Assessing index insurance: conceptual approach and empirical illustration from Burkina Faso*

Thomas Barré, Quentin Stoeffler and Michael Carter †

January 2016

Abstract

Index-insurance is a growing area of interest among researchers and policy-makers for protecting poor households against adverse shocks and promoting their productive investments. However, the quality of index-insurance products is rarely assessed, and there is currently a lack of consensus regarding the criteria of index-insurance quality. Unlike conventional insurance products, index-insurance products explicitly include a positive probability of not receiving a payout in case of shock. Thus, a bad quality index-insurance may actually reduce household well-being. Besides, insurance quality is difficult to assess for households, even ex-post, in the short term. In this article, we suggest a simple measure of index-insurance quality. The measure is based on the index-insurance target of income stabilization, and takes into account desirable properties of insurance products. We provide a first graphical illustration of the measure using simulated index-insurance products. For these simulated products, our quality measure helps rank products and discriminate low quality insurances. Then, the article gives an empirical application of the measure to an area-yield insurance product implemented in Burkina Faso. We use farmer group-level data from the cotton company and individual-level data which we collected before implementation of the insurance project. Our measure shows that despite a relatively good performance of the insurance product at the group level, quality of the product is disappointing at the individual level due to high idiosyncratic risk levels and high prices. These results suggest that greater attention has to be paid to index-insurance quality to enable products to reach their objective of income stabilization for vulnerable farmers. Specifically, insurance design should systematically employ index-insurance quality measures based on farmer outcomes– such as the one developed in this paper– to discard low quality insurance products, choose the best performing available alternative, and price the product in a way which makes farmers better-off than without insurance.

[NOTE: PRELIMINARY AND INCOMPLETE]

*Acknowledgements: [Wouter Gelade, Catherine Guirkinger, Karim Pare... GAN meetings/conf...]

†Department of Agricultural and Resource Economics, University of California, Davis

1 Introduction

Risk has been widely acknowledged as playing a major adverse role for households in the developing world, harming individuals when shocks occur and threatening livelihoods in the long-run [World Bank, 2014, Fafchamps, 2003]. Receiving a shock can have life-long consequences, and the simple threat of uninsured risk pushes household to under-invest in high-return activities [Dercon, 2004, Hoddinott and Kinsey, 2001, Zimmerman and Carter, 2003]. In this context, index insurance is increasingly considered as a promising tool for providing cost-effective protection to poor and vulnerable households. However in general, demand for index insurance products in developing countries has been disappointingly low [De Bock and Gelade, 2012]. "Behavioral" explanations have been suggested to explain low take-up: aversion for uncertainty about the insurance product itself, cash and credit constraint, etc. [Karlan and Morduch, 2010, Binswanger-Mkhize, 2012, Carter and Serfilippi, 2015]. However, the lack of quality of specific insurance products is as likely to prevent rational and well-informed households from purchasing these products [Clarke, 2011]. In this paper, we develop simple measures for assessing and comparing index insurance products, in order to design (better) insurance products and rule out contracts which do not increase household well-being. This approach can be used to assess any risk alleviating intervention, including infrastructure development or cash transfers. We illustrate how our measures can be applied to an area-yield cotton insurance in Burkina Faso, and show the challenges of designing efficient protection against shocks when a large share of the risk is idiosyncratic in nature.

From a developmental agenda point of view, the value of insurance derives from the increase in present and future well-being provided to farmers.[1] A high quality insurance will make payments to an individual farmer when she receives a shock, thus providing her adequate protection against risk. As a consequence, farmers will avoid costly coping strategies and short-and long-term negative consequences of shocks [Janzen and Carter, 2013]. In addition, they will have greater opportunities to realize risky but highly profitable investments, thus increasing their income in the long run [Elabed and Carter, 2015, Jensen et al., 2014a]. These positive impacts can be determined through an ex-post impact evaluation. However, assessing the quality of an insurance product ex-ante offers unique advantages. First, it allows a better design of a given insurance product before selling a potentially poor quality (or even harmful) contract to farmers. Second, it can assess the theoretical "core economic value" of an insurance, regardless of the quality of the implementation and the magnitude of the take-up.[2] Third, a positive impact on household investments can still be observed even in the case of a contract which fails to adequately protect farmers. This situation is of course positive in the absence of shocks, but extremely dangerous and problematic in general [ref of that ex]. In this article, we adopt a client-centric approach and provide an ex-ante assessment of the value of the protection offered to farmers.

From the farmer's perspective, protection means that she will receive indemnity payments commensurate with the losses she experiences. While this correlation between losses and indemnities is (supposed to be) perfect in the case of traditional insurance products relying on claims and field verifications, index-insurance is based on an index which does not depend directly on the losses encountered by a particular household. Thus, index-insurance includes a "basis risk" for insured households: there is a positive probability to receive payments when no loss occurs, or to receive no payments when the household does receive a shock. The latter possibility is arguably the most damaging, as households are left worse-off than without insurance: they receive a shock, did not obtain any payout, and paid an insurance premium [Clarke, 2011]. The lack of insurance quality has several sources, which are not strictly restricted to the prediction power of the index used (e.g. yield predictions from rainfall data). These sources include: index quality, geographic scale, product design, type of shocks covered, price, etc. The measures proposed in this paper aim at taking into account all these different components of insurance quality, by observing the correlation between insurance payments and actual deviations from expected yields obtained by farmers. Also, by decomposing different components of insurance quality, the analysis conducted in this paper enables us to identify whether lack of quality of an insurance

[1] Obviously, the objective may be different for a private insurance or reinsurance company.

[2] This is particularly important for the insurance product studied in this article, where implementation issues reduce the possibility to observe ex-post impact on investments.

product derives from its design and price, and thus can be improved; or if, on the other hand, lack of quality is due to high idiosyncratic risk, thus creating low potential for an index-insurance to protect individual households efficiently.

Index-insurance is a relatively new area of interest, especially in developping countries. A growing number of studies have been interested in the developments of the index-insurance sector, which has quickly evolved from small rainfall-based pilots to larger programs relying on area-yield or on new technologies such as satellite indices [De Bock and Gelade, 2012, Chantarat et al., 2013, Barnett et al., 2008]. Most of the original studies have focused on insurance demand, which has often been disappointingly low [Binswanger-Mkhize, 2012, Cole et al., 2013, Winters et al., 2010]. However, lack of demand can be driven by several confounding factors such as: lack of understanding of the product; lack of trust in the insurance scheme; cash or credit constraint; high prices; basis risk; existing alternative protections; etc. To unpack this low demand, recent studies have started to look into another, possibly essential factor: the core quality of insurance products. In particular, Clarke [2011] has shown that in case of a positive probability of uncompensated loss, demand from fully informed rational individuals with high risk aversion could be null since households are made worse-off with insurance than without insurance in some states of the world. In a companion paper, Clarke et al. [2012] proposes to set a minimum standard in terms of expected indmnities in case of catastrophic loss, and shows that a rainfall-based index insurance program in India barely meets this minimum standard. Looking at a case study of satellite-based dought insurance for herders in Northern Kenya, Jensen et al. [2014b] show that while this insurance product protects efficiently against covariate shocks, it leaves at least 62% of the total risk uninsured when taking into account idiosyncratic shocks . A number of recent studies promote the use of existing theories of decision under uncertainty (Expected Utility theory or Cumulative Prospect theory) to compute farmers'willingness-to-pay for insurance and assess the potential welfare gains and commercial viability (share of the population willing to buy insurance at a commercial price) [Flatnes, 2015, Elabed et al., 2013, Jensen et al., 2014b]. One of our contributions to this nascent literature is to develop measures which allow a fair comparison between different insurance options, focusing on the income target that the insurance aims at protecting. We also suggest meaningful (and more exigent) minimum standards in order to rule out products which do not make households significantly better-off.. In doing so, our objective is double: helping improve the design of insurance contracts and select the best products when several alternatives are available; and rule out products which should not be sold to farmers.

The next section of this article presents the measures suggested for assessing index insurance. The third section describes the insurance product and the data used for the empirical application of our measures, as well as our empirical specifications. The fourth section shows the results in terms of insurance products comparisons and minimum standards, while the last section concludes.

2 Measures for assessing index insurance

2.1 Protecting agricultural income

The debate on index-insurance quality focuses too often on the exact definition of "basis risk": for instance, should basis risk include all agricultural shocks, or only those covered by the insurance? In other words, does an insurance fail if farmers field are destructed by locust, but the insurance only covers against drought? While such a product may not fail from a legal point of view, a development approach should focus on farmers present and future well-being. For farmers, a good insurance product provides them with affordable protection against agricultural shocks so that it stabilizes their current and future income.

Based on this definition, index insurance quality arises from several factors. First, the quality of the index itself matters: an index needs to accurately identify covariate shocks in a given area. For instance a satellite index needs to be strongly correlated with actual area-yields. Second, the geographic scale of an index affects insurance quality: a weather station may not capture rainfall well for plots which are far away, or area-yields may correlate poorly with plot-yields when the area is too wide. Third, if an insurance protects against a specific shock (e.g. drought), that type of shock has to be a substantial source of loss compared to losses from other type of shocks (e.g. animal damage). Fourth, covariate shocks can have an heterogeneous impact on different households, depending on their assets (high fields vs. low fields), production function (use of input, etc.) and on their coping strategies. Fifth, the value of an insurance depends crucially on the level of covariate risk compared to idiosyncratic risk at the aggregate level. If most of the risk is idiosyncratic, an insurance protecting against covariate shocks has little value for farmers. Last but not least, the price of a product may be too high compared to the protection provided. The objective of this section is to develop simple measures which take into account all the different sources of lack of quality in index insurance.

For that reason, we try to move beyond the debate on the definition of "basis risk" and we start from the final objective of index-insurance: stabilizing agricultural income. The end objective is actually to stabilize lifetime consumption- or in other words, current consumption and assets. However, the effect of shocks (i.e. agricultural income variation) is not directly observed on consumption, because households implement coping strategies. These strategies include working more, borrowing, eating less, selling assets, engaging in sex work, migrating to gold mines, and other damaging options [Bird and Prowse, 2009, Collins, 2004]. While these coping strategies are clearly painful in current times and threaten future livelihoods, they make us unable to observe shocks by looking at current consumption directly. Consequently, our measures of insurance quality focus on agricultural income.[3]

2.2 Measures of index-insurance quality

A few measures of index-insurance quality have been suggested in the literature. In particular, the simplest measure, which only captures an element of basis risk, is the "probability of obtaining an insurance payment in case of catastrophic loss" (PP):

$$PP = prob(I_i > 0)\,|\,y_i < \tau \tag{1}$$

where I_i is the amount of payout received by individual i, y_i is yields of individual i, and is the threshold under which losses are considered as "catastrophic", for instance 80% of yield historical average (see section 3 for more details on the calibration of loss parameters). The attractive simplicity of the PP measure must be balanced with some important flaws in its design. Indeed, the PP measure would give a perfect score to a product that only pays a minimal payout every time a loss happens and has a very high premium cost.

Clarke [2011]extends this insurance quality measure by taking into account the amount of insurance payment received and the price of the insurance, and calls this measure "basis risk ratio" (BRR):

$$BRR = \frac{E(I_i|y_i < \tau)}{mE(I_i)} = \frac{\text{Expected payout in case of catastrophic loss}}{\text{Commercial premium}} \tag{2}$$

where m is a loading factor applied to the actuarially fair premium (AFP) $E(I_i)$, so that $m * E(I_i)$ represents the commercial premium. A BRR measure inferior to 1 means that, on average, during a bad year, insurance payments don't even cover the premium paid, leaving farmers worse off in bad years. Such low quality index should not be allowed for commercialization.[4] The BRR quality measure has the advantage of being intuitively understood and gives clear guidelines in terms of minimum standards. However, it presents some restrictions, including: (i) the BRR only considers one state of the world ("catastrophic loss") without taking into account the level of loss: for instance, it makes no difference between insurance payments when losses are small (e.g. 35%) or when losses reach 100%

[3]To translate agricultural production into agricultural income, one need to use agricultural prices. For the goal of assessing index-insurance, it is useful to employ a single price over the period studied, to focus on agricultural production. This is not a strong issue in our empirical application, as farm-gate prices are set by the cotton company, and relatively stable over the period.

[4]Clarke [2011] shows that any risk averse expected utility maximizers individual should not purchase the insurance if $BRR < 1$ for all levels of loss .

and farmers need insurance payouts the most; (ii) it averages "expected payouts" and thus does not differentiate between insurance products which pay a given amount certainly (100% of the time) or products which pay twice this amount only 50% of the time; (iii) because it only signals insurance products that harm farmers in bad years, the BRR measure sets an extremely low minimum standard, letting "bad" products pass the minimum requirements; it does not tell if a product actually fulfills its income stabilization objective.

Several recent articles build on existing theories of decision under risk to investigate if index insurance can be welfare improving at its actuarially fair price and compute the maximum markup that can be charged without nullifying these welfare gains [Flatnes, 2015, Elabed et al., 2013, Jensen et al., 2014b]. There exists a debate on the appropriate theory and functional forms that could be used here (McIntosh & Sadoulet). Comparing these possible approaches goes beyond the objective of the present paper. Instead, we will simply highlight the fact that minimum standard based on the willingness-to-pay can only reject products that harm farmers. We believe this is a very low requirement, leading to outcomes which are not realistic or desirable for farmers.x

Assume that a farmer faces some uncertainty in the amount of cotton she will harvest every season. Suppose she has 20% chance to obtain a low income of \$150 and 80% chance to receive a high income of \$750. Her expected income is \$630. Suppose our farmer is risk averse with an utility function of her income U(Y) = $\frac{1}{1-\alpha}Y^{1-\alpha}$ with $0 \leq \alpha \leq 2$. A risk neutral individual ($\alpha = 0$) is indifferent between this risky situation and a scenario where she would receive her average income every year, not matter the state of nature. However, a risk averse individual ($\alpha = 2$) would be willing to give up about 35% of her expected income in order to fully stabilize her income. This high willingness to pay is unrealistic, especially in the case of Bukinabe cotton farmers who already live under the national poverty line in normal years and can fall under the chronic poverty line if a disaster occurs. Giving up 35% of expected income is not a realistic option for these farmers even though it would not be welfare reducing from the perspective of expected utility theory. In comparison, an insurance contract that guarantees that income after indemnities cannot fall below 80% of the expected income would only cost 12% of our famer's income during good years. Willingness-to-pay minimum standards set the bar too low.

The measure suggested in this article is based on the developmental objective of index-insurance of stabilizing agricultural income and is constructed so that it verifies that losses are commensurate with losses, penalizes larger failures of the index, accounts for the likelihood of shocks of different intensity and compares apples to apples by normalizing the premium rate of the different contracts.

Let$R(y)$ be agricultural income from yields y without any insurance (all parameters are expressed in share of average revenue). We define an insurance income stabilization level R^*, which is the income target guaranteed, for each individual, by a "perfect insurance" offered at its actuarially fair price (AFP).[5] Thus, the perfect insurance makes payments $I^p = R^* + AFP - R(y)$ when $R(y) - AFP < R^*$, and does not pay anything when $R(y) - AFP \geqslant R^*$. As a result, individuals obtain $R^p(y) = max(R^*, R(y) - AFP)$ with the perfect insurance.

The index-insurance, on the other hand, makes payments $I^i(Z)$ depending on the external index Z and the farmer's revenue is $R^i(y, Z) = R(y) + I^i(Z) - CMP$, where CMP is the commercial premium. This insurance is not perfect, which means that farmers do not necessarilyreceive the full amount R^* in case of a negative shock: typically, they receive less in case of low levels of yields (this has been called "false negative"). In addition, there may be areas of the yield distribution where the insurance overshoots its target and provides more than R^*(overshooting in the absence of shocks has been called "false positive"). Finally, the commercial premium is typically higher than the AFP. Consequently, there is an income gap between the insurance target and actual income obtained by insured farmers. We propose an Income Stabilisation Score (ISS) to measure the quality of a given insurance contract given the income target:

$$ISS_\alpha = \int E[\big(max(R^* - R^i_{adj}(y, Z); 0) + 1\big)^\alpha - 1|y].d\Phi(y) \quad (3)$$

[5] This level is chosen normatively, similar to the "catastrophic loss" level in Clarke [2011].

where $\Phi(y)$ is the yield distribution and $R^i_{adj}(y,Z)$ is the revenue with index-insurance adjusted for the difference in prices with the perfect contract: $R^i_{adj}(y,Z) = R(y) + (I^i(Z)/CMP) * AFP - AFP$. The measure is basically the expected gap between the perfect insurance R^* (which reaches only and fully the income target) and the adjusted revenue with index-insurance $R^i_{adj}(y,Z)$ when we only consider false negatives ($R^* > R^i_{adj}(y,Z)$).[6] In addition, we weight index failures by α in the spirit of the Foster et al. [1984] (FGT) indices, in order to penalize products when realized farmer income after indemnity payment is further away from the target.

In order to make interpretation easier, we further normalize the quality measure by dividing it by the value of ISS without insurance, so that we obtain the share of the income stabilization objective that is actually fulfilled.[7] We call this measure the Normalized Income Stabilisation Score (NISS):

$$NISS_\alpha = 1 - \frac{\int E[\Big(max(R^* - R^i_{adj}(y,Z);0)+1\Big)^\alpha - 1|y].d\Phi(y)}{\int (max(R^* - R(y);0)+1)^\alpha - 1.d\Phi(y)} \tag{4}$$

The denumerator of $NISS_\alpha$ is simply ISS_αcomputed in the absence of insurance. Consequently,$NISS_\alpha$is equal to 0 in the absence of insurance, and is equal to 1 for a perfect insurance. For an index-insurance which is better than the absence of insurance, $NISS_\alpha$would be between 0 and 1. However, an index-insurance which performs worse than the absence of insurance (for instance, when farmers pay high premiums with little probability to receive payments in case of shocks) would have a negative $NISS_\alpha$.

The NISS includes some key features: (i) Appropriate payment level: for each loss level, the measure must compare the realized payments against an income stabilization target; (ii) Penalize larger gaps: it can be argued that insurance failures matter more when they leave farmers with extremely low incomes. We introduce a weight $\alpha \geq 1$ that puts more weight on larger failures of index insurance when α increases; (iii) Probabilistic sensitivity: index failures that happen often must be given a larger weight; (iv) Insurance premium: protection against risk comes at a price and we include it by normalizing the price of the contracts compared.

The next sub-section shows how these properties are satisfied through numerical simulations of index-insurance contracts, and through graphical illustrations.

2.3 Simulations and graphs

We use three simulated index-insurance contracts to illustrate how the measures presented above allow us to compare index-insurance quality of different products. Farmers yields are simulated, and three insurance contracts are designed. Levels of payments are adjusted so that all contracts have the same price as the perfect insurance (which is the AFP). Contract A is the worst contract: its insurance payments correlate poorly with farmers' losses (low yields) compared to other contracts. Contracts B and C perform better than contract A, as they unambiguously provide higher average payments in case of loss. However, it is unclear which one of contract B or contract C performs better: contract B performs better in case of extremely low yields (rare event but very detrimental to farmers), while contract C performs better in case of losses which happen more often but are not as high.

Figure 1 graphically illustrates index PP: it shows for contracts A, B and C the probability of receiving an insurance payment as a function of the level of yields (in proportion of historical average). The level of catastrophic loss is reached when yields are below 80% of historical average (i.e. $\tau = 0.8$), and the dotted curve shows the yield probability density function. The solid line represents a perfect insurance contract, which pays with probability 1 when yields are below 80% of the historical average, and probability 0 otherwise. The dashed lines represent the probability of payment of the insurance

[6] It can be argued that "false positives" (overshooting when yields are high) brings direct negative value for farmers, as it decreases confidence in the product and generates a feeling of unfairness. However, we consider that directly penalizing overshooting would result in double counting. Indeed, overshooting means that for a given premium, farmers pay for unecessary protection. As compared to an alternative insurance at the same price, this means a lower protection in other states of the world, and thus is already accounted by QG_α.

[7] We also substract 1 in order to obtain a measure which increases when quality improves.

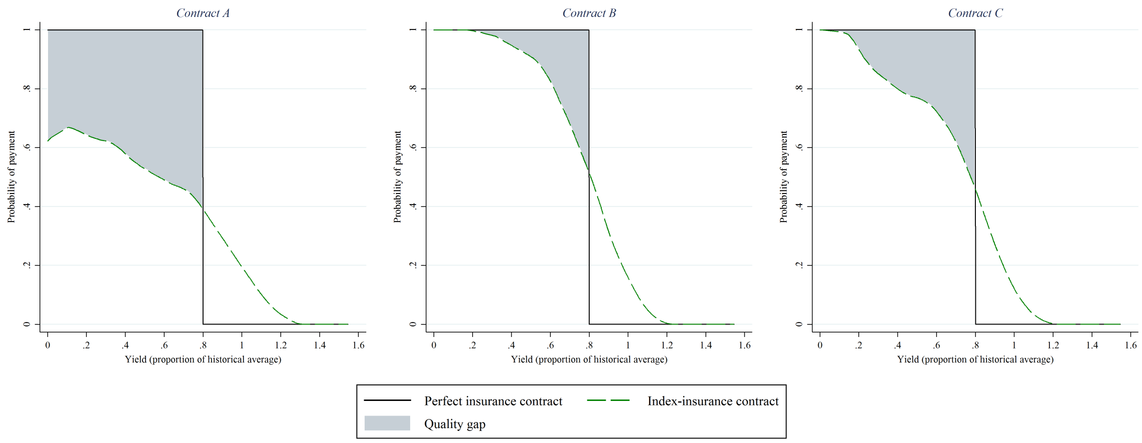

Figure 1: Quality Gap Measure #1: Probability of Payment

for each level of yields, and the grey area represents the quality gap between the perfect contract and the index-insurance contract. The grey area (and greater distance to the perfect contract) is greater for contract A compared to contract B and C, which indicates that contract A performs poorly in terms of providing payments when losses occur. The difference between contract B and contract C is more ambiguous, since the grey area is smaller for very low yields (very rare events) with contract B, but larger when yields are closer to the 80% threshold (ordinary event), compared to contract C. Table 1 shows the coresponding values of the PP measure where we can see that contract A is clearly dominated by contract B and C. Even if these latter contracts are closer in terms of probability of payment in case of loss, contract B outperforms contract C in this dimension.

Figure 2 illustrates the BRR index, showing the expected indemnity per unit paid in premium for each insurance contract. Similar to previous graphs, contract A performs poorly, providing low payments in case of low yields. This results in a large quality gap compared to the perfect contract. The difference between contracts B and C becomes more obvious in Figure 2: in case of extremely low yields, payments over premium are lower for contract C (a ratio of approximatively 5 when yields are 0) compared to contract B (a ratio of approximatively 7 when yields are 0). On the other hand, the quality gap between index-insurance and the perfect insurance is smaller for contract C around the 0.8 threshold, i.e. for levels of yields which happen more often. The respective BRR measures for the three contracts in Table 1 show that contract A is still largely dominated, and contract B is still prefered to contract C despite its larger failures in case of extreme losses.

Finally, Figure 3 represents the ISS/NISS measures, i.e. the expected revenue as a function of yields under different insurance status. The solid line and the blue dashed line still represent the perfect and index-insurance contracts respectively, and the grey 45 degree line represent the absence of insurance. The grey area is the gap between the index-insurance and the perfect insurance, but in these graphs, it represents the revenue which the index-insurance fails to stabilize. More precisely,the grey area is the amount of loss which is expected to be not covered by the insurance when the blue dashed line is below the solid line on the left side of the graph: expected insurance payments are not sufficient because of false negative, which results in lower revenues than the income stabilization target. On the right of the graph, insurance payments are too high because of false positives: farmers are receiving funds in the good state of the world, which does not correspond to the insurance objective. Again, these graphs show that contract A performs poorly, as indicated by its larger grey area. The poor performance of contract C in case of extremely low yields is obvious, but when one considers the probabilistic distribution of the level of yields, it is not obvious whether contract B or C dominates.

In table 1, the actual value of the measures are computed for the three simulated contracts. For

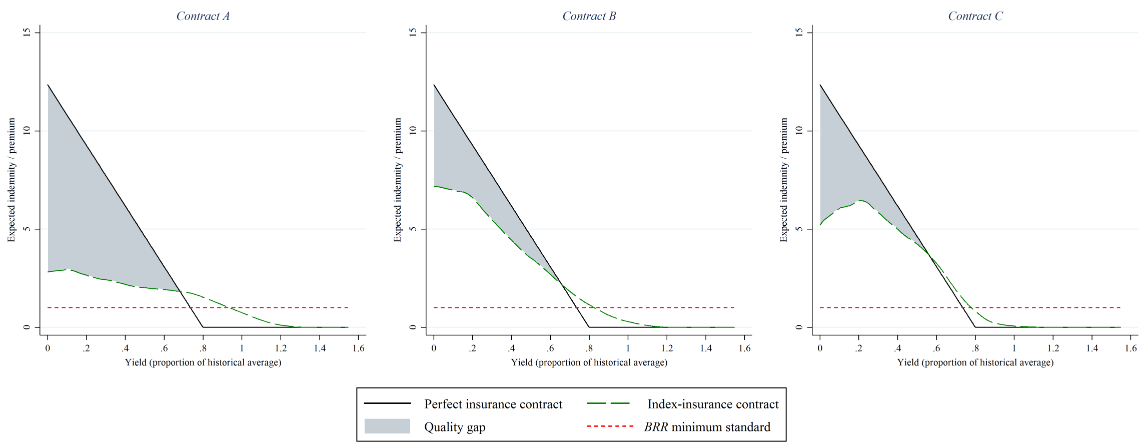

Figure 2: Quality Gap Measure #2: Basis Risk Ratio

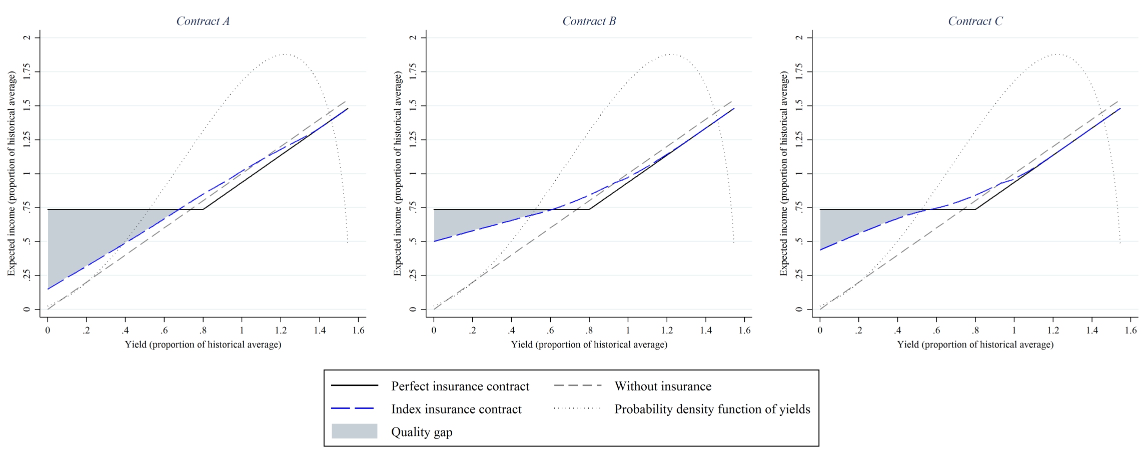

Figure 3: Quality Gap Measure #3: Income Target

	Contract A	Contract B	Contract C	No insurance	Perfect insurance
PP	0.514	0.842	0.771	0	1
BRR	2.044	3.335	3.544	0	3.768
$NISS_1$	0.327	0.795	0.818	0	1
$NISS_2$	0.451	0.932	0.922	0	1
$NISS_3$	0.551	0.977	0.965	0	1
$NISS_4$	0.635	0.992	.985	0	1

Table 1: Contract Quality Measures

all quality measures, the graphical intuition is confirmed: contract A performs poorly compared to contract B and C. However, some of the measures go further and show that contract A is not necessarily prefered to no insurance, as $NISS_\alpha$ is negative with contract A for $\alpha \geq 2$. This is due to the fact that with index-insurance, farmers can possibly be worse-off when they have a loss, pay a premium, but receive no insurance payments [Clarke, 2011]. When α is high, $NISS_\alpha$ penalizes more these events obtain when farmers obtain a revenue lower than the revenue without insurance.

In addition, the quality measures allow us to rank contracts B and C, whereas this ranking was left ambiguous by looking at the graphs. According to all measures except BRR, contract B outperforms contract C. This means that contract C poor performance for extremely low yields (compared to contract B) is not compensated by its better performance around the 0.8 threshold. BRR still ranks contract C above contract B because it is only based on the expected amount received. $NISS_\alpha$ on the other hand, considers each individual farmer's situation. Thus, it penalizes more larger gaps such as the gap observed when yields are close to 0 with contract C, especially when $.\alpha$ is high.

The rest of this article goes beyond these simple simulations and applies the approaches described above to data from an index-insurance implemented in Burkina Faso.

3 Empirical application: cotton insurance in Burkina Faso

3.1 Project & insurance design

Burkinabe farmers activity is characterized by the pervasiveness of risk and income variability. Farmers face many sources of risk, in particular drought and floods, but also localized shocks such as pest, diseases, fire, animal damage (livestock, elephants, etc.) or other production shocks (loss of labor, etc.).[8] Numerous studies have demonstrated that shocks affect both farmers consumption, which is far from being fully smoothed, and production choices [Fafchamps et al., 1998, Carter, 1997, Kazianga and Udry, 2006, Stoeffler, 2015]. For that reason, index-insurance has been promoted as a promising instrument for improving household capacity to cope with risk, smooth consumption and invest in productive assets and activities in Burkina Faso [Berg et al., 2009, Carter and Lybbert, 2012, Sakurai and Reardon, 1997].

In this context, an index-insurance product for cotton farmers has been piloted in Burkina Faso since 2013, building on the experience in neighboring Mali.[9] As in other West-African countries, the cotton sector in Burkina Faso is organized by cotton companies which have local monopolies on the purchase of cotton production. In the areas studied in this paper, Sofitex (the parastatal company) provides input on credit to farmer groups, *Groupes de Producteurs de Coton* (GPCs). Sofitex then buys the entire cotton production from GPCs and uses this production to perfectly enforce credit reimbursement, using cotton production as a collateral. If an individual farmer does not produce enough cotton to reimburse her individual part of the group credit, other farmers reimburse her part

[8] In our sample, about 70% of the farmers received at least one negative shock on their cotton or cereal plots in the last 12 months.

[9] The similar, promising insurance product piloted in Mali had to be discontinued in 2012 due to political instability [Elabed et al., 2013, Elabed and Carter, 2015].

this year and she has to pay them back with her cotton production the following year. This is known to generate tensions in GPCs related to individual debts and the group monitoring, which the insurance has the potential to ease if it manages to provide sufficient coverage to individual households [Gelade, 2015]. The structure of the cotton market makes Burkina Faso (and other West-African countries) an ideal setting to set up an area-yield product. Indeed, the cotton company has good information on surface planted, monitors the technical process and weights the cotton production, so that it can offer an area-yield contract with no additional costs and limited risk of moral hazard.

The actual insurance contract was offered in our research area by Planet Guarantee in collaboration with the cotton producer union (UNPCB) and Sofitex since 2014. The insurance contract proposed to Burkinabe cotton farmers has five important features:

1)It is an area-yield index contract where the area considered is the GPC, a group of 10 to 80 cotton farmers inside a village.

2) Similar to the product offered in Mali, the insurance has a double-trigger mechanism [Elabed et al., 2013]. The first trigger is set up at the GPC level: if cotton yields are below a certain threshold, the first condition for insurance payments is met. The second trigger is set up at the "neighborhood" level: the yields of neighboring GPCs also have to be below a (different, higher) threshold to meet the second condition for insurance payments. Only if both conditions are met will a GPC receive insurance payouts. The second trigger is supposed to prevent possible moral hazard at the GPC level, in case GPC members would decide to coordinate to obtain low yields in order to get insurance payments. While the neighborhood trigger is set higher than the GPC trigger to take into account the possibility of localized shocks, it introduces some basis risk since a GPC can have a shock but obtain no payouts if neighboring GPCs are not affected by the shock.[10]

3) GPCs have been grouped in four categories based on their historical average yield from 2000 to 2014.

4) For each group, a distribution of yields was fit and insurance thresholds were set at the 20th, 8th and 4th percentile. At the 20th percentile, i.e. at yields obtained every 5 year on average, GPCs receive a "small payout" of 11,200 FCFA per hectare, which is equivalent to the premium cost (or about 20 USD). At the 8th percentile, GPCs receive a "medium payout" of 34,000 FCFA per ha. Finally, at the 4th percentile (i.e. once every 25 year on average) GPCs receive a "large payout" of 90,000 FCFA per ha. This large payout is roughly equivalent to the credit cost of the input bundle per hectare.

5)This product was not offered as the actuarially fair price computed by the design team. Mark-ups for the commercial partners were added, resulting in a commercial premium three times higher than the actuarially fair price.

We decompose these five characteristics in section 4 to see how theey impact contract quality.

3.2 Data

The data used for our assessment of the Burkina Faso area-yield cotton insurance come from two sources. First, we use the Sofitex data which were used to design the insurance contract. This dataset includes information at the GPC level (not individual level) on production and yields for each GPC in the Sofitex regions of Hounde and Dedougou (main cotton producing regions in the country). This information corresponds to surface actually cultivated by farmers in the group and cotton production weighted (and paid to farmers) by Sofitex each year from the 2000-01 season to the 2013-14 season.[11] This dataset includes all the 704 GPCs in Hounde.

The second data source comes from surveys conducted for the impact evaluation (IE) of the pilot index-insurance program in 2014 (baseline) and 2015 (follow-up) in Hounde. In 2014, information was collected on cotton production for 2013-14, as well as retrospective information from 2008-09 to 2012-13

[10]Based on our GPC sample, the neighborhood trigger would have prevented 20% of the insurance payments to GPCs with yields below the GPC threshold since 2006, compared to a contract with a single GPC trigger.

[11]By construction, this area-yield index is more strongly correlated with the yields of the larger farmers inside each GPC. An index computing the average yield accross farmers of a same GPC would have more value for the average farmer of the GPC.

(including surface cultivated, production and yields). In 2015, information was collected for the 2014-15 season on the same variables. In total, the two surveys provide a panel dataset of 7 seasons.[12] While retrospective information is usually less accurate than current information, the structure of cotton production in Burkina Faso is likely to make this information more reliable than in other contexts.[13] Questionnaire were administered with tablets and simultaneous checks (and corrections) performed. Our sample is composed of 1,015 cotton farming households, selected from the lists of 80 GPCs (about 13 households per GPC).[14] Attrition is low: only 4 households (0.4%) were not found in 2015.

The Sofitex dataset and the household level dataset are combined in the analysis: the Sofitex dataset provides a greater number of GPCs and years, but the household level data is needed to compute our farmer-centric measure of index insurance quality.

3.3 Empirical specifications

In order to build an empirical evaluation of the index-insurance contract offered to cotton farmers in Burkina-Faso, we simulate a database of shocks (idiosyncratic and covariate) based on the empirical distributions observed in our historical data.

3.3.1 Estimating shocks

The first step for measuring index-insurance quality empirically consists in estimating the shocks received by farmers.In that analysis, we make two fundamental assumptions: i) we assume that the data used for insurance design are stationary. Hence, shocks at the individual or GPC levels can be defined as deviations from their respective mean yield; ii) we assume multiplicative shocks at the GPC and individual levels. Thus, the cotton yield obtained by farmer$i$ in GPCv at time $t(y_{ivt})$ can be decomposed into a product of the GPC-level average yield in GPC v (y_v,), the relative productivity of farmer i compared to the average farmer of his GPC v (u_{vi},) and the GPC-level and idiosyncratic shocks (e_{vt} and e_{ivt}):

$$y_{ivt} = y_v \times e_{vt} \times u_{vi} \times e_{ivt} \tag{5}$$

If we had long enough (stationary) time series for each farmer, we could compute the components of equation (5) as deviations around the appropriate empirical means. Unfortunately, we only have eight years of GPC-level data from Sofitex (2006-2013) and seven years of individual yield data (2008-2014). To overcome this limitation, we develop a two-stage estimation strategy that builds on the panel structure of our data. First, we exploit the panel structure of the GPC-level yield data (y_{vt}) obtained from Sofitex and estimate equation (6) where we model the GPC-level average yield y_v as a fixed effect for each GPC:

$$y_{vt} = y_v \times e_{vt} \tag{6}$$

Second, we exploit our individual-level data and define an idiosyncratic shock for a given year as a deviation from the GPC-level yield for this year that is not explained by the individual's normal deviation from her GPC's average yield. If we had a perfect correlation between the GPC-level yields and the individual-level yields, we could merge our survey data with the data obtained from Sofitex, and the GPC-year specific average computed using survey data would be equal to the Sofitex recorded yield. However, Figure (4) shows that survey (recall) data tend to overestimate the average yield in bad years and underestimate it in good years.

[12] In 2014-15, 40 of the 80 GPCs have been offered the insurance contract, and 17 have purchased it. However, there was no impact on cotton production, most likely because the insurance was sold too late in the season.

[13] Indeed, surface is planned in advance, inputs are requested accordingly, and production is weighted and paid at a relatively stable (and national) price. Some records (although not always accurate) exist of these transactions and some households used them to support their memory.

[14] About 95% of the households of the research area are cotton farmers, i.e. are members of a GPC.

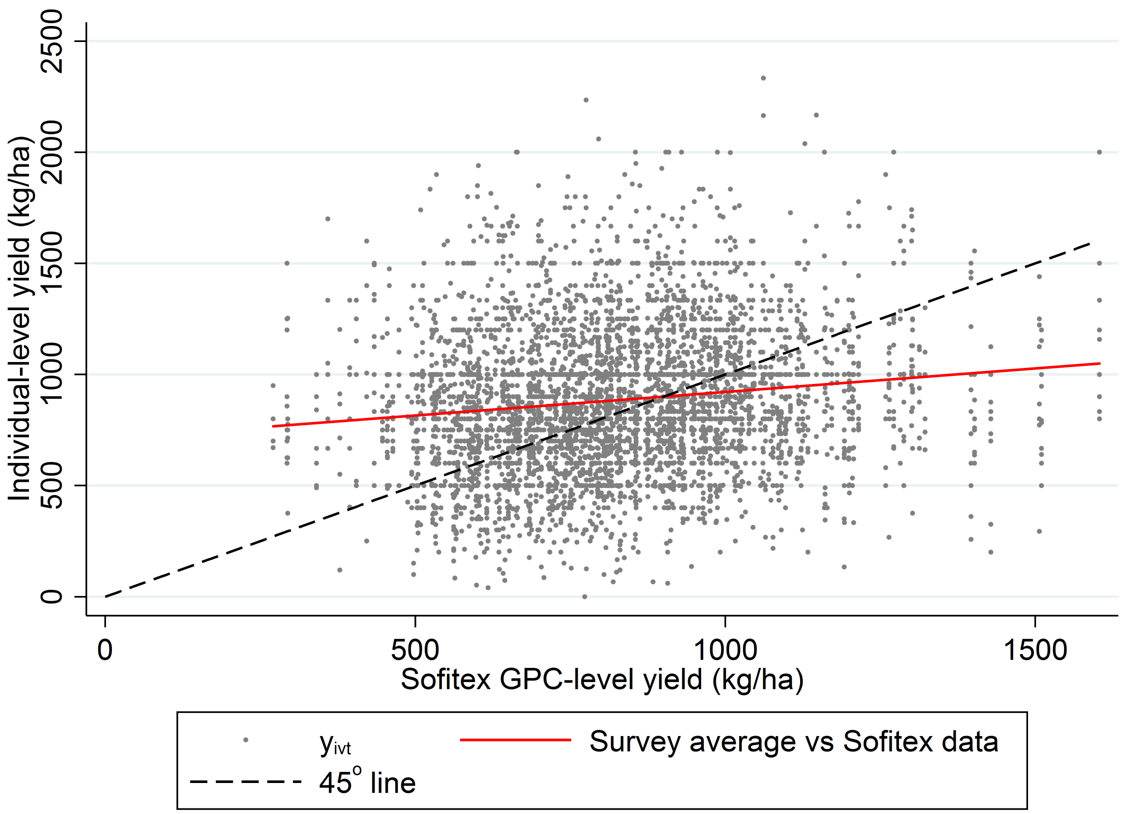

Figure 4: Sampling and memory bias in recall data yield

We account for this bias in our analysis by controlling for the severity of the covariate shock as well as the number of years elapsed between the harvest and the survey. Hence, we estimate idiosyncratic shocks as the residuals e_{ivt} of equation (7):

$$ln(y_{ivt}/y_{vt}) = Z_{ivt} + ln(u_{iv}) + ln(e_{ivt}) \tag{7}$$

where Z_{ivt} is a set of control variables for memory bias.

Based on our estimation results presented in Figure 5, it appears that idiosyncratic shocks and covariate shocks have very similar distributions, with a slightly larger proportion of extreme negative idiosyncratic shocks. Hence, we expect that a perfect insurance contract could not cover the full production risk for these cotton producers.

3.3.2 Shock simulations

In the empirical quality assessment, we use the estimated distributions of covariate and idiosyncratic shocks to simulate 1,000 covariate shocks and 1,000 idiosyncratic shocks, creating a full database of a million observations. These simulations, based on the observed group-level and individual-level yields, are generated to increase the size of the historical sample of shocks. Indeed, by definition, extreme events are rare. For that reason, performing non-parametric estimations and measuring index-insurance quality from simulations is more straightforward than conducting these analyses with the historical data. We use these simulated data to analyze the quality of the area-yield index product offered to Burkinabe farmers.

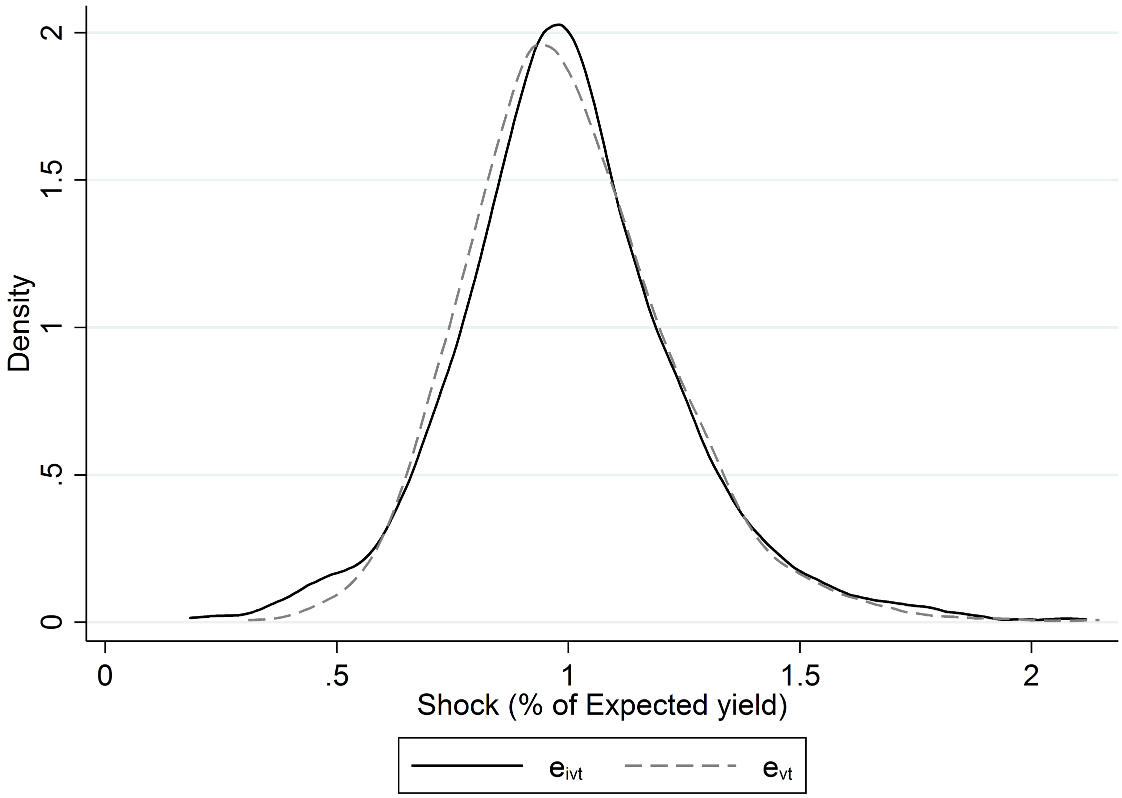

Figure 5: Empirical distribution functions of shocks

4 Results: index insurance quality in Burkina Faso

This section analyzes the quality of the index insurance contract offered to Burkine farmers, based on the empirically simulated yield data. This contract has five characteristics that we want to evaluate here:

1) It is an area-yield index at the GPC level that cannot cover idiosyncratic risk.

2) GPCs with similar average yield are grouped together and offered the same triggers and premium.

3) Indemnities are paid only when two thresholds are triggered simultaneously: one at the GPC level and one at the neighborhood level

4) A commercial mark-up is added to the Actuarially Fair Premium

5) There exist three levels of indemnity payments (see section 3.1)

We investigate the quality of a contract including these characteristics by comparison against a perfect contract covering the individual farmer against all types of shocks. This perfect contract pays indemnities such that the farmer's income never falls below 75% of his expected income. The actuarially fair price of such contract equals 5% of the expected income, so that insurance payouts are perfectly equivalent to the loss experienced when yields fall below 80%. In order to simplify comparisons between contracts, indemnities of all the contracts that follow are adjusted so that they all have the same price (which is the actuarially fair price of the perfect contract). The next subsection assesses the quality of a simple area-yield contract defined at the GPC level only. The following subsections show how each characteristic of the commercial contract (i.e. departure from the perfect area-yield contract) impact index insurance quality. Results for different degrees of sensitivity to shocks (α) are summarized in

Table 2.

	No insurance	Perfect	Area-yield	Mean categories	Markup	Double trigger	Lump-sum	All features
	0	1	0.456	0.128	0.456	0.364	0.501	0.501
BRR	0	3.551	2.99	2.099	0.996	2.99	2.941	0.98
$NISS_1$	0	1	0.235	0.234	-0.003	0.142	0.23	-0.017
$NISS_2$	0	1	0.251	0.248	0.043	0.137	0.251	0.022
$NISS_3$	0	1	0.237	0.234	0.052	0.109	0.241	0.029
$NISS_4$	0	1	0.204	0.202	0.033	0.066	0.213	0.012

Table 2: Result summary (Normalized Quality Gap)

4.1 Standard area-yield contract

We first assess the value of a simple area-yield index contract against the perfect contract. This area-yield contract is designed to perfectly cover covariate shocks at the GPC level. We have seen above that a large share of the risk is idiosyncratic in nature. Hence, it should not be expected from an area-yield contract to fulfil more than 50% of our income stabilization objective. This somewhat pessimistic observation can be seen in Figure 6: the grey short-dash line corresponds to the case without insurance; the solid black line represents the perfect contract which stabilizes income at 0.75 of average income; and the blue long-dash line shows the average performance of a perfect area-yield contract . The dotted line draws the probability density function of shocks.

Figure 6 shows that on average an area-yield contract tends to pay too often (or too much) in the neighborhood of normal production conditions, thus creating some "false positive" events. On the other hand (and arguably more importantly), it fails to completely cover losses in the case of extreme events. This mismatch between farmers' income stablization needs and insurance payments results in an area-yield insurance contract that can only fulfill 23.5% of the income stabilization objective of an individual with linear loss aversion (α=1). In the case of moderately loss averse individuals (α=2), this area-yield insurance contract has a little more value, reaching a quality measure of 25.2%. However, this area-yield index loses some of its value when we tend towards farmers with high loss aversion (α>2). Indeed, for these farmers, facing a shock without receiving indemnities is extremely painful, and offsets the gains from the indemnities received.

4.2 Grouping GPCs by average yield category

The commercial contract offered to Burkinabe farmers was designed so that the insurance premium, indemnity levels and probabilities of payment are the same for every GPC among each of the four categories (see section 3.1). Hence, it would have been necessary to compute GPC-specific thresholds to achieve this objective. However, having a specific threshold for each GPC was considered unrealistic for commercial purposes, and it was decided to group GPCs in four categories, based on their historical average yields. This grouping necessarily creates some inequalities between GPCs (and consequently, between farmers). Indeed, if the GPC's true average yield is below that of its category, the insurance will pay more often, because normal yields may be considered as shocks by the insurance. Consequently, the GPC's expected indemnity can be even larger than the premium paid. On the opposite, a GPC with average yields higher than the average of its category has a very low probability to receive a payment because the thresholds are too low compared to its historical yields. We analyze the importance of this grouping in a given category by looking at the average quality of the index inside a category of GPCs. We use a k-means clustering approach to group our GPCs in different categories based on their average yields.[15] We follow Makles [2012] and compute the optimal k-means clustering by increasing

[15]In the actual contract, GPCs were arbitrarily separated in four categories, but the final contract includes a fifth category where GPCs with high "burn rate" are placed (in an ad hoc manner). We proceed more systematically by

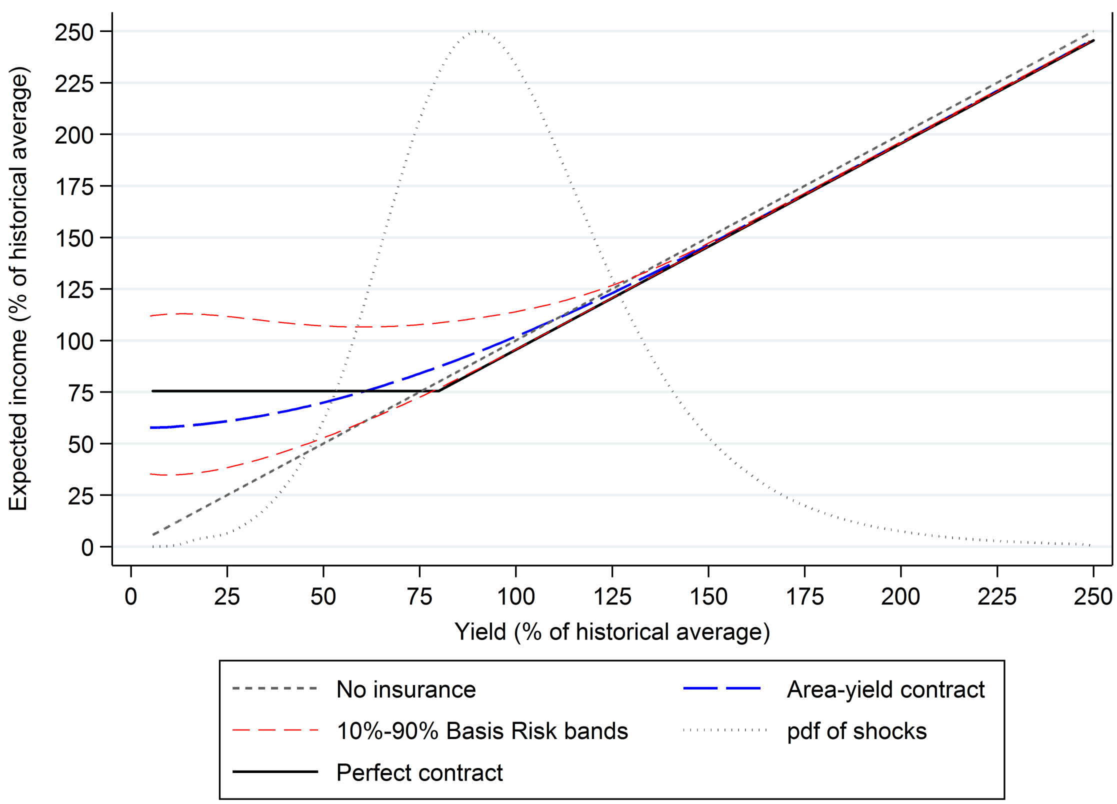

Figure 6: Simple area-yield contract

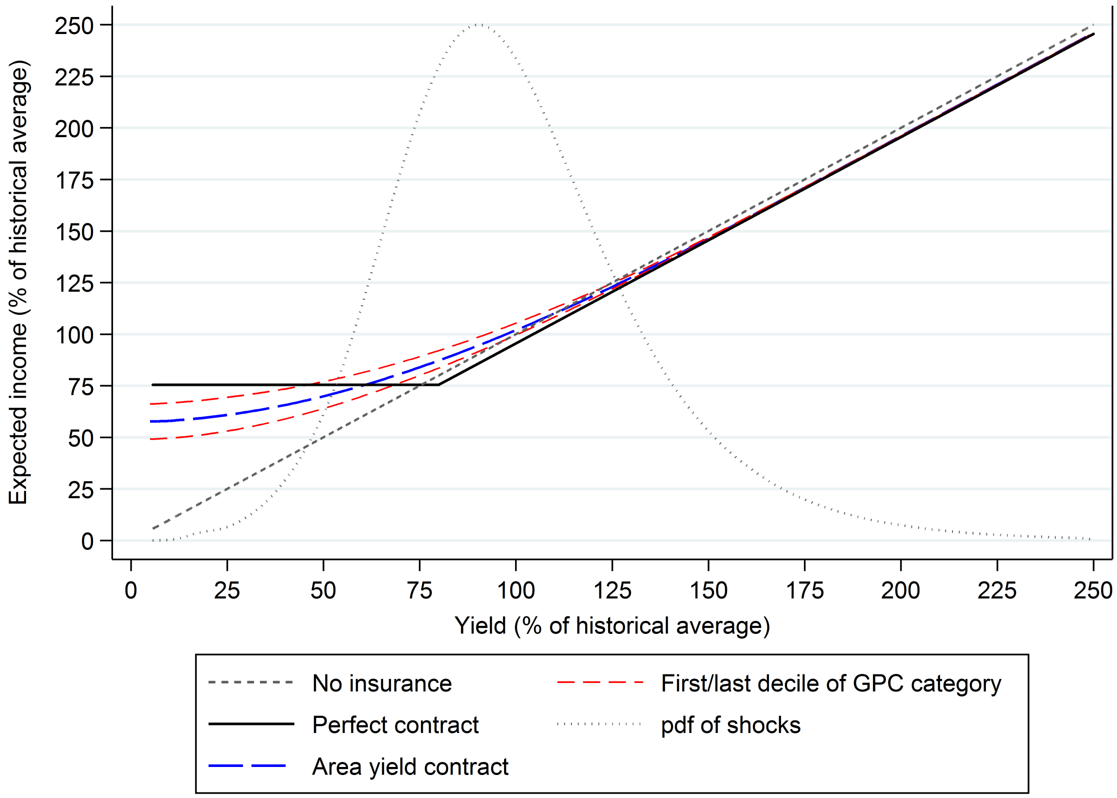

Figure 7: Grouping by GPC category

the number of clusters until the grouping can explain at least 90% of the total variance observed or adding one more group does not significantly increase the share of explained variance. In doing so, we create five categories, in which a GPC can be more or less productive than its category average. In each category however, a GPC has less than 1% chance to have an average yield lower than 80% or higher than 120% of its category's average yield. This grouping appears to have a limited impact on the quality of the insurance contract. Figure 7 shows that within a category, low performing GPCs can benefit from the high likelihood of getting a payout, but the gains for these low performing GPCs (first decile) are compensated by the losses for the high performing GPCs (last decile) who would pay the premium every year but rarely receive a payout. While the overall quality measure of the area yield index insurance decreases slightly when GPCs are grouped, it must be kept in mind that some important differences between GPCs are hidden: for $\alpha = 1$, the quality index increases to .42 for the first decile of the GPC mean distribution while it drops to .04 for the last decile (compared to an average quality measure of 0.234).

4.3 Double-trigger

The actual Burkina Faso contract employs a double-trigger design to address moral hazard concerns. Indeed, a GPC is a fairly small group of farmers who could coordinate to reduce their yields and receive undue indemnities if the index was solely based on the GPC's yield records. This potential manipulation of the index is avoided in the insurance contract offered to cotton farmers in Burkina

employing k-means clustering.

Faso where triggering the indemnity threshold at the GPC level is not enough to receive indemnities: a second threshold was set at the "neighborhood" level, where a neighborhood is a group of GPCs that are considered close enough to experience the same climatic conditions (see 3.1). Constructing retrospective payouts from GPC-level data suggests that the double trigger cancels out 20% of the indemnity payments triggered at the GPC level. Introducing this additional probability of failure of the index reduces the quality of the contract significantly to 0.142 for farmers with linear loss aversion (=1) and even further for other types of farmers. However, it must be kept in mind that this double-trigger design is actually an improvement compared to the other commercially viable option based on a single trigger defined only at the neighborhood level [Elabed et al., 2013]. Indeed, insurance companies would refuse a contract solely based on GPC data because of the risk of index manipulation.

4.4 Lump-sum payments

In a perfect contract, insurance payments would always perfectly match farmers' losses so that indemnity payments would be a linear function of predicted losses. In practice, it is often simpler to implement a lump-sum contract. The contract offered to Burkinabe farmers was set up so that farmers would have a 20% chance to get an indemnity at least as high as their insurance premium (11,200 FCFA), 8% chance to receive at least 34,000 FCFA and a 4% chance to receive a high payment of 90,000 FCFA (see 3.1). This payment structure could influence the quality of an index contract because of the mismatch between losses and indemnities that it creates. According to our analysis however, this lump-sum payment structure does not dramatically influence the quality of the index insurance.

4.5 Commercial mark-up

An insurance project always involves several partners (calculating agents, insurance company, reinsurance company, brokers, etc.) generating additional costs beyond the pure price of risk and the fixed cost of developing the index (which is often covered by public funding). In the case of the cotton project in Burkina-Faso, these additional costs increased the insurance premium threefold compared to the actuarially fair price. Our simulations indicate that such high mark-up causes the potential welfare gains of index insurance to vanish, for any level of loss aversion. This observation raises serious doubts regarding the viability of index insurance when priced at a commercial rate. Finding ways to cut these additional costs seems necessary for achieving the objective of keeping index insurance valuable for farmers.[16]

4.6 Combining the different features of the cotton index insurance contract

When we look at the quality of the final product including the different features detailed above, we observe that the quality measure of our index contract is close to zero or even welfare reducing for farmers with linear loss aversion ($\alpha = 1$). We showed above that the two features that heavily impact the quality of the proposed contract are the double-trigger design and the high commercial mark-up put on the premium. It appears that the most important feature is- by far- the commercial price of the contract which multiplies the actuarially fair premium threefold. We computed the maximum mark-up that can be applied to this contract to make farmers indifferent between buying the insurance or not. We found that for farmers with linear loss aversion ($\alpha = 1$), the maximum mark-up equals 2.67. While reasonably high, this maximum mark-up is still much lower than the mark-up applied by the commercial partners of the Burkina-Faso cotton program, resulting in a low demand in unsubsidized groups.

Indeed, the predictions of our analyses have been confirmed in the field where demand for this contract was close to zero when the contract was not highly subsidized.

[16] While public subsidies are often used to keep premiums lower than their commercial price, this solution does not necessarily address the root causes of the high prices. A policy option considered to avoid the "uncertainty loading" applied by reinsurance companies is to set up a public reinsurance facility [Carter, 2013].

4.7 Normative minimum standard

The quality measures presented so far allow us to rank several potential insurance contracts and make sure that the best of them is not welfare reducing. But a contract that does not hurt farmers is still far from a contract that can help them manage risk. In this section, we introduce a normative minimum standard that can be applied to index insurance contracts. We propose to set a threshold below which insured farmers' income cannot fall (on average) to keep these farmers financially safe. The idea is that on average, the index must at least guaranty that farmers can repay their input expenses.[17] In our data, the standard input package for 1ha of cotton costs 90,000 FCFA, the price of cotton is about 935 FCFA per kg and the average yield is 900kg of cotton per ha. Hence, the average farmer needs at least 43% of his expected yield to repay his input loan. We want to make sure that when yields are below 43% of expected yields, insurance indemnities compensate so that on average farmers' income does not fall below 43% of the expected income. Formally, we compute the expected income with insurance given that yields are below 43% of expected yields. A contract that fails to guarantee such minimal level of income could not be considered as a valuable financial protection tool since insured farmers often be left in the incapacity to repay their input loan.

We apply this criterion to the different contracts presented above and, if a particular contract passes this minimum standard requirement at its actuarially fair price, we compute the maximum mark-up that can be charged so that the willingness-to-pay remains positive (i.e. the product is not welfare reducing) and the minimum standard criterion is fulfilled.

Our empirical results show that while a simple area yield contract defined at the GPC level and priced at its actuarially fair price would pass the minimum requirement test, increasing the expected income of insured farmers to 58% of expected income when yields fall below 43% of expected yield. The commercial contract offered to Burkinabe farmers (combining the double-trigger, mean groups and lump-sum payments) also performs well when priced at the actuarially fair price, offering an expected income equal to 53% of expected income when yields fall below 43% of expected yields. The drop in quality appears to be mainly due to the lump-sum structure of indemnity payments.

Unfortunately, the high markup charged on this contract reduces the expected income with insurance to 38% of the expected income without insurance when yields fall below 43% of expected yields. This is only marginally better than without insurance in which case the expected income given that yields are below 43% of expected yields equals 34% of the expected income. According to our estimations, the markup cannot be higher than 1.8 if we want to ensure that the commercial contract does not fall below the minimum standard criterion. In comparison, the simple area-yield contract could support a maximum markup of 2.15.

5 Conclusion

This article demonstrates in different way the pressing need for rigorous ex-ante assessment of index-insurance products, and suggest simple measures to conduct that assessment. Index-insurances are complex products. Many factors affect their quality, in particular index prediction quality, idiosyncratic risk and price. Measuring insurance quality requires to take into account several dimensions: severity of the loss, frequency of the events, individual outcomes (as opposed to average outcomes) and insurance payouts compared to insurance premium. To account for these dimensions, we developed a farmer-centric approach, using as a starting point the stabilization target of the insurance product. Using simple simulations, we show how a measure based on that approach allows us to distinguish good and bad contracts, and to rank contracts when none of them is, apparently, unambiguously superior.

Applying this measure to data from Burkina Faso, we show that even a well-designed, area-yield insurance contract present serious quality issues. There, the lack of quality is mostly due to high idiosyncratic risk, and at commercial price, to very high premiums. Because the contract functions well

[17] This is particularly relevant in our context, since not repaying a loan is associated with high social costs within the joint-liability credit group [Gelade, 2015].

at the group level, these results emphasize the need for index-insurance assessment at the individual level. This requires appropriate measures of quality as well as individual level data. Focusing on insurance value for individual farmers is also an opportunity to move from a debate on the exact definition of basis risk towards a focus on insurance quality more consistent with the development objectives of index-insurance.

These findings have three main implications for research and policy partners involved in promoting index-insurance in developing countries. First, the results suggest that some index-insurance products should not be sold to farmers. Indeed, if products do not meet minimum quality standards, they do not make farmers better-off. This is obvious from the simulations which we conducted, as the quality measure is worse for the bad index-insurance contract than for the "no-insurance" contract. Because the commercial premium is also very high- and because of high idiosyncratic risk- the Burkina Faso index-insurance contract is also worse than the absence of insurance when sold at commercial price.

The second implication regards the issue of contract design, selection and pricing. Different contract options are usually considered during the design phase. These options have to be assessed at the individual level, by carefully measuring the relationship between farmer losses and insurance payouts. The measures suggested in this article offer a simple way to conduct such comparisons and to select the best contract among those available, as they indicate which trade-offs (e.g.: more costly area-yield vs. more affordable rainfall index) make farmers better-off. Finally, the measures developed here allow us to observe the impact of the high loading factors applied by insurance and re-insurance companies, and to know when products should not be sold to farmers if the resulting premium generates an unacceptable level of quality gap.

The third implication regards the index-insurance research agenda. While program assessment is increasingly conducted through rigorous impact evaluations, this paper shows that greater attention should be paid to program quality in general, and in particular for complex interventions such as index-insurance projects. Some aspects of program quality, such as the core economic value of index-insurance for farmers, can- and should- be assessed ex-ante. If program quality is too low, an impact should not be expected- and if an impact is found nonetheless, it may actually put farmers in dangerous situations (unknown to them). For index-insurance, ex-ante assessments imply an investment in data collection in order to obtain long time series, required to rigorously assess index-insurance product in time.

Altogether, these three implications mean that more effort has to be put into product design, selection and assessment. This outlines a research and policy agenda which, only then, will allow us to know whether or not index-insurance has the actual potential to protect farmers and improve their well-being.

References

Barry J Barnett, Christopher B Barrett, and Jerry R Skees. Poverty traps and index-based risk transfer products. *World Development*, 36(10):1766–1785, 2008.

Alexis Berg, Philippe Quirion, and Benjamin Sultan. Weather-index drought insurance in burkina-faso: assessment of its potential interest to farmers. *Weather, Climate, and Society*, 1(1):71–84, 2009.

Hans P Binswanger-Mkhize. Is there too much hype about index-based agricultural insurance? *Journal of Development Studies*, 48(2):187–200, 2012.

Kate Bird and Martin Prowse. Vulnerability, poverty and coping in zimbabwe. *Chronic Poverty Research Centre Working Paper*, (136), 2009.

Elabed Ghada Carter, Michael and Elena Serfilippi. Behavioral economic insights on index insurance design. *Agricultural Finance Review*, 75(1):8–18, 2015.

Michael Carter. Sharing the risk and the uncertainty: Public-private reinsurance partnerships for viable agricultural insurance markets. *I4 Index Insurance Innovation Initiative Brief*, 1, 2013.

Michael R Carter. Environment, technology, and the social articulation of risk in west african agriculture. *Economic development and cultural change*, 45(3):557–590, 1997.

Michael R Carter and Travis J Lybbert. Consumption versus asset smoothing: testing the implications of poverty trap theory in burkina faso. *Journal of Development Economics*, 99(2):255–264, 2012.

Sommarat Chantarat, Andrew G Mude, Christopher B Barrett, and Michael R Carter. Designing index-based livestock insurance for managing asset risk in northern kenya. *Journal of Risk and Insurance*, 80(1):205–237, 2013.

Daniel Clarke. *A theory of rational demand for index insurance*. Department of Economics, University of Oxford, 2011.

Daniel Clarke, Olivier Mahul, Kolli N Rao, and Niraj Verma. Weather based crop insurance in india. *World Bank Policy Research Working Paper*, (5985), 2012.

Shawn Cole, Xavier Giné, Jeremy Tobacman, Robert Townsend, Petia Topalova, and James Vickery. Barriers to household risk management: evidence from india. *American economic journal. Applied economics*, 5(1):104, 2013.

Greg Collins. The coping strategies index (csi) baseline survey: World food programme (wfp) assisted refugees in western tanzania. Technical report, WFP Tanzania, 2004.

Ombeline De Bock and Wouter Gelade. The demand for microinsurance: A literature review. *Microinsurance Innovation Facility, Research Paper*, (26), 2012.

Stefan Dercon. Growth and shocks: evidence from rural ethiopia. *Journal of Development Economics*, 74(2):309–329, 2004.

Ghada Elabed and Michael Carter. Ex-ante impacts of agricultural insurance: Evidence from a field experiment in mali. 2015.

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

Marcel Fafchamps. *Rural poverty, risk and development*. Edward Elgar Publishing, 2003.

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

Michael Flatnes, Jon Einar Carter. Fail-safe index insurance without the cost: a satellite based conditional audit approach. Working paper, 2015.

James Foster, Joel Greer, and Erik Thorbecke. A class of decomposable poverty measures. *Econometrica: Journal of the Econometric Society*, pages 761–766, 1984.

John Hoddinott and Bill Kinsey. Child growth in the time of drought. *Oxford Bulletin of Economics and statistics*, 63(4):409–436, 2001.

Sarah A Janzen and Michael R Carter. After the drought: The impact of microinsurance on consumption smoothing and asset protection. Technical report, National Bureau of Economic Research, 2013.

Nathaniel Duane Jensen, Christopher B Barrett, and Andrew Mude. Index insurance and cash transfers: A comparative analysis from northern kenya. *Available at SSRN 2547660*, 2014a.

Nathaniel Duane Jensen, Andrew Mude, and Christopher B Barrett. How basis risk and spatiotemporal adverse selection influence demand for index insurance: Evidence from northern kenya. *Available at SSRN 2475187*, 2014b.

Dean Karlan and Jonathan Morduch. *Handbook of Development Economics*, volume 5, chapter Access to Finance, page Chapter 2. Amsterdam: NorthHolland, 2010.

Harounan Kazianga and Christopher Udry. Consumption smoothing? livestock, insurance and drought in rural burkina faso. *Journal of Development Economics*, 79(2):413–446, 2006.

Anna Makles. Stata tip 110: How to get the optimal k-means cluster solution. *Stata Journal*, 12(2): 347–351, 2012.

Takeshi Sakurai and Thomas Reardon. Potential demand for drought insurance in burkina faso and its determinants. *American Journal of Agricultural Economics*, 79(4):1193–1207, 1997.

Quentin Stoeffler. Crop portfolio choices in burkina faso. *Journal of developing areas*, 2015.

Paul Winters, Lina Salazar, and Alessandro Maffioli. Designing impact evaluations for agricultural projects. Technical report, Inter-American Development Bank, Washington, DC, 2010.

World Bank. World development report 2014: Risk and opportunity-managing risk for development, 2014.

Frederick J Zimmerman and Michael R Carter. Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints. *Journal of Development Economics*, 71(2):233–260, 2003.