Feasibility Study on Agricultural Index Insurance in Nepal:

Preliminary Final Report

Thomas Barre, Michael Carter & Jisang Yu

University of California, Davis

October 2015

INTRODUCTION

The objective of the feasibility study came in two parts:

- Determine where (which activities/regions/etc.) the development impacts of insurance-based risk management strategies would be maximized; and,
- Assess whether there is appropriate and available data that is correlated with the identified agricultural outcomes (based on identified activities/regions/etc.) such that it can be used to create an effective index for an insurance product.

In the initial concept note for this feasibility study, I4/BASIS presented the following diagram to explain the approach we would take to explore the feasibility of an agricultural index insurance contract that would have the potential for development impact. The vertical axis of the diagram measures the potential development impacts of region/activity. The horizontal axis measures the quality of the index insurance that can be designed.

Figure 1: The "Sweet Spot"

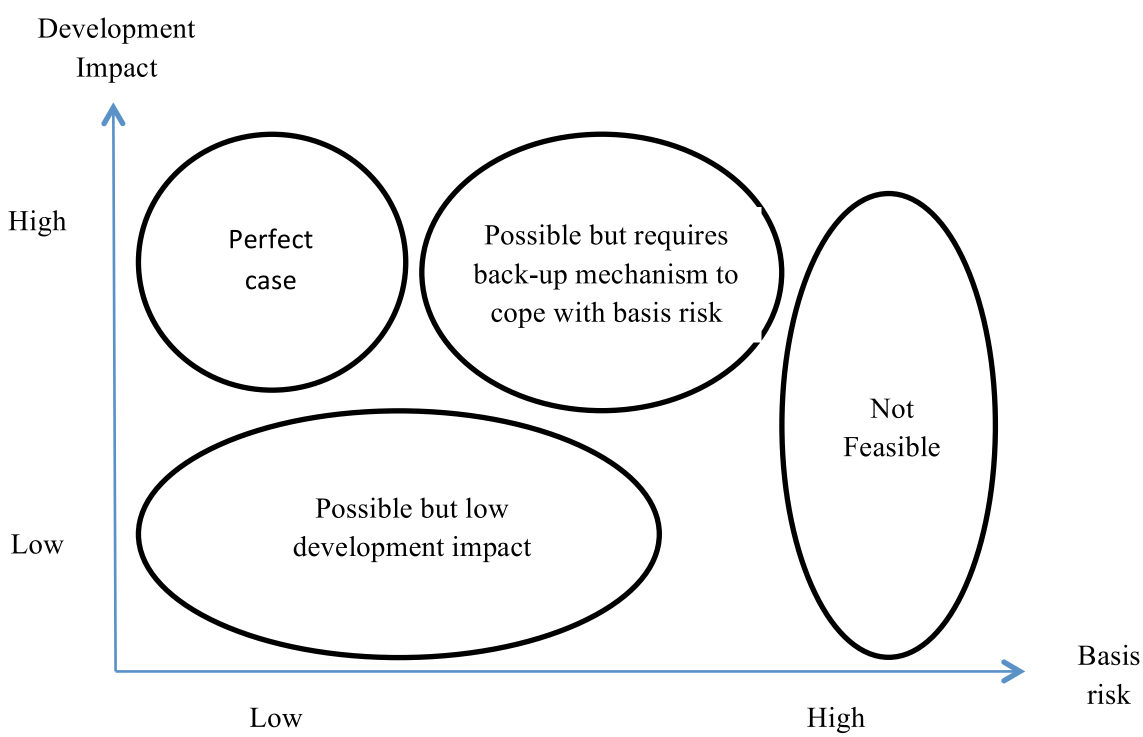

Our goal was to identify geographic areas and commodities that have both high potential for significant development impacts, as well as areas where there is data that can be effectively used to predict farmer outcomes and design an effective index insurance contract. In Stage 1 of the feasibility study we looked across a broad variety of commodities and areas and create a short list of those commodities/areas where insurance can have large impacts by crowding in new investment and prudential risk-taking by small-scale farms. Based on our initial work under this activity, we determined that insurance for rice in the Terai had the greatest potential for landing us the "perfect case" sweet spot in the upper left corner of the diagram.

Since that time, we have devoted ourselves to seeing if we can design an index insurance contract of sufficient quality (low basis risk) to get us into that "sweet spot" region. In Stage 2, we have closely examined rice in the Terai to determine whether an effective index insurance contract can be designed. Unfortunately, our work to date on contract design and contract quality measurement returns a pessimistic conclusion on the feasibility of an effective index insurance contract. In visual terms we appear to be well to the right of the perfect case "sweet spot".

The remainder of this report is organized into the following sections:

1. Finding the "Sweet Spot": Lessons from Tanzania
2. Data Collection
3. Testing and Design of the Index for an Insurance Product
4. Next Steps.

SECTION 1: FINDING THE "SWEET SPOT"

Lessons from Tanzania

In Tanzania, I4/BASIS researchers, in partnership with a leading geospatial software engineering firm (Vencore), developed a potential solution to the traditional challenges of index insurance design by using satellite information at a high resolution, coupled with better modeling approaches. This approach has potential to reduce basis risk and improve index insurance performance. The following were essential for these advances:

1. High spatial resolution: With individual resolution of 500 meters by 500 meters, readings could be fine-tuned to improve estimates of yields.
2. Satellite measures: These sophisticated measures allow for tracking of plant growth and better estimation of yields. More information on the different measures available is available in Appendix A.
3. Crop masking: Researchers and partners developed a crop-masking model that can differentiate within the satellite images what is the intended insured crop (rice) vs. other non-crop vegetation.
4. Planting date detection algorithm: This algorithm can detect when a field is flooded for planting. This helps to identify the "start" date and then to identify the best "end" date for the contract. In practice, the planting date detection did not add significant value to the contract.

Figure 2: The Study Area

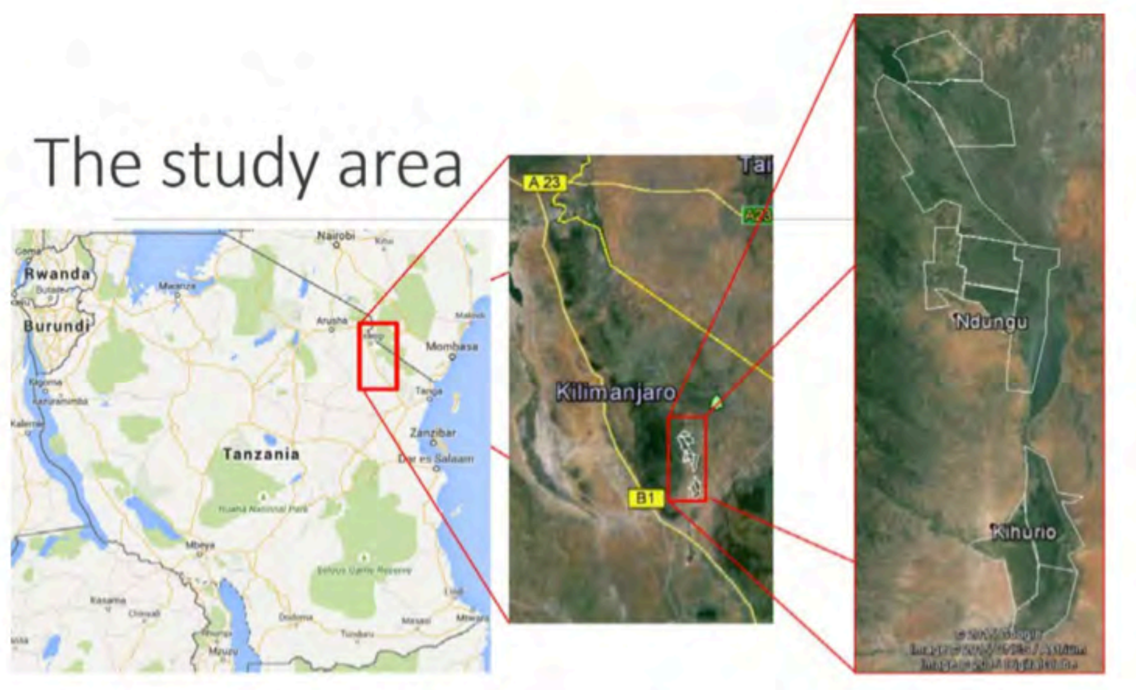

In our experience in Tanzania, these satellite measures were shown to have a strong correlation with yields, such that they had a strong potential to serve as an effective index in an insurance product. This correlation can be seen in the Figure 3.

Figure 3: Predicted vs. Actual Yields in Makindube, Tanzania

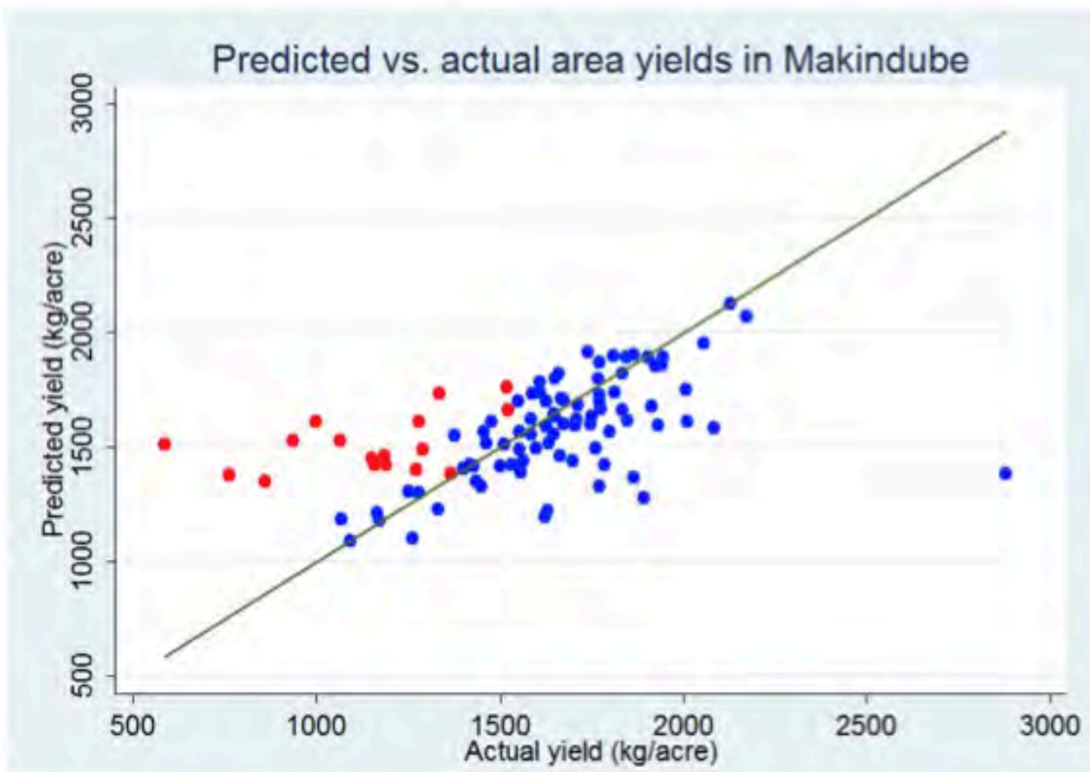

However, in Figure 3 (above), the red dots indicate the cases when the actual farmer losses are below average, so perhaps an insurable event, and the error between predicted yield and actual yield is greater than 5 percent. These red dots are causes for concern. This is because farmers experienced losses, and the extent of these losses were underestimated by 5 percent or more.

From this we see that even with a strongly correlated satellite index, there remains some possibility for basis risk. However, with the addition of a contract clause (Figure 4), which allows farmers to request an on-the-ground audit when the satellite data fails to correctly measure yields, cost-effective and risk-reducing insurance plans can be implemented.

Figure 4: Satellite Index with "Fail Safe" Audit Contract Clause

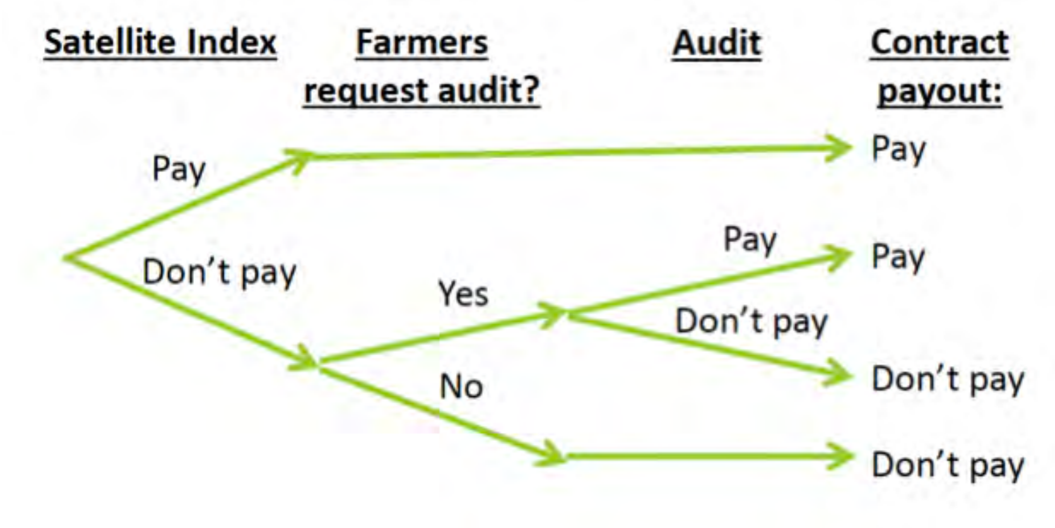

An area yield is statistically the best predictor of individual farmer yields, so this would be the most accurate index for an insurance product and would provide the best protection to farmers. It would, however, be the most expensive option, as well, as adequate yield information can be quite costly to collect. A satellite-based contract would be much cheaper to administer, but is less accurate in predicting the outcomes of individual farmers, so offers less protection to the farmer.

The satellite-based index with a "fail safe" audit clause in the contract provides a third option between area yield and satellite, both in terms of expense of contract administration and with regard to accuracy (with regard to paying farmers when they actually experience an insurable loss). Based on our experience in Tanzania, a farmer-requested audit would occur an estimated 17 percent of the time. This provides an effective option between an area yield index and a satellite-based index alone, both in terms of expense of the contract and in terms of overall protection for the farmer.

There is also evidence that a satellite contract with this "fail safe" audit option will be valued by farmers, as can be seen in Figure 5, below. This figure indicates that farmers prefer a satellite-based contract with the "fail-safe" audit clause over the other options.

Figure 5: Demand for Different Index Insurance Contracts in Tanzania

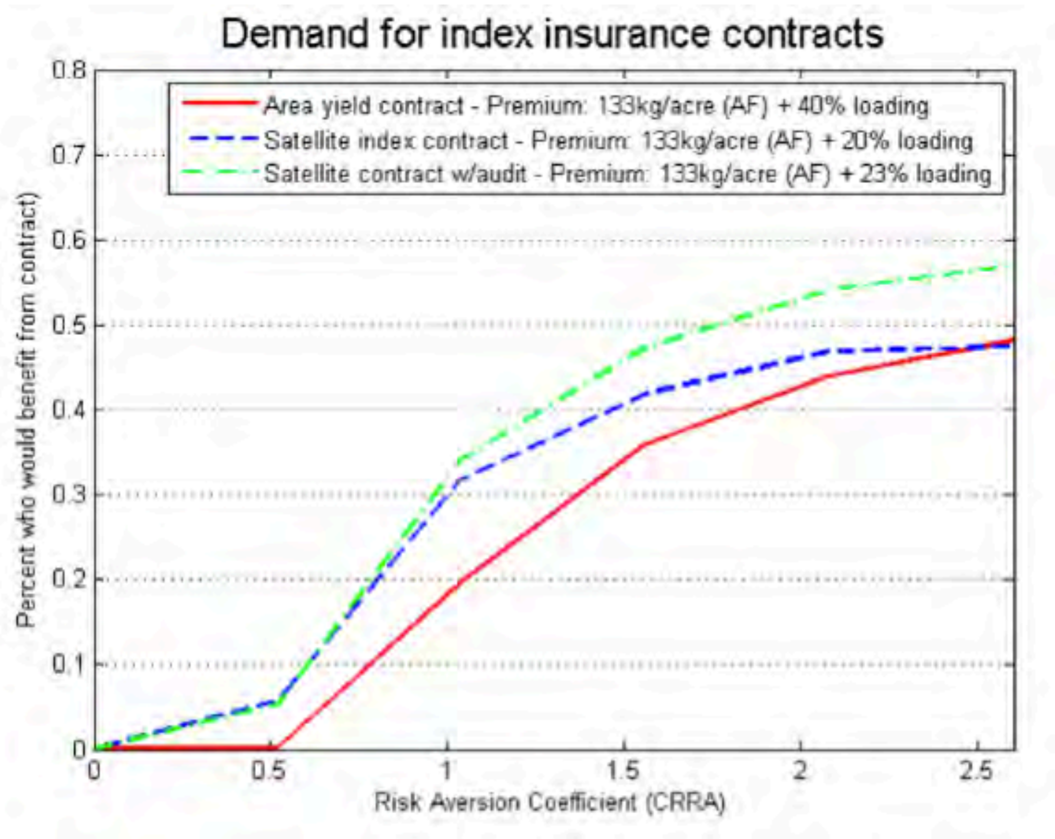

By reaching out to Vencore to access unique expertise, I4/BASIS researchers were able to use the algorithms produced from this partnership to implement fruitful innovations, lowering insurance costs and creating better protections for at-risk farmers. This work from Tanzania helped significantly in the development of a more accurate index for use with an insurance product. This "fail-safe" contract design encourages insured famers to confidently invest in more productive agricultural practices that ultimately create pathways out of poverty.

It was this "sweet spot" of development potential and contract design feasibility that I4/BASIS researchers sought out in Nepal.

SECTION 2: DATA COLLECTION

Historical Rice Yield Survey in Western Terai Region, Nepal

To test the feasibility of replicating the success in Tanzania, I4/BASIS researchers first had to gather gather disaggregated data from household surveys, village/district statistics, processing industry, etc. The objective is to build a dataset of historical yield that covers a large number of villages and many years. This dataset is of particular importance because it was used to compare the proposed indices with historical yield. Because high quality disaggregated and geo-referenced data is essential for the development of the index, I4/BASIS researchers had to collect recall yield data with a particular attention to the losses experienced during very bad years.

This field survey was conducted because existing data, including the Agricultural Census and Nepal Living Standard Survey, do not provide enough rice yield data (especially time-series). BASIS researchers designed a recall survey to gather historical rice yield data for the rice farmers in Western Terai region.

In March 2015, BASIS researchers worked with a local survey firm (IDA: Interdisciplinary Analaysts) to conduct surveys across 50 VDCs in 8 districts (Banke, Bardiya, Dang, Kailai, Kanchanpur, Kapilbastu, Nawalparasi, Rupandehi). Researchers surveyed 500 farmers, 10 farmers in each VDC (5 farmers in each Tole, 2 Toles in each VDC). For each group of 5 farmers, researchers organized a focus group survey followed by individual interviews of the participants to the focus group. Each group was asked to provide historical information on shocks on rice production from 2001 to 2014.

Then, each individual farmer was asked to provide details about his rice plot size and quantity produced from 2001 to 2014. This survey allowed BASIS researchers to build a database of historical yield of about 560 VDC-level observations over fourteen years (14 years and 50 VDCs with some gaps).

Figure 6: Map of the VDCs visited for the recall data survey

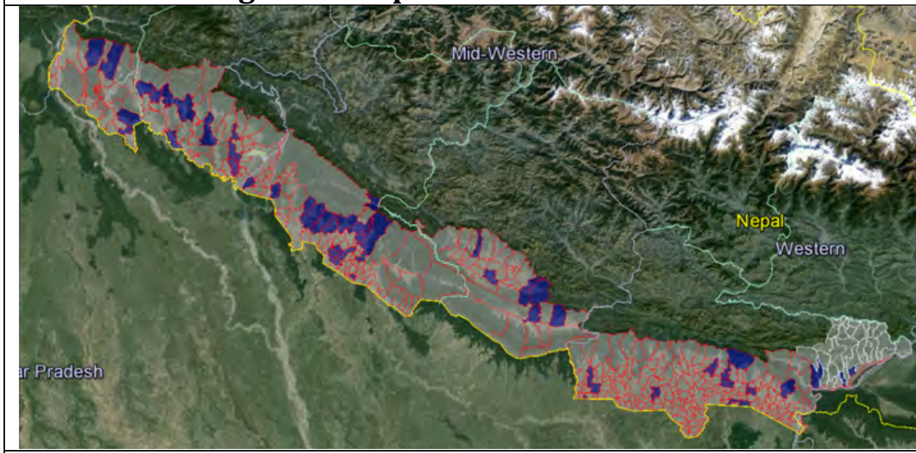

VDCs visited are highlighted in blue

In addition to the rice yield historical data, BASIS researchers also collected the GPS coordinates of the rice cluster where the interviewed farmers cultivate rice. Production inside these rice clusters is expected to be very homogenous because of the collective nature of some of the significant steps in the growing season, such as flooding. Furthermore, because these plots are very close to each other, a drought or flood in the region should impact the farmers of the same cluster in the same way.

Figure 7: Map of VDC

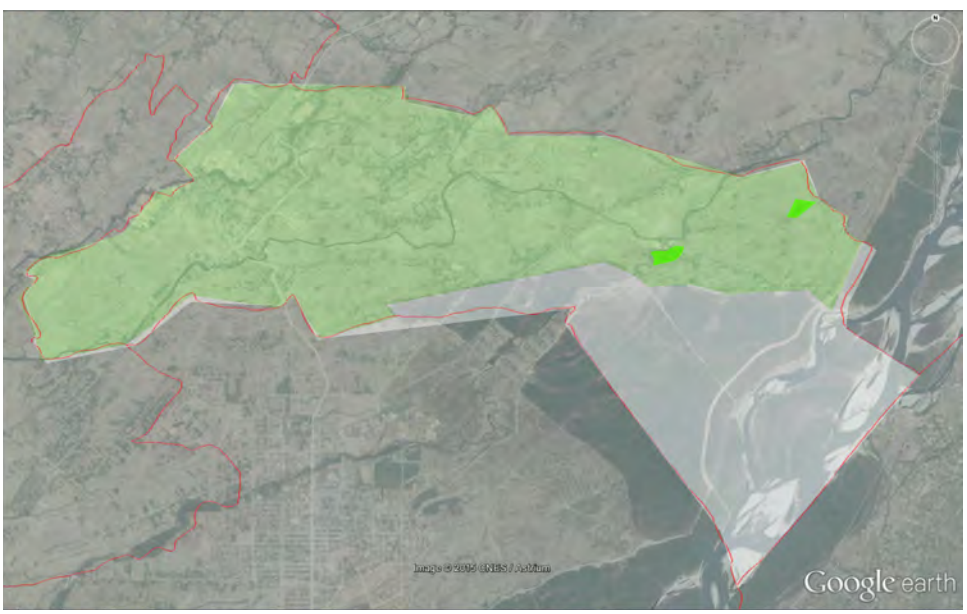

The above image shows a VDC area of about 2600 ha, of which rice fields are shown through crop masking to be approximately 1,750 ha (the light green area). The two brighter green areas highlights "rice clusters" (rice cluster 1 = 4.7 ha, rice cluster 2 = 3.5

ha). The “rice clusters” indicate areas where flooding of the area occurs collectively for all farmers in that area.

Satellite Data

Because the collection of rice yields can be very expensive, BASIS researchers investigated the possibility to substitute an area-yield contract with a satellite-based index. The satellite data used in this study are available for free form NASA's website. We only use the data from MODIS, which provides indices with different time and spatial resolution, going from daily/250m pixels to 16 days/1km depending on the index we look at. The indices we can study are: NDVI, EVI, LAI, LSWI, FPAR and GPP. This assessment is based on the work already done to construct an effective index in Tanzania.

When we try to correlate the satellite data to yield, it is important to only use satellite data from pixels that actually correspond to fields of the crop considered. If we want to study rice yields, pixels corresponding to urban area for example have no value. In the case of Western Terai in Nepal, the monsoon season is such that rice is almost the only cereal grown during the summer. Hence, almost every field can be classified as rice, and we only need to distinguish cropland from other types of cover (forest, city, water). This selection of the relevant pixels is implemented based on the amount of light that is reflected/absorbed throughout the year. Only cropland shows a strong cyclical pattern in the amount of light absorbed/reflected (measured using NDVI) following the cropping seasons. The GoogleEarth image bellow shows the performance of this crop masking in the Far-Western Terai. The dark areas are identified as "non-rice". We can see that the forests are indeed masked out.

Figure 8: Crow Masking in the Far Western Terai

Google Earth image of the Far Western Terai	Crop mask for the Far Western Terai

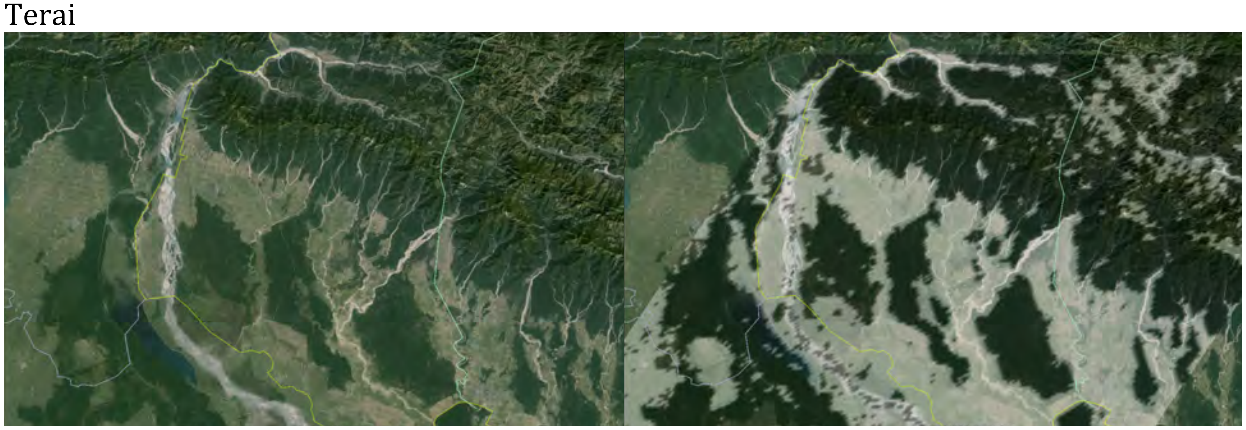

SECTION 3: DESIGN AND TEST OF THE INDEX FOR AN INSURANCE PRODUCT

One of the key challenges to the development of an index insurance product is to assess whether there is available data that can indeed produce an index that effectively predicts average yields. As we've discussed previously, the task of identifying a measure that can be used as an effective index for an insurance product is critical. "Basis risk" or the risk that the product will not accurately trigger when farmers experience a covered loss, is of particular importance because it will determine the value of index insurance for farmers. If the index is only moderately able to measure negative shocks, and therefore does not trigger a payout every time the farmer experiences significant losses from the insured shock, then the farmer will continue costly risk mitigation strategies such as underinvestment in productive technologies.

3.1: WOULD AN AREA YIELD INDEX WORK?

In theory, an area yield-based index insurance product should be best able to predict farmer outcomes, as it is based on actual yield information rather than an outside indicator (such as rainfall or satellite-based measures). To assess the potential for an index-based insurance to work, we began by examining a theoretical area yield-based measure. By assessing this model, we can better understand the upward limit of what an effectively correlated index-based insurance product can offer in terms of protection against risk for the insurance buyers.

To try this, we took the yield data we had leading up to 2014 and fit a relationship between individual yield and the "area yield" upon which an index could be best. Then, using these measures, we estimated theoretical payouts for 2014. This allowed us to compare what would have been paid out using this predictive model to the actual experiences of farmers in 2014.

This model (Figure 9) indicates that, for 2014, an area yield contract would have issued insurance payouts for individual farmers who experienced a loss (a below average year) 66% of the time that individual farmer experienced a loss. For 34% of the time, an individual experiencing a loss is likely experiencing a loss due to some idiosyncratic cause, and they will not receive a payout from the insurance.

Figure 9: Covariate vs. Idiosyncratic Shock

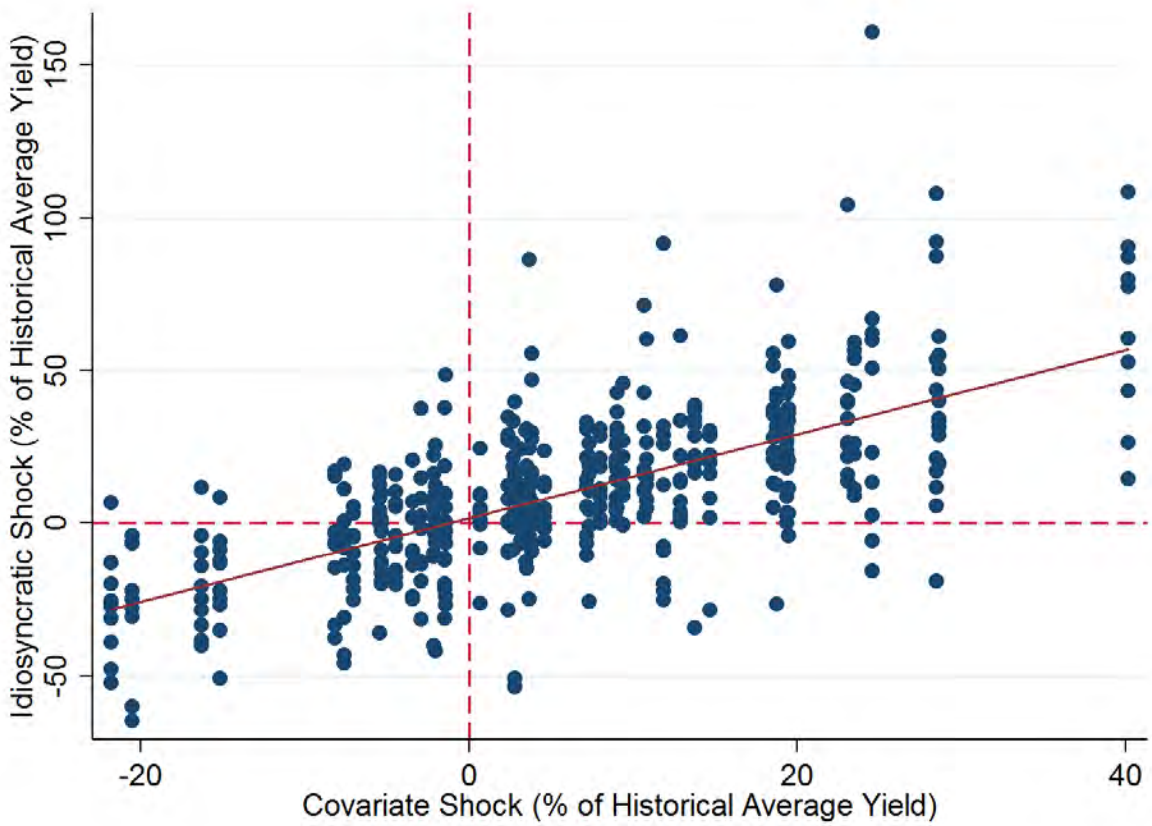

Unfortunately, the amount of yield information required to construct an area-yield based index insurance contract is currently unavailable in Nepal, and would be extremely costly to collect. As such, we turn to other measures, including satellite-based indices, to assess whether these less costly options can produce a reliable index.

If area yield-based index insurance is predicted to correctly payout in case of a loss 66% percent of the time, this is roughly the best we could imagine any index performing. In reality, most outside (non-yield) measures are unlikely to achieve this level of accuracy, as they are only indirectly correlated with yield.

3.2: SELECTION OF A SATELLITE MEASURE

The remote sensing literature and our previous work in Tanzania tell us that GPP (Gross Primary Production) is a good candidate for that purpose. Indeed, while other indices often focus on one attribute of ground conditions (rainfall, temperature, soil moisture) or photosynthesis (NDVI, EVI, LAI), GPP estimates plants biomass production, which is more closely related to grain production. Our previous study in Tanzania indicates that it is indeed a better predictor of yields.

In this study, I4/BASIS tested whether these promising solutions explored originally in Tanzania could accurately predict VDC-level yields in Nepal.

3.3: DETREND THE DATA

The rice yield data we collected show an increase of about 1.2% in rice yield per year on average. It is important to account for this trend in yield in our analysis for two reasons:

1. If we simply look at yield shocks around the overall mean while the data are positively trending, the early years are automatically labeled as "bad years" and the recent years as "good years". Instead, what we want to look at is if the growth in yield in a given year was higher or lower than 1.2% (the normal growth). If the growth in a given year is only .5%, it is a "bad year" because the normal growth should have been 1.2%.

2. It is difficult to say if we expect the satellite data to be trending too or not and in which direction. Consider the example of NDVI or Leaf area Index: the modern rice varieties could be smaller than the traditional varieties, so that the plant spends more energy into grain production rather than in the production of leaves. If we observe fewer leaves, NDVI and LAI decrease, while grain production increases. So we could find a negative correlation between rice yields and these indices. But then it becomes difficult to distinguish the trend in modern varieties adoption (creating a downward slopping trend in NDVI) from a weather shock that hurts plants, making NDVI drop. So we need to detrend the satellite data as well, and focus on shocks around the trend.

3.4: THE RESULTS

For Nepal, we computed the “shocks” (deviation from average yield years) in our VDC-level yield data and compare with the prediction from the satellites. If the satellite data can accurately predict yield, we would see the data points cluster around the 45-degree dotted line, as we did in our experience in Tanzania (page 5, Figure 3). Unfortunately, we observe that our remote sensing data do not predict yield shocks accurately, indicating that the data would not work as an effective index for an insurance product (Figure 10).

Figure 10: GPP Satellite-based Yield Predictions in Nepal

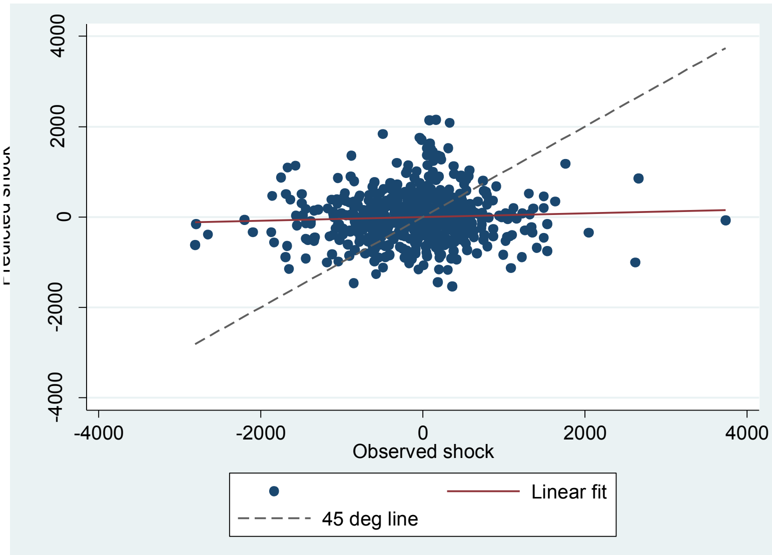

3.5: THE CHALLENGES THAT REMAIN

There are several possible explanations for this mismatch between GPP predictions and yield outcomes. One potential explanation is that the pervasive cloud coverage negatively impacts the quality of our index. In Figure 11 below, we compare the GPP series (green) and the GPP data quality (blue). The quality of the GPP data systematically decreases during the monsoon season because of cloud cover.

Figure 11: Quality of GPP Data

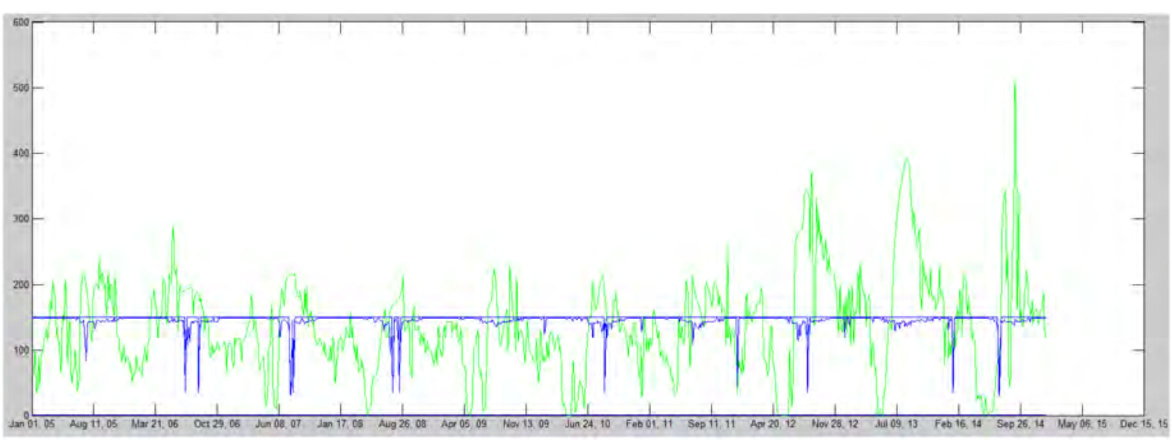

We attempted to overcome the challenges presented by cloud cover, but were unable to find a satisfactory solution. Simpler indices, such as NDVI, did not perform any better. It is possible, however, that part of the problem is that much of the yield data used is recall data, which could create “noise” around the data because of inaccuracies in farmer recall of yields and shocks.

SECTION 4: OTHER OPTIONS

Though our initial assessment based on what worked in previous assessments did not prove promising for Nepal, there are a handful of other options that can be explored. Alternatively, we may conclude that a risk management instrument like index insurance may not be feasible in Nepal given current technological constraints.

Tweaking the insurance design

There are two additional potential solutions to these design challenges that I4/BASIS has been working on. The first tested whether other, less cloud sensitive satellite measures are adequately correlated with farmer yields to serve as an index in an insurance product (these measures are included in Appendix A). The second potential solution would involved "downscaling" the insurance zone. I4/BASIS researchers tested whether "downscaling" or reducing the size of the insurance zone could improve the correlation of the satellite data with farmer yields. This could be the case if the areas within these zones are more heterogeneous than originally thought.

After testing these potential alternative solutions, none of the results were promising. We can provide additional specific information on what was tested if that would be helpful.

Consider a rainfall-based insurance contract

We examined the potential for a rainfall-based index insurance product, for which there are additional details in Appendix B. Our analysis indicated that a rainfall-based index insurance product in infeasible in this environment both because of the high spatial disbursement between weather stations and the poor correlation between rainfall measured at the weather stations and the experience of farmers even a relatively modest distance away from the closest weather station.

Consider an area yield contract

Because the yield data did indicate that at least roughly 39 percent of losses experienced by individual farmers were covariate shocks, there is potential for a moderately effective insurance contract using area yield to be effective. Average area yield would clearly be highly correlated with individual farmer losses and therefore have high predictive power, but the data needed for such a contract could be prohibitively costly to collect.

Though this yield data is not currently collected, this kind of data collection for risk management instruments could be donor supported as part of a development agenda. This public support of the necessary data could make an area yield contract financially feasible.

Our preliminary analysis of the potential for an area-yield based index contract was not promising (this analysis is detailed in Appendix C). We tested hypothetical yield-based index insurance, and the index did not prove effective in predicting individual farmer losses. In addition, we explored how much farmers would be willing to pay for such an insurance product, and it falls below that which is typically considered feasible for these kinds of products.

Conclude an effective index insurance contract is not feasible under current circumstances

It is possible that, for this given situation with currently available technologies, there are no available solutions to create an index insurance product that has the potential for development impact and responsibly protects farmers against risk. Moving forward with an irresponsible product that does not accurately predict farmer losses could be catastrophic, both for the individual farmers who purchase such an insurance product and for the long-term insurance market prospects in the country.

APPENDIX A: BRIEF DESCRIPTIONS OF SATELLITE INDICES

The satellite-based indices used in this feasibility study are extracted from NASA's MODIS (Moderate Resolution Imaging Spectroradiometer) program. MODIS' instruments capture data in 36 spectral bands ranging in wavelength from 400 nm to 14,400 nm and at varying spatial resolutions (2 bands at 250 m, 5 bands at 500 m and 29 bands at 1 km).

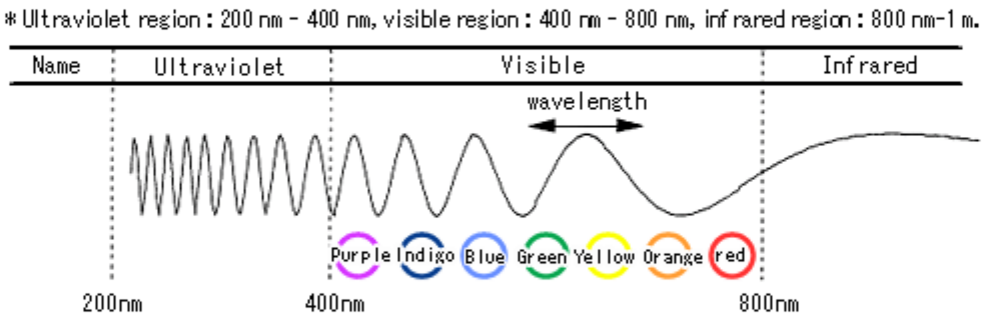

Each spectral band carries some specific information about our environment[1]. Focusing on specific wavelengths or combining them together allows us to compute vegetation indices that capture different dimensions of plants' health.

- The wavelength of 450-520 nm (blue-green) are short wavelength of light and penetrate better than the other bands so it is often the band of choice for aquatic ecosystems. It is used to monitor sediment in water, mapping coral reefs, and water depth. Unfortunately short wavelength blue light is scattered more than the other bands. For this reason it is rarely used for "pretty picture" type images.

- The wavelength of 520-600 nm (green) has similar qualities and matches the wavelength for the green we see when looking at vegetation.

- Since vegetation absorbs nearly all red light (it is sometimes called the chlorophyll absorption band) the wavelength of 630-690 nm (red) can be useful for distinguishing between vegetation and soil and in monitoring vegetation health.

- Since water absorbs nearly all light at wavelength 760-900 nm (near infrared), water bodies appear very dark. This contrasts with bright reflectance for soil and vegetation so it is a good band for defining the water/land interface.

- The band 1,550-1,750 nm (mid-infrared) is very sensitive to moisture and is therefore used to monitor vegetation and soil moisture. It is also good at differentiating between clouds and snow.
- The 10,400-12,500 nm band (thermal infrared) is a thermal band, which means it can be used to measure surface temperature. This is primarily used for geological applications but it is sometime

[1] The information on wavelength here is obtained from Landsat Spectral Band Information guide: http://ibis.geog.ubc.ca/courses/geob373/lectures/Handouts/CBC_landsat_band_info.pdf

used to measure plant heat stress. This is also used to differentiate clouds from bright soils since clouds tend to be very cold.

The different bands can be combined to derive several vegetation indices that could serve as indices in satellite-based index insurance:

- ❖ The most common index used to monitor vegetation health is the NDVI (Normalized Difference Vegetation Index). The NDVI is a normalized transform of the ratio of the near infrared reflectance to the red reflectance;

$$NDVI = \frac{NIR - RED}{NIR + RED}.$$

As it combines red and infrared wavelength, this index can easily distinguish vegetation from dirt or water. Using the red (chlorophyll) band, it can also measure the amount of light that is being absorbed by plants, giving us a measure of plants' health. One advantage of this index is that it is relatively simple minimizing some noises from the correlation across bands and effects of variations in irradiance, clouds and cloud shadows, sun and view angles, topography and atmospheric attenuation. However it can be insensitive to vegetation variations over certain land cover conditions; in particular, NDVI is known for its inability to track vegetation in areas with high biomass because the index saturates.

- ❖ The limitations of NDVI can be overcome using the EVI (Enhanced Vegetation Index). This index improves on NDVI by adding the blue wavelength into the equation:

$$EVI = 2.5 * \frac{NIR - RED}{NIR + 6 * RED - 7.5 * BLUE + 1}.$$

The introduction of the blue band helps to solve the insensitivity to vegetation variations in the high-biomass area. However this index does not work well over bright regions (heavy clouds, and snow/ice) because of the blue band.

Other indices tracking plants' activity are the fraction of Photosynthetically active radiation (fPAR), the Leaf Area Index (LAI) and the Gross Primary Production index (GPP). While these indices might better capture the different dimensions of crop production, they are only available at the 1km spatial resolution, while NVDI and EVI can be computed at a 250m resolution.

- ❖ Photosynthetically active radiation (PAR) is the spectral range from 400-700nm that is used by plants in photosynthesis. The fraction of PAR (fPAR) measures the portion of PAR used by plants. fPAR is often used in ecosystem models because it plays a key role in exchanges of energy between the surface of the earth and the atmosphere. It is an important parameter in measuring biomass production because vegetation development is related to the rate at which radiant energy is absorbed by vegetation. Precipitation and temperature are two of the major factors that determine the proportion of PAR absorbed by plants.

❖ The leaf Area Index (LAI) is the one-sided green leaf area per unit ground surface area (m2 / m2) in broadleaf canopies. LAI can be used to predict photosynthetic primary production. Indeed, an inverse exponential relation between LAI and light interception, which is linearly proportional to the primary production rate, has been established, making LAI a good candidate for index insurance applications.

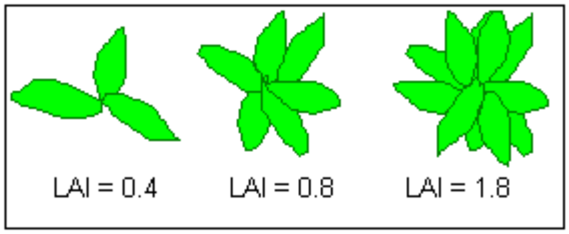

A potential difficulty with the indices presented so far is that they focus on plants' green matter while our focus is on cereal production. Several modern rice varieties are semi-dwarf varieties that are chosen because they proved to be highly resistant to a variety of pests and diseases and produced the slender rice grain preferred in many countries. In addition, semi-dwarf varieties mature quickly – in 105 days instead of the 130 days to 170 days for traditional varieties. Hence, the introduction of modern rice varieties can lead us to observe the combination of a downward sloping trend in the Leaf area Index over time together with a steady increase in yield. In this case, distinguishing the normal trend in the use of smaller rice varieties from a weather shock causing plants to stop their growth might prove extremely difficult.

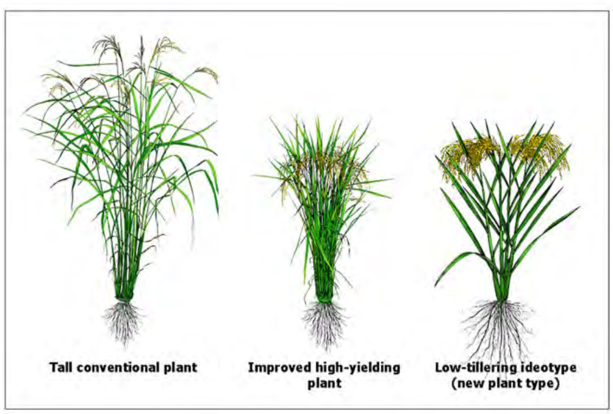

❖ The Gross Primary Production index is based on the fPAR combined with signal from the shortwave infrared band (SWIR - wavelengths from 900 to 1700 nm):

$$GPP = \varepsilon * 0.45 * SWIR * fPAR$$

where ε captures the effect of daily minimum temperature and daylight average vapor pressure deficit.

Because the GPP index is an approximation of true biomass, it not only measures the production of leaves, but also the production of panicles, i.e. the quantity of rice that can be harvested. In particular, the plant growth cycle (below) shows that most of the biomass production during the early stages of growth is concentrated in leaf production. However, during the second half of the crop cycle most of the energy spent by the plant is devoted to panicle production. It is thus possible to estimate rice yields using the amount of biomass produced during the second half of the crop cycle.

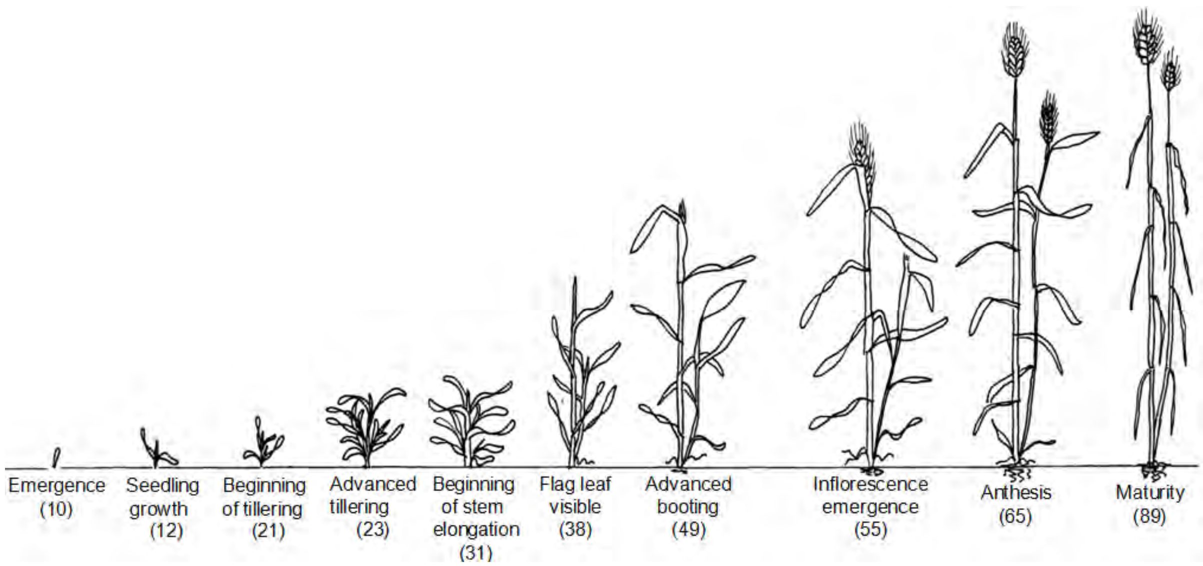

An alternative approach to these vegetation indices is to focus on the environment, and the resources available to plants. In particular, water availability is a key parameter in rice production.

- The Land Surface Water Index (LSWI or NDWI (Normalized Difference Water Index)) can detect a key event in rice production in Nepal: the onset of the monsoon. The LSWI is sensitive to the first substantial monsoon rainfall and the subsequent increase in the vegetation and soil liquid water content.

$$LSWI = \frac{NIR - SWIR}{NIR + SWIR}.$$

  This index would however not be able to capture a drop in rice yield due to cold temperature during the reproductive stage or other non-monsoon related events.

APPENDIX B: ANALYSIS OF POTENTIAL FOR RAINFALL-BASED INSURANCE

To start, we analyzed the spatial availability of weather stations in Nepal. The map below indicates the spacing of weather stations (indicated with black & white circles in the map below), as well as district boundaries (thick red boundaries), VDC boundaries (light red boundaries) and surveyed rice clusters (white areas). To give an idea of relative distances, the distance between two weather stations (indicated by the green line) is 30 kilometers.

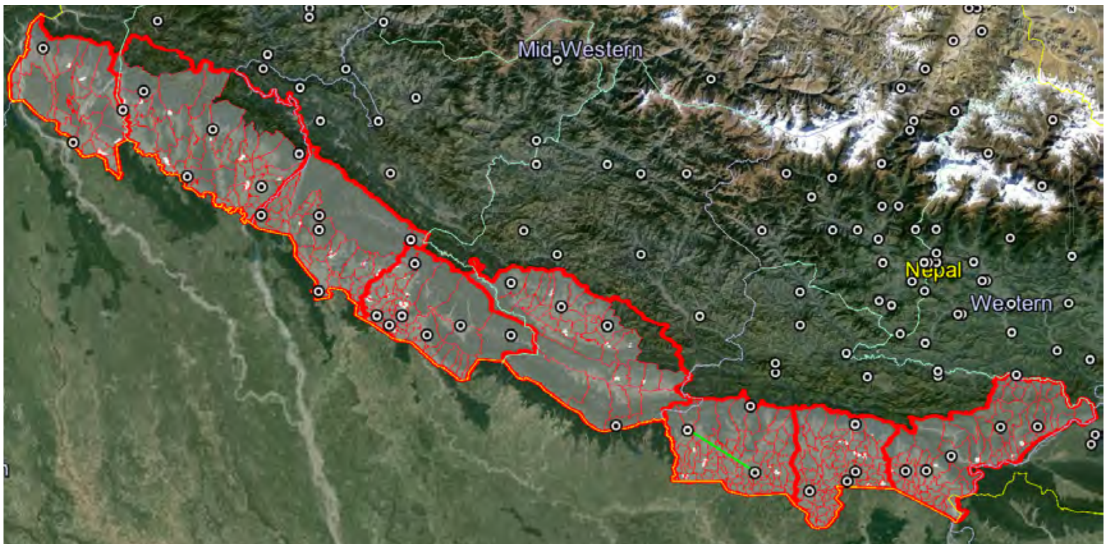

There are 36 functioning stations in the region (one was closed in 1974 in Kailali). Recent literature on rainfall index insurance suggests that we need a density of rainfall stations of at least one station every 10km to get reasonable estimates of rainfall between stations. This condition is not met in Nepal. Furthermore, we expect worse results in Nepal because of the proximity of the mountains, which have significant impact on the climate in the Terai. The climate in Nepal, as a result, is less homogenous than in other studied cases.

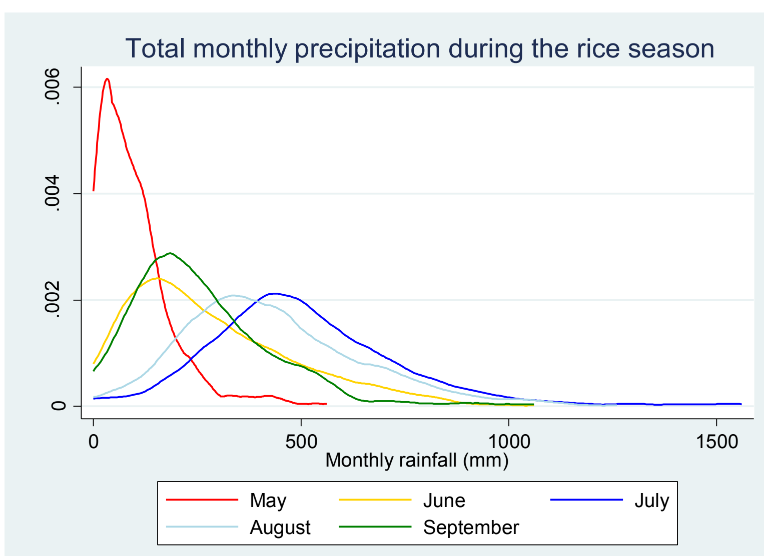

Analysis of correlation coefficients versus distance

Daily data

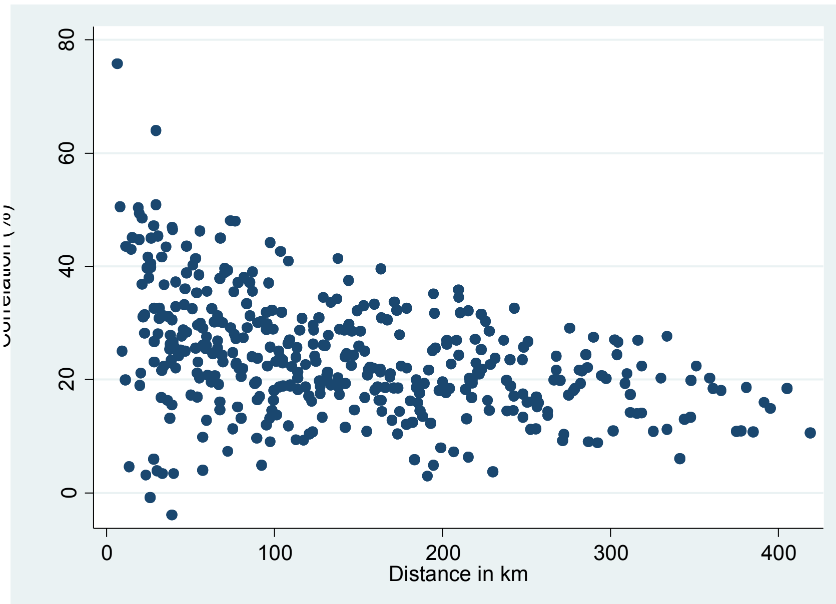

Weekly data

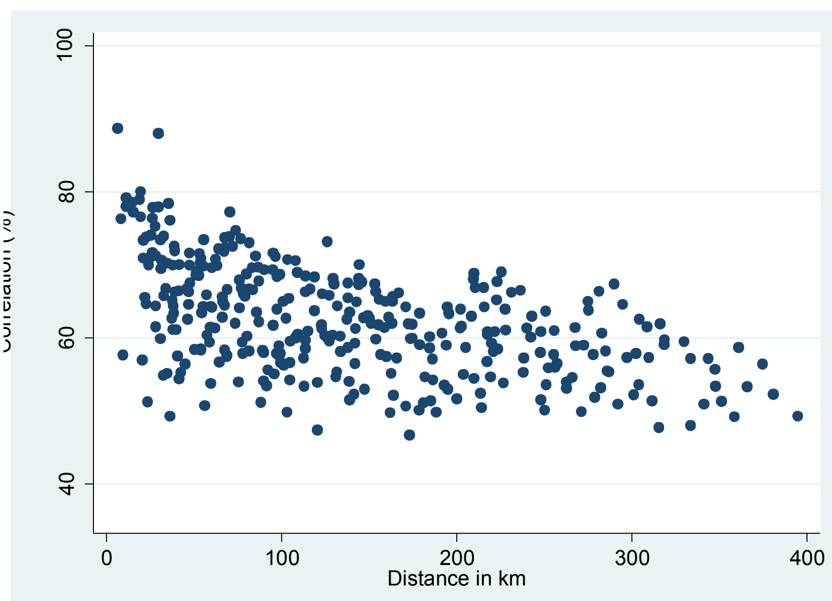

Monthly data

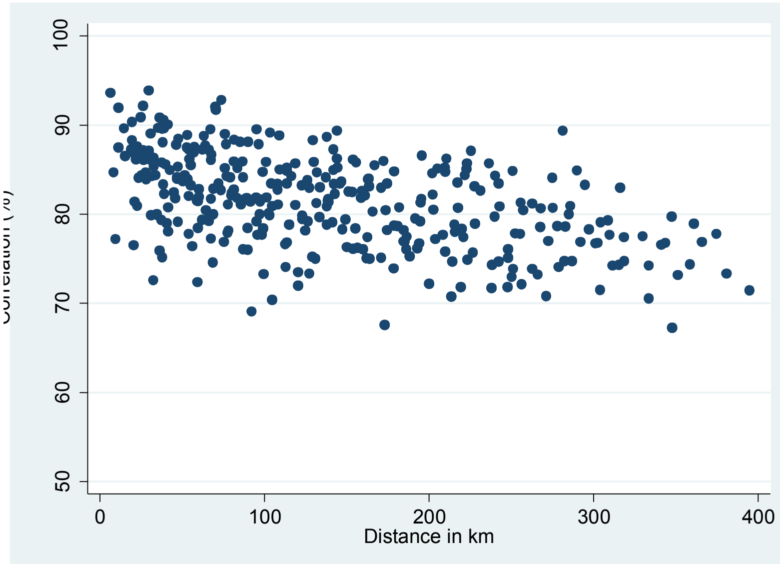

In the case of Nepal, the closest weather stations are 6 km, with most more widely spaced than that. Even in the best-case scenario (which would be one weather station every 6 km) farmers would experience a variation from the predicted rainfall at the nearest weather station of approximately 80-95 percent. In order to design a rainfall index, we would require this correlation to be superior to 90% in all cases; otherwise the insurance based on this index would not provide adequate risk protection to farmers.

Furthermore, the geography of Nepal is such that rainfall in the mountains can have important effects in the Terai region. Hence, floods in Terai might not be related to rainfall in the Terai; and drought in the Terai could be due to lack of rainfall in the mountains. This leaves us with little hope that a satellite-based rainfall index could offer reasonable coverage to rice farmers in Nepal.

APPENDIX C: ANALYSIS OF POTENTIAL FOR AREA YIELD INSURANCE

In section 3.1, we briefly examined a theoretical area yield-based insurance. Based on our retrospective survey data, about 26% of the shocks (computed as deviations from the individual's historical mean yield) faced by individuals are VDC-level shocks that could be insured under an area-yield contract.

Note that our VDC level yield measure is the average yield of 10 farmers from two toles with 5 farmers per tole. (A ward is a group of toles, and a VDC is a group of wards).

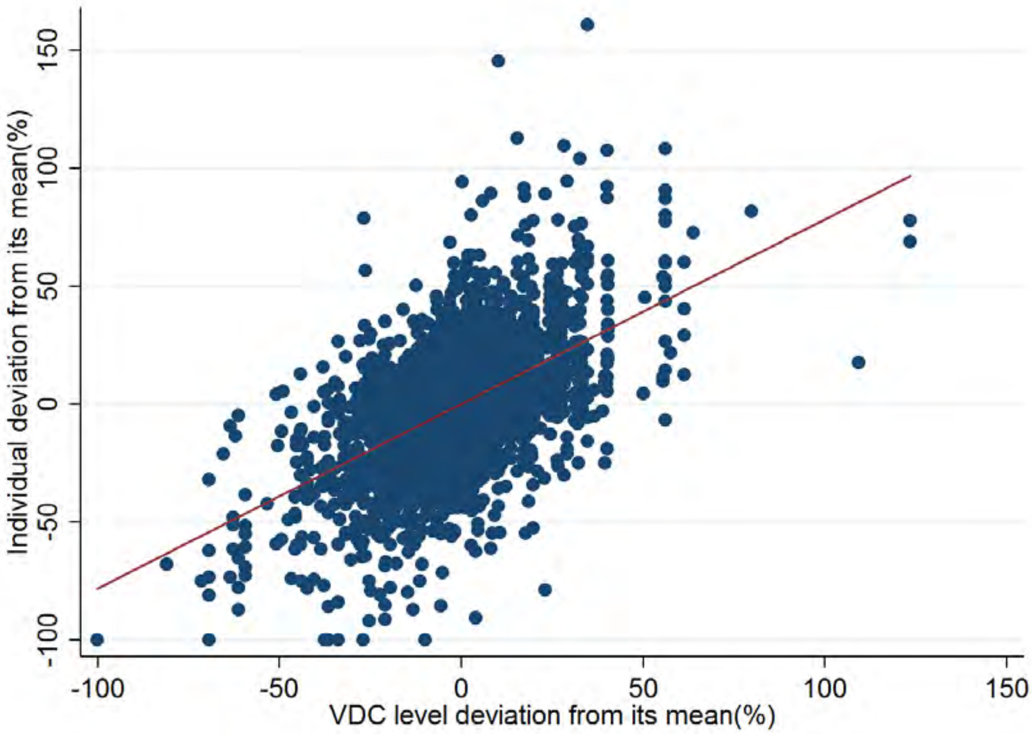

Figure 1 Scatter Plot of Individual Deviation and VDC Level Deviation

Using our recall data from 2001 to 2013, we computed a VDC-level historical average. This historical average is then compared to actual VDC-level yield data for 2014 and theoretical payouts are computed for the year 2014 if actual average yields are lower than the computed historical yield.

This model indicates that, for 2014, 66% of the farmers who did experience a loss in 2014 would have received a payout under a contract that pays systematically when VDC-level yields fall below average.

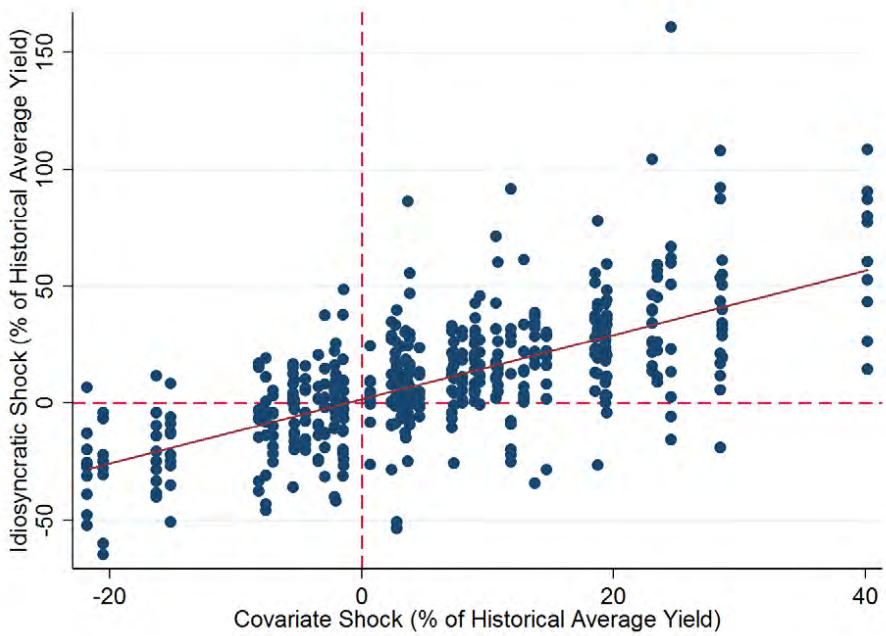

Figure 2 Covariate and Idiosyncratic Shocks for 2014

In this appendix, we complement our examination of section 3.1 by providing some analysis of the maximum price that farmers are willing to pay for this insurance product. In other words, we compute the price at which farmers would be indifferent between having no insurance and purchasing an insurance product. Such equilibrium price is called a "reservation price".

First, we estimate the reservation prices for an area-yield insurance assuming that farmers within a VDC all get the same yields; there are no idiosyncratic shocks. This provides an upper bound of willingness-to-pay for a VDC-level area yield insurance.

Then, we estimate the reservation prices for a theoretical area yield insurance for individual farmers who do experience idiosyncratic shocks.

A.1. Reservation Price under No Individual Yield Deviation from VDC Yields

Contract Description

Consider a VDC yield insurance product that pays whenever the VDC yield is below its historical average and where indemnities are proportional to the drop in VDC-level yields. With such insurance product, and in the absence of idiosyncratic shocks, farmers would always have yields above or equal to their historical averages. Figure 3 shows this case for our sample from the retrospective survey. The red dots represent the case without insurance. The blue dots represent the case with insurance where indemnity payments guaranty that farmers always receive at least 100% of their historical average.

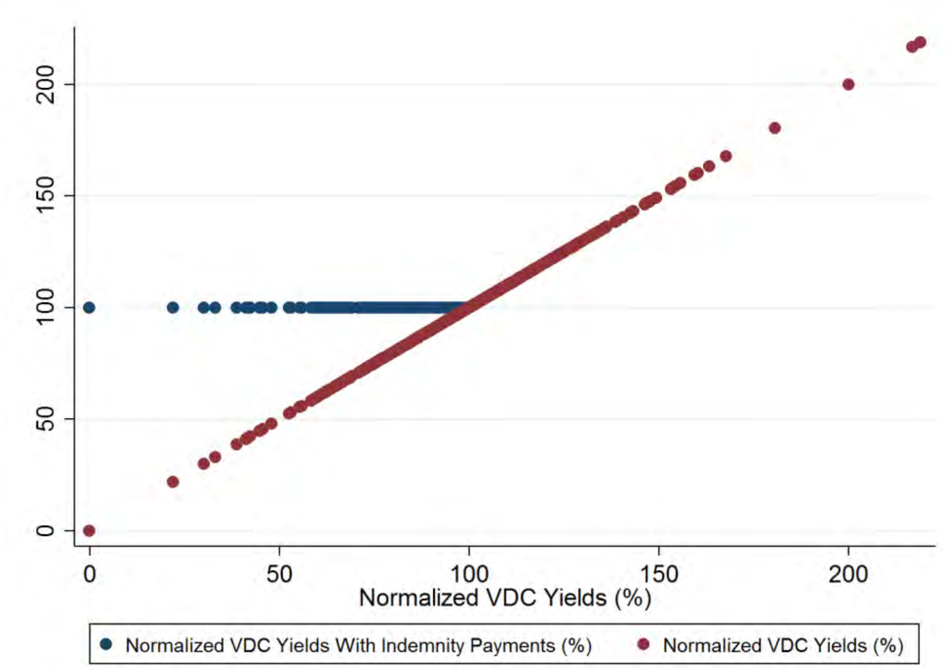

Figure 3 Indemnity Payment Structure

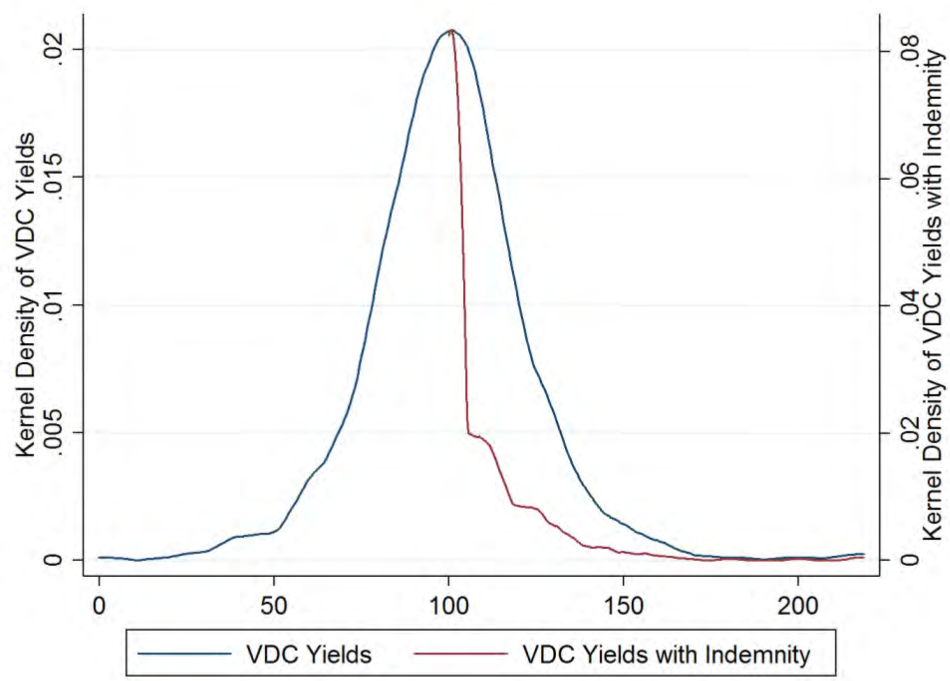

Figure 4 Distributions of Yields with and without Indemnity Payments

Reservation Price

The contract description above ignores the question of the price of such contract. In this section, we will compare the "reservation price" (the maximum price farmers are willing to pay for such contract) against the actuarially fair price (the minimum possible price such that over the long term, the sum of indemnities received equals the sum of premiums paid).

The reservation price is linked to the degree of risk aversion people have. Indeed, someone who is risk neutral would necessarily have a reservation price equal to the actuarially fair price. However, someone with a greater level of risk aversion could be willing to pay a much higher price because he fears low yield years much more than he enjoys high yield years.

Figure 5 reports the average reservation prices across VDCs for various levels of risk aversion (0 = risk neutral; 3 = highly risk averse). We see that risk neutral farmers have a reservation price equal to the actuarially fair price while highly risk averse individuals are ready to pay a premium 54% higher. We can foresee a difficulty in marketing area-yield insurance under these conditions since commercial premiums in the real world are often higher than 150% the actuarially fair price. Only extremely highly risk averse rice farmers would buy index insurance at such price in Nepal.

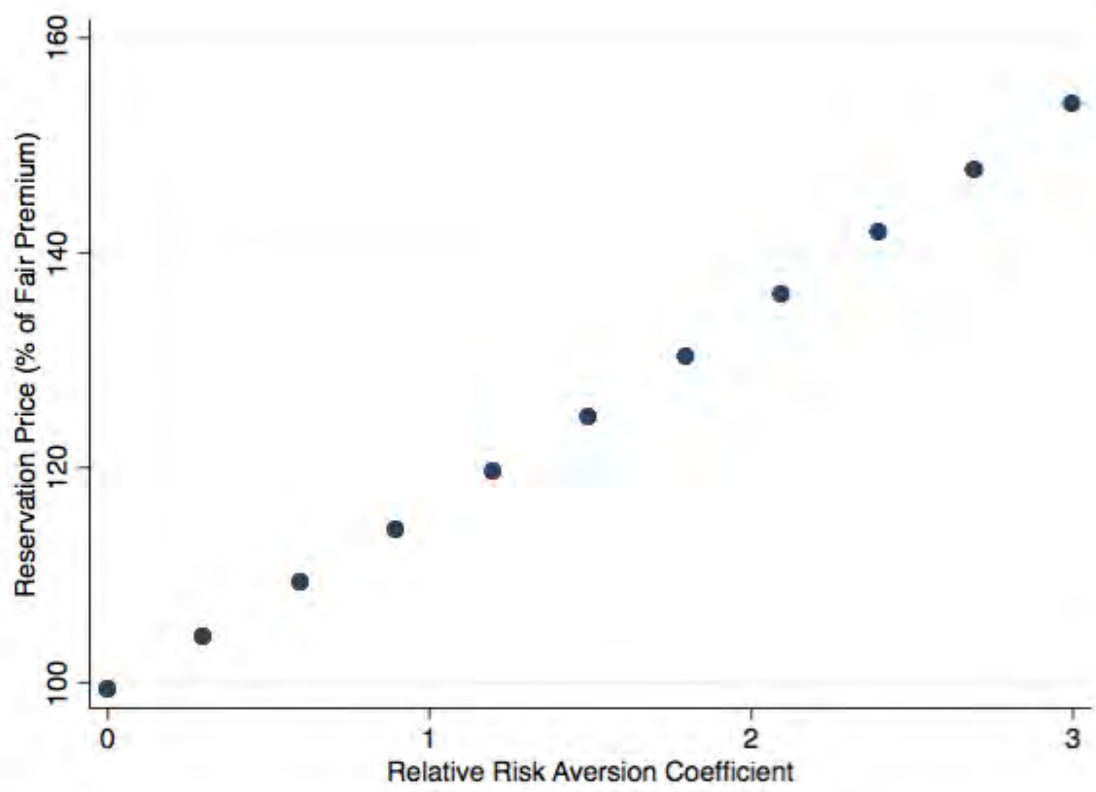

Figure 5 Average Reservation Prices as % of Actuarially Fair Premium

A.2. Reservation Price under Individual Deviations from VDC Yields

The study above assumes the absence of idiosyncratic shocks so that every farmer inside a VDC always receives the same yield as his neighbors. In the real world, we observe a lot of variations inside a VDC, and such variations (we call these “idiosyncratic shocks”) reduce the value of an area yield contract for individual farmers. The purpose of this section is to measure how much value is lost because of these idiosyncratic shocks.

In order to answer this question, we use the same contract we used before, that triggers indemnity payments whenever the VDC yields are lower than their historical average computed on the data from 2001 to 2013. But now the farmer not only suffers from these VDC-level shocks, he can also experience individual “idiosyncratic” shocks that make his own yields differ from the average yield in his VDC.

In this case, it is possible that the farmer suffer a loss while his VDC-level yields are above average. Hence, this farmer needs to receive indemnities, but the are-yield index does not trigger. It is also possible that the farmer enjoys a good year compared to his own historical average, but his VDC-level yields are below average so that he receives an indemnity that he does not need. Figure 6 compares farmers with and without insurance when we introduce idiosyncratic shocks. We can observe that the VDC area-yield insurance no longer guarantees individual yields at the 100% of their historical average. When the individual farmer loses his entire production, he can only expect to receive 60% of his normal income.

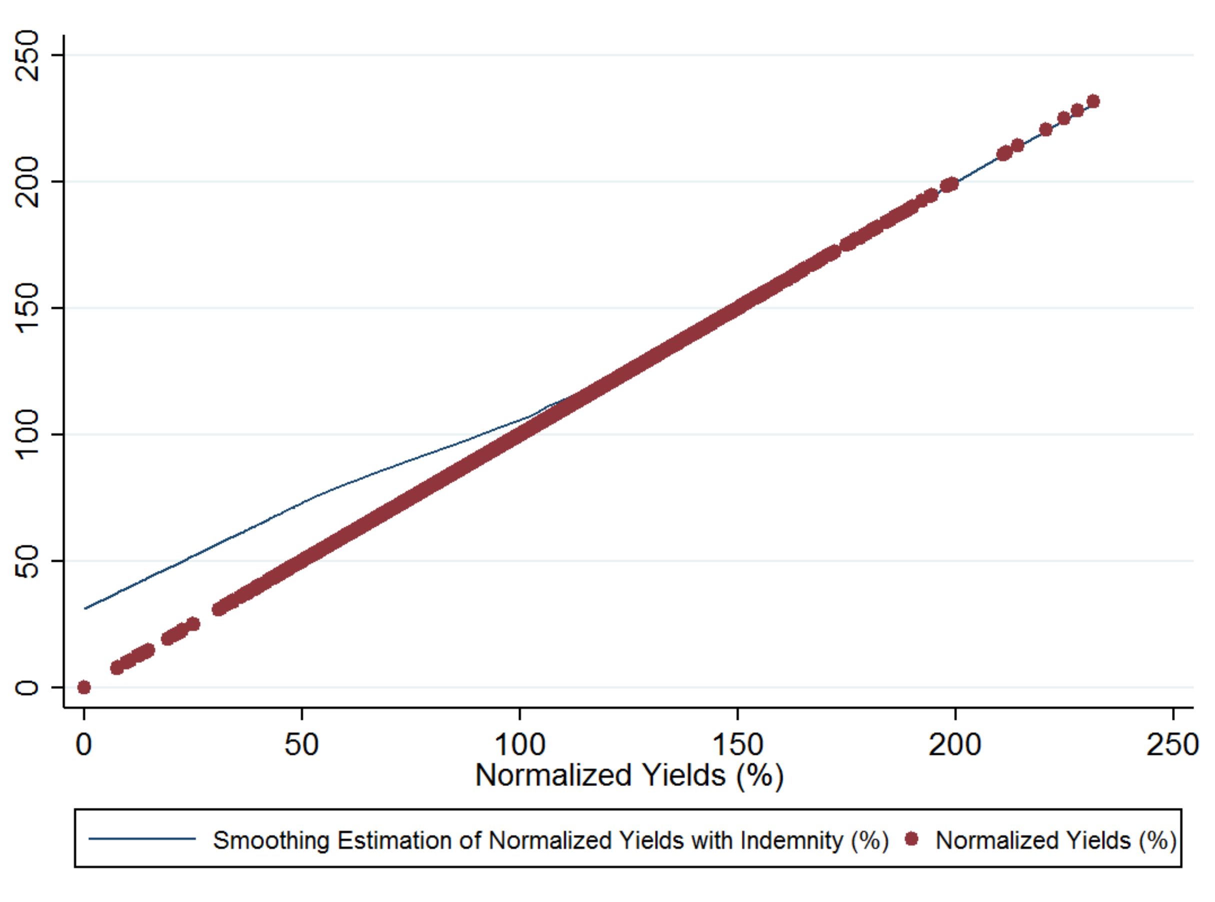

Figure 6 Indemnity Payment Structure

Reservation Price

Given the additional noise introduced by idiosyncratic shocks, we expect farmers' reservation prices to be lower than what was found in the previous section. We redo our comparison of these reservation prices against the actuarially fair price and Figure 8 reports the average reservation prices across individuals for each level of risk aversion. We can see that as the level of risk aversion increases the reservation price increases. For a risk neutral individual, the reservation price still equals the actuarially fair price. However, for a highly risk-averse individual, the reservation price is now equal to 145% of the actuarially fair price. This drop in the willingness-to-pay reduces further the chances that an area-yield contract could be successfully implemented for rice farmers in Nepal, as commercial products typically are at least 150% or more of the actually fair price.

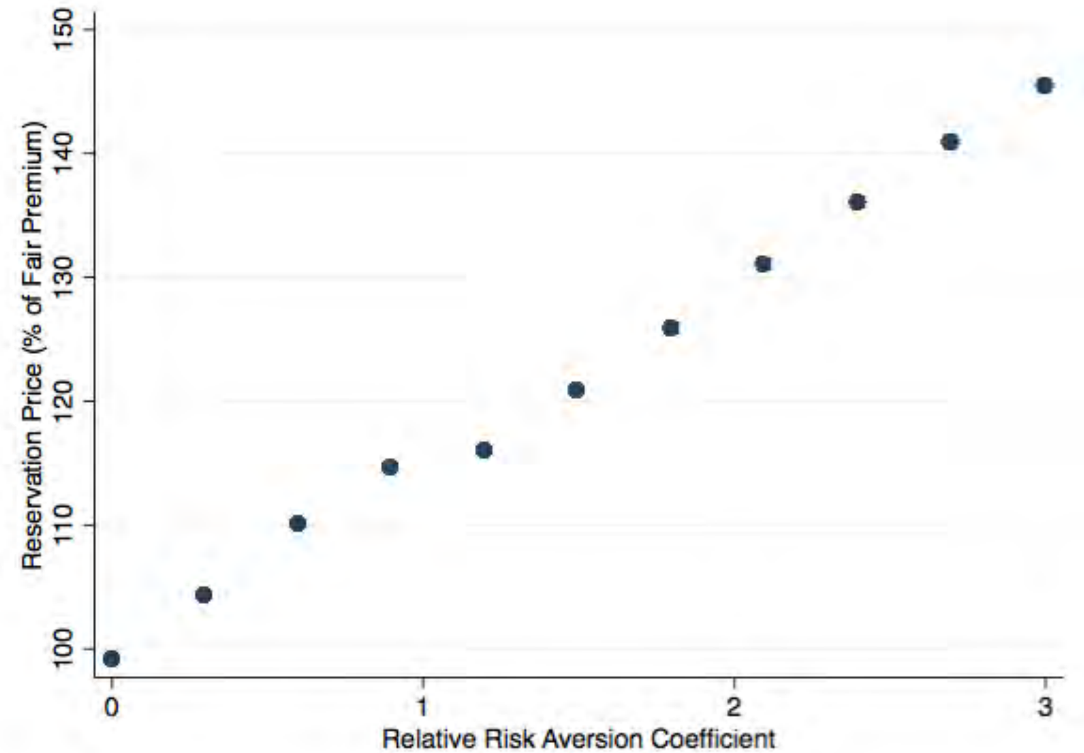

Figure 8 Average Reservation Price