Rainfall Insurance for Midwest Crop Production

Selected Paper for AAEA Annual Meetings 2001

Dmitry V. Vedenov and Mario J. Miranda *

Correspondence to: Dmitry Vedenov, Dept. of AED Economics,
2120 Fyffe Rd., Rm. 103, Columbus, OH 43210, e-mail: vedenov.2@osu.edu

May 7, 2001

Abstract

The paper discusses a methodology for design and pricing of index insurance contracts for crop production. The methodology heavily relies on establishing a relationship between the index and yields in order to evaluate the contract performance in hedging farmers' risk. However, analysis of yield/rainfall data series for Iowa corn and Kansas wheat fail to produce a reliable and meaningful relationship which can be used uniformly across several counties and/or crop producing districts. Further research is needed as to applicability of rainfall insurance to specific crop/region combinations.

Keywords: risk management, weather derivatives, index contracts, crop insurance

---

Rainfall Insurance for Midwest Crop Production

Introduction

Agricultural production has always been a risky endeavor. Farmers constantly have to deal with unfavorable weather conditions, variability in prices of inputs and outputs, livestock disease outbreaks, pests, etc. The uncertainty of future incomes complicates both short-term production decisions and long-term planning (e.g. expansion of production or capital investments in machinery and equipment). It also renders lending institutions less willing to provide loans to farmers, since the probability of default is relatively high. Although some forms of self-insurance may be available to farmers (e.g. crop diversification or intertemporal income transfers), these have certain limitations and ultimately reduce farm profits in the long term.

Of all the risk factors affecting agricultural production, especially crop production, weather is typically the most significant. Weather phenomena are hard to predict (at least in the beginning of the growing season) and even harder to mitigate against. Moreover, since unfavorable weather conditions such as floods or droughts often affect large areas, the risks faced by different producers are correlated. The latter presents a stumbling block to traditional insurance, which is designed to pool a large number of small unrelated risks rather than handle widespread simultaneous (catastrophic) losses.

In a perfect world, risk sharing would be complete and the cost of ceding would be close to the actuarially-fair rate. High levels of catastrophe risk retention by retail insurance and high reinsurance rates are evidence of incomplete or imperfect markets. Applied to crop insurance, market imperfection translates into unacceptably high rates or even complete lack of risk protection for farmers.

Traditionally, unwillingness or inability of insurance markets to provide risk management mechanisms for agricultural crop production has prompted many governments to step in with various forms of support of agricultural producers (subsidized loans, price-support programs, etc.). Unfortunately, the government support programs are often inefficient and come at high social cost (Skees, Hazell, and Miranda).

An emerging trend in recent years has been to develop new financial instruments (catastrophe options, catastrophe bonds) which would allow insurers to securitize correlated risks and circumvent the limitations of traditional insurance market. In application to agricultural insurance, the innovations include area-yield insurance program and various exchange-traded area-yield contracts. A characteristic feature of these instruments is that their payoff depends on values of a specially designed measure, or *index*, related to the risk being hedged against.

The main advantage of indices is that they can be measured objectively and do not depend on individual actions of market participants. Transparency of index-based contracts along with the fact that they are uncorrelated with the traditional financial instruments (stocks, bonds, etc.) make them attractive to investors outside of insurance industry. Insurers benefit by getting access to additional funds that can be used to indemnify large simultaneous losses caused by natural disasters.

In the same way, using index-based instruments in agricultural risk management (e.g. rainfall insurance, heating degree-days contracts, area-yield contracts, etc.) allows to circumvent problems faced by traditional insurance and provide farmers with an efficient hedge of weather-related risks. However, a design of an index-based insurance contract or financial instrument requires answering very important questions, such as what variable to use as an index, how to structure the indemnity schedule, how to price a contract, how to sell it, and whom to sell it to.

General Approach to Design and Pricing of Index Contracts

In order to create an efficient index insurance contract for agricultural crop production, the following principles should be adhered to. First, the contract should have a relatively simple and transparent structure. An individual buying the contract should not be confused as to when and under what circumstances the contract will pay off.

Second, the index should be easily observable and measurable on a regular basis. Further, there should be enough historical observations of the index in order to derive its distribution and relation to the actual risk being insured.

Third, using an index to compute indemnities rather than actual losses inevitably introduces some basis risk. Therefore, careful consideration should be given as to what variable (or combination of variables) to use as an index in order to minimize the basis risk to the extent possible.

Fourth, given the traditional accounting practices of insurance companies, the premium rate and the risk exposure borne by the company issuing the contract should be clear.

While an index contract may be structured in many different ways with different coverage layers and provisions, we will focus our attention on a specific class of elementary contract. Specifically, an elementary contract pays an indemnity $f(\tilde{\iota})$ conditional on realization of the index $\tilde{\iota}$ according to the following schedule (figure 1)

$$
f(\tilde{\iota}) = \begin{cases} x, & \text{if } \tilde{\iota} \leq \lambda i^*; \\ x\dfrac{i^* - \tilde{\iota}}{(1 - \lambda)i^*}, & \text{if } \lambda i^* < \tilde{\iota} \leq i^*; \\ 0, & \text{if } i^* < \tilde{\iota}, \end{cases} \tag{1}
$$

In other words, the contract pays whenever the index $\tilde{\imath}$ falls below a specified trigger i^*, with the indemnity proportional to the difference between the index and the trigger. The maximum indemnity $\$x$ is paid whenever the index falls below a critical value λi^*, $0 \le \lambda \le 1$.

Although developed independently, this setup is somewhat similar to that used by Martin, Barnett, and Cobble. The main differences are that we emphasize the insurance aspect of the contract and specify the payment structure so that indemnities are paid for rain deficiency rather than excess rainfall. This approach seems to be more suitable for crop production where drought negatively affects the harvest-time yield.

Elementary contracts in (1) are convenient for analysis, yet offer enough flexibility to construct more complicated instruments. Combining elementary contracts with different triggers i^*, limit parameters λ, and maximum liabilities x, one can recreate or otherwise approximate more complicated, multi-layered indemnification schedules that may provide efficient risk protection whenever expected losses are not linearly related to index. The elementary contract also contains the simple "all-or-nothing" contract as a special case. Specifically, when $\lambda = 1$, the contract pays the maximum indemnity if the index falls below the trigger level i^*, but pays nothing otherwise.

In order to specify a particular contract, we need to impose three conditions on the contract parameters i^*, λ, and x. These conditions can be chosen so that the contract has pre-specified properties. For the purposes of our analysis, it will be convenient to further standardize the contracts under consideration by requiring them all to have an expected indemnity, or pure premium, of $\$1$. From the buyer's standpoint this is a convenient normalization, since it allows him to see readily how much protection he can buy for $\$1$ in pure premium by inspecting the indemnity

schedule. The normalization, however, is not restrictive, since one can achieve any coverage level desired simply by buy the necessary number of $1 contracts.

We shall refer to an elementary contract with a pure premium of $1 as a *standard contract*. For the standard contract, the parameters must be chosen so that

$$1 = \mathbf{E}_{\widetilde{\iota}}\, f(\widetilde{\iota}\,; i^*, \lambda, x) = x \left(\int_0^{\lambda i^*} h(i)di + \int_{\lambda i^*}^{i^*} \frac{i^* - i}{(1-\lambda)i^*} h(i)di \right). \tag{2}$$

where $h(i)$ is the probability density function of the index $\widetilde{\iota}$.

Insurance companies usually design contracts based on a premium rate, which is the ratio between the premium and largest risk on line. Because a standard contract has a pure premium of $1, the premium rate is simply equal to

$$\pi = \frac{1}{x}, \tag{3}$$

where x is the maximum liability. In other words, fixing the premium rate of a standard contract immediately fixes its maximum liability.

Thus, a standard contract can be uniquely identified by its premium rate π and limit parameter λ. Condition (3) fixes the maximum liability x given the premium rate, and condition (2) implicitly fixes the trigger i^* given the limit parameter λ.

One of our goals is to develop a systematic approach to designing standard index contracts that are optimal, in some sense, for potential buyers. Let us look at the standard index contract from a buyer's standpoint. Assume that the individual's income $\widetilde{r}$ subject to an index-related risk can be expressed as a function g of the index $\widetilde{\iota}$ and an independent random shock ε

$$\widetilde{r} = g(\widetilde{\iota}\,; \varepsilon), \tag{4}$$

5

where ε essentially represents the basis risk.

In case of agricultural insurance, we assume that the buyer is a farmer involved in agricultural production, i.e. growing crops. The farmer's income, or net revenue $\widetilde{r}$ can be calculated as $\widetilde{r} = \widetilde{p} \cdot \widetilde{y} - c$, where $\widetilde{p}$ is the harvest time price, $\widetilde{y}$ is the crop yield, and c is the total production cost. Depending on the particular situation, the risk faced by the farmer may be caused by either uncertainty about yields or prices.

Further, assume that the individual has some target income level r^* he wishes to protect. The target income may be some fraction (e.g. 75%) of the expected income $\overline{r} = \mathbf{E}\, g(\widetilde{\iota}; \varepsilon)$ or some other income level (e.g. the level at which the individual breaks even). Income lower than the target is considered to be a loss

$$\widetilde{L} = \max\{0, r^* - g(\widetilde{\iota}; \varepsilon)\}.$$

If the individual buys N standard contracts defined by the parameters π and λ and conditions (2) and (3), his total loss *with contracts* is then

$$\widetilde{L}_c = \max\left\{0, r^* - [g(\widetilde{\iota}; \varepsilon) + N f(\widetilde{\iota}; \pi, \lambda) - N]\right\}.$$

We assume that the individual tries to avoid the downside loss at all states of nature and therefore wants to minimize his total expected root-mean square (RMS) loss. Hence, the optimal number N^* of standard contracts as well as the optimal parameters π and λ can be determined as a solution to the optimization problem

$$\min_{\{N, \pi, \lambda\}} \mathbf{E}_\varepsilon \left(\int_0^\infty [\max\{0, r^* - [g(i; \varepsilon) + N f(i; \pi, \lambda) - N]\}]^2\, h(i) di \right)^{1/2}. \qquad (5)$$

Given the above considerations, the process of designing and pricing of a standard

index contract for a representative buyer can be outlined as follows.

First, given the distribution of the underlying index $\widetilde{\iota}$ and relation between the index and buyer's risk exposure $\widetilde{L}$ (determined by the function g and the target income level r^*), the optimal number N^* of standard contracts and the contract parameters are determined by solving (5). The buyer is then offered a composite contract consisting of N^* standard contracts with the specific parameters i^*, λ, and x.

The standard contract structure provides a convenient basis for comparing contracts with different premium rates and triggers. Along with specifying the risk an insurance company undertakes, it also gives the buyer an opportunity to determine which contract provides the best coverage for the same price. The RMS loss measure is a transparent selection criterion, which allows rank-ordering of various contracts available on the market.

The composite contract does not necessarily need to consist of N^* identical contracts. In principle, one can combine contracts that differ in structure. For instance, we might consider a combination of contracts with different triggers and limit parameters to obtain coverage over several risk layers simultaneously. However, finding an optimal combination of multiple contracts of different structures renders the optimization problem more difficult by increasing the number of variables to be chosen. Alternatively, a more complicated contract may be first constructed as a weighted average of several standard contracts of different structure (in order to preserve the fixed premium rate and unit price) and then the composite contract may be constructed according to (5).

Note that instead of finding an optimal number of contracts N^*, the problem may also be reformulated in terms of the amount of money the buyer is prepared to pay

7

for insurance. In this case, the number of contracts is fixed, and the optimization problem in (5) becomes a condition on only the parameters λ and π. While fixing the commitment of the buyer may not result in the best available risk hedging, such an analysis may be important if buyer operates under tight budget constraints and cannot afford the best available insurance coverage. However, the pre-specified number of contracts may not necessarily be optimal in the sense that the buyer sometimes can achieve a greater risk reduction by actually *reducing* his commitment level.

Basis Risk

No matter how accurately the relation in (4) is estimated, there are ultimately some risks that cannot be hedged against by using an index contract. In case of crop production, harvest at a particular farm depends on a variety of factors, such as the soil moisture at planting time, the amount of rainfall, the temperature patterns during the growing season, and the amount of fertilizers in the soil. While one index, e.g., amount of rainfall, may account for several risk factors, there is always some risk attributed to other influential factors that are not correlated with the index. In other words, unless the index exactly reflects the risk exposure of the buyer, there is always some basis risk present. The latter may have several components, one caused by a nondeterministic relation between the index and targeted variable (random shock ε in (4)), as well as those caused by other factors.

The major issue in specifying the details of the contract is the trade-off between transparency of the contract structure and amount of basis risk. More specifically, there are three types of basis risk we need to be concerned about — temporal, spatial, and crop-specific.

Temporal Component

The sensitivity of yield to climatic conditions varies over the stages of growth. A typical growth cycle can be divided into four phenological periods, *viz.* germination, bloom, development, and maturity. Across the periods, weather parameters such as rainfall, amount of sunshine, etc., have different effects on the prospective yield. Too much rain during germination may slow down the overall plant growth, while drought during the bloom and/or maturity may prevent crops from realizing their full potential. Historical weather patterns are also different during each period.

As an extreme case, one can create several different contracts for each of the phenological periods, and thus reduce the temporal basis risk to minimum. However such an arrangement would involve added transaction costs for both sellers and buyers (marketing, monitoring, etc.). Alternatively, an aggregate index may be created based on amount of rainfall or sunshine during the whole growing season. The transaction cost involved in marketing such a contract would be lower than for several distinct contracts. However, the basis risk embodied by the single contract would be larger and the index variable may be less transparent.

Spatial Component

Weather patterns differ across locations within the same region. Measurements at one station may track precipitation level or temperature at nearby farms very closely, but diverge considerably from observations at farms located farther away. Thus, a contract based on measurements at one station would bear very little basis risk for some farmers, but perform very poorly for others. As before, an ideal solution to the problem would be to create a separate contract for each station in the region. However, this may result in dozens of distinct contracts marketed within a single

crop producing district or even a county.

A more reasonable approach is either to choose a central (in some sense) station as a reference point for the entire region, or create a weighted average of the observations at different stations. The latter may not necessarily be a more efficient alternative, since computing the average would require coordination of data collection from several different locations. In addition, the beginning of the growing season may vary slightly from one part of the region to another. Therefore, the temporal component of the basis risk may also be involved and further complicate the matter.

Crop-Specific Component

Different crops vary in their sensitivity to rainfall, duration of the growing season, and planting times. Obviously, a single contract cannot provide an optimal protection for all crops. Therefore, either the contracts should be specifically tailored for each crop, or a series of contracts should be developed for the whole season providing different levels of protection for different periods within the season. Once again, the crop-specific component of the basis risk may interact with temporal and spatial components in a complex way.

Ultimately, the decision as to how much basis risk to sacrifice for contract transparency is up to the insurance company issuing the contract and the buyer purchasing it. An efficient market will eventually ensure that only the contracts which provide an optimal basis risk/transparency combination survive and trade actively.

Implementation Issues

The suggested methodology assumes that the index distribution $h(i)$ and the relation between the index and income $g(\tilde{\iota}; \varepsilon)$ are known. In practice, however, only particular realizations of $\tilde{\iota}$ and $\tilde{r}$ are often available. In this case, the distribution of

the index may be estimated by using either nonparametric techniques (e.g. kernel-smoothing) or by fitting the observed series by one of the standard distributions (e.g. by the maximum likelihood method). The function g also has to be estimated from the available data, with the appropriate functional form chosen based on either agronomic or statistical considerations. Both estimations may present a challenge in a practical situation.

Deriving the probability density function of the index is an easier task to handle, since observations on weather patterns are available for most locations in the US oftentimes for more than a hundred years (NCDC databases). For the rainfall, the conventional assumption is to use the gamma-distribution (Martin, Barnett, and Coble). The probability density function for this distribution is

$$f(x \,|\, \alpha, \beta) = \begin{cases} \dfrac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x} & \text{for } x > 0, \\[2mm] 0 & \text{for } x \leq 0 \end{cases}$$

(DeGroot). The distribution accounts well for the stylized facts about rainfall patterns such as nonnegativity, skewedness to the left, and possibility of events higher than those observed historically (this would not be the case, for example, with beta-distribution).

The relationship between yields and rainfall and/or other weather parameters turns out to be more elusive. The agronomic literature generally suggests that the relationship, if exists, is represented by a quadratic function of rainfall and possibly temperature during the growth period. In particular, the models derived by Teigen and Thomas (the most comprehensive study of the topic we found) for corn yields in ten Midwestern states included linear, quadratic, and cross-product terms for both precipitation and temperature for the months of May through September. The rele-

vant independent variables were then selected for each state based on the regression analysis.

While the models reported 90% and above goodness-of-fit for most states, an independent analysis of the NASS yield data series revealed that 50 to 70 percent of the variation was explained by the liner trend variable included in all models. Since the trend represents a systematic change in yields due to improved technologies and agricultural practices, this is obviously not a risk component and thus does not need to be hedged against. If the linear trend variable is removed from the models, the goodness of fit drops dramatically thus raising the question of model applicability.

Martin, Barnett, and Coble also adopt a quadratic model derived in an earlier study based on observations of several test plots in Mississippi. However, they do not provide actual data on yield realizations and how those are scattered around the fitted function.

Nevertheless, we attempted to implement a similar approach to establish a relationship between yield and weather variables for Iowa corn and Kansas wheat. In both cases, a crop-producing district was selected randomly and then yields were regressed against various combinations of weather variables (precipitation and temperature) for each county in the district. The yields were detrended to account for technology changes. The rainfall and temperature data were considered in two variants: absolute monthly data and deviations from the long-term monthly averages (both types of data available from NCDC). Unfortunately, the analysis failed to show consistent patterns of relationships between yields and a specific set of variables. Even the visual analysis of scatter diagrams of yields vs. rainfall and/or temperature showed little or no systematic dependence, with high yields corresponding to both high and low rainfall and or temperature and vice versa.

The adjusted R^2 rarely exceeded 30% for selected counties, and no single model performed consistently well even in two adjacent counties. In addition, the models often had to include such esoteric combinations of variables as the product of August rainfall and temperature along with July rainfall squared but without any linear terms. Clearly, such combinations can hardly suit as an index, which needs to have a relatively simple structure transparent for both farmers and exchange brokers. In addition, failure of any single index to performed uniformly well within the crop producing district implies that each county would need its own contract based on its own index. This again undermines the idea of simplicity and makes the contracts much harder to market at an exchange or as an investment instrument.

Conclusion

Index insurance contracts may provide an alternative way to hedge the risk of agricultural production. The contracts are more efficient and may provide an access to financial resources usually unavailable to traditional insurance. The suggested methodology allows one to design and price index insurance contracts based on actuarial requirements. It also allows a buyer to determine the optimal number of contracts he needs to buy in order to obtain the best coverage.

However, in order to bring the methodology beyond purely academic exercise, one needs to establish a relationship between the index and the risk being hedged against. In the cases of Iowa corn and Kansas wheat, attempts to establish such a relationship between yields and weather parameters did not produce consistent results. Further analysis of crop/weather variable combinations may result in more clear understanding of the contract applicability.

Figure 1: Indemnity Payments of a Standard Contract

References

DeGroot, M. H. *Probability and Statistics, 2nd Ed.* Reading MA: Addison-Wesley, 1989.

Martin, S.W., B.J. Barnett, and K.H. Coble. "Developing and Pricing Precipitation Insurance." *J. Agr. Res. Econ.* 2001: forthcoming.

Skees, J., P. Hazell, and M. Miranda. *New Approaches to Public/Private Crop Yield Insurance.* To be published by The World Bank, Washington, DC.

Teigen, L.D. and M. Thomas. *Weather and Yield, 1950–1994.* ERS Staff Paper No. 9527, 1995.