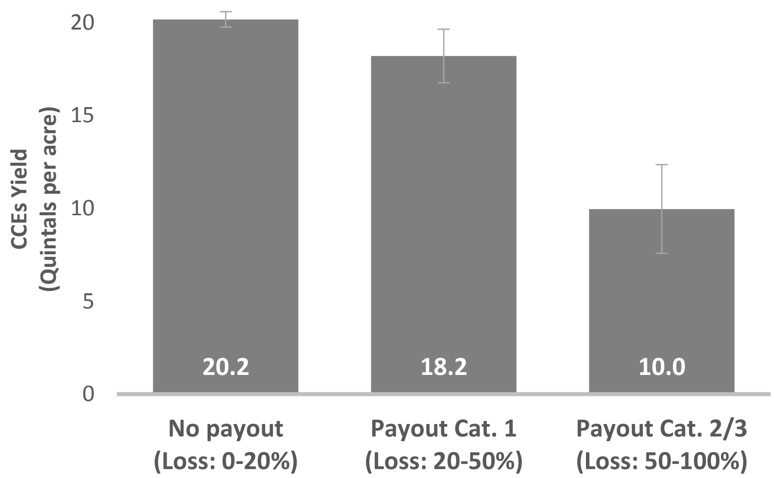

Panel B. Average yields at different WBI payout categories

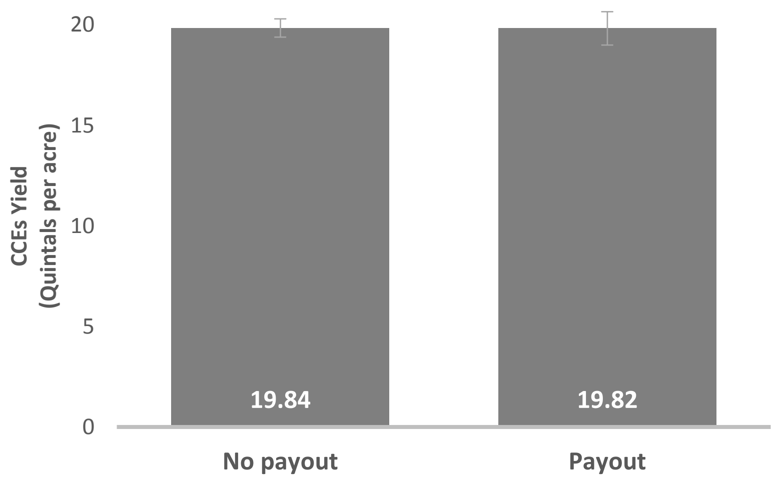

In the case of WBI, in contrast, the yields from farmers for whom WBI triggered a payout are virtually indistinguishable from those for whom no payout was triggered. This is indicative of a very large degree of overall basis risk in weather-index based insurance, where actual losses do not correspond with insurance payouts. To be accurate, this result is only applicable to the specific weather-index based product implemented in this project, which was however carefully designed to reflect the perceptions of both farmers and expert wheat agronomists in the study region. Moreover, as seen in Table 2, most WBI payouts were very low, with the maximum payout at just under 30% of the insured amount, indicative of the absence of a severe weather shock affecting wheat. This leaves room for the possibility of the weather-index product performing better after more extreme weather shocks. Nevertheless, in the light of the limited risk coverage of WBI in relation to the more comprehensive PBI coverage illustrated in Figure 3, the preliminary evidence seems to consistently point in the direction of reduced basis risk in PBI.

An alternative set of products to weather-index based insurance, which have been gaining popularity in the past decade, are the ones based on satellite imagery (Stanimirova et al., 2013). An increasing number of insurance projects in the developing world rely on this system of loss identification, and the relatively recent flourishing of very high temporal and spatial resolution micro-satellite products seems to have pushed this trend even further and led to some proponents to argue that this technology alone can tackle most of the drawbacks seen in weather-index based products.

Promising as this technology may be, satellite-based insurance products nevertheless suffer from many disadvantages. For instance, cloud cover may not allow for an image to be captured to estimate vegetation indices, which can get particularly problematic in rainy seasons when no images may become available over several consecutive weeks. Moreover, estimation of yields through proxies such as vegetation or weather indices are subject to considerable measurement error, introducing basis risk into any insurance product that relies on these. In addition, crop yields depend on the complex interplay of several weather factors and other elements, an aspect which cannot be captured through proxy indices. In terms of spatial resolution, publicly available satellites such as MODIS or HLS have resolutions that are too coarse for yield estimation at the (typically-small) size of plots ubiquitous in developing regions, and higher-resolution data available from micro-satellites is generally quite costly and imposes large storage and computational burdens which add up to the costs even further. Finally, even if in the near future insurance products from micro-satellites may become more attainable through increased competition and the further development of more efficient processing and storage technologies, for these data to be useful to small farmers at a large scale, detailed geo-referenced cadasters or other ways of delineating the insured plots would be needed in order to estimate individual yields.

Despite requiring a much larger involvement by the farmer (which, as mentioned, can however help to increase a farmer’s ownership in the product and thus his demand for it), PBI can tackle many of the issues above. By placing ‘eyes on the ground’, smartphone pictures can derive many similar indices to those used in satellite products and can contribute with a wealth of additional information only visible at ground-level, such as the standing of the crop or the presence of specific pests, diseases, or other subtle features indicating damage by hail or suboptimal temperatures. In addition, advanced machine learning techniques can exploit all these features and potentially provide a much more accurate estimate of yield loss.[5] In sum, as much as PBI does not constitute a substitute for satellite insurance products, it does provide an attractive alternative in those cases best suited to its own strengths, and, more generally, can stand as a valuable complement to the increasingly popular remote sensing products.

Lastly, a common approach for determining crop losses, area-yield index insurance (AYI), relies on the estimation of the average yield across a broader geographic area for individual payouts. In order to compare the AYI and PBI approaches, we combine the yield data from the crop-cutting experiments together with the expert loss assessment data and simulate the proportion of farmers who would have received payouts under different types of products: AYI only (triggering when average yields measured for a random sample of farmers in a given area drop more than 20% below normal yields); PBI only (triggering when visible damage is either more than 20%, in case of a 'lenient policy', or 50%, in case of a 'strict policy', as assessed by subject matter experts); and a product that combines AYI and PBI. Table 3 presents average proportions of farmers receiving payouts and standard deviations from 10,000 simulations where, depending on the product type, each simulation randomly selects four farmers within a cluster of two nearby villages (weather station level) or a cluster of all villages within a district to determine the area-yield index for that cluster.

[5] Interestingly, some algorithms have been developed to count, for instance, the number of grains in wheat or the number of fruits in trees, coming much closer to a direct estimation of yields than proxy methods.

Table 3: Simulations of area-yield index and picture-based insurance performance

	Probability of receiving payout							
	All farmers		Has >=50% loss		Has 20-50% loss		Has < 20% loss	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
A. Area-yield index (AYI)								
Weather station level	0.079	0.034	0.510	0.298	0.202	0.080	0.050	0.025
District level	0.069	0.106	0.247	0.409	0.097	0.149	0.059	0.090
B. Picture-based insurance (PBI)								
- Lenient policy (pay if >= 20% loss)	0.097	0.009	0.750	0.079	0.073	0.027	0.071	0.009
- Strict policy (pay if >= 50% loss)	0.029	0.005	0.599	0.087	0.000	0.000	0.007	0.003
C. AYI + PBI: Lenient policy (>= 20%)								
Weather station level	0.150	0.021	0.794	0.080	0.255	0.067	0.112	0.023
District level	0.151	0.083	0.771	0.094	0.150	0.137	0.122	0.079
D. AYI + PBI: Strict policy (>= 50%)								
Weather station level	0.091	0.025	0.734	0.106	0.202	0.080	0.052	0.023
District level	0.091	0.094	0.666	0.157	0.097	0.149	0.064	0.087
Number of observations in total*	357		14		33		310	
Number of weather stations	25		5		17		25	
Number of districts	6		2		6		6	

Notes: Mean and standard deviation based on a Monte Carlo simulation with 10,000 replications and 4 CCEs per geographical unit (weather station level or district level). We are not simulating area-yield indices at the village level due to a limited number of observations in villages. *Observations that are randomly selected for inclusion in the CCEs in a simulation are dropped from the payout analyses for that simulation in order to avoid mechanical correlations between the CCE yields and insurance payouts, which we would not avoid to occur in the actual implementation given that for one village with more than 100 farmers there are typically 4 CCEs.

The first two columns show, under each alternative product, the probability of receiving a payout including all farmers. Panel A shows that the AYI product would have triggered for, on average, 7.9 percent of farmers if yields were measured in a cluster of nearby villages (weather station level), and for 6.9 percent of farmers if measured at the district level. To also assess the degree of basis risk, the remaining columns distinguish between farmers with different levels of actual damage as measured through CCEs. As would be expected, measuring yields for a cluster of nearby villages—although costlier and logistically more cumbersome—minimizes basis risk compared to measuring yields at the district level due to spatial correlation in yields. At the weather station level, the area-yield index identifies about half of all farmers with severe damage (50% or more), and about one fifth of all farmers with moderate damage (20-50%), which is twice the proportion identified through the district-level index.

Panel B in Table 3 shows that PBI suffers from both advantages and disadvantages compared with the AYI product used in the simulations. On one hand, without triggering significantly more often on average, the lenient PBI policy triggers significantly more often than AYI for farmers with severe damage; and the stricter PBI policy triggers at least as often as AYI, but at significantly lower cost, as it rarely triggers for other farmers. On the other hand, while outperforming AYI in identifying farmers with severe damage, PBI does not help distinguish farmers with moderate damage from farmers with less or zero damage.

In reality, PBI does not need to be compared against potential alternatives, but rather as a complement when offered in combination with existing alternatives. Smallholder farmers can benefit from an ecosystem of insurance products available for them, that can cater to their individual preferences and characteristics and that can best tackle the nature of production risks in a given geographic area and for a given crop. In this regard, we believe PBI has the potential to serve as a top-up component to other insurance products that can help reduce their basis risk while retaining some of the cost advantages of more traditional index schemes. This was indeed the approach followed in the project, where PBI was offered as a top-up component to a weather-index based insurance product, which helped provide a more comprehensive protection against hazards that would otherwise not be visible through smartphone pictures. A similar scenario could be applicable to satellite or AYI products, where farmer appeals to non-triggering indices could be dealt with by relying on smartphone pictures taken from the ground under the PBI protocol.

In this regard, panels C and D of Table 3 illustrate the advantages from combining an area-yield insurance product with PBI. In Panel C, combining the lenient PBI policy with AYI reduces downside basis risk compared with AYI, but also increases the proportion of farmers that receive payouts while not experiencing damage, leading to upside basis risk and higher costs of the insurance policy. Finally, the stricter policy in Panel D increases the overall proportion of farmers receiving payouts only slightly compared to the AYI products in Panel A, maintaining the weather station level AYI-PBI combination as a viable option. Combined, these findings indicate that using pictures for loss assessment in combination with AYI can substantially reduce the downside basis risk in AYI products observed in the simulations, without significant increases in costs.

5. Conclusions

Picture-Based Crop Insurance (PBI) is a new approach to improve smallholder farmers' access to affordable but high-quality insurance. By leveraging increasing smartphone ownership among smallholder farmers and relying on automated image processing techniques, the goal of PBI is to combine key advantages of index insurance—fast and inexpensive claims processing—with those of indemnity insurance—low basis risk and easy-to-understand products. To our best knowledge, the feasibility of this approach has never been evaluated systematically, and this study is a first step in that direction.

Based on a first pilot year, we find that (a) farmers are able—at large—to follow picture-taking protocols; (b) agronomists and farmers agree that the most important risks in wheat production can be visible in pictures, and expert loss assessments are indeed able to detect such damage in the pictures; and (c) picture-based insurance can provide payouts which are better correlated with yields than those from weather-index insurance and seems to offer considerable advantages to other common index products such as satellite-based or area-yield-based insurance.

In related research, we will study PBI sustainability considerations such as the extent of moral hazard observed during this first season and its dynamics over time, including product design aspects that can help to limit this moral hazard issue; analyze differences in willingness to pay for picture-based insurance versus weather-index based insurance; test to what extent there is adverse selection and how to overcome it; and further develop image processing algorithms that will allow automate loss assessment.

Importantly, this approach is not exclusively reserved to areas with sufficient smartphone penetration. An equivalent insurance model could be achieved by relying on village representatives who could be provided with an inexpensive Android smartphone (when one is not already available) and requested to visit every insured plot a few times a week in order to capture the corresponding repeat picture. This representative could also serve as distribution channel and as a key link with the insurance company, in exchange for a commission on premiums.

All in all, PBI offers a promising alternative to existing insurance products for poor farmers, with the potential to bring about important changes in the way that insurance is offered to smallholders in rural areas of the developing world.

References

Barrett, C., and J. McPeak, 2006. "Poverty Traps and Safety Nets." In *Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke,* edited by A. de Janvry and R. Kanbur, 131–154. New York: Springer.

Berhane, G., S. Dercon, R.V. Hill, and A. Taffesse. 2015. "Formal and Informal Insurance: Experimental Evidence from Ethiopia." Paper presented at International Conference of Agricultural Economists, Milan, August 8–14.

Bhushan, C., and V. Kumar, 2017. "Pradhan Mantri Fasal Bima Yojana: An Assessment." Centre for Science and Environment, New Delhi.

Cai, J., 2013. "The Impact of Insurance Provision on Households' Production and Financial Decisions," MPRA Paper 46864, University Library of Munich, Germany.

Cai, H., Y. Chen, H. Fang, L. Zhou, 2009. "Microinsurance, Trust and Economic Development: Evidence from a Randomized Natural Field Experiment," NBER Working Papers 15396.

Carter, M.R., C.B. Barrett, S. Boucher, S. Chantarat, F. Galarza, J.G. McPeak, A.G. Mude, and C. Trivelli, 2008. "Insuring the Never Before Insured: Explaining Index Insurance through Financial Education Games." BASIS Briefs, University of Wisconsin, Madison.

Carter, M.R., Cheng, L. and Sarris, A., 2016. "Where and how index insurance can boost the adoption of improved agricultural technologies." *Journal of Development Economics* 118(1): 59-71.

Ceballos, F. and B. Kramer, 2018. "Big brother might be watching you: Testing for moral hazard in imagery-based crop insurance." Mimeo.

Chantarat, S., A.G. Mude, C.B. Barrett, and M.R. Carter, 2013. "Designing Index-Based Livestock Insurance for Managing Asset Risk in Northern Kenya." *Journal of Risk and Insurance* 80(1): 205-237.

Cole, S.A., X. Giné, J. Tobacman, P.B. Topalova, R.M. Townsend, and J.I. Vickery, 2013. "Barriers to Household Risk Management: Evidence from India." *American Economic Journal: Applied Economics* 5 (1): 104–135.

Cole, S.A., X. Giné, J.I. Vickery, 2017. "How Does Risk Management Influence Production Decisions? Evidence from a Field Experiment." *The Review of Financial Studies* 30(6): 1935–1970.

de Janvry, A., V. Dequiedt, and E. Sadoulet. 2014. "The Demand for Insurance against Common Shocks." *Journal of Development Economics* 106: 227–238

Dercon, S., R.V. Hill, D. Clarke, I. Outes-Leon, and A. S. Taffesse. 2014. "Offering Rainfall Insurance to Informal Insurance Groups: Evidence from a Field Experiment in Ethiopia." *Journal of Development Economics* 106: 132–143.

Dercon, S., and J. Hoddinott, 2004. "Health, Shocks, and Poverty Persistence." In *Insurance against Poverty,* edited by S. Dercon, 123–136. Oxford, UK: Oxford University Press;

Elabed, G., M.F. Bellemare, M.R. Carter, and C. Guirkinger, 2013. "Managing Basis Risk with Multi-Scale Index Insurance." *Agricultural Economics* 44: 419–431.

Hazell, P., C. Pomareda, and A. Valdes. 1986. *Crop Insurance for Agricultural Development: Issues and Experience.* Baltimore: Johns Hopkins University Press.

Hill, R.V., L.M. Robles, and F. Ceballos. 2016. "Demand for a Simple Weather Insurance Product in India: Theory and Evidence." *American Journal of Agricultural Economics* 98 (4): 1250–1270.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry, 2014. "Agricultural Decisions after Relaxing Credit and Risk Constraints." *The Quarterly Journal of Economics 129*(2): 597-652.

Matul, M., A. Dalal, O. De Bock, and W. Gelade. 2013. "Why People Do Not Buy Microinsurance and What We Can Do about It." Briefing Note 17. Geneva: Microinsurance Innovation Facility.

Mobarak, A.M., and M. Rosenzweig, 2012. "Selling Formal Insurance to the Informally Insured." Working Paper No. 97, Economic Growth Center Discussion Paper No. 1007, Yale University.

Richardson, A.D., K. Hufkens, T. Milliman, D.M. Aubrecht, M. Chen, J.M. Gray, M.R. Johnston, T.F. Keenan, S.T. Klosterman, M. Kosmala, E.K. Melaas, M.A. Friedl, and S. Frolking, 2017. "Tracking Vegetation Phenology across Diverse North American Biomes using PhenoCam Imagery." Scientific Data.

Stanimirova, R., H. Greatrex, R. Diro, G. McCarney, J. Sharoff, B. Mann, A.L. D'Agostino, M. Rogers-Martinez, S. Blakeley, C. Small, and P. Ceccato, 2013. "Using Satellites to Make Index Insurance Scalable." Final IRI Report to the ILO Micro-Insurance Innovation Facility.