

Ceballos, Francisco; Kramer, Berber; Robles, Miguel

Article

The feasibility of picture-based insurance (PBI): Smartphone pictures for affordable crop insurance

Development Engineering

Provided in Cooperation with:

Elsevier

Suggested Citation: Ceballos, Francisco; Kramer, Berber; Robles, Miguel (2019) : The feasibility of picture-based insurance (PBI): Smartphone pictures for affordable crop insurance, Development Engineering, ISSN 2352-7285, Elsevier, Amsterdam, Vol. 4, pp. 1-14, <https://doi.org/10.1016/j.deveng.2019.100042>

This Version is available at:

<https://hdl.handle.net/10419/242300>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>



The feasibility of picture-based insurance (PBI): Smartphone pictures for affordable crop insurance

Francisco Ceballos^{a,*}, Berber Kramer^a, Miguel Robles^b

^a Markets, Trade and Institutions Division, International Food Policy Research Institute (IFPRI), 1201 Eye Street, NW, Washington, DC, 20006, USA

^b Office of the Executive Director - EDS08, World Bank, Washington, DC, USA

ARTICLE INFO

Keywords:

Risk and insurance

Mobile technology

Basis risk

IndiaJEL classification:

G220

O13

O16

Q14

ABSTRACT

Smallholder farmers are increasingly exposed to weather extremes but lack access to affordable insurance products for catastrophic crop damage. This paper analyzes the feasibility of Picture-Based Insurance (PBI) as a low-cost tool to improve coverage. PBI verifies insurance claims using smartphone pictures of insured plots, taken by farmers themselves, to minimize asymmetric information and costs of claims verification, while reducing basis risk compared to index-based insurance. A pilot implementation in the rice-wheat belt of India speaks to PBI being a feasible and valuable innovation to reduce downside basis risk in index insurance: nearly two-thirds of trained farmers took at least four pictures (roughly one per growth stage), which was considered sufficient for loss assessment; severe damage was visible from smartphone pictures in 71 percent of affected sites; and this was a significant improvement over alternative index-based products, which identified severe damage in at most 34 percent of affected sites.

1. Introduction

Climate change is increasingly exposing smallholder farmers to natural hazards such as drought, heat, excess rainfall, hail, and pests and diseases (Porter et al., 2014), while the supply of reliable indemnity insurance coverage against weather extremes remains limited. Indemnity insurance premiums are unaffordable because the amounts that smallholder farmers seek to insure are small relative to the transaction costs associated with providing insurance and because asymmetric information between farmers and insurance providers can lead to adverse selection and moral hazard (Hazell et al., 1986). This leaves smallholder farmers' livelihoods vulnerable to extreme weather shocks, limiting risk averse farmers' ability to invest in productivity-enhancing technologies and hampering investments in the production of profitable crops (Barrett and McPeak, 2006; Cai, 2013; Cai et al., 2009; Cole et al., 2017; Dercon and Hoddinott, 2004; Karlan et al., 2014; Mobarak and Rosenzweig, 2012).

In the past few decades, various index-based insurance products have been piloted as a potential solution to the high transaction costs and information asymmetry problems that challenge indemnity

insurance. Index-based insurance pays out according to a pre-determined index, which proxies for losses resulting from weather and other catastrophic events. By determining insurance payouts through an objective index such as the amount of rainfall or the average temperature, insurance providers eliminate asymmetric information and do not need to send claim adjusters to assess damage on individual fields, reducing the cost of claim verification and the time until claim settlement. Yet, demand for such index-based insurance products has been low, even when offered at subsidized premiums, due to limited trust in insurance providers, a lack of understanding of these insurance products, and high levels of basis risk, meaning that indices and associated insurance payouts often do not correlate well with plot-level damage (Cole et al., 2013; Hill et al., 2016; Matul et al., 2013; Mobarak and Rosenzweig, 2012).

This paper describes and analyzes the feasibility of a new approach to overcome these challenges: Picture-Based Insurance (PBI). PBI provides insurance coverage for damage detected from a time-lapse of the insured crop, built from both pre- and post-damage georeferenced pictures that farmers take themselves using regular, low-cost smartphones. This approach allows farmers to reliably document losses from

This work was undertaken as part of the CGIAR Research Program on Policies, Institutions, and Markets (PIM) led by the International Food Policy Research Institute (IFPRI). Funding support for this study was provided by the CGIAR Research Programs on Policies, Institutions, and Markets (PIM) and Climate Change, Agriculture, and Food Security (CCAFS), the CGIAR Platform for Big Data in Agriculture, and the International Initiative for Impact Evaluation (3ie). The opinions expressed here belong to the authors, and do not necessarily reflect those of PIM, CCAFS, 3ie, IFPRI, or CGIAR.

* Corresponding author.

E-mail address: F.Ceballos@cgiar.org (F. Ceballos).

<https://doi.org/10.1016/j.deveng.2019.100042>

Received 13 July 2018; Received in revised form 16 April 2019; Accepted 23 May 2019

Available online 25 May 2019

2352-7285/ © 2019 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

natural calamities, while providing evidence that the affected crop was managed appropriately, helping reduce information asymmetries and costs of plot-level loss verification that have challenged traditional indemnity insurance. By being participatory and tangible, and by delivering plot-level assessments of damage, PBI has the potential to reduce basis risk and improve trust and understanding, key challenges for index-based insurance.¹ As such, PBI is designed to combine key advantages of both index-based insurance—timely compensation without expensive loss assessments—and indemnity insurance—minimum basis risk and a tangible product.

PBI works as follows. A farmer installs an application into his smartphone, takes an initial picture of the site he wishes to enroll, and documents how the crop develops in this site by taking repeat pictures throughout the growing season. The app facilitates this task through geotags (warning the farmer if taking the repeat picture at a different location than the initial picture) and visual aids that ensure an almost-identical view frame throughout the season (displaying a mildly-transparent version of the initial picture that allows the farmer to align each repeat picture to permanent background features in the landscape). All pictures are automatically uploaded to a server, with the farmer unable to manipulate them in the phone. At the end of the season, experts evaluate the time-lapse of the plot and estimate a percentage of crop damage. The insurance company uses these loss assessments to verify claims that trigger payouts to the farmer.

PBI is a novel concept that has not been tested before in a systematic way. In this paper, we therefore address key knowledge gaps around the technical feasibility of PBI, both in terms of farmers complying with the PBI protocol and of pictures serving as a medium to identify crop damage. Given that the product was provided for free during this initial season, we are unable to inform questions about the demand for PBI were it to be offered commercially. Our formative evaluation of PBI, targeting 750 smallholder wheat farmers in Haryana and Punjab—two states in northwest India—, aims to answer three research questions: (i) to what extent do farmers comply with a PBI protocol that requires them to regularly upload pictures of their plots, and how is this linked to traditional determinants of technology acceptance such as age, education, and caste; (ii) to what extent is damage visible in smartphone camera data, that is, can experts identify a reduction in productivity from smallholders' pictures of insured crops; and (iii) does PBI reduce downside basis risk compared with conventional index-based insurance products?

Overall, the results speak to PBI being a feasible and valuable option to complement existing index-based insurance. Compliance was relatively good: close to two-thirds (63 percent) of trained farmers uploaded at least four repeat pictures throughout the season (roughly one picture per growth stage and considered sufficient to assess losses). Using objective site-level yield measurements, we compare basis risk across different insurance products—a crucial yet unusual undertaking in the index insurance literature, with [Jensen et al. \(2016\)](#) and [Jensen et al. \(2018\)](#) being two notable exceptions (albeit focused on livestock). PBI was unable to identify sites with moderate degrees of damage, but did reduce severe cases of downside basis risk: picture-based loss assessments identified 71.4 percent of sites with severe crop losses, outperforming both weather index-based insurance (which could not cover the perils that were causing severe losses) and a simulated area-yield index (which identified on average only 34.4 percent of sites with severe crop losses). We conclude that PBI offers a promising option to reduce downside basis risk in insurance products for smallholder

farmers.

This innovation comes in a timely manner, as it takes advantage of increased smartphone ownership and penetration of low-cost mobile internet services among smallholder farmers. PBI further builds upon recent advances in image processing; particularly on applications that use digital repeat photography for near-surface remote sensing (the PhenoCam project, [Richardson, et al., 2017](#)). The approach thereby relates to the use of high temporal- and spatial-resolution satellite imagery to identify losses at finer scales for smallholder farmers ([Stanimirova et al., 2013](#); [Chantararat et al., 2013](#); [Carter et al., 2008](#)). Satellite-based damage estimation is subject to shortcomings, and PBI can stand as a valuable tool to complement existing products.² Moreover, PBI can help adding an extra layer of protection on top of a traditional index product, in the form of fail-safe indices or gap insurance, which allow for audit-based payouts if an index does not trigger in an area where a sufficient proportion of farmers claim to have suffered losses ([Berhane et al., 2015](#); [Flatnes and Carter, 2015](#)). By documenting pre-damage and pre-audit crop conditions, PBI can help operationalize gap insurance and reduce basis risk.

2. Context and procedures

In this section, we describe the study context and procedures, focusing first on the study region and sampling procedures, followed by a description of the insurance products that were tested as part of the study, including the PBI product and a weather index-based product that was used for comparison purposes. The final part of this section describes the study procedures.

2.1. Study context and sampling

The study was conducted for wheat grown during the Rabi (winter) season in the states of Haryana and Punjab. These states are the second and third largest wheat-producing states in India and play a critical role in India's food grain supply.³ Although yields in these two states have traditionally been among the highest in the country, and although most farmers have access to irrigation, wheat yields have stagnated, and are increasingly exposed to extreme weather events including excess rains and heatwaves due to climate change. The formative evaluation targeted this region in part due to this increasing exposure to weather risks, and in part due to near-universal ownership of smartphones among farmers, with smartphone penetration still gaining momentum in other parts of India.

We targeted six districts (three from Punjab, two from the west of Haryana and one from the northeast of Haryana) for which the underwriter of the insurance products, HDFC Ergo General Insurance

² Loss estimation at the field level can be very expensive (since sufficiently high-resolution satellite data tends to be provided by private companies) and faces limitations, such as cloud cover, large computational storage and processing costs, and poor availability of georeferenced cadasters to accurately identify insured plots. Moreover, in most settings, basis risk is still present, due to, for instance, measurement error or intercropping practices. PBI can tackle many of the issues above. By placing 'eyes on the ground', smartphone pictures can provide a wealth of additional information visible only at ground-level, such as the standing of the crop or the presence of specific pests, diseases, and other subtle features indicating damage by hail or suboptimal temperatures ([Hufkens et al., 2019](#)).

³ Wheat was selected because nearly all farmers in the selected districts grow this crop, whereas there is more heterogeneity in the production of other crops. Moreover, because agriculture is largely irrigated in Haryana and Punjab, the main risk for other crops—a drought or late onset of the monsoon—does not affect production as much as in states with rain-fed agriculture. Wheat is, however, considered a relatively safe crop in these states, which may reduce the need for insurance and thus reduce both take-up and compliance with the picture-taking protocol.

¹ PBI can reduce basis risk by assessing losses at the plot level as opposed to a distant weather station (reducing spatial basis risk); by covering visible damage, including some pests and diseases, and lodging, as opposed to only weather-related events (reducing design basis risk); and by following farmer-specific timing in terms of planting time and the individual seed's crop cycle instead of the average risk profile for the crop (reducing temporal basis risk).

Limited, could source rainfall and temperature data from weather stations, to allow for comparisons of PBI and weather index-based insurance. From a list of available weather stations, we randomly selected 25 stations stratified by district, with the number of weather stations per district (ranging from three to eight) being proportional to district size. Subsequently, we randomly selected two rural villages within a radius of 5 km from each weather station (to limit geographical or 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).⁴ This resulted in a study sample of 50 villages.

In each village, we listed 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 for a baseline survey. These farmers were equally distributed across three categories: operating less than five acres, five to ten acres, and ten to fifteen acres of farmland. In doing so, we oversampled relatively smaller farmers, instead of constructing a sample of farmers representative for our study area, motivated by external validity considerations: representative farmer populations in other states of India typically have smaller landholdings than farmers in Haryana and Punjab. In addition, as a smartphone picture can capture only a limited area, the PBI approach appears more relevant for farmers with smaller landholdings.

2.2. Insurance products

All farmers in the baseline survey were offered insurance, free of charge, for one acre of wheat grown during the Rabi 2016/17 season (spanning November through April). In all 50 study villages, the product included a standard weather index-based insurance component (WBI), which triggered payouts in case of unseasonal rains or above-normal temperatures between February and April (around flowering and harvest time). The indices relied on daily minimum temperature and rainfall collected at the nearest weather station and were developed based on focus group discussions with farmers, key informant interviews with local wheat agronomists, and indices designed by past governments. For each index, small payouts were triggered once the index exceeded a strike value, and payouts were linearly increasing in the index until reaching an exit value. For index levels at or above the exit value, farmers would receive the total sum insured. Payments were made for either the rainfall or the temperature index, whichever triggered the highest payout. The maximum payout was 13,000 Indian Rupees (Rs.) or 200 US dollars per acre. This amount was based on the average total production costs for one acre of wheat, including labor, as determined during the initial focus group discussions.

For every weather station, we randomly selected one of the two villages—or 25 villages in total—to receive *in addition* to the standard WBI component a picture-based insurance component (PBI), providing coverage for visible damage during the Rabi season.⁵ Farmers from these 25 PBI villages were informed that to determine payouts, independent experts would inspect their pictures for visible damage due to risks beyond their control. This procedure was meant to be as transparent as possible to participating farmers, who could perceive the insurer to downplay losses, particularly given the novelty of the

product. Farmers were told that damage below 20 percent would not trigger payouts; damage between 20 and 50 percent would trigger a payout of Rs. 3900; damage between 50 and 75 percent would trigger a payout of Rs. 7800; and damage above 75 percent would trigger the maximum PBI payout of Rs. 13,000. Farmers would receive a payout for either the WBI or the PBI component, whichever triggered a higher payout. In the remaining 25 WBI villages, farmers received payouts for only the WBI component.⁶

2.3. Procedures

During July and August 2016, we conducted a baseline survey among the 15 selected farmers in each of the 50 study villages. In October 2016, we invited these farmers to village sessions, introduced as a training on agricultural risk management, in which they were informed that they would receive agricultural insurance for one acre of wheat grown during the upcoming Rabi growing season. While only farmers from the 25 PBI villages were insured for damage visible in their smartphone pictures under PBI, *all* farmers were told that their WBI coverage would be conditional on following the picture-taking protocol. Project staff enrolled these farmers by downloading a dedicated smartphone application named WheatCam, designed to implement the PBI procedures; by taking an initial picture within the app of a randomly selected field (i.e. a *site*); and by training farmers on how to take repeat pictures of this site.

The protocol requested farmers to regularly capture repeat pictures of the same portion of the enrolled site throughout the entire growing season.⁷ These pictures were ideally taken between 10am and 2pm to maintain appropriate and comparable lighting levels across all images, and pictures had to be taken from the same spot, pointing at the same direction every time. WheatCam facilitated this task through geotags, used to issue warnings if the farmer was taking a repeat picture at a different location than the initial picture, and by displaying the initial

⁶ The research project paid for the premiums. The cost of the stand-alone WBI product was Rs. 3133 (incl. taxes) in Punjab and the northeastern Haryana district, and Rs. 3149 in the two western Haryana districts. The premium was relatively high, at 24 percent of the insured sum, for two reasons. First, given that the trigger was set to the 70th percentile of historical index values, both the excess rainfall index and the above-normal temperature index were designed to trigger small payouts once every three or four years. Second, this non-subsidized premium included a relatively high loading factor of approximately 50 percent. Adding PBI coverage increased the cost of the insurance product by approximately Rs. 630 (or 20 percent) to Rs. 3760 in the western Haryana and Punjab districts and Rs. 3779 in the northeastern Haryana districts. This increase in cost reflects the increased probability of payouts for risks not covered under WBI, including damage due to lodging, hailstorms, pests and diseases, and wild animals. Further, underwriting the PBI component was relatively expensive due to a lack of historical yield data to assess risk at the individual farmer level. Additional costs associated with PBI coverage, notably the data management and loss assessments by independent experts, were not included in the insurance premium but borne by the project. At current expert consultancy rates of Rs. 35,000 per month and assuming an assessment takes about 10 min per claim, one expert would be able to verify 960 claims per month, resulting in an approximate cost of Rs. 36.50 per assessment, or—assuming only 20 percent of farmers file a claim and three experts review each case—an estimated Rs. 21.90 per policy (around 0.34 U.S. dollars); a negligible 0.17 percent of the sum insured.

⁷ Initially, we set the number of required pictures to three per week because a larger number of pictures would benefit the training of image processing algorithms to automate damage assessment. At the same time, loss assessment experts indicated that losses could be quantified from irregular, infrequent pictures showing the development of the wheat plant at a few different growth stages (to verify normal crop growth) in addition to a few pictures at the time of damage and immediately before harvest (to quantify the nature and level of the loss). Because of this, combined with technical problems in the initial roll-out of the WheatCam app, we decided to consider farmers for insurance payouts as long as they had taken at least 2 pictures throughout the season.

⁴ Weather stations with fewer than five such villages were excluded from the sampling frame.

⁵ We randomized the type of insurance product offered to farmers in order to test whether PBI affects farmer behavior. Lower input application or higher yields in the group of farmers insured under PBI would indicate the presence of moral hazard. Ceballos and Kramer (2018a) discuss these analyses in detail and find no evidence of moral hazard, possibly due to the systematic picture-taking protocol designed to reduce information asymmetries.




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

Note: This figure shows the visual aids used to ensure a fixed view frame when taking repeat pictures at a farmer's site. Panel A shows the "ghost" image, consisting of a mildly transparent version of the initial image that allowed the farmer to align static features in the landscape across pictures. Panel B shows the auxiliary pole, serving as a tripod to maintain a fixed position for the phone, and the reference pole, serving as a fixed reference point in the plot.

picture as a "ghost" image (a mildly transparent image): a visual aid allowing the farmer to align static features in the landscape (such as distant trees or structures and reference poles in the field) with those same elements in the initial picture, thus ensuring an almost-identical view frame throughout the season (panel A of Fig. 1). To further standardize the time-lapse, WheatCam applied a fixed white balance, keeping in-camera RGB ratios constant. Pictures were uploaded to a server automatically and it was not possible to upload pictures taken outside WheatCam, eliminating possibilities for editing pictures. Farmers could reach out to project staff throughout the season for troubleshooting in case they encountered any problems with the app or protocol.

All farmers who agreed to take pictures were given 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 (panel B of Fig. 1). Further, due to app compatibility issues with older Android versions in farmers' phones at the launch of the project, all farmers who had agreed to take pictures were provided with a low-cost Android smartphone and a data plan of Rs. 258 per month to upload the pictures, conditional on following the protocol.⁸

At the end of the season, six wheat experts evaluated the pictures and estimated a percentage of crop damage for each site with at least two repeat pictures. Each time-lapse was randomly assigned to three different experts. Assessments were first done by experts individually and the median assessment was used to determine insurance payouts. In case of large disagreement between individual assessments, a final damage estimate was agreed upon through consensus to determine insurance payouts. Assessments were anonymous with no access to the farmer's personal details or type of insurance coverage. For farmers with at least 20 percent estimated damage, claims were submitted to the insurance company, which issued payments directly into farmers' bank accounts, based on the experts' loss assessment. Although disbursing payouts in the study season took around four months, this was mainly due to administrative delays; without those delays, it would have been possible to disburse payouts within one month from harvest.

⁸ The poles will not be provided moving forward. Pictures of insured sites were showing enough trees and other permanent structures in the background to align repeat pictures and verify that pictures were always taken at the same location. In addition, the provision of smartphones and data plans will be discontinued in the future, as farmers often report having a data plan and that this is not a crucial implementation requirement.

3. Data

In this section, we describe the primary data sources used in the analyses below. First, we use self-reported data collected through baseline and endline surveys, conducted with all available farmers during August 2016 and April 2017, respectively. The baseline survey inquired about an array of farm and household characteristics, including plot characteristics, cultivation practices, input use and agricultural technology adoption in previous seasons, household composition, income and risk perceptions. The endline survey gathered data for the Rabi 2016/17 season, including cultivation practices, input use and self-reported wheat yields, perceptions about the insurance product received, and experiences with the WheatCam app.

Second, as an objective measure of wheat yields, we use crop cutting exercises (CCEs) conducted during the last month of the wheat growing season at all sites in which a farmer had taken at least two pictures over the season. The field team informed farmers of the CCEs only the day before their visit, so that farmers would not adjust their behavior in anticipation of these validation visits for picture-based loss assessments. During this visit, the field team identified two separate sections of one square meter at each site, both visible in the time-lapse of pictures: 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 cut, threshed, and the resulting grain was weighed. We use the average weight from the two square meters from the same site to determine a final yield estimate for a given site.⁹ The CCEs were not used to determine insurance payouts.

Finally, the analyses use the expert assessments of visible damage in the wheat pictures. Once the Rabi season was over and the time-lapse of pictures had been processed and cleaned, each crop site was individually reviewed by three wheat experts. For each site, the experts would assess whether the crop was damaged. If the crop was damaged, they would also indicate the loss percentage, cause of damage, and the

⁹ Using CCEs is one of the most accurate ways available for estimating plot-level yields. An improved way for estimating plot-level yields, considered to be the "gold standard", is to harvest the entire plot, which was logistically and financially infeasible in our project. See Lobell et al. (2018) for a complete discussion. Although it is common agronomic practice to sample larger areas than one square meter to improve the precision of yield estimates, we find very high correlation between yield estimates for the samples to the left and to the right of the reference pole, and the two measures almost always overlap when discretizing yields into categories. This indicates that the implemented procedure was sufficiently precise for our objective.

percentage of visible damage that was due to unavoidable hazards or due to mismanagement by the farmer. Finally, for the few 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.

Appendix Table A1 summarizes baseline characteristics across different groups of study farmers: all 736 farmers who completed a baseline interview; the 467 farmers who actively participated in the project and took at least two pictures throughout the season, thus qualifying for expert loss assessments; and the 357 farmers for whom we also conducted a CCE. Attrition in the CCEs is largely due to farmers having already harvested before the team reached a village, or the crop not yet having ripened at the time the team reached a village. Active participation in the project appears unrelated to observable farmer, household, or farm characteristics, and also differences between farmers with and without data from CCEs are small and statistically insignificant. Thus, we find no significant evidence of attrition bias in terms of observable farm and household characteristics in the samples used for the analyses.

4. Results

This section reports findings regarding three key knowledge gaps concerning the feasibility of the picture-based insurance approach. First, we analyze compliance with the picture-taking protocol. Second, we assess whether damage can be quantified accurately from smartphone camera data or not. Third, we analyze to what extent PBI reduces basis risk compared with alternative index insurance approaches.

4.1. Compliance with the picture-taking protocol

A first prerequisite for PBI to be feasible is that there is enough image data available at the time of loss assessment to determine the damage (if any) suffered by the crop. For this, farmers need to take pictures of their fields regularly. Out of the full sample of 736 farmers, 548 farmers participated in the village sessions and were trained on the insurance products. Other farmers were unable to attend or were not interested in the project. We consider this as a determinant of insurance take-up, and treat the group of 548 trained farmers as the appropriate benchmark to study compliance with the picture-taking protocol. Among this sample, 122 farmers (22.3 percent) did not take any repeat picture, meaning that they did not qualify for loss assessment.

Panel A of **Fig. 2** presents the distribution of the total number of pictures taken per farmer, conditional on taking at least one repeat picture. Close to two-thirds (345 farmers or 63 percent) of all trained farmers took at least four pictures throughout the season—or roughly one picture per growth stage. While loss assessment can potentially be done with as little as two pictures (one picture early in the season to verify normal growth and one picture close to harvest to assess any damage present in the crop), having four pictures allows to monitor crop growth throughout the season and to identify the growth stage (Hufkens et al., 2019), and we hence use this as our preferred compliance measure. Finally, 150 farmers (27.4 percent of trained farmers) took pictures twice a month or more, resulting in a high-quality time lapse that is being used to develop image processing algorithms for automated loss assessment.

Note: This figure shows different statistics of farmer's picture-taking activity throughout the Rabi 2016/17 season. Panel A shows a frequency histogram indicating the number of farmers that took a certain number of pictures through the season, conditional on taking at least one repeat picture. Panel B plots the number of farmers that took at least one picture per calendar week across the season.

Panel B 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

Panel A. Total number of pictures taken
by farmer




Fig. 2. Picture-taking activity.

Note: This figure shows different statistics of farmer's picture-taking activity throughout the Rabi 2016/17 season. Panel A shows a frequency histogram indicating the number of farmers that took a certain number of pictures through the season, conditional on taking at least one repeat picture. Panel B plots the number of farmers that took at least one picture per calendar week across the season.

period.¹⁰ Thus, while not all farmers followed the requested protocol, nearly two-thirds of the trained sample submitted at least one picture per growth stage, and a sustained participation was observed over time. Notably, WheatCam did not have built-in reminders nor other mechanisms such as picture-based advisories or pest and disease detection services that could have helped encourage further compliance with the picture-taking protocol.

Finally, we analyze whether farmers' ability and willingness to take pictures for insurance purposes depends on observable farmer characteristics, using picture-taking activity as a proxy for product engagement. Because smartphones could be a relatively unfamiliar technology, and technology acceptance could vary across demographic or socioeconomic dimensions, it is important to analyze whether PBI is more inclusive for specific segments of the population. To that end, **Table 1** assesses the relationship between trained farmers' characteristics and picture-taking behaviors. Column (1) focuses on the extensive margin, that is, whether a farmer took at least one repeat picture throughout the season. Column (2) shows the results from Tobit regressions of the number of pictures taken. Finally, Column (3) analyzes our compliance measure, using a dummy variable taking the value of 1 when a farmer took at least four repeat pictures throughout the season (approximately one picture per growth stage) and 0 otherwise.

In Column (1), none of the variables predict whether a farmer took at least one repeat picture, whereas in Column (2), belonging to a lower caste is a strong determinant for reduced picture-taking in the intensive margin. Interestingly, farmers sowing a smaller portion of land and farmers who do not own the insured plot took more pictures on average. While landless farmers are normally excluded from the crop insurance market, the present study explicitly allowed them to obtain coverage, perhaps inducing higher participation among this sample.

¹⁰ In the beginning of the season, participation remained limited because wheat had not started growing yet and because of technical challenges with WheatCam. In initial versions of the app, GPS restrictions (imposed to prevent tampering) and frequent crashes were a challenge. These issues were resolved in later versions of the app.

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




Fig. 2. (continued)

Farmers whose plots are located farther from their homes tended to take more pictures. Although we do not have data to test this hypothesis, this could be related to more established routines for visiting their plots. Finally, in terms of age, we expected higher technology acceptance and hence higher participation among the youngest tercile of farmers in our sample (those with an age below 30 years) but observe similar participation across all age terciles, indicating that age is not a key factor for adoption of this particular technology.

In sum, certain characteristics related to technology acceptance in other contexts, such as caste and potential access to insurance substitutes, are also important for compliance with PBI protocols. Other characteristics, such as a farmer's age, education, landholdings, progressiveness, or smartphone experience, are not significant determinants of participation.¹¹ Our initial concern that such farmers would be unwilling to engage with an innovative product through a relative unfamiliar technology appears unfounded in the context of this study.

4.2. Do pictures capture damage?

A next question is whether damage arising from different types of hazards is indeed visible in a smartphone picture. To address this question, we will first present qualitative results and descriptive statistics for the expert loss assessments, followed by a discussion of the agreement across experts. Finally, we will analyze the association between expert loss assessments and yields measured through CCEs. From an insurance perspective, it is important that the loss assessment identifies plots with severe damage. Thus, as our main measure of (downside) basis risk, we will analyze the proportion of sites with arguably severe damage for which PBI would have triggered payouts. We will show that the expert loss assessments would have triggered insurance payouts for more than 70 percent of sites with severe damage.

Our protocol required farmers to take overview pictures of insured plots, taken at a distance such that a large fraction of the plot was visible. We opted for this protocol since close-up pictures could become

subject to tampering too easily. However, damage might be less visible in overview pictures than in close-up ones, raising the question whether these pictures can indeed capture damage events. Initial conversations with local wheat agronomists indicated that overview pictures would be able to capture most—though not all—hazards and hence eliminate basis risk. Certain events such as lodging (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 heat stress during the grain filling stage would be more difficult to identify. Endline survey data indicate that farmers have similar perceptions (Appendix Figure A1), suggesting that PBI is well suited for minimizing basis risk, at least in the study context.

Table 2 provides descriptive statistics for the expert loss assessments (and associated insurance payouts). Panel A shows the percentage of cases for which the assessments triggered a payout, together with the different PBI payout categories. Based on experts' visual inspection of the pictures, 9 percent of farmers experienced a loss above 20 percent—thus triggering a payout. The average PBI payout, conditional on triggering, was Rs. 5,200, with most cases assessed to have between 20 and 50 percent damage.

Fig. 3 presents a box plot of the expert loss assessments for total damage, including damage both due and not due to mismanagement, ordered by the median assessment within a site. The figure reveals a few interesting patterns. For low levels of damage (median damage below 20 percent), we observe high levels of agreement between experts, with most assessments falling within the same damage category. For sites with higher levels of visible damage (median damage above 20 percent), we naturally observe more disagreement regarding the exact level of damage. Most experts nonetheless agree on 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 can identify crop losses from direct visual inspection of pictures.

Our main measure for downside basis risk—the proportion of farmers with severe damage for whom PBI would have triggered an insurance payout—relates these loss assessments to the CCEs conducted closely before harvest. Fig. 4 shows a scatterplot, mapping measured yields on the vertical axis against the final expert loss assessment for the same field (that is, the joint expert consensus for sites with a lot of disagreement and median assessment for the rest).¹² On the right

¹¹ Interestingly, picture-taking behavior does not seem to have been affected by the expectation of losses either. If a farmer were to expect little or no losses in his plot, it would be reasonable to exert less effort in taking pictures. However, we find this not to be the case: farmers with little or no losses took a very similar (and statistically-indistinguishable) number of pictures than those with severe losses. This is reasonable to an extent, since a farmer cannot rule out the possibility of his crops being damaged until very close to harvest, creating dynamic incentives to take pictures as a safeguard measure. We thank an anonymous reviewer for pointing this out.

¹² Ideally, we would convert measured yields into measured damage, as a percentage of a farmer's attainable yield, but this requires several assumptions, including an estimate of what the yield would have been in the absence of

Table 1
Factors related to picture-taking behavior and compliance.

	(1)	(2)	(3)
	Took repeat pictures	Number of repeat pictures	Took at least four pictures
	OLS	Tobit	OLS
Landholdings (HAs)	−0.001 (0.004)	−0.201 (0.246)	0.001 (0.005)
Age is under 30 years	−0.008 (0.047)	−0.121 (2.645)	0.032 (0.049)
Age is over 50 years	0.062 (0.057)	1.847 (2.566)	0.045 (0.051)
Highest level of education	0.014 (0.012)	0.087 (0.535)	0.005 (0.011)
Belongs to sched./OB caste	−0.148 (0.095)	−8.768** (3.843)	−0.129 (0.094)
Perception of yield variability	0.002 (0.011)	−0.414 (0.634)	−0.008 (0.013)
Household size	−0.008 (0.008)	−0.201 (0.357)	−0.006 (0.005)
Takes pictures on phone often/very often	−0.064 (0.045)	−2.397 (2.825)	−0.067 (0.055)
Has network signal often/very often	0.011 (0.044)	0.338 (3.556)	0.004 (0.050)
Ever used laser land leveller	−0.058 (0.050)	−1.357 (3.053)	−0.043 (0.049)
Wheat yield Rabi 2015/16	0.007 (0.010)	0.831 (0.597)	0.014 (0.012)
Share of income from crops	0.088 (0.127)	4.201 (8.007)	0.010 (0.138)
Share of crop income from wheat	0.121 (0.165)	2.557 (8.489)	0.235 (0.178)
Fraction of land planned to be sown with wheat	−0.227 (0.151)	−28.704* (15.265)	−0.410* (0.216)
Owens insured plot	0.022 (0.075)	−9.409* (4.866)	−0.076 (0.074)
Distance from plot to home (in minutes)	−0.000 (0.001)	0.085* (0.044)	0.001 (0.001)
Mean of dep. Variable	0.726	14.14	0.631
Observations	539	539	539
R-squared	0.118		0.123

Note: This table shows ordinary least squares and Tobit regressions for the determinants of various measures of picture-taking activity. Standard errors, clustered at the village level, are in parentheses. We also control for a constant, weather station fixed effects, and dummy variables to indicate PBI villages and villages where payments were conditional on not burning the previous season's crop residue (this condition was orthogonal to the treatment under consideration). We do not present coefficients for these variables here. Columns (1) and (2) are conditional on a farmer taking an initial picture. The variables "Owens insured plot" and "Distance from plot to home (in minutes)" were imputed for 13 missing observations using the mean value of the observed sample. A dummy to account for this imputing is controlled for but not reported. ***p < 0.01, **p < 0.05, *p < 0.10.

vertical axis, distinguishing between three yield ranges that roughly echo the average farmer's PBI payout categories, we indicate—by yield range—the proportion of farmers for which the PBI product would have triggered a payout. Although the correlation is not perfect, the figure shows a negative relationship between damage estimated through pictures and CCE yields, consistent with our hypothesis that picture-based loss assessments capture crop losses.

From an insurance perspective, it is important that the loss assessment identifies plots with severe damage, that is, sites with very low yields (below, for instance, 10 quintals per acre, corresponding to 50 percent

(footnote continued)

damage, which is unavailable given the endogenous, farmer-specific component of yields. We therefore assess the relationship between assessed damage and measured yields as opposed to measured damage.

Table 2
Loss assessments and insurance payouts.

	Mean	Std. Dev.	Median	N
<i>Panel A: Triggering of indices</i>				
PBI index triggered (%)	9.0	—	—	412
Slightly damaged: 20–50% (payout Rs. 3,900)	6.3	—	—	412
Severely damaged: 50–75% (payout Rs. 7,800)	2.4	—	—	412
Fully damaged: 75–100% (payout Rs. 13,000)	0.3	—	—	412
WBI index triggered (%)	22.8	—	—	412
<i>Panel B: Payout if index triggered (in Rs.)</i>				
PBI payout	5,200	2,188	3,900	37
WBI payout	2,307	585	2,259	94

Note: This table shows summary statistics of weather-based insurance (WBI) and picture-based insurance (PBI) payouts across all study villages. The table includes only the 412 farmers for which a loss assessment based on their pictures was completed. In WBI only villages, PBI payouts include the hypothetical payouts that would have been made based on experts' loss assessments of pictures if farmers would have been insured under PBI.

damage for the average farmer). We therefore use the proportion of sites with arguably severe damage for which PBI would have triggered payouts as our measure of (downside) basis risk. Experts generally identify severe damage when it exists. PBI would have triggered for 71.4 percent of sites with yields lower than 10 quintals per acre. Upon further inspection, of the four sites with assessed damage under 20 percent (in red in the left-bottom corner), three did not show any visible damage in the pictures, and one showed visible damage that was categorized as mismanagement of the crop under water-logging conditions.

Assessments were less accurate for cases with moderate damage: while triggering payments for 7.1 percent of farmers with normal yields (above 16 quintals per acre), experts assessed losses of at least 20 percent for only 6.1 percent of farmers with moderate damage (yields between 10 and 16 quintals per acre). Conditional on triggering a payment, average payouts were nonetheless substantially lower for farmers in the moderate and low/no damage categories, somewhat mitigating the apparent misclassifications incurred under PBI.

4.3. Does PBI reduce basis risk compared to index insurance?

In this subsection, we compare the effectiveness of PBI and alternative index products in terms of their ability to provide comprehensive coverage with minimal basis risk. We compare PBI with both the WBI product and a simulated area-yield insurance (AYI) product. AYI pays out according to the average yield estimated from a limited number of crop-cutting exercises (CCEs) in a given geographic area (such as a village or block in the case of India). The comparison with such a product is highly relevant. In India alone, the Pradhan Mantri Fasal Bima Yojana (PMFBY, or Prime Minister's National Crop Insurance Scheme), which was launched in 2015/16 and covered more than 40 million farmers during the Kharif 2017 season, relies mostly on this type of loss assessment mechanism (Bhushan and Kumar, 2017).

Table 2 summarized the proportion of farmers for whom the WBI product triggered, and the average payouts in case the product triggered. In contrast to the 9% of farmers for whom the PBI product triggered, WBI triggered payouts for a significantly higher 22.8% of farmers, an implausibly high number given that no widespread damage due to excess rainfall or extreme heat was reported in the Rabi 2016/17 season. Table 2 Panel B shows that the average payout for WBI was quite lower than that for PBI.¹³ This reflects the fact that WBI tended to

¹³ The maximum WBI payout was Rs. 3,382, indicating that the index was far from the product's exit trigger at which it would have paid the full coverage amount of Rs. 13,000. PBI payouts, on the other hand, include one case with a full payout and five cases where the intermediate payout (Rs. 7500) was issued.




Fig. 3. Individual expert loss assessments.

Note: The figure shows the dispersion of individual experts' damage assessments across different levels of median damage assessment at a site (where for each crop site there are three expert assessments). A median damage assessment category may contain multiple sites with the same median assessment.




Fig. 4. Yields from crop-cutting exercises (CCEs) and expert loss assessments.

Note: This figure shows a scatterplot of (a) wheat yields captured during CCEs conducted immediately before harvest and (b) the median damage assessment from wheat experts based solely on the time-lapse of pictures from the entire season. Each observation corresponds to one single plot. † indicates the probability of receiving a payout in each actual yield category; and the average payout is shown within brackets.

trigger at relatively low index values, around levels at which substantial damage to the crop is uncommon, potentially inducing numerous instances of false positives. On the other hand, as Fig. 4 indicates, PBI identified farmers with substantial damage, naturally resulting in higher payouts.

We did not implement an AYI product but can simulate payouts under this approach, using the yield data collected through CCEs right before harvest. Specifically, we randomly select four farmers for each (a) district or (b) cluster of two nearby villages (located within 5 km from the same weather station), construct an area-yield index based on their average yields, and assume payouts are triggered (to all farmers in the district or cluster) when this area-yield index is below 16 quintals per acre (20 percent or more below normal yields). We repeat this exercise 10,000 times and report the relevant average values across these iterations. In a way, this is a best-case scenario for AYI, given that the data used to determine the index and the data used to validate this index are originating from the same dataset, which is rarely the case in practice.

Fig. 5 summarizes the proportion of farmers for whom the PBI and WBI products triggered and the average simulated proportion of farmers for whom the district-level and village cluster-level AYI products would have triggered. As in Fig. 4, we disaggregate farmers into three yield categories, roughly echoing the damage categories used for

PBI payouts. Horizontal lines indicate the proportion of farmers receiving a payout for a given product across the three yield categories. For each product, the left bar is our main measure of downside basis risk: the proportion of farmers with severe damage for whom a product triggers a payout. The middle bar indicates to what extent a product accurately triggers payouts for farmers with moderate levels of damage, and finally, the right bar provides a measure of upside basis risk: the proportion of farmers with normal yields for whom the insurance product nonetheless triggers a payout.

First, when comparing PBI with WBI, the contrast is striking. WBI fails to identify farmers with severe damage altogether, mainly because damage arose from hail storms and lodging events which WBI was not designed to cover. For farmers with moderate damage, WBI was more likely to trigger than the PBI product, making small insurance payouts to 36.4 percent of farmers. Worryingly, however, for farmers with low or no damage, WBI triggered significantly more often than PBI, issuing small payouts for 24.2 percent of farmers with yields above 16 quintals per acre, which is indicative of upside basis risk in the WBI product. Of course, these findings are only applicable to the specific WBI product implemented in this project and to the specific study season under consideration, which was not characterized by severe covariate weather shocks affecting wheat. However, comparable high levels of basis risk have been found across similar WBI products in the literature (see




Fig. 5. Payouts from different insurance products by crop-cutting yields.

Note: This figure shows the probability of receiving a payout from four alternative insurance products for three groups of farmers based on their yield during the Rabi 2016/17 season. The horizontal blue bars represent the probability of receiving a payout for a given product across all farmers. The vertical grey bars indicate confidence intervals at the 95 percent level.

Morsink et al., 2016), suggesting that our finding may generalize to other seasons.¹⁴

Second, the simulated AYI products also suffer from high levels of basis risk; the district-level product triggers payouts for, on average, 22.1 percent of the farmers with severe damage; 8.9 percent of the farmers with moderate damage; and 5.5 percent of the farmers with no or minimal damage. The cluster-level product (encompassing two nearby villages), which is closer to the operational guidelines followed in India's national crop insurance scheme (the PMFBY), performs slightly better, triggering payouts for 34.4 percent of farmers with severe damage, 13.1 percent of farmers with moderate damage, and 3.6 percent of farmers with no/minimal damage. As would be expected due to spatial correlation in yields, measuring yields for a cluster of nearby villages—although costlier and logistically more cumbersome—reduces basis risk compared to measuring district-level yields, but still fails to trigger for the majority of farmers with severe damage. Product performance also varies widely across simulations, as indicated by wide confidence intervals.

PBI suffers from both advantages and disadvantages compared with these products. While PBI appears well-suited to identify farmers with severe damage, it does not help distinguish farmers with moderate damage from farmers with less or no damage. Thus, instead of testing PBI as a standalone product against potential alternatives, there is potential for it to complement existing index products, acting as a top-up component that can identify severe localized damage. Such a scheme also introduces the possibility for the traditional index product to capture non-visible crop damage, a task for which PBI is inadequate. Appendix 1 simulates the advantages (in terms of reduced basis risk) from bundling PBI with existing index products.

All in all, the findings above indicate that PBI can bring about significant improvements in insurance coverage for farmers experiencing severe damage. Although PBI fails to identify cases with moderate damage, it does trigger payouts for 71.4 percent of those with severe damage. This is a substantial improvement over both WBI and simulated AYI products, which identify on average at best 34.4 percent of these cases. PBI can hence be a valuable complement to include as a fail-safe trigger in traditional index products.

5. Conclusions

Picture-Based Insurance (PBI) is a new approach to improve smallholder farmers' access to affordable but high-quality crop 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 comply with picture-taking protocols and send in smartphone pictures of their crops on a regular basis; (b) expert loss assessments are able to detect, using these smartphone pictures, the majority of cases with severe damage; and (c) PBI reduces downside basis risk compared to both weather index-based insurance and area-yield insurance.

With the insurance product provided for free during this initial season, we are unable to inform questions about the levels of take-up that would be observed were PBI to be offered commercially.¹⁵ We focus instead on assessing the technical feasibility of such an approach for insuring agricultural risks, both in terms of farmers complying with the PBI protocol and of pictures serving as a medium to identify damage in the crop. In future research, it is important to study PBI sustainability considerations under commercial premium rates, including the question whether shifting costs associated with claims verification to the client—by requesting farmers to take pictures—reduces take-up in settings with already low demand for insurance. In a follow-up study, having to take pictures did not crowd out demand and no evidence was found of adverse selection or moral hazard (Ceballos and Kramer, 2018a, 2018b). Nonetheless, given the challenges that donors, policy-makers and insurance providers have faced in scaling index insurance, further research around the sustainability of PBI would not be misplaced.

Moving forward, scaling-up of the PBI approach could take a few forms. First, an insurance product could be implemented featuring two

¹⁴ Even in years with catastrophic damage from the perils aimed to be covered through the weather index, WBI may suffer from fundamental design issues such as the specific choice of index, temporal coverage, and heterogeneity in exposure to these risks, resulting in a large degrees of design basis risk even at the covariate level.

¹⁵ The free nature of insurance coverage means that we have observed an upper bound of take-up. We may however have observed a lower bound for compliance, conditional on already having insurance, by making insurance available free of charge, through three channels: (i) selection into the sample of those with lower valuation and thus lower willingness to comply; (ii) farmers not feeling the “sunk costs” of acquiring insurance that could induce higher compliance in a commercial setting; and (iii) a zero price signaling low quality, thus discouraging compliance (see Cohen and Dupas, 2010).

layers: (i) a standard low-cost index, for instance a weather index or a coarse area-yield index; and (ii) damage estimates from visual inspection of pictures by experts. To make this approach scalable, the main difference with the approach that we followed would be for time-lapses of pictures to be assessed only in case the first layer does not trigger a significant payout, and only for farmers who file an insurance claim for visible damage. The pictures and engagement with farmers could generate the ground truth data needed to further strengthen the quality of the standard low-cost index, so that not only farmers with the PBI layer benefit.

Second, combining ground pictures and claims data from the initial seasons, together with weather data, georeferenced yield data, and satellite observations of insured plots, algorithms can be developed for automated claims processing, which would further decrease cost and improve the speed of damage assessments. For instance, vegetation and texture indices are already being derived from the pictures throughout the season, and used to approximate growth stages (Hufkens et al., 2019). Insurers could use this to dynamically adjust their weather indices, based on the productivity loss from a weather event predicted for the observed growth stage, reducing temporal basis risk (Dalhaus et al., 2018). Moreover, while the development of machine learning algorithms to detect visible damage requires a large amount of data to train reliable models, functional models could be developed after a few seasons, boosting the scalability of the PBI approach (whether in combination with other index products or as a standalone instrument).

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 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. Moreover, future efforts could make communication two-way, potentially in partnership with telecoms or financial institutions trying to increase their market share among smallholder farmers and concentrate on bundling picture-based insurance with agro-advisory and pest detection services to make the benefits of taking pictures more salient to farmers. This more holistic risk management system could also help minimize any potential picture tampering as it would be in the best interest of the farmer to provide accurate information on the crop status.

It is worth noting that the implementation costs under PBI versus either WBI and AYI are of a different nature. On one hand, WBI and AYI products have a fixed implementation cost at the regional level, introducing an implicit trade-off between basis risk reduction and cost: while a geographically-narrower index may reduce the level of basis

risk, it can do so only at the expense of increasing the per-cluster implementation cost. Such a tradeoff is not present in PBI, where the quality of coverage of the insurance product is unaffected by increasing cluster size. On the other hand, using pictures as a method for loss verification carries a constant variable cost, both in terms of the time needed for experts to assess losses for each individual site (which may reduce with time, once algorithms for automatic loss assessment are developed), and in terms of the cost for farmers, including time spent taking repeat pictures and costs of cellular data to upload these pictures to the server (which could potentially be overcome by including PBI within the umbrella of a subsidized national crop insurance scheme). Given this variable nature of costs associated with PBI, this proposed solution may be better-suited to riskier cash crops such as fruits or vegetables, which are not grown on a sufficiently large scale to justify the fixed costs of implementing an area-yield insurance scheme.

In conclusion, smallholder farmers can benefit from access to an ecosystem of insurance products, catering to their individual preferences and characteristics, and tackling the nature of production risks in a given geographic area and for a given crop. In this regard, PBI is a promising concept to complement existing insurance products, serving as an additional layer to protect against extreme damage and reduce basis risk, while at the same time retaining the cost advantages of more traditional index schemes. It has the potential to bring about important changes in how insurance is offered to smallholders in rural areas of the developing world.

Declarations of interest

None.

Acknowledgements

We gratefully acknowledge Braulio Britos and Matthew Krupoff for excellent research assistance; Azad Mishra and Siddhesh Karekar from HDFC for pioneering the implementation of PBI; the Borlaug Institute for South Asia for data collection; and Pramod Aggarwal, Koen Hufkens, Michael Mann, and Eli Melaas for their guidance. We received valuable suggestions from seminar participants at Göttingen University and participants of the Platform for Agricultural Risk Management (PARM) K-Sharing and Learning Workshop at IFAD; IPA's 3rd Annual Researcher Gathering on Financial Inclusion and Social Protection at Northwestern University; the CEAR Academic Pre-Conference in Microinsurance in Lima, Peru; the Regional Dialogue on The Role of Agricultural Research in the Design, Implementation, and Evaluation of Agricultural Insurance at the BRAC Centre in Dhaka, Bangladesh; and ILO's workshop on agricultural insurance in Accra, Ghana.

Appendix. A1 Descriptive statistics

Table A1
Descriptive statistics and attrition

	(1)	(2)	(3)	(4)	(5)	(6)
	Completed base-line	Qualified for loss assessment (LA)	Difference (2)–(1)	Completed LA and CCE	Difference (4)–(1)	Difference (4)–(2)
PBI village	0.503 (0.071)	0.467 (0.074)	–0.036	0.462 (0.076)	–0.041	–0.005
Age (in years)	39.143 (0.722)	39.364 (0.759)	0.221	39.423 (0.814)	0.280	0.059
Completed tertiary education	0.438 (0.019)	0.443 (0.022)	0.006	0.420 (0.024)	–0.017	–0.023

(continued on next page)

Table A1 (continued)

	(1)	(2)	(3)	(4)	(5)	(6)
	Completed base- line	Qualified for loss assessment (LA)	Difference (2)–(1)	Completed LA and CCE	Difference (4)–(1)	Difference (4)–(2)
Belongs to sched./OB caste	0.102 (0.029)	0.077 (0.027)	–0.025	0.067 (0.028)	–0.035	–0.010
Landholdings (hectares)	8.845 (0.181)	8.937 (0.212)	0.092	8.990 (0.251)	0.144	0.052
Household size	6.240 (0.119)	6.156 (0.140)	–0.084	6.132 (0.157)	–0.109	–0.025
Perception of yield variability	3.993 (0.088)	4.073 (0.094)	0.080	4.073 (0.096)	0.080	0.000
Share of income from crops	0.835 (0.012)	0.842 (0.013)	0.008	0.845 (0.015)	0.010	0.002
Share of crop income from wheat	0.376 (0.010)	0.383 (0.012)	0.008	0.389 (0.013)	0.013	0.005
Fraction of land planned to be sown with wheat	0.961 (0.007)	0.965 (0.007)	0.003	0.969 (0.006)	0.008	0.004
Wheat yield Rabi 2015/16	19.570 (0.156)	19.556 (0.178)	–0.015	19.568 (0.196)	–0.002	0.012
Ever used laser land leveler	0.707 (0.038)	0.700 (0.043)	–0.008	0.713 (0.046)	0.006	0.014
Distance from plot to home (minutes)	15.900 (1.186)	15.777 (1.385)	–0.123	16.545 (1.586)	0.645	0.768
Owns insured plot	0.938 (0.011)	0.941 (0.012)	0.004	0.936 (0.015)	–0.002	–0.006
Takes pictures on phone often/very often	0.774 (0.025)	0.752 (0.026)	–0.023	0.754 (0.029)	–0.021	0.002
Has network signal often/very often	0.755 (0.030)	0.756 (0.032)	0.000	0.754 (0.034)	–0.002	–0.002
Number of farmers	736	467		357		

Note: This table shows the mean value of baseline farmer characteristics across different sub-groups of study farmers. Column 1 includes all farmers who completed a baseline interview, column 2 those farmers who qualified for expert loss assessments, and column 4 those for whom both loss assessment and crop cutting exercises were conducted. Columns 3, 5, and 6 show the results of tests of equality of means between every pair of groups, where none of the differences are statistically significant. ***p < 0.01, **p < 0.05, *p < 0.10.

A2 Simulations of an area-yield index

This appendix describes a simulation exercise to illustrate the advantages of bundling PBI with, alternatively, WBI and AYI products. Table A2 summarizes the results from such an approach. We first consider a lenient policy that triggers payouts when expert assessments indicate that the farmer experienced at least 20 percent damage. In Panel A, combining this lenient PBI policy with AYI reduces downside basis risk compared with the standalone AYI products, by increasing the proportion of farmers with less than 10 quintals per acre receiving payouts from less than 40 percent to approximately 75 percent. At the same time, it increases the proportion of farmers that receive payouts while not experiencing damage, leading to upside basis risk and higher costs of the insurance policy. However, it is worth noting that in identifying farmers with severe damage, the district-level product now performs as well as the product measuring yields at the village-cluster level, and the number of farmers with moderate damage that would receive payouts under the combined district-level product is almost as high as when AYI is offered at the village level. In other words, at a low additional cost, PBI reduces downside basis risk, potentially realizing cost savings by reducing the number of CCEs required for AYI loss indemnification.

Compared with introducing PBI in an AYI product, combining PBI with the WBI product results in a similar proportion of farmers with severe damage receiving payouts, while increasing the probability of payouts among farmers with moderate damage. However, under such a product, 30 percent of farmers without damage would receive payouts, due to the high degree of upside basis risk in the WBI product.

In Panel B, we consider bundling with a stricter PBI policy that pays out only in case experts identify more than 50 percent damage in the pictures. This product, when combined with AYI products, increases the overall proportion of farmers receiving payouts only slightly compared to the standalone AYI products in Fig. 5, limiting the additional costs of providing PBI, while still significantly increasing the proportion of farmers with severe damage receiving payouts. The combined village-level product would have made payouts to 66.3 percent of farmers with severe damage, to 13.1 percent of farmers with moderate damage, and to only 3.9 percent of farmers with limited or no damage. 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.

Table A2
Bundling PBI with index insurance products

	Probability of receiving a payout						Overall proportion of payouts
	Yields < 10 quintals/acre		Yields 10–16 quintals/acre		Yields > 16 quintals/acre		
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.	
A. WBI/AYI + Lenient PBI							
Weather index-based insurance (WBI)	0.714		0.394		0.300		0.325
Area-yield: District level	0.766	0.087	0.142	0.125	0.119	0.074	0.146
Area-yield: Village cluster level	0.744	0.035	0.179	0.043	0.101	0.016	0.133
B. WBI/AYI + Strict PBI							
Weather index-based insurance (WBI)	0.571		0.364		0.248		0.271
Area-yield: District level	0.657	0.143	0.089	0.136	0.060	0.082	0.086
Area-yield: Village cluster level	0.663	0.069	0.131	0.052	0.039	0.017	0.072
Number of observations in total**	14		33		310		357
Number of weather stations	5		17		25		25
Number of districts	2		6		6		6

Notes: * Cost estimates between standalone AYI and PBI are not comparable since they depend on statistical estimation of expected payouts under each insurance system, for which we do not count with data. Instead, the cost estimates under PBI reflect the additional costs for, respectively, loss assessments and additional expected payouts for idiosyncratic events (for which we take as representative Rabi, 2016-17 season payouts). Mean and standard deviation based on a 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.

A3 Farmers' perceptions of wheat hazards




Fig. A1. Farmers' Perceptions of Wheat Hazards.

Note: This figure shows farmers perceptions around wheat hazards, captured during the baseline survey. Panel A shows answers to questions on whether different hazards to wheat would be visible in overview pictures (taken from a distance of approximately 5–15 meters). Panel B presents the extent to which farmers are concerned about different hazards (self-reported at baseline) and the average rate of occurrence of the hazard during the Rabi 2016/17 (self-reported at endline). To measure the former, we asked farmers during the baseline to divide tokens between different hazards, with more tokens being allocated to hazards that worried the farmer more.

Note: This figure shows farmers perceptions around wheat hazards, captured during the baseline survey. Panel A shows answers to questions on whether different hazards to wheat would be visible in overview pictures (taken from a distance of approximately 5–15 m). Panel B presents the extent to which farmers are concerned about different hazards (self-reported at baseline) and the average rate of occurrence of the hazard during the Rabi 2016/17 (self-reported at endline). To measure the former, we asked farmers during the baseline to divide tokens between different hazards, with more tokens being allocated to hazards that worried the farmer more.

References

- Barrett, C., McPeak, J., 2006. Poverty traps and safety nets. In: In: de Janvry, A., Kanbur, R. (Eds.), *Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke*, vols. 131–154 Springer, New York.
- Berhane, G., Dercon, S., Hill, R.V., Taffesse, A., 2015. Formal and informal insurance:

- experimental evidence from Ethiopia. In: Paper Presented at International Conference of Agricultural Economists, Milan, August 8–14.
- Bhushan, C., Kumar, V., 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. University Library of Munich, Germany MPRA Paper 46864.
- Cai, H., Chen, Y., Fang, H., Zhou, L., 2009. Microinsurance, trust and economic

- development: evidence from a randomized natural field experiment. NBER Working Papers 15396.
- Carter, M.R., Barrett, C.B., Boucher, S., Chantarat, S., Galarza, F., McPeak, J.G., Mude, A.G., Trivelli, C., 2008. Insuring the Never before Insured: Explaining Index Insurance through Financial Education Games. University of Wisconsin, Madison BASIS Briefs.
- Ceballos, F., Kramer, B., 2018a. Big Brother Might Be Watching You: Testing for Moral Hazard in Picture-Based Crop Insurance. Mimeo.
- Ceballos, F., Kramer, B., 2018b. How Attractive Is Picture-Based Insurance? Willingness to Pay and Testing for Adverse Selection. Mimeo.
- Chantarat, S., Mude, A.G., Barrett, C.B., Carter, M.R., 2013. "Designing index-based livestock insurance for managing asset risk in northern Kenya. *J. Risk Insur.* 80 (1), 205–237.
- Cohen, J.P., Dupas, 2010. Free distribution or cost-sharing? Evidence from a randomized malaria prevention experiment. *Q. J. Econ.* 125 (1), 1–45.
- Cole, S.A., Giné, X., Tobacman, J., Topalova, P.B., Townsend, R.M., Vickery, J.L., 2013. Barriers to household risk management: evidence from India. *Am. Econ. J. Appl. Econ.* 5 (1), 104–135.
- Cole, S.A., Giné, X., Vickery, J.L., 2017. How does risk management influence production decisions? Evidence from a field experiment. *Rev. Financ. Stud.* 30 (6), 1935–1970.
- Dalhaus, T., Musshoff, O., Finger, R., 2018. Phenology information contributes to reduce temporal basis risk in agricultural weather index insurance. *Sci. Rep.* 8 (46), 1–10.
- Dercon, S., Hoddinott, J., 2004. Health, shocks, and poverty persistence. In: Dercon, S. (Ed.), *In Insurance Against Poverty*. Oxford University Press, Oxford, UK, pp. 123–136.
- Flatnes, J.E., Carter, M.R., 2015. Fail-Safe Index Insurance without the Cost: A Satellite Based Conditional Audit Approach. Mimeo.
- Hazell, P., Pomareda, C., Valdes, A., 1986. *Crop Insurance for Agricultural Development: Issues and Experience*. Johns Hopkins University Press, Baltimore.
- Hill, R.V., Robles, L.M., Ceballos, F., 2016. Demand for a simple weather insurance product in India: theory and evidence. *Am. J. Agric. Econ.* 98 (4), 1250–1270.
- Hufkens, K., Melaas, E.K., Mann, M.L., Foster, T., Ceballos, F., Robles, M., Kramer, B., 2019. Monitoring crop phenology using a smartphone based near-surface remote sensing approach. *Agric. For. Meteorol.* 265 (15), 327–337.
- Jensen, N.D., Barrett, C.B., Mude, A.G., 2016. Index insurance quality and basis risk: evidence from northern Kenya. *Am. J. Agric. Econ.* 98 (5), 1450–1469.
- Jensen, N.D., Mude, A.G., Barrett, C.B., 2018. How basis risk and spatiotemporal adverse selection influence demand for index insurance: evidence from northern Kenya. *Food Policy* 74, 172–198.
- Karlan, D., Osei, R., Osei-Akoto, I., Udry, C., 2014. Agricultural decisions after relaxing credit and risk constraints. *Q. J. Econ.* 129 (2), 597–652.
- Lobell, D.B., Azzari, G., Burke, M., Gourlay, S., Jin, Z., Kilic, T., Murray, S., 2018. Eyes in the Sky, Boots on the Ground: Assessing Satellite- and Ground-Based Approaches to Crop Yield Measurement and Analysis in Uganda. World Bank Policy Research Working Paper No. 8374.
- Morsink, K., Clarke, D., Mapfumo, S., 2016. How to Measure whether Index Insurance Provides Reliable Protection. World Bank Policy Research Working Paper No. 7744.
- Matul, M., Dalal, A., De Bock, O., Gelade, W., 2013. Why People Do Not Buy Microinsurance and what We Can Do about it. vol. 17 Microinsurance Innovation Facility, Geneva Briefing Note.
- Mobarak, A.M., Rosenzweig, M., 2012. Selling formal insurance to the informally insured. Yale University " Working Paper No. 97, Economic Growth Center Discussion Paper No. 1007.
- Porter, J.R., Xie, L., Challinor, A., Cochrane, K., Howden, S., Iqbal, M., Lobell, D., Travasso, M., Netra Chhetri, N., Garrett, K., 2014. Food security and food production systems. In: Field, C.B., Barros, V.R., Dokken, D.J., Mach, K.J., Mastrandrea, M.D. (Eds.), *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Working Group II Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, pp. 485–533 (Chapter 7).
- Richardson, A.D., Hufkens, K., Milliman, T., Aubrecht, D.M., Chen, M., Gray, J.M., Johnston, M.R., Keenan, T.F., Klosterman, S.T., Kosmala, M., Melaas, E.K., Friedl, M.A., Frolking, S., 2017. Tracking Vegetation Phenology across Diverse North American Biomes Using PhenoCam Imagery. Scientific Data.
- Stanimirova, R., Greatrex, H., Diro, R., McCarney, G., Sharoff, J., Mann, B., D'Agostino, A.L., Rogers-Martinez, M., Blakeley, S., Small, C., Ceccato, P., 2013. Using satellites to make index insurance scalable. Final IRI Report to the ILO Micro-Insurance Innovation Facility.