WPS 3392

DOCUMENT OF THE WORLD BANK

Risk Modeling Concepts Relating to the
Design and Rating of Agricultural Insurance Contracts

Barry K. Goodwin

North Carolina State University

Olivier Mahul

The World Bank

Abstract

Goodwin and Mahul identify the key issues and concerns that arise in the design and rating of crop yield insurance plans, with a particular emphasis on production risk modeling. The authors show how the availability of data shapes the insurance scheme and the ratemaking procedures. Relying on the U.S. experience and recent developments in statistics and econometrics, they review risk modeling concepts and provide technical guidelines in the development of crop insurance plans. Finally, they show how these risk modeling techniques can be extended to price risk in order to develop crop revenue insurance schemes.

Executive summary

Crop insurance is one mechanism for the management of the risks associated with random yield shocks once all cost-effective risk mitigation strategies have been implemented. Many developing and developed countries have well-established public-private crop insurance programs. However, their overall financial experience has been very expensive for the taxpayers and, in many instances, far from popular with farmers. The poor financial performance is usually due to the inability or the unwillingness of the government to charge adequate premiums to farmers. Beyond the interactions with the government, which is tempted to use crop insurance programs as a social vehicle for subsidy to the agricultural sector, the lack of sound ratemaking procedures is one of the key reasons for their poor financial performance in several developing countries, like India.

Since the late 1990s, due to dwindling government subsidies to agricultural producers in developing countries, there has been a renewed interest among development lenders and policy makers in using agricultural insurance as an important risk management tool. The financial viability of agricultural insurance programs relies on the design and the rating of actuarially sound insurance contracts in order to give private insurance companies incentives to offer market-based products and/or to assess the fiscal consequences of government subsidy programs. In India, for example, the government is looking for international technical assistance to devise and implement an actuarially sound ratemaking process for the National Agricultural Insurance Scheme.

The purpose of this paper is to provide Bank staff and policy makers involved in agricultural insurance programs in developing countries with an overview of the latest developments in the modeling of crop yield risk and, to some extent, crop revenue risk, and to discuss how these modeling concepts affect the design and the rating of crop insurance. The overarching concern with such models is the derivation of the parameters needed to adequately measure risk and thus design and rate viable crop insurance contracts. Though the procedure commonly used to measure yield risk (i.e., collecting data and estimating parameters of the yield distribution) seems straightforward, a number of important concerns are germane to the process. Indeed, these concerns underlie actuarial science and underwriting considerations.

This paper draws some lessons from the U.S. experience. The U.S. crop insurance program is a joint effort of the Federal government and private industry. The insurance scheme, known as Multiple Peril Crop insurance (MPCI), has been offered to the U.S. farmers for more than three decades. In crop year 2003, the MPCI program provided coverage on 217 million acres (almost 76%) of U.S. cropland, insured USD 40.6 billion in crops, generated a total premium of USD 3.4 billion (of which USD 2.0 billion are premium subsidies) and distributed USD 3.2 billion of indemnity payments. The Risk Management Agency of the U.S. Department of Agriculture has developed a well accepted ratemaking procedure for the MPCI program.

This paper lays out recommendations on the design and rating of viable agricultural insurance policies. It provides Bank staff and policymakers with two sets of guiding principles for the development of viable agricultural insurance programs in developing countries.

The first set of guiding principles refers to some key issues that should be addressed in devising agricultural insurance contracts.

- Insurance contract design is driven by the data availability. Individual farm data are almost always missing or unreliable. Consequently, aggregate data can often be used either to develop individual insurance coverage, where the individual indemnity depends on the individual loss, or index-based coverage, where the individual indemnity is based on some indices correlated with the individual loss (e.g., area yields, weather parameters).

- Insurance coverage should be offered at the whole farm level in order to stabilize the farmer's overall agricultural revenue. This global coverage focuses on losses that cannot be mitigated through diversification and avoids fraudulent claims (e.g., yield shifting across fields).

- All-peril coverage should be subject to sufficiently high deductibles and coinsurance and to effective loss adjustment process in order to efficiently control moral hazard problems.

- Multi-year contracts should be promoted because they protect insurers against inter-temporal adverse selection (i.e., farmers seek insurance only when they expect a bad year) and they offer more financial stability to the emerging agricultural insurance market.

- Crop revenue insurance, covering both yield and price downside variations, should be considered as an extension of crop yield insurance.

The second set of guiding principles pertains to crop risk modeling, which is a key step in insurance contract design and ratemaking procedures, particularly when historical loss data are limited or missing. Its main task is to get information about the stochastic nature of random yields. Ultimately, the objective is to derive the estimated crop yield stochastic distribution function.

- The modeling approach (e.g., parametric or non-parametric statistical methods) and the underlying assumptions (e.g., family of stochastic distributions) are of fundamental importance as they affect directly the estimated yields, the estimated losses and thus the insurance premium rates.

- Non-parametric methods should be preferred when large datasets are available because they impose no minimal structure on the estimated distribution. On the contrary, parametric methods are more preferable when the data set is small.

- Crop yield distributions are usually skewed to the right, with yields close to the maximum yield observed more frequently than very low yields. The normality assumption is thus usually rejected.

- When historical losses are available, rating procedures should be based on historical loss cost ratios. The ratemaking methodology should rely on several steps, such as loss adjustment to a common level of coverage, derivation of unloaded base rates, base rate loading, capping rate changes and updating.

- When historical losses are not available, the rating methodology should rely on simulated losses derived from statistical crop yield models.

- Crop revenue insurance rating requires innovative procedures that are available through recent experience on the U.S. crop insurance market.

The guiding principles outlined in this paper aim at devising viable crop insurance products. They are based on a sound actuarial approach. Experience shows that such a rigorous approach may be missing in government-sponsored crop insurance programs, because the high social welfare content imposed by the government may not be compatible with these insurance guidelines. For example, premiums rates derived under these actuarial principles may be considered as too high to meet social objectives. As a consequence, a prerequisite in the implementation of these guiding principles is to make a clear distinction between social insurance, which cannot be achieved through market-based products, and market-based insurance, which should rely on the above principles.

1. Introduction

The management of crop production risks is an issue of fundamental importance to agricultural economies. Because of the random nature of production conditions (e.g., weather, pests, diseases), agricultural producers face an array of risks that may influence their level of output per acre from year to year. Management of such yield risks has long been an important issue for producers as well as for policy makers.

Crop insurance is one mechanism for the management of the risks associated with random yield shocks once all cost-effective risk mitigation strategies have been implemented. Many countries have well-established public and/or private crop insurance programs, although most, if not all, public multi-peril schemes are not sustainable without heavy government subsidies. In the United States, for example, federally subsidized crop insurance has been in existence for over 60 years.[1] Such insurance is currently available for over 75 different crops and livestock products, with premium subsidies representing on average 60% of total premiums in 2003 (RMA 2004).[2]

In its most fundamental form, a crop insurance plan will pay producers an indemnity in the event that their yields fall below a pre-determined level. Construction of such a seemingly simple contract requires representation of a number of important parameters. Accurate measurement of such parameters may be quite complex, especially in cases where limited knowledge of the risks or levels of protection being provided are available. The challenges associated with accurately measuring the parameters that determine liability, premiums, indemnities and other components of a crop insurance plan are often complicated and may require the application of rather complex actuarial methods, models, and assumptions in order to design and rate viable insurance contracts.

The purpose of this paper is to provide Bank staff and policy makers involved in crop insurance programs in developing countries with an overview of the latest developments in the modeling of crop yield risk and, to some extent, crop revenue risk, and to discuss how these modeling concepts affect the design and the rating of crop insurance. The overarching concern with such models is the derivation of the parameters needed to adequately measure risk and thus design and rate viable crop insurance contracts. Though the procedure commonly used to measure yield risk (i.e., collecting data and estimating parameters of the yield distribution) seems straightforward, a number of important concerns are germane to the process. Indeed, these concerns underlie actuarial science and underwriting considerations. Based on the latest developments in agricultural risk modeling and on the U.S. experience, this paper reviews each of these concerns and identifies possible solutions and recommendations that may be commonly used to address each issue.

This paper is organized as follows. Practical considerations in devising crop insurance schemes are presented in Section 2. Section 3 describes how production risk can be modeled using parametric and non-parametric methods. Section 4 examines how risk

---

[1] Goodwin and Smith (1995) review the history and operation of the U.S. crop insurance program. The European Commission (2001) provides a description of crop insurance programs in European countries, Canada and Japan. FAO (1991) describes several crop insurance programs in developing countries (e.g., Chile, Cyprus, Mauritius, Philippines).
[2] Subsidies are applied to a base premium rate of 1.075, as required by law.

modeling can be used to develop crop insurance contracts. Ratemaking procedures for existing insurance schemes and new schemes are detailed in Section 5. Section 6 examines how to extend coverage to address price risk through crop revenue insurance. Finally, the key issues are summarized in the conclusions.

2. Practical Considerations in Devising Crop Insurance

The theory underlying the construction of insurance contracts is generally straightforward. However, there are a number of issues that are relevant to any practical application of these principles. Indeed, it is attention to such issues that makes up much of actuarial sciences. In this section, we discuss these issues and various approaches used in the practical development of crop insurance tools.

2.1. Data Availability

The feasibility of any particular crop insurance plan is often influenced by the availability of yield data. Detailed time series of individual farm data of the sort needed in measuring yield risks are usually rare, even in cases where farm records are maintained. In the case of the U.S. crop insurance program, only 4-10 years of yield data are used in determining the yield guarantees. Seldom, if ever, do substantially rich yield records exist at the individual farm level that would permit the detailed rating of crop insurance contracts at the individual level. This suggests that the development of rates and levels of protection for individual farmers may be difficult and that developers will likely need to give consideration to the use of aggregate data. The use of such aggregate data in devising crop insurance contracts may follow two distinct avenues: individual insurance coverage, and index-based insurance coverage.[3]

Individual insurance coverage

One may wish to use aggregate data to make inferences about the risks facing individual farmers. This is very common in the development of individual plans. For example, under the U.S. crop insurance program, twenty years of the loss cost experience at the county level is used to obtain a measure of the typical level of risk for farmers in the county.[4] Rates at the individual level are adjusted in accordance with how the individual farm's average yield compares to the county average.

Other approaches have been used to combine individual farm yield data with aggregate data. The developers of the U.S. Income Protection (IP) program, a revenue insurance scheme, combine county yield data with individual farm yields to make use of the longer county series and thus improve the amount of yield information used in developing the IP contracts. In the case of the IP program, an attempt is made to distinguish yield variability specific to the farm from yield variability experienced by the entire county. Under provisions of the U.S. federal crop insurance program, farmers lacking a yield history are assigned a yield that is based upon the county average yield (in

---

[3] A conceptual discussion on the role of information in insurance contract design is provided in the Appendix.
[4] Loss cost is given by the ratio of indemnities to total liability and represents a measure of the expected loss rate which, in turn, is a measure of the actuarially fair premium rate. The ratemaking procedures are discussed in Section 6.

some cases, a proportion of the average yield is used). Other approaches to combining aggregate and disaggregate data are possible. Most approaches adopt a measure of the relationship between farm level yields and yields at a more aggregate level such as the county. For example, one may use regression methods to estimate the relationship between a limited number of farm-level yields and county yields and then use this relationship along with a longer series of county levels to infer a pseudo series of farm-level yields. In a related way, aggregate county yields are often used to approximate missing farm-level yields. In the U.S. crop insurance program, farms that lack a yield history from which to construct a measure of their expected farm-level yield are assigned a proportion of the county average yield in place of the missing data. For example, if a farmer can prove that they have never produced the crop before, they are eligible for 100% of the county average. Farmers that are not new producers but that are unable to furnish yield histories are assigned 60% of the county average yield.

Research by the Risk Management Agency (RMA) suggests that individual average yields below the county average tend to be more variable than those above the county average. This inverse relationship is incorporated into the rating of individual contracts by considering an index formed by the ratio of the farm average yield to the county average yield. This accounts for the fact that county average yields are variable across counties and thus adjustments to rates are made on the basis of how an individual farm may compare to the county average. Hence, rates are higher for growers with below average yields and lower for growers with above average yields.

Index-based insurance coverage

An alternative approach to devising crop insurance plans when individual data are lacking is to base coverage on some aggregate index that conveys information about (e.g., that are correlated with) the individual losses. Examples include area-yield insurance, which bases coverage and indemnities on aggregate yields of a given geographical area, and weather-based insurance, which is based on an index reflecting the weather experience (e.g., rainfall) at a particular weather station. These index-based plans have substantial advantages in that aggregate data are generally more reliable and longer histories of aggregate experience are more likely to be available.

A potential disadvantage associated with such plans relates to the fact that indemnity payments may not be perfectly correlated with the individual loss and thus there may be yield basis risk. Consequently, an individual farmer may suffer a loss and not be eligible for an indemnity payment. The degree of correlation between an individual farmer's yields and the index that forms the basis of the area-wide plan is a key factor determining the extent to which an individual's yield risks are covered by the insurance plan.

This raises a substantial distinction between actuarial soundness and effectiveness of the plan for protecting against individual risks. An index for which reliable data are abundantly available may allow one to construct very accurate measures of risk and thus accurate premium rates. In terms of actuarial soundness, being able to accurately measure the index and model its uncertainty is the primary concern. However, the effectiveness of the index-based insurance plan and its likely acceptance by farmers will be determined by the extent to which the index reflects their individual risks. For a farmer with yields that are unrelated to the index, the index-based plan will provide little protection for their

yield risks. However, if a farmer's yields are perfectly correlated with the index, the index-based plan will provide perfect coverage for their yield risks.

However, index-based plans may offer some advantages. A greater availability of data will likely allow insurance premium rates and other parameters of the contract to be measured with greater accuracy. This limits the problems associated with adverse selection. Index-based plans also lessen concerns of moral hazard. If an individual farmer is so small as to be unable to affect the overall index, there is no potential for moral hazard since the agent's actions cannot influence the likelihood that indemnities will be paid. In addition, reflecting the fact that such index-based plans are much cheaper to administer since individual producer records are not needed and loss adjustment need only be done at the aggregate level, index-based plans can generally be offered at a lower cost.

It may also be the case that there is a lag in measuring the index, generating a delay in settlement of claims. Under the Indian crop insurance program, for example, the delay is close to one year. It is caused by the time taken to carry out crop cutting experiments and, in some cases, delayed flows of funds from the state government. The delays prevent the restoration of credit eligibility of farmers for the following crop season (World Bank 2003). Likewise, indemnity payments under the U.S. area yield plan, are not made until the following year due to the lag associated with the measurement of area yields.

2.2. *Specific Peril versus Multiple Peril versus All Peril Coverage*

An important distinction in the nature of coverage offered under alternative crop insurance plans pertains to the types of hazards that are insured against. Some forms of insurance are specific peril, meaning that only losses from a pre-specified peril are insurable. A prominent example is hail insurance, which will pay only for those losses that are caused by hail damages. Fire insurance is another example. Alternatively, crop insurance may provide indemnity payments when losses are realized as a result of multiple causes of loss. As an example, crops are commonly insured against losses caused only by both fire and hail. It is more common for such multiple peril policies to list causes of loss that are not insurable. That is, all causes of loss except for a specific list of perils will trigger indemnity payments. Finally, some agricultural insurance plans provide all peril coverage. For example, the U.S. federal crop insurance program is designed to insure yields of farm producers over an entire growing season on an all risk basis. The primary cause of loss is weather, either for a single identifiable event or over an extended period. This program provides protection against all losses except the following: (1) negligence, mismanagement, or wrongdoing; (2) failure to follow recognized good farming practices; (3) losses caused by water contained in any public dam or reservoir; and (4) failure to carry out good irrigation practices for irrigated crops.

A fundamental issue that is key to the question of whether coverage can be single or multiple peril involves the extent to which the loss adjustment process is able to adequately determine the causes of loss. As was noted in the case of the U.S. crop insurance program, losses occurring due to improper actions of the producer are generally not insurable. Such exclusions are closely related to the issue of moral hazard. Insurance against the adverse actions of farmers will naturally work against the soundness of the insurance program. Monitoring of insured agents' actions and adequate loss adjustment is

an issue that arises in almost all lines of insurance. Adequate training of loss adjusters is a critical component needed in the development of a sound insurance program. It is also important that any conflicts of interests between loss adjusters and the insured agents be identified and remedied.

Actuarial considerations must, of course, make a distinction between measuring all peril risk and measuring the risks associated with a specific peril. It is often argued that it is easier to measure the risks associated with a specific peril than to attempt to measure all risks, seen and unseen, that might be expected to cause yield losses. This may be true in cases where exogenous information about risks is available. For example, in the case of hail, historical weather records may allow actuaries to accurately measure the risks associated with hail damages in any particular area. Specific peril lines of insurance may also be provided against events that are very low probability and non-systemic. For example, fire may be a rare event and may not affect large numbers of producers. In such cases, adequate measures of risk may be obtained from considering a large, cross-sectional pool of producers.

Specific peril insurance may also rely on historical yield data to measure risks. In such a case, actuaries must be able to identify the exact causes of loss associated with historical yield shortfalls. The methods discussed below for detrending yields can be extended in a straightforward manner to "condition" on non-insured causes of loss. For example, consider a case where factors associated with yield shortfalls can be segmented into a single hazard that is covered (e.g., hail), denoted by X, and hazards that are not covered, denoted as Z. One would run a regression of yields on Z, thereby conditioning out the effects of those factors affecting yields that are not covered by the insurance. The residuals from such a regression represent yield shocks that are associated with the covered factor X, as well as random, unmeasured factors. Thus, the overall approach to modeling the risks associated with a specific peril from historical yield data is no different than what is commonly undertaken to detrend data. The only distinction involves the choice of conditioning factors included in the regression.

2.3. Unit Insured

Agricultural insurance is an integral part of risk management. It deals with residual risks that cannot be prevented through cost-effective preventive measures. Field diversification is a standard risk management technique that allows farmers to reduce risk. On the demand side, farmers would prefer to cover their assets at risk by purchasing an "umbrella" contract that guarantees a minimum revenue at the farm level. However, premium subsidies and information asymmetries would give them incentives to insure each field, plot or crop separately. On the supply side, such an umbrella coverage would correspond to the primary objective of crop insurance, i.e., stabilizing the agricultural revenue at the farm level.

Indemnity payments received on a disaggregated set of units are at least as large as what would be received if all the units were aggregated. Intuitively, if yields are not perfectly correlated across individual units, it is likely that shortfalls for some units may be offset by higher than expected yields on others, such that the aggregate yield does not qualify for a payment while the individual units would qualify if insured individually.

However, actuarially sound premium rates would be expected to fall as the insurance unit is defined over more and more individual farm units.

The fact that yield risk tends to fall as the unit of observation increases reflects factors that are associated with the correlation of yields across units. Spatial diversification may be possible if yield correlation is less than perfect. Individual units may be separated by a considerable distance. Individual farm units may have their own idiosyncratic factors affecting yields. For example, the soil quality on one unit may be much higher than on another.

Multi-field contracts

In the case of the U.S. crop insurance program, an across-the-board discount of 10% is given to farmers that insure their individual fields growing the same crop (which are called "optional units") as a whole (which is called the "basic unit").

However, under the optional units, it may be difficult for the insurer to verify the field from which insured production originated. To the extent that it is difficult to monitor and verify production from individual fields, producers may have an incentive to "shift" production across individual fields. For example, if two fields had an identical guarantee of 100 bushels per acre and each field produced 105 bushels, no indemnity would be paid. If, however, the producer reports that one field generated 125 bushels and the other generated 85 bushels, indemnities would be triggered. This point involves actions on the part of producers that are obviously fraudulent. The extent to which such actions actually take place in the U.S. crop insurance program is unclear, though unpublished research has suggested that there is evidence of such yield switching in some areas and for some crops.[5]

When the interest is in addressing whole farm risk or risk at some aggregate level, policies based on production across all relevant fields should be considered. In addition, the potential for fraudulent claims through yield shifting across fields may exist when coverage is provided at the individual field level. This is essentially a monitoring and yield verification problem and thus its relevance depends upon the extent to which adjusters can verify yields on individual fields.

Multi-crop contracts

Most contracts are developed for a single crop. However, there may be some value to provide coverage over multiple crops due to the diversification effect, as this effect applies over multiple fields for a same crop. If yields across crops are not perfectly correlated, overall risk will be lower when multiple crops are combined into a single contract and, therefore, such contracts may be offered at a lower price. For example, shortfalls in corn yields may be offset by good soybean production on a single farm. In such a case, corn insurance would trigger an indemnity payment while soybeans would not. However, when total revenues from the two crops are combined, no indemnity is triggered since the below average corn yields are offset by above average soybean yields.

---

[5] Unpublished research at Montana State University has provided evidence that may be consistent with such switching.

Rating such multi-crop plans presents a number of challenges since one must have adequate measures of the correlation of yields for different crops on a single farm. Measurement of the degree of correlation of yields for different crops on a single farm is difficult and may require strong assumptions about the degree of correlation. Other factors such as the distance between fields for alternative crops may be relevant to the degree of correlation. The authors are aware of only two insurance policies available on a multi-crop basis in the U.S. The whole farm option is available for corn and soybean producers under the Revenue Assurance plan. In this case, the revenue guarantee and indemnity trigger are based upon the total revenue generated by both corn and soybeans. Substantial discounts are available to growers that insure their crops as a whole farm unit. The second example is the Adjusted Gross Revenue (AGR) plan of insurance. The coverage and indemnities are based on the farm's adjusted gross revenue as reported on the farm's tax return form. This plan covers revenues from all sources, including livestock commodities.

2.4. Multi-Year Contracts

Most crop insurance contracts are of a single year in duration. This reflects the seasonal nature of crop production and the fact that production decisions are typically made over an annual growing period. The same general arguments regarding the variability reduction that may be possible from combining yields across crops also applies to the case of combining yields over time. If one bases coverage on the value of production over multiple years, it may be possible to lower the cost of insurance since poor production in one year may be offset by above average production in another. In light of the fact that many farmers are credit constrained and dependent upon short term loans, multi-year coverage may not be a viable alternative.

A related issue pertains to the potential for offering discounts to producers that have strong performance from year to year through a bonus/malus arrangement. Producers that do not submit a claim over a few years may reveal that they are of a lower risk to the insurer than was thought when they came into the program. In such a case, it may be preferable to lower their premium rates, thus making use of the revealed information about their risk. This simply involves using their performance information to fine-tune rates in an attempt to make them more accurate. This intuition also works in the opposite direction; it may be to the insurer's advantage to raise the rates of agents that frequently submit claims. Commercial lines of insurance typically undertake such actions in an attempt to make premium rates more accurate.

3. Production Risk Modeling

3.1. Time Trend

Crop yields have realized substantial changes in recent years as improved production techniques have been adopted. Conceptually, this means that the data-generating process underlying yield realizations is not stable but rather changes over time. Thus, one cannot combine yields observed over different periods of time. For example, corn yields observed in the 1970s are clearly not comparable to those observed today.

To address the problem of structural changes in yields observed over time, a variety

of methods for "detrending" yield data have been adopted. Generally, one uses least squares regression techniques to account for deterministic trends that have moved yields up over time. Deviations from this trend can then be added to current yields to produce a series of yields that can be compared over time. Trends may be linear or nonlinear (the latter is more common) and a variety of estimation approaches have been used to account for such trends. These univariate methods include autoregressive-moving average models, splines, and local nonparametric smoothing techniques.[6]

A closely related issue that one must address when detrending yields involves the second moment of the detrended yield data. In particular, is the variance of the yields constant or does it vary as the level of yields changes? Goodwin and Ker (1998) evaluate this issue using parametric and nonparametric tests for heteroscedasticity. Their results indicate that the standard deviation of yields tended to be proportional to the average yield. This suggests a somewhat different approach to normalizing yields.

Assume that the following functional relationship is appropriate for detrending a temporal series of crop yields and has been estimated for a series of yields ranging from t=1970, …, 2003:

$$y_t = X_t \beta + e_t,$$

where X_t represents some linear or nonlinear function of time. Estimates of such a relationship will yield trend-predicted yields (y_t) and deviations from the trend (e_t). If one wants to normalize yields to a 2003 level and one believes the magnitude of the errors is not affected by the level of yields, one would add all of the residuals to the 2003 yield prediction, such that:

(normalized yield)$_t$ = y_{2003} + e_t.

However, if one believed that the deviations from the trend tended to be proportional to the level of yields, one might consider constructing normalized yields as:

(normalized yield)$_t$ = $y_{2003}(1+e_t/y_t)$.

Both approaches have been adopted in the literature though the latter probably has more empirical support. One must balance the complexity of empirical techniques that may be more appropriate against practical considerations associated with modeling what is often a very large collection of yield trends.[7]

We have not assumed that yield trends are linear in nature. In many cases, the exact functional form for trend effects is unknown and thus nonlinear functions including higher-ordered polynomials, splines, and nonparametric regression procedures may be used to represent time trends. In the case of the U.S. group risk (GRP) program, a linear

---

[6] For a detailed discussion of this literature and references to specific applications, see Goodwin and Ker (2002).

[7] A number of other subtle issues underlie the normalization of yields. Departures from normality, which often underlie considerations of modeling yield risk, may suggest that least squares techniques are less than optimal since least squares tends to make detrended yields (residuals from the OLS regression) more symmetric than the population error distribution. This issue is discussed in detail by Goodwin and Ker (1998), who conclude that a quadratic loss function of the sort that underlies least squares regression may be merited on economic grounds. One could also argue that any suspected heteroscedasticity should be modeled in the detrending equation directly rather than accounting for it after estimation.

spline with a single knot point is used to detrend county-level yields. Of course, it is always possible to improve the fit of any regression by adding additional terms (or in the case of nonparametric regression, by lowering the bandwidth). Put differently, deviations from trend can always be made smaller by increasing the order of a polynomial regression of yields on time. One must take care to avoid overfitting when estimating trends and other deterministic factors. One approach to guarding against overfitting when detrending yield data is to choose the function on the basis of *out-of-sample* predictive power rather than in-sample performance. One might fit the regression model on one portion of the sample and predict the values of the other portion. Alternatively, one might drop out each observation, fit the model, and predict the omitted observations, an approach known as cross-validation. Semiparametric and nonparametric methods are well-known for their overfitting problems and thus care should be used when modeling trend or other deterministic factors.

3.2. Spatial Correlation

Variations in crop yields tend to be driven by factors that typically affect a large area, such as weather and pest infestations. This raises the issue of systemic (i.e. covariate) risk or, put differently, spatial correlation of yields.[8] The fact that average yields tend to be subject to substantial correlation across space suggests that standard appeals to normality based upon central limit theorem results are likely to be invalid. Goodwin (2001) demonstrates that yield correlation tends to die off slowly over space. His results also suggest that the degree of correlation may be state-dependent, i.e., yields tend to be more highly correlated over space as extreme yield events (e.g., drought) occur.

The issue of spatial yield correlation has many other important implications for the construction of yield insurance contracts. Systemic risk may suggest that the potential for catastrophic losses is higher than what one might reveal in a short sample of data. The degree to which risks in a large insurance pool can be diversified is diminished as yields become more spatially correlated.

In addition to spatial correlation, it is possible that the effects of exogenous factors that underlie realized yields may persist beyond a single production cycle. For example, the effects of a severe drought on crop yields in one year may persist into the next year as soil moisture levels remain below average. Likewise, the outbreak of a disease or pest infestation may affect yields over multiple growing seasons, suggesting that low yields this year may be followed by low yields again next year. The potential for such persistence is an agronomic issue and is likely to be crop- and region-specific. To the extent that the effects of weather or other exogenous factors persist across years, it may be necessary to include forms of autoregressive or moving average effects in any detrending equation. In doing so, care must be exercised in using the results since one may want predictions of future yields to account for any such effects but may not want to condition out the effects of such factors when detrending yields, since deviations from trend for indemnifiable reasons should not be accounted for when modeling deviations from trend. Evidence regarding the extent to which yield shortfalls may persist beyond a single growing cycle is unclear, though some authors have obtained results that suggest such persistence may indeed occur (Goodwin and Ker 1998).

---

[8] This correlation tends to decrease as there are elevation variations (e.g., hills, valleys).

3.3. *Yield Modeling: Parametric vs Nonparametric Methods*

Inherent in the design of an insurance contract is a mechanism for determining the probability of loss and the expected level of loss when losses occur. More formally, one is generally interested in a measure of the probability density function (pdf) underlying the event or events that trigger losses. Thus, the concept of modeling yield risk for the purposes of designing and rating a crop insurance contract is fully analogous to modeling the probability distribution for the crop yield in question.

At its most fundamental level, the task of modeling yield risk involves obtaining estimates of parameters or descriptions of patterns depicting the stochastic nature of random yields. In formal terms, one is interested in measuring the probability density or distribution function of yields. Approaches to modeling yield densities and distributions fall into two broad areas. The first involves estimation of the parameters of a parametric distribution. The second general approach encompasses the wide variety of nonparametric methods that are commonly used to measure or approximate yield distributions. Alternative semi-parametric methods may combine elements of both parametric and nonparametric methods (by, for example, using a high-ordered polynomial series expansion to approximate an unknown function). In what follows, we provide an overview of parametric and nonparametric methods commonly used in estimating crop yield densities.

Parametric methods

Parametric approaches to the estimation of probability densities generally involve using an observed series of yield realizations to estimate specific parameters that describe a probability density or distribution function. Examples include the normal distribution, a two parameter distribution that is completely described by its mean and variance. Given a set of independent and identically distributed (iid) yield realizations (a critical assumption for yield data, which is discussed below) one can estimate the parameters of the distribution using maximum likelihood or method of moments estimation procedures. Of course, such an approach to modeling yields assumes that one knows *a priori* the correct parametric family underlying the yield data generating process. In reality, there is usually very little to guide one in choosing a specific parametric distribution, though certain characteristics of yield distributions (such as negative skewness) tend to suggest specific candidates.

A number of specific parametric distributions have been considered for modeling crop yield distributions. A case can be made for the appropriateness of the normal distribution on the basis of the central limit theorem (CLT), since a degree of aggregation is almost always present in crop yields such that one is measuring average yields. However, a problem arises when considering central limit theorem results for crop yields because of the systemic (covariate) component. The risks affecting agricultural yields are typically of a systemic nature in that they arise from weather, pests, or other natural phenomena that tend to affect large geographic areas. The central limit theorem, in its most basic form, is based upon the independence of identically distributed individual random draws, a situation that is unlikely to occur for crop yields. The spatial correlation that affects crop yields and results in systemic risks suggests that simple versions of the CLT may be invalid when invoked for crop yields. Weakest versions of the CLT that

allow for dependence and non identical distributions exist, though requirements that such dependence decays over space limits their suitability for crop yield modeling.

Box 1. Basic statistical measures of risk

Consider a sample of n observations, y_i, for i=1 to n. The mean (or average) is the value that is

expected to occur from a draw from the distribution, i.e., $\mu = \frac{1}{n}\sum_{i=1}^{n} y_i$.

The second moment about the mean is measured using the variance, which describes the amount

of variation about the mean exhibited by the distribution, $\sigma^2 = \frac{1}{n}\sum_{i=1}^{n}(y_i - \mu)^2$

Skewness is determined by the third moment about the mean, $\frac{1}{n\sigma^3}\sum_{i=1}^{n}(y_i - \mu)^3$. It is a measure

of the extent to which a distribution is pulled to the right (negative skewness) or to the left (positive skewness). A completely symmetric distribution has zero skewness.
Kurtosis refers to the mass that is present in the tails of the distribution and is often defined

relative to the normal distribution (which has a kurtosis of 3), $\frac{1}{n\sigma^4}\sum_{i=1}^{n}(y_i - \mu)^4 - 3$. Positive

values of this metric are said to indicate excess kurtosis and densities are said to exhibit leptokurtosis (for heavy tails) and platykurtosis (for lighter tails).
Source: authors.

Many researchers have observed that crop yields tend to be negatively skewed, with yields near the maximum being observed more frequently than yields near the minimum. In this light, a frequent choice for modeling yield distributions is the four-parameter beta distribution. The beta distribution can accommodate the skewness so commonly observed for crop yields and can assume a variety of shapes. This distribution does suffer from a number of shortcomings, however. Although the distribution has four parameters (two shape parameters and a maximum and minimum possible yield), it is commonly applied with only the shape parameters being estimated. The maximum and minimum yields are generally set by assumption or in some ad hoc manner.[9]

Other parametric distributions that have been used to model crop yields include the Weibull, the log-normal, the gamma, the logistic, versions of the Burr distribution, and mixtures of parametric distributions. These parametric distributions vary in terms of their flexibility and ability to capture intrinsic properties of various crop yields and thus they differ in terms of their appropriateness for modeling crop yield densities. For example, the log-normal distribution imposes positive skewness on the distribution, a characteristic that is not typically expected for crop yields.

---

[9] Ker and Coble (2003) show that shape parameter estimates may be very sensitive to assumptions about the maximum and minimum possible yields. They also present evidence suggesting that estimation of three or four parameter beta distributions may be less efficient than nonparametric or semi-parametric alternatives.

In summary, intuition about the nature of crop production (i.e., the fact that biological constraints limit the maximum possible yield, with yields close to the maximum being observed more frequently than very low yields) suggests that yield distributions are likely to be negatively skewed. Likewise, much of the existing empirical evidence about crop yield distributions has confirmed the prevalence of negative skewness. Several distributions are compatible with negative skewness, including the beta, gamma, and Burr distributions. Our recommendation would be that such distributions should be the first choice in cases where a parametric distribution must be adopted.

Nonparametric methods

An alternative approach to modeling crop yield distributions is to adopt a nonparametric type of estimator. The simplest estimator is the histogram. The range of the data is divided into bins of a certain width and the yield observations falling into each bin are counted. An obvious problem with such an approach involves the determination of the appropriate bin width and placement of bins. Alternative choices can have a substantial impact on the shape of the estimated density, especially in cases of small sample sizes.

In cases where one is working with large dataset to estimate insurance premium rates or other parameters of an insurance contract, a simple nonparametric approach may involve simply counting the number of positive events that occur in a large sample (to gain a measure of the probability of the event) or measuring the average level of loss (to gain a measure of the expected loss). Such a method has distinct advantages in that no assumptions regarding the underlying probability distribution are needed. This empirical approach is commonly referred to as using "empirical estimates" of premium rates or expected loss, and is used in rating the U.S. federal crop insurance program.

The obvious shortcoming of this empirical approach is its reliance on a large sample and the potential difficulty in measuring rare events. For example, using simple binomial probabilities, the likelihood of observing a 1 in 100 year event one or more times in a sample of 100 years of data is only 0.63. When one is interested in measuring the risks associated with rare, catastrophic events, using empirical likelihood measures is problematic. In this light, such empirical measures are often combined with other approaches such as catastrophic loading or risk pooling to measure catastrophic risks and the likelihood of other rare events.

An alternative approach to measuring yield risks involves the application of nonparametric kernel methods (see Box 2). A shortcoming associated with kernel density estimation techniques pertains to their lack of structure relative to parametric methods. By its very nature, a nonparametric technique imposes minimal structure on the estimated distribution. If one knows the appropriate parametric structure (i.e., the proper distribution function), it is more efficient from a statistical perspective to impose this structure when estimating the distribution. However, an inappropriate choice of a parametric distribution may result in biased estimates of the distribution and thus inaccurate insurance premium rates. This tradeoff between efficiency and bias is omnipresent in considerations of the proper specification of econometric or statistical models and one may be willing to accept biases if the gains in efficiency are large (and vice versa).

15

Box 2. Nonparametric kernel methods

The kernel density estimator places an individual kernel at each observation from the density of interest. The area under these kernels at each point is summed to give an estimate of the density at that point. The area under a properly defined kernel must integrate to one. Many specific functions, including properly defined probability density functions (pdf), are suitable as kernels and research has demonstrated that the estimated density is not sensitive to the choice of kernel function in moderately sized samples.

More important than the specific function chosen for the kernel is the width of the kernel. This width, which corresponds to the variance of the density if a pdf is used for the kernel, is called the "bandwidth" parameter. It is fully analogous to the bin width chosen in the construction of a histogram. A variety of methods have been developed for choosing an optimal bandwidth. The simplest involves a "rule of thumb," which Silverman (1986) has shown to perform well as the data deviate from normality. Alternative methods including various plug-in approaches and cross-validation procedures may offer advantages as data depart radically from normality. Nonparametric kernel density estimation techniques were applied to rate area-wide crop insurance contracts by Goodwin and Ker (1998). Ker and Goodwin (2000) considered a number of extensions to the nonparametric approaches, including methods that allowed the bandwidth parameter to vary across different points of the density and empirical Bayes type methods that utilized aggregate data to construct priors for use in the estimation of nonparametric densities.

Source: authors.

Therefore, a tradeoff exists between the efficiency associated with imposing structure in the estimation of risk versus the potential for biases when one imposes an inappropriate parametric structure. Rarely does the analyst know *a priori* the correct parametric form. Nonparametric procedures, while offering substantially more flexibility, necessarily require more observations in order to obtain reliable probabilistic estimates.

Precise recommendations regarding which approach to use are difficult to make since the approach is likely to depend on the specific characteristics of the problem. In general, if one has strong prior knowledge of the statistical distribution that is likely to describe yields or if one is working with small data sets (of less than 30-40 observations), a parametric approach is likely to be preferred. In cases where a large amount of data is available and one is unsure of the appropriate distribution, a nonparametric approach offers advantages in terms of its flexibility.

4. Using Risk Models to Develop Crop Insurance Contracts

We consider the simple case of a yield insurance contract that pays indemnities at a predetermined, fixed price if realized yields fall beneath some threshold that defines a guarantee. Such a contract has characterized the U.S. multiple peril crop insurance program over most of its existence. Two fundamental parameters are inherent in such an insurance contract. First, one must establish the guarantee that determines the conditions under which indemnity payments will be paid. The yield guarantee establishes total liability (i.e., the maximum possible indemnity or, equivalently, the amount of indemnity paid in the event of a total loss). Second, one must establish the appropriate price (premium) that should be charged for the coverage offered under the contract. An error in either parameter can be costly for the insurer and/or the insured.

16

Both the premium, which reflects the likelihood and expected level of loss, and the insurance guarantee, which is generally set to reflect the expected yield, are parameters that depend upon the underlying yield distribution. Thus, accuracy in the determination of insurance contract parameters is intimately linked to an accurate representation of the distribution of yields. Most crop insurance contracts establish the yield guarantee to reflect a proportion of the expected yield. In the U.S. crop insurance program, growers are able to select to insure from 50-85% of their expected yields. Expected yields are generally measured by taking an average of previous realized yields. In the case of the U.S. program, expected yields are the simple average of the preceding 4-10 years of yields.

The insurance premium is usually expressed as a "premium rate" which represents the dollars paid in premium for each dollar of liability (i.e., as a proportion or percentage of total liability). Abstracting from administrative and program costs, an actuarially fair premium will equal the expected insured loss (expected indemnities). Thus, the insurance premium rate will be expressed as the ratio of expected loss to total liability. Note that the price at which indemnities (yield losses) will be paid is irrelevant to the risk of the policy (assuming that risk is not endogenous to price) and thus to the premium rate. Total premiums and total liability will move proportionally as price is changed and thus their ratio, the premium rate, will be unchanged by price changes.

Consider an insurance contract that will pay indemnities anytime yields fall beneath a certain proportion λ of the expected insured yield μ. Expected insured loss will be given by the product of the probability that a loss will be realized times the expected loss, given that a loss occurs. In other words,

Expected insured loss $(y) = E\max[\lambda\mu - y, 0] = \text{Prob}[y < \lambda\mu][\lambda\mu - E(y|y < \lambda\mu)]$.

Assuming one has an adequate representation of the yield density $f(y)$, this is given by:

$$\int_{-\infty}^{\lambda\mu} f(y)dy(\lambda\mu - \int_{-\infty}^{\lambda\mu} yf(y)dy / \int_{-\infty}^{\lambda\mu} f(y)dy), \text{ where } \mu = \int_{-\infty}^{\infty} yf(y)dy.$$

This discussion demonstrates the importance of adequate measurement of the probability distribution function for yields. All key parameters of the insurance contract are related to such a measurement of yield risk. Though straightforward in theory, a number of complicating issues may arise when this is put into practice with actual yield data.

In order to illustrate these concepts, a set of 21 years of regional yield data for groundnuts in an important producing country was assembled. No time trend was found. The average yield is 751.3 kg/ha, the standard deviation is 304.2 kg/ha and the skewness is slightly negative (-0.14). Maximum likelihood procedures were used to estimate parameters of a beta density (the minimum possible yield was fixed at zero while the maximum possible yield was chosen to maximize the likelihood function). In addition, a normal distribution was used to model insurance contract parameters (using the mean and variance calculated from the 21 yield observations). Finally, nonparametric kernel density

estimation techniques were used to estimate the yield density.[10] It should be noted that this sample is too small to derive reliable estimates of a density function and thus it is only used here to illustrate the method.

As shown on Figure 1, the three alternative density estimates are very similar to one another, though subtle differences do exist. The beta distribution is constrained to lie above zero while no such constraint is imposed on the normal density and the kernel estimate. We work only with that part of the densities that lie above zero (i.e., insurance contract estimates are taken from truncated densities).

We consider insurance contracts providing a λ percent insurance coverage, with $\lambda = 50\%$, 65%, 75% and 85%, i.e., that will pay indemnities for any yield realization that is below λ % of the average (375.5 kg/ha, 488.2 kg/ha, 563.3 kg/ha and 638.4 kg/ha, respectively). Estimates of relevant yield insurance parameters are presented in Table 1. In spite of the visual similarity of the alternative densities, substantial differences are apparent for the estimated insurance parameters.

The probability of loss is higher under the kernel estimation when the coverage is 50% and 65%. It is higher under the Beta distribution for larger coverage levels. This is the direct consequence of the "bump" in the lower tail of the kernel density function (see Figure 1). The expected loss conditional on the occurrence of a loss is higher under the kernel estimation. However, the difference with the parametric (normal and beta) estimations decreases as the insurance coverage increases. The premium rate, expressed in percent of liability, varies accordingly. For 50 percent coverage, it is 3.50% under the Beta distribution, 4.19% under the Normal distribution and 6.01% under the kernel estimation. For 85% coverage, the premium rate varies from 11.47 (Normal distribution) to 12.91% (kernel estimation).

---

[10] The nonparametric estimates used a bandwidth chosen to minimize the approximate mean integrated square error under the assumption that the true distribution is normal.

Figure 1. Groundnut density estimates: Normal, Beta, and Kernel

Curves: Normal (Mu=751.33 Sigma=304.23)   Beta(Theta=0 Scale=2038 a=3.44 b=5.96)
          Kernel (c=0.79)
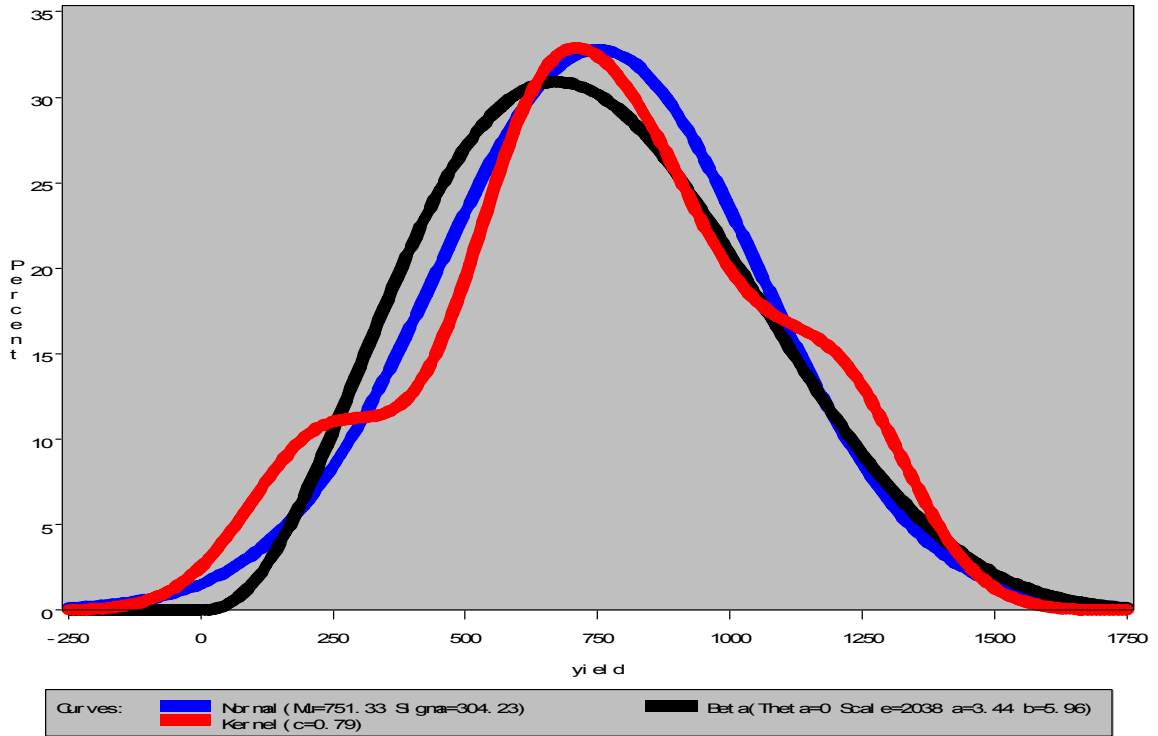
Table 1. Groundnut Crop Insurance Parameter Estimates

Value	Beta	Normal	Kernel
Probability of Loss (%)			
50 percent coverage	11.23	10.85	13.78
65 percent coverage	12.76	19.37	20.41
75 percent coverage	30.20	26.85	27.03
85 percent coverage	39.27	35.55	35.62
Expected Loss | Loss Occurs (Kg)			
50 percent coverage	98.06	138.46	163.90
65 percent coverage	137.25	167.92	203.33
75 percent coverage	165.21	191.46	218.64
85 percent coverage	192.81	208.73	231.43
Premium Rate (%)			
50 percent coverage	3.50	4.19	6.01
65 percent coverage	6.16	6.72	8.50
75 percent coverage	8.77	8.89	10.49
85 percent coverage	11.79	11.47	12.91

Source: Authors.

This example illustrated the application of models of yield risk to a simple problem of rating an all-peril yield insurance contract. Essential to constructing accurate insurance contract parameters is an adequate measure of the risks underlying random yields. Conceptually, these risks are represented by yield density estimates. As discussed previously, a number of different parametric and nonparametric techniques are available to estimate yield densities. Alternative approaches each have relative merits and shortcomings. Our application to groundnut yield data demonstrates that insurance parameters such as premium rates are likely to be quite sensitive to the assumptions made in modeling yield distributions.

5. Crop Insurance Ratemaking Procedures

Two ratemaking procedures can be considered, depending on the previous existence of the insurance product. If the insurance scheme has been offered for several years (at least 15-20 years), as it is the case many crops covered by the U.S. Multi-peril crop insurance (MPCI) program, then historical indemnity and liability data can be used to assess the premium rate through the loss cost ratio. The main advantage of this procedure is that changes in farming practice once the farmer is insured (for each level of coverage), i.e., moral hazard, is captured in the expected value of future costs. When a new insurance product is launched, such historical data are not available and thus the ratemaking procedure relies on the yield data available and on simulated losses.

5.1. *Rating an Existing Insurance Scheme*

It is first noteworthy that there is no single ratemaking approach that will apply to all insurance contracts. Insurance companies have derived their own ratemaking procedures. The case of U.S. MPCI scheme is detailed here to illustrate how public-private multi-peril crop insurance contracts are rated. The involvement of the Federal government in the MPCI program creates a social insurance program which operates on different principles than a privately underwritten insurance market. For example, insurers are required to accept all applicants. The U.S. Department of Agriculture (USDA) is mandated by law to set a base premium rate with an expected loss ratio of 107.5 percent. Subsidies are applied to this rate. All base rates are either produced or approved by the Risk Management Agency (RMA) and companies are not allowed to load or modified the RMA-produced rate. As a branch of the federal government, the USDA fully subsidizes the administrative delivery costs associated with this program. Companies are thus compensated for delivery expense. Currently, Congress has authorized an expense reimbursement of 24.5% of the premium. However, the actual expenses of the crop insurance industry may exceed the amount of the expense reimbursement (Deloitte & Touche LLP 204). These costs are excluded from the rate paid by the producers.

Crop insurance ratemaking is a particular case of property and casualty (P&C) insurance ratemaking. Unlike standard P&C contract, MPCI coverage is not triggered by an event. Instead, the indemnity payment is determined after the crop is harvested. As for many other insurance coverages, the ratemaking process for MPCI has evolved over time. Josephson et al. (2000) and Schnapp et al. (2000) provide a detailed analysis of the current MPCI ratemaking procedures. The ratemaking process relies on two important

values: liability and indemnity. Liability is a measure of the insurer's exposure to loss. It represents the total insured value of the crop, defined as

Liability = acres planted x APH yield x selected coverage level x base price x price election.

The APH yield, based on the producer's actual production history, represents the producer's normal yield defined as the 4-10 year average yield. The coverage level represents the producer's deductible, with a coverage level of 75% meaning that indemnities are paid only if the actual yield is less than 75% of the APH yield. Coverage levels are usually available between 50% and 85%, in 5% increments. The base price represents the price of the crop at the start of the growing season, as established by RMA. The producer may choose to insure production at any level (i.e., price election) from 60% to 100% of the base price of the crop.

Indemnity is the amount paid under the MPCI coverage for a producer suffering a covered loss. Indemnity is paid when the value of production is less than the liability purchased. The amount of indemnity is thus equal to the difference between the liability purchased and the value of production, i.e.,

Indemnity = acres planted x base price x price election x max[APH yield x selected coverage level - actual yield,0].

The premium paid by a producer is thus derived as

Premium = liability x rate x adjustment factor.

Box 3 illustrates the steps in determining the liability, premium and loss.

Box 3. Multiple peril crop insurance: an illustration

Step 1. Liability

Acres planted	500 acres
Farmer's actual production history	140 bushels/acre
Coverage level selection	75%
Base price	USD 3
Price election	100%
Liability	USD 157,500 = 500 x 140 x 75% x 3 x 100%

Step 2. Premium

Liability	USD 157,500
Premium rate	3%
Premium	USD 4,725 = 157,500 x 3%
Premium subsidy	40%
Producer paid premium	USD 2,835 = 4,725 x (1 – 40%)

Step 3. Loss

Actual yield	95 bushel/acre
Value of production	USD 142,500 = 95 x 500 x 3 x 100%
Indemnity payment	USD 15,000 = 157,500 – 142,500

The producer expects to harvest 70,000 bushels and insures 52,500 bushels. The actual production is 47,500 bushels. The indemnity is thus equal to 5,000 bushels evaluated at the base price.

Sources: Authors, from Schnapp et all (2000).

The objective of the ratemaking process is to derive the premium rate that is used in the above formula. Fundamentally, a premium rate is an estimate of the expected value of future costs. One of its largest components is typically the provision for losses. This component is derived from the expected loss component, which is represented by the loss cost ratio (LCR), i.e., indemnity divided by liability. The LCR is a measure of loss per unit of exposure. Because loss costs vary between crops and locations, the MPCI procedure establishes rates from each crop and each county. This empirical approach based on observed losses typically precludes from the use of individual data because it is rare that a single insured is sufficiently large to allow the insurer to derive the LCR solely from the insured's loss history. Consequently, the aggregate experience of a group of similar risks is considered in developing premium rates. The main five steps of the MPCI procedure are detailed below (Josephson et al. 2000)

Adjusting loss and exposure to a common coverage level

MPCI is offered at various coverage levels, generally from 50% to 85%. In order to use extensively the historical data, the first step in the ratemaking process is to adjust the data to a common coverage level (CCL) set at 65%.

The adjustment of the liability simply requires one to multiply the current liability by the ratio of the common coverage level to the CCL. For example, to adjust the liability from 75% to 65%, one multiplies the current liability by the ratio 0.65/0.75.

The adjustment of the indemnity is less simple and two cases need to be considered. For coverage levels higher than the CCL, the adjusted indemnity is equal to the unadjusted indemnity minus the reduction in liability, for all actual yields (in percentage of APH yield) lower than the CCL. The adjusted indemnity is zero for other actual yields. For coverage levels lower than the CCL, the adjustment procedure is more difficult because only the individual insured loss is known. For example, adjusting from a 50 percent coverage level to a 65 percent coverage level is problematic because actual yields higher than 50% of the AHP yields did not generate indemnity payments. RMA has developed an adjustment procedure based on an interpolation between minimum and maximum bounds.

An alternative adjustment process would be to treat each coverage level separately and to develop rates for each of them. While this segregation process may capture the differences in loss experience for producers selecting different coverage levels, it can result is less stability and predictability of expected losses due to the greater statistical variation. The adjustment process used by RMA is considered actuarially sound by several experts (see, for example, Josephson et al 2000).

Derivation of county unloaded base rates

The adjusted LCR is derived from the adjusted indemnity and the adjusted liability, for each year. It reflects the percentage of liability that would have been paid to producers if they had all purchased 65% coverage insurance.

The first issue is about the number of years to be used in the LCR projection. The year to year variability in loss costs and the long term trend have to be taken into account. The MPCI process uses years from 1975. The long-term average LCR is simply based on

the average LCR over the period 1975 and subsequent, with equal weighting of all years. An alternative to equal weighting would be liability weighted LCR.

Regarding excess loss adjustment, RMA has developed a procedure in order to reduce the impact of a single year on the average loss cost of each county. Under this procedure, the adjusted average LCR for any single year is capped at the 80[th] percentile LCR of all years. Excess indemnity is then the amount of indemnity that is excluded from the capped LCR. However, this 80/20 rule may be too evasive for some crops/states to capture the impact of catastrophes and, therefore, a threshold varying by state and by crop should be considered.

The final predictive value of the loss experience of a specific county is based on the LCR of this county, and the simple circle LCR that is the average of surroundings counties' LCRs weighted by each county's total adjusted liability. Weighting together different LCR estimates is part of the credibility process (see Box 4). The MPCI procedure uses the credibility formula $Z = \sqrt{\min(1, P/271)} \times 60\%$, where P is the number of insurance policies indemnified. A county with 271 or more policies indemnified will receive the maximum credibility of 60%. The remaining is assigned to the simple circle LCR. The county unloaded rate is thus defined as the weighted sum of the county's LCR and the simple circle LCR. It should be noticed that the number of claims, while being a common measure of credibility for property insurance ratemaking, has no clear justification. In addition, this number may need to vary, by state or region, to reflect the expected claim frequency.

Box 4. Credibility

Credibility is a measure of the predictive value of the loss experience. The credibility process is the weighting together of different estimates to come up with a combined estimate. The general credibility formula is

I=Z x I_1 + (100% - Z) x I_2

where   I is the item that is estimated (e.g., county LCR)

            I_1 is an indicator of the item experience (e.g., county average LCR)

            Z is the credibility attached to I_1

            I_2 is an alternative indicator (e.g., simple circle LCR).

Credibility Z ranges from 100% (full credibility) to 0% (no credibility). For an item with credibility less than 100%, an appropriate item to receive the remaining wight must be identified.

Sources: Authors, from Josephson et al. (2000).

Base rate loading

The county unloaded rate is an estimate of the expected capped LCR for the forthcoming year. It is then adjusted by several factors.

First, a disaster reserve factor, equal to 13.64% of the unloaded rate, is added. This reserve factor is intended to meet the Congressional requirements to maintain "reasonable reserves". RMA ha defined a reasonable reserve as an amount sufficient to achieve financial adequacy over a 10 year period at a 85% confidence level.

Second, a state excess load adds a provision for the excess indemnity excluded from the LCR in the unload rate calculation. It is equal to the state excess indemnity (i.e., the sum of the county excess indemnities at the state level) divided by the state adjusted liability (i.e., the sum of the county adjusted liability), with a minimum of 1% (cupping) and a maximum of 5% (capping). The excess indemnity for each county is thus pooled for the entire state and then spread back to each county. The capping process is actuarially justified by the fact that catastrophic events lack statistical credibility.

Third, a prevented planting load is added as a provision for losses due to crops never planted because of external factors which are not directly related to yield loss.

Finally, the rate loaded by the previous components is loaded by the unit division load, which takes into account that indemnity is not computed for an entire farm but rather for each division of farm.

Capping rate changes

The implied base rate derived above is divided by the current base rate to produce a normalized loss ratio. This loss ratio is compared to a pre-determined schedule of rate changes centered at the mandated target loss ratio of 107.5%. The current base rate is then adjusted to meet this target loss ratio, if necessary. However, public policy considerations force the rate change to range between –5% and 10% of the current base rate. The imposition of such limits, based on political considerations, is to stabilize premiums paid by producers from year to year.

Update

The rates developed to this point have used aggregate data. They are then adjusted to take into account individual characteristics: farming practices, coverage level and average yield. Studies have shown that producers with lower APH yields also have significantly higher variability of yields. For this reason, different Rate Classes are established for producers with different APH yields. Finally, other expenses and capital costs are added.

APH premium rates are calculated by multiplying the premium rate at 65 percent coverage level by a rate relativity factor that is the same for all crops and all regions. For example, to find a 75 percent APH rate, one multiplies the 65 percent premium rate by 1.53. Babcock, Hart and Hayes (2002) show that these constant rate relativities may be inconsistent with the underlying probability distribution of yields and therefore premium rates may not be actuarially sound. They show that farmers who plant crops in region where the 65 percent pure premium rate is above 3 percent face unsubsidized premiums that are too high at the 80 percent coverage level. This discrepancy may explain why 65 percent coverage was the most popular coverage level in higher-risk areas.

5.2. Rating a New Insurance Scheme

In the absence of historical indemnity data, indemnity payments of a new crop insurance scheme can be simulated from the yield data. As discussed in Section 3, the data generating process of yields can be estimated through parametric or non-parametric methods, depending on the data available and a priori beliefs. The distribution function of indemnity payments is thus derived from the distribution function of yields. However, this procedure does not capture the impact of insurance on farming practices, contrary to the ratemaking procedure described earlier. In other words, the yields observed are those

of uninsured farmers, but they may have been different if these farmers were insured. This is the standard problem of moral hazard.

In this simple form, the insurance premium can be decomposed as

Premium = pure premium + reserve load + administrative cost + return on equity.

We focus here on the first two components, i.e., the pure premium and the reserve load. Expressed in percentage of the liability, the pure premium is the expected LCR. The reserve load is the cost of reserves the insurer must set aside in order to pay unexpected losses, with a predefined degree of confidence (e.g., 99%).[11] The reserves, or minimum capital requirements, are derived from the probability distribution of insured losses. The reserve hold is thus the opportunity cost of holding these reserves. The two components of the insurance premium are illustrated using the previous example on groundnut crop insurance (see Table 1). The groundnut yield density is estimated using the non-parametric kernel procedure. The reserves are built to pay losses caused by events with a 100 year return period. The opportunity cost of capital is 5%. Table 2 shows the pure premium rate (i.e., expected LCR), the minimum capital requirements (MCR) in percent of liability and the reserve load in percent of liability, for different levels of insurance coverage.

Table 2. Pure premium, MCR reserve load, in percent of liability - groundnut crop insurance scheme.

Insurance coverage	Pure premium	MCR	Reserve load
50 percent coverage	6.01%	50.82%	2.54%
65 percent coverage	8.50%	62.17%	3.11%
75 percent coverage	10.49%	67.21%	3.36%
85 percent coverage	12.91%	71.07%	3.55%

Source: Authors.

The reserve load can represent a significant fraction of the premium rate. In the above example, it represents 42% (resp. 37%, 32% and 28%) of the pure premium rate, under 50 (resp. 65, 75, 85) percent coverage. The lower the coverage, the higher is the reserve load as a percentage of the pure premium rate. In other words, under catastrophic coverage, the gross premium rate becomes more and more driven by the reserve load.

---

[11] Part of these unexpected losses are usually transferred through reinsurance treaties.

6. Extending Coverage to Price Risk: Crop Revenue Insurance

Analysis of agricultural risk is usually concerned with the effects of yield and price uncertainty on revenues. Revenue insurance plans account for the risks associated with uncertain yields and prices. They are used to guarantee producers a minimum level of gross revenues. A number of revenue insurance products have been offered to the U.S. farmers since the late 1990s and have grown greatly in prominence, surpassing standard multiple peril contracts in terms of total liability in many areas for many crops.[12] These plans had their beginnings in crop insurance legislation in 1994 that mandated development of a cost of production plan. A private company offered an insurance scheme that would pay indemnities at harvest time prices rather than at the predicted price, which was determined prior to planting. This product was developed into a revenue insurance plan known as Crop Revenue Coverage (CRC). Indemnities under the CRC plan could be triggered by low prices, low yields, or a combination of the two. In addition, lost yields would be valued at the higher of the predicted or harvest time prices. CRC was the first form of revenue insurance and accounted for a large proportion of liability after it was introduced. Alternative forms of crop insurance were also introduced in the late 1990s. Income Protection (IP) provided a true form of revenue insurance where guarantees were based upon the product of predicted prices and yields. Revenue Assurance (RA) also provided revenue coverage. RA was later expanded to include a "harvest price option" that operated much like CRC in that lost yields were paid at the higher of the predicted or harvest time price. This form of coverage has proven to be quite popular in the U.S.

A key issue is the modeling of revenue risk. Joint distributions and accounting for correlation between two random variables must be addressed. Precise measures of yield risk (the yield density) and price risk (the price density) are first needed. However, the measure of these marginal densities is still incomplete for capturing revenue risk, as yields and prices are rarely independent. To the extent that yields are spatially correlated, yield shortfalls on an individual farm are likely to be correlated with higher prices, meaning that yields and prices tend to be negatively correlated. Any consideration of revenue risk that utilizes measures of price and yield risk must also consider the degree of correlation between yields and prices. In the case of a negative correlation between yields and prices, a "natural hedge" exists such that revenue shortfalls associated with low yields are naturally offset by higher prices. This implies lower risk and premiums for revenue coverage than would be obtained from independently insuring price and yield. Measurement of correlation may not be straightforward since annual data are used (since contracts are annual) and thus data may be limited. Correlation is also location-specific. The empirical evidence suggests that price and yield will be much more highly correlated in major growing regions than is the case as one moves outside of such regions. The U.S. Revenue Assurance plan adopted an innovative parametric assumption when modeling revenue risks. Insurance rates for a particular level of coverage from the standard yield crop insurance program are used as a base measure of risks. A beta distribution is assumed to be appropriate for yields. Using assumptions about the minimum and maximum possible yields, the parameters of the beta function are first estimated. This

---

[12] In 2003, revenue insurance products represented about 60% of total crop insurance premium volume and about 55% of total acres insureds (RMA 2004).

26

yields a marginal distribution for yields at the farm level. The options market provides a direct means of measuring the marginal distribution of prices. However, it is again necessary to assume a particular parametric distribution. A near universal assumption regarding commodity prices is lognormality.[13] Revenue risk is then simulated from the marginal yield and price distribution and assumptions about the price-yield correlation (see Box 5).

Box 5. Simulating crop revenue

With knowledge of the marginal distributions and degree of correlation of prices and yields, it is possible to simulate the joint distribution in order to obtain simulated revenues. Two procedures have been utilized in the literature to accomplish random sampling of correlated random variables from specified marginal distributions.

The first of these procedures is often called the "weighted linear combination" approach, developed by Johnson and Tenenbein (1981). This approach essentially involves a translation scheme whereby functions of uncorrelated random draws from two different marginals (say for price and yield) are combined to form a linear combination which yields bivariate distributions with the appropriate degree of correlation. Essentially, this procedure is analogous to taking random draws from a known distribution (such as a normal), constructing a weighted linear combination that will result in variables with a known degree of correlation, translating these to a uniform using the cumulative distribution function (cdf) and then using these uniform random deviates in the inverse cdf for each variable to generate random draws from the desired marginals. A limitation associated with this approache is that it is fundamentally parametric; a priori knowledge of the appropriate marginal distributions and the degree of correlation is necessary in order to implement the procedures. Thus, the aforementioned limitations associated with parametric modeling of price and yield distributions are again relevant. Measurement of correlation is especially important as the joint distribution will be significantly influenced by the extent of correlation among the individual variables. The Johnson and Tenebein (1981) procedure is currently used to draw correlated yields (from a beta marginal distribution) and prices (from a log-normal marginal distribution) in constructing rates for crop revenue insurance. These rates are quite sensitive to assumptions regarding the degree of correlation between yields and prices. In particular, rates fall substantially as this correlation increases and thus accurate rates are critically dependent upon measurement of this correlation. Unfortunately, the amount of data available for measuring this correlation is often quite limited, and thus measures of this correlation may be inaccurate. A more fundamental limitation of this method involves the fact that, although the techniques are able to match first and second moments, including covariances, there are no assurances that higher ordered "cross-moments" are matched.

The second procedure is the method of copulas. A copula is a function that relates marginal distributions to a joint distribution function. It may hold some promise for measuring risk in multivariate relationships. Frees and Valdez (1998) discuss techniques for identifying copulas for a sample of data.

Source: Authors, from Goodwin and Ker (2002).

The design and rating of such contracts presents a number of other challenges that are unique to revenue plans. For example, Mahul (2003) and Mahul and Wright (2003) examine the design of optimal revenue insurance contracts and show how crop revenue

---

[13] The Black-Scholes option pricing formula implies lognormality for prices.

insurance, crop yield insurance and commodities futures and options tend to be complementary in efficient risk financing strategies.

7. Conclusions

A comprehensive review of the data that are available and credible is a fundamental step in the design and rating of any insurance scheme. Individual coverage is typically preferred by farmers. However, the data needed to derive accurate insurance rates for individual farmers are almost never available. In the event that consistent data are lacking at the individual level, two options are available. One may wish to use aggregate data in an attempt to measure individual contract parameters. An alternative approach is to consider coverage based upon observable indices that reflect as close as possible the loss incurred by the farmer. Examples include area-yield insurance and weather-based insurance.

The availability and credibility of data are also a central issue in production risk modeling. In cases where a significant amount of reliable data is available, one may choose to adopt nonparametric approaches to modeling yield risks. More common, however, is the case where data are limited. In such cases, parametric methods may be preferable. The choice essentially involves a tradeoff between bias and efficiency.

Ratemaking procedures are also based on data availability. When historical losses are available, rating procedures are based on the historical loss cost ratios. In contrast, the rating of a new insurance scheme in the absence of historical losses usually relies on simulated losses derived from crop yield models.

This paper has laid out recommendations on the design and rating of viable agricultural insurance policies. They rely on recent statistical and actuarial developments as well as empirical work based on the U.S. experience. They provide Bank staff and policymakers with technical guiding principles for the development of viable agricultural insurance programs in developing countries.

Finally, prior to any considerations regarding the contract design, one must investigate the demand for the insurance plan, i.e., willingness and ability to pay on the part of those targeted by the insurance plan. This issue has not been addressed in this paper and is left to subsequent work. An actuarially sound crop insurance scheme will fail if farmers have limited interest (e.g., because of risk management alternatives) and/or limited financial resources. This raises the tradeoff between actuarial soundness and effectiveness of the insurance plan for protecting against individual risks.

References

Babcock, B., C. Hart and D. Hayes. 2002. "Crop Insurance Rates and the Laws of Probability." Center for Agricultural and rural Development, Iowa State University, working paper 02-WP 298.

Deloitte & Touche LLP. 2004. "Federal Crop Insurance Program: Profitability and Effectiveness Analysis." Prepared on Behalf of the National Crop Insurance Services, Inc.

Doherty, N. and O. Mahul. 2001. "Moral Hazard in Insurance Securitization," Insurance Project Workshop, National Bureau of Economic Research, Boston MA, (February 2001).

European Commission. 2001. "Risk Management Tools for EU Agriculture with a Special Focus on Insurance," European Commission Working Document.

FAO. 1991. *Strategies for Crop Insurance Planning*, FAO Agricultural Services Bulletin.

Frees. E.W. and E.A. Valdez. 1998. "Understanding Relationships Using Copulas," *North American Actuarial Journal* 2: 1-25.

Goodwin, B. K. 2001. "Problems with Market Insurance in Agriculture," *American Journal of Agricultural Economics* 83(3):643-49

Goodwin, B. K. 1993. "An Empirical Analysis of the Demand for Multiple Peril Crop Insurance," *American Journal of Agricultural Economics* 75(2): 425-434.

Goodwin, B.K. and A.P. Ker. 1998. "Nonparametric Estimation of Crop Yield Distributions: Implications for Rating Group-Risk (GRP) Crop Insurance Contracts." *American Journal of Agricultural Economics* 80(1):139-153.

Goodwin, B. K. and A. P. Ker. 2002. "Modeling Price and Yield Risk," *A Comprehensive Assessment of the Role of Risk in US Agriculture.* (Norwell MA: Kluwer), p.289-323.

Goodwin, B. K. and V. H. Smith. 1996. *The Economics of Crop Insurance and Disaster Relief.* Washington, DC: American Enterprise Institute Press.

Holmström, B. 1982. "Moral Hazard in Teams," *The Bell Journal of Economics* (13): 324-340.

Johnson, M. and A. Tenenbein. 1981. "A Bivariate Distribution Family with Specified Marginals," *Journal of the American Statistical Association* 76(1):198-201.

Josephson, G.R., R. B. Lord, and C. W. Mitchell. 2000. "Actuarial Documentation of Multiple Peril Crop Insurance Rate Making Procedures," Report of Milliman and Robertson, Inc. prepared for the Risk Management Agency.

Ker, A.P. and B.K. Goodwin. 2000. "Nonparametric Estimation of Crop Insurance Rates Revisited." *American Journal of Agricultural Economics* 83 (2):463:478.

Ker, A. P. and K. H. Coble. 2003. "Modeling Conditional Yield Densities," *American Journal of Agricultural Economics*, 85(2):291-304.

Mahul, O. 2002. "Hedging in Futures and Options Markets with Basis Risk," *Journal of Futures Markets* 22(1): 59-72.

Mahul, O. 2003. "Hedging Price Risk in the Presence of Crop Yield and Revenue Insurance," *European Review of Agricultural Economics* 30 (1): 1-23.

Mahul, O. and B.D. Wright. 2003. "Designing Optimal Crop Revenue Insurance, *American Journal of Agricultural Economics*," 85(3): 580-589.

Silverman, B. 1986. *Density Estimation for Statistics and Data Analysis*. New York: Chapman and Hall.

Schnapp F., J. Driscoll, T. Zacharias, and R. Josephson. 2000. "Ratemaking Considerations for Multiple Peril Crop Insurance," Report prepared for the Risk Management Agency.

Appendix. Information and Insurance Contract Design: Conceptual Issues

Informational Asymmetries

If the risks associated with individual insuring agents is heterogeneous and the insurer is unable to accurately discern each individual's level of risk, individual insurance premium rates (the price of insurance, expressed as a percentage of liability) will not be accurate. If agents are overcharged for insurance relative to their level of risk, they will be less likely to insure. Likewise, if high-risk agents are undercharged for their risk, they will be more likely to buy insurance. Thus, if the price of insurance is based upon some average level of risk, participation in the insurance pool will be distorted in favor of the higher risk individuals. Such a situation is commonly called "adverse selection" since the insurance pool will tend to be comprised of agents with higher risks. Adverse selection is one important consequence of inaccurate premium rates.

When one considers the dynamic operation of an insurance program over time, the consequences of adverse selection become clearer. The insurer will lose money since premiums are inadequate to cover the indemnities being paid out. The insurer will likely respond by raising the overall level of premiums charged for insurance. However, this action causes the insurance pool to be made up of a smaller number of riskier individuals since lower risk agents will be priced out of the insurance plan. The risk of the pool rises and the plan loses even more. Such a cycle continues until the program fails entirely. This process is commonly referred to as the "Death Spiral of Adverse Selection." Goodwin (2001) provided empirical evidence suggesting that such adverse selection and its consequences were present in the U.S. crop insurance program.

A closely related problem in the design and operation of insurance programs involves the extent to which the insurer is able to monitor the actions of insuring agents. If agents undertake actions that alter their risks of collecting indemnity payments and if the insurer is unable to monitor and adequately price such actions, rates may not reflect the risks of loss and the program may again lose money. A situation where risks are endogenous to the insurance decision is commonly referred to as "moral hazard." Such moral hazard effects are present in almost every insurance transaction, though careful rating and monitoring as well as careful judgment in adjusting for losses are key to minimizing the effects of moral hazard. Goodwin and Smith (1996) demonstrated that substantial moral hazard effects are present in the U.S. crop insurance program. In particular, they found that farmers buying crop insurance tended to use less fertilizer and chemicals on their wheat crop than did those farmers that did not insure.

A number of other problems that are indirectly related to the measurement of risks are also important to the design and operation of crop insurance programs. The potential for fraud (e.g., claiming losses when none occurred or over-reporting existing losses) is often an issue important to any insurance program. The accuracy of the loss adjustment process and the extent to which losses can be verified to have occurred because of hazards that are covered under the plan are important aspects of any insurance program. The marketing and servicing aspects of an insurance program may be important to the reception of the program by potential insureds and the degree of participation that the program generates. These issues are secondary to the key issue addressed in this paper.

They are, however, important to the overall success of any insurance program.

Informative signals

The ultimate objective of insurance is to provide indemnity payment that fit the individual losses. Under most of the existing crop insurance schemes, the indemnity schedule is based on the farmer's individual yield. This tailored insurance scheme is exposed to all types of informational asymmetries (moral hazard, fraud, adverse selection). These potential sources of uninsurability are managed through the introduction of deductibles and/or coinsurance rates. However, they create large administrative costs that lead to high insurance premiums.

Because of the lack of individual data, alternative insurance contracts have been implemented where the indemnity is triggered by an index that is correlated with, but not identical to, the individual loss. For instance, the area-based index is used as a proxy of the average yield in a given geographical area. It has been implemented in India since 1985 and in the USA since 1994. More recently, weather-based indices have been suggested as alternative triggers.[14]

There is a growing debate among experts about the feasibility and the performance of such contracts. In our opinion, the critical issue is not about the insurance contract design but about the information conveyed by the index. In other words, the primary step is to analyze how valuable is the information conveyed by the signal on the individual loss and what is its cost. The efficient index or set of indices that should be used in an insurance contract will be the result of this cost-benefit analysis. In order to set the limits of this analysis, consider the following two indices: an individual index, perfectly corrected with the individual loss but available at prohibitive costs, and an aggregate index perfectly uncorrelated with the individual loss and available for free. The latter is obviously worthless in the design of an insurance policy because it brings no information on the individual loss. The former is perfectly informative of the individual loss but its costs are so high that the insurance premium becomes unaffordable. In both cases, these indices cannot be used in the design of an insurance policy.

The concept of informative signal is broader than the concept of correlation. There may exist a non-linear relationship between the individual loss and the index and the concept of correlation, based on an implicit linear relationship, cannot capture it. This is particularly the case for a weather-based index that usually affects the individual loss in a non-linear manner. The value of the informative index relies on its ability to explain the variability of the individual loss. The remaining part of the individual loss variability that cannot be explained by the signal is the basis risk.

Some index characteristics are listed in Table A1 for three different types of triggers: individual trigger, area-based trigger and parametric (e.g., weather-based) trigger. Each of these characteristics generates costs for the insured farmer. If the trigger is not easily comprehensible, the farmer will have to spend time and money to understand it. This criteria is particularly relevant in developing countries where the fraction of illiterate people may be high. The presence of (uninsurable) basis risk is costly for the risk-averse

---

[14] The Bank has assessed the feasibility of weather-based index insurance in several developing countries (Morocco, Tunisia, Nicaragua, Mexico…).

farmer. Delays in settlement may create opportunity costs. Informational asymmetries (moral hazard and adverse selection) are also costly because they cause agency costs (and thus higher premiums) and/or the existence of a minimum deductible and/or coinsurance forces the farmer to retain a fraction of the loss. Unreliable data, index manipulation and administrative costs create additional costs that are passed to the insured farmers through a higher premium. Table A1 clearly shows that there is not a single index that complies with all the above-mentioned characteristics. An individual trigger is not exposed to basis risk and is easily comprehensible, but it can be manipulated and subject to moral hazard and high administrative costs. On the contrary, a parametric trigger, like a rainfall index, is not exposed to delays in settlement or high administrative costs, but it faces basis risk.[15] As a consequence, an effective index is the result of the multiple tradeoffs.

Table A1. Index effectiveness

	Indemnity trigger	Area-based trigger	Parametric trigger
Easily comprehensible	Yes	Often	Yes
Basis risk	No	Yes	Yes
Delays in settlement	Sometimes	Yes	Rarely
Moral hazard/adverse selection	Yes	Sometimes	No
Reliable data	Rarely	Rarely	Often
Index manipulation	Yes	Yes	Sometimes
Administrative costs	High	Low	Very low

The result of these tradeoff may lead to different choices depending on the insured's characteristics, and specifically his/her aversion towards risks. This is illustrated in Figure A1. The individual preferences are mapped in the mean-variance model. Point A represents the initial position of the farmer when no insurance is available. His expected profit is high, but the variance of his profit is also high. When an individual insurance policy is available, the farmer can reduce the variability of his profit by purchasing insurance but, because insurance is sold costly (the insurance premium is higher than the actuarially fair premium), this involves a loss in expected profit. The curve A-C represents the efficient frontier under the individual insurance policy. Notice that full insurance is not available, i.e., point C leaves the farmer exposed to basis risk (e.g., because of the presence of a minimum deductible in order to deal with moral hazard). The curve A-B is the efficient frontier under an index-based insurance contract. Because such a contract is less costly than an individual policy, so the premium is, and thus any decrease in the variance of profit generates a lower loss in expected profit. As a consequence, the curve A-B is above the curve A-C. However, because the index is less

---

[15] Such an index can also be manipulated by a third party.

informative of the individual loss, the remaining basis risk is larger. Consequently, a low risk-averse farmer will choose the index-based insurance policy (point E), while a high risk-averse farmer will purchase an individual insurance contract (point D).

Figure A1. Tradeoff between basis risk and transaction costs
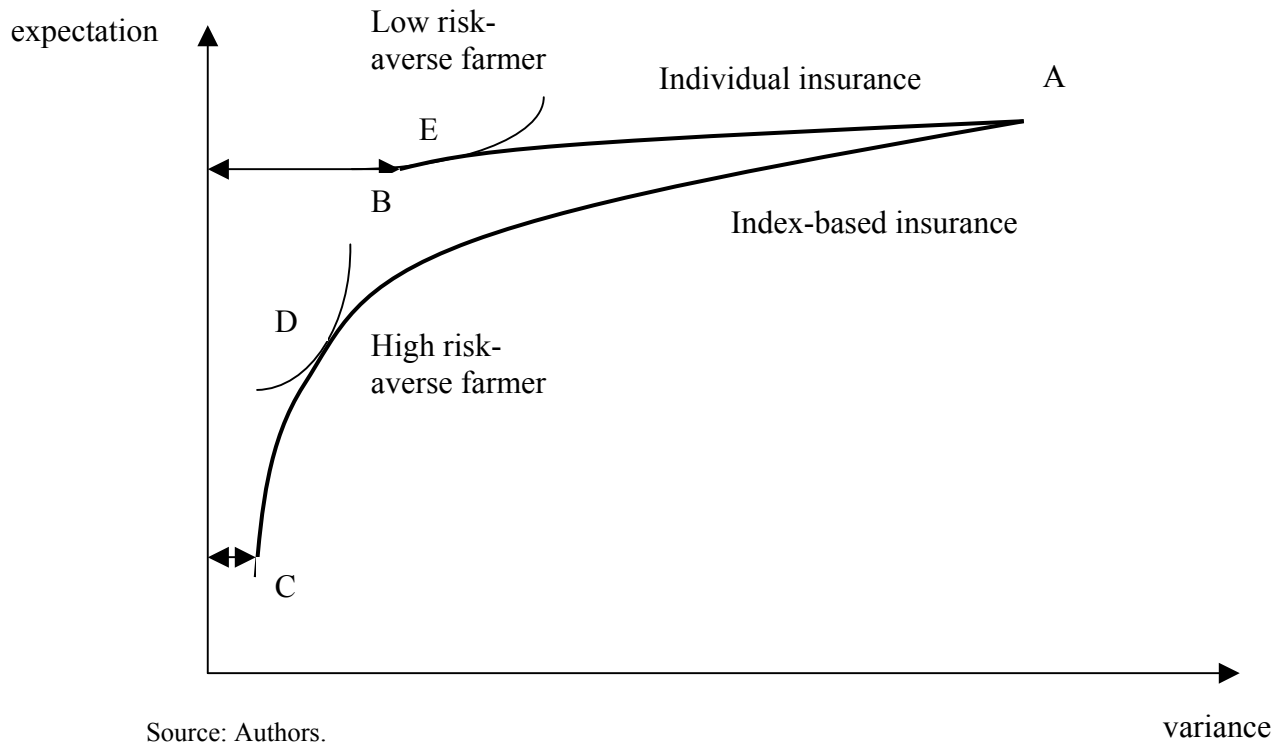
Source: Authors.

Given informative signals, an optimal insurance contract can be designed. It critically depends on the relationship between the individual crop yield and the vector of proxies taken into account in the indemnity schedule. Such a relationship can be captured by a pseudo production loss function where the output is the individual yield and the inputs are the informative signals (see Box A1).

Box A1. Optimal index-based insurance contract design

The design of an efficient insurance contract design relies first on the characterization of the relationship between the individual yield and the proxies available for the indemnity schedule. For example, if the indemnity function is based on rainfall and temperature, then one must first define the pseudo-production function where the rainfall level and the temperature levels are the main inputs. This thus gives the estimated individual yield function. Formally, define y_i the individual crop yield loss, **x** the vector of indices (rainfall levels, temperature…) and ε_i the individual error term (yield basis risk). From an appropriate characterization of the pseudo-

34

production function $f_i(.)$, we have $y_i = f_i(\mathrm{x}, \varepsilon_i)$. Mahul (2002) shows that, when the insurance premium depends only on the actuarial value of the policy, the optimal insurance contract design is given by

$$ I_i(\mathrm{x}) = \left[1 + \Psi \sigma_{\varepsilon_i}^2\right] \max\left(X_i - Ef_i(\mathrm{x}, \widetilde{\varepsilon}_i), 0\right), $$

where $\sigma_{\varepsilon_i}^2$ is the variance of the individual residual, X_i is the strike and Ψ is an index reflecting the farmer's behavior toward risk (called prudence). Indemnities are thus paid when the expected yield conditional on the observed indices is lower than the strike.

This expression is interesting in several points. First, the cost of insurance will decrease the yield guarantee but will not affect (at least in the primary effect) the coverage under the yield guarantee. Second, when positive, the indemnity depends on the conditional expected yield $Ef_i(x, \widetilde{\varepsilon}_i)$, the variability of the individual yield basis risk and the farmer's aversion to this error term. The higher the variance of the basis risk and/or the aversion toward this variance, the higher then marginal insurance coverage. This provides a rationale for the well-known argument stating that in the presence of an (uninsurable) basis risk over-indemnification is optimal. If the farmer preferences are based only the first two moments of the yield distribution and/or there is no yield basis risk, then the optimal indemnity schedule, when positive, is equal to the expected estimated yield, under the strike.

However, this optimal insurance contract design based on multiple indices may be difficult to implement because of the potentially complex estimated pseudo-production function. The efficient insurance indemnity should thus rely on the tradeoff between the simplicity of the indemnity function and its accuracy.

Sources: Authors from Mahul (2002).

The presence of basis risk is usually considered as one of the main obstacles that makes the index-based insurance product unattractive for farmers. However, the concept of basis risk needs to be discussed in the context of production risk. Basis risk initially refers to the difference between the spot price of asset to be hedged and the futures price of contract used (Hull 2000). Price basis risk is thus the remaining part of variability of the spot asset that cannot be explained by the variability of the futures price. It is usually captured using regression models that are linear in the price basis risk. Yield basis risk can be defined as the yield variability that cannot be explained by the variability of informative signals. This type of basis risk has several sources: (i) heterogeneity among crop yields, (ii) systemic risks not captured by the proxies (iii) misspecification of the pseudo-production function, (iv) errors in yield measurement. The first source of basis risk is common to all index-based contracts and it due to heterogeneous characteristics among farmers (e.g., soil, farming practices). This risk is likely to be independent among farmers and thus it may be pooled among a large number of farmers. The second source is related to parametric index based insurance. For example, under rainfall insurance, the rainfall index captures the systemic loss caused by excessive or deficient rainfall but it does not capture other systemic sources of risk (temperature, wind…) that also affect farmers at the same time.[16] Contrary to basis risk (i), this source of basis risk is systemic

---

[16] Part of the variability can be captured by the index if the index is correlated with the other sources of risk.

and thus it cannot be pooled among farmers. It is noteworthy that area yield index are not exposed to this type of basis risk because, by definition, area yield captures all the sources of systemic risk. Basis risk (iii) is due to an inappropriate characterization of the relationship between crop yield and the (informative) indices. Finally, basis risk (iv) is common to all individual or index-based crop yield insurance contracts. It is caused by the difficulty to properly assess crop yields. Other sources of potential basis risk are uncertainty in the settlement of claims, risk of manipulation of the data, etc.

Combining individual and index-based contracts

Following the development and the implementation of index-based contracts in the late 1990s, there has been a growing interest in new insurance products that could wrap individual insurance around the index-based insurance. The objective of such contracts would be to offer a coverage against the yield basis risk. They are called gap insurance contract, basis risk insurance contract or "wrap-around" insurance contract. It is first noteworthy that this issue is different of that examined in the preceding sub-section because the individual loss is now allowed to be included in the insurance indemnity schedule. In fact, this issue is not new in the economic literature and it was first examined by Holmstrom (1979) to deal with moral hazard in a principal-agent model. He shows that additional signals will only be included in the optimal risk-sharing contract if they are informative of the agent's hidden action. More recently, Doherty and Mahul (2000) reconsider this problem and propose to make a distinction between signals that are directly informative of the agent's hidden action and signals that are directly informative of the loss (e.g., correlated with the individual loss). Area-based or weather-based indices are examples of signal directly informative of the individual loss.

The optimal risk-financing strategy relying on index-based contract and individual policy will be the result of the tradeoffs between the (marginal) insurance/hedging costs. While the individual policy is not exposed to basis risk, it should be more costly than the index-based contracts because of, e.g., the agency costs to deal with moral hazard, loss adjustment, and the substantially greater record keeping burden associated with individual coverage.[17] In this case, the optimal risk financing decisions are based on a two-step procedure. First, the index-based contract is design to cover the component of the individual loss that is correlated with the indices. In the second step, the remaining risk is covered with the individual policy. The optimal design of both individual and index-based contracts depends on technical constraints that are discussed in Box A2. It is of interest to notice that the optimal individual contract design displays a contingent strike. The higher the index-based indemnity is, the lower the contingent strike under the individual insurance policy.

Box A2. Efficient index-based and gap insurance contracts

We consider the pseudo-production function $y_i = f_i(\mathbf{x}, \varepsilon_i)$, where y_i is the individual crop yield, $\mathbf{x}$ is the vector of indices and ε_i is the yield basis risk. The farmer can cover a loss in crop yields through an index-based contract under which the indemnity function depends on indices $\mathbf{x}$

---

[17] Obviously, if the (marginal) cost of individual insurance is always lower than the (marginal) cost of index-based contracts, only the individual coverage is purchased.

and/or an individual insurance contract under which the indemnity function can be based on both the individual yield and the indices. We assume that the cost of one unit of coverage is more expensive under the individual policy than under the index-based contract because of, e.g., the presence of agency costs in the individual policy. The optimal risk financing strategy of the farmer is a two-step strategy. Formally, the indemnity function of the index-based contracts is

$$I_i(\mathrm{x}) = \max\big(X_i - Ef_i(\mathrm{x}, \widetilde{\varepsilon}_i), 0\big),$$

where X_i is the strike.

In the second step, the remaining risk, $y_i + I(x)$, is (partly) covered using the gap insurance policy. Two cases are considered. When the indemnity payments must be non-negative, the optimal indemnity schedule is

$$J_i(y_i, \mathrm{x}) = \max\big(Y_i - I_i(\mathrm{x}) - y_i, 0\big),$$

where Y_i is the strike. This contract thus displays a strike $\big[Y_i - I_i(\mathrm{x})\big]$ that is contingent on the realization of the indices. If the indemnity payments can be negative, i.e., the farmer may face an *ex post* premium adjustment based on the realization of the index, then the indemnity schedule is

$$J_i(y_i, \mathrm{x}) = \max\big(Y_i - y_i, 0\big) - I_i(\mathrm{x}).$$

Source: Authors.

This analytical discussion on the optimal combination of insurance/hedging instruments raises an additional issue regarding the constraints on the design of the individual insurance policy. This indemnity schedule of the individual contract can depend either on both the individual yield and the indices or on individual yield only. In the first case, the individual contract acts as a gap insurance policy dealing only with the remaining part of the individual yield not covered by the index-based contract. In the second case, such a distinction is not possible and thus both individual and index-based contracts deal with the systemic component of the risk captured by the indices. Allowing gap insurance is obviously a more efficient strategy because it relies on an efficient decomposition of the individual yield.