InsuranceSim: Selecting Crop Insurance for Your Farm

December, 2002
Terry Kastens, K-State Ag. Economics (785-532-5866; tkastens@agecon.ksu.edu)
Kevin Dhuyvetter, K-State Ag. Economics (785-532-3527; kdhuyvet@agecon.ksu.edu)
G.A. (Art) Barnaby, K-State Ag. Economics (785-532-1515; abarnaby@agecon.ksu.edu)

This paper accompanies the *InsuranceSim.xls* spreadsheet that is used to assess the risk and profitability associated with different multiple-peril crop insurance products for your farm. The current version of *InsuranceSim* handles wheat, corn, milo (grain sorghum), and soybeans. For those interested, *InsuranceSim* also can assess the impact of combining futures hedging or forward contracting with crop insurance, as well as limited futures options and minimum price contracting strategies. Essentially, *InsuranceSim* simulates crop yield and price outcomes to help with the crop insurance purchase decision at the time it is typically made – near planting time or just before the crop insurance sign-up deadline. This paper and the associated spreadsheet can be found at the Dhuyvetter/Kastens webpage, www.agecon.ksu.edu/kdhuyvetter. Readers in a hurry to use *InsuranceSim* can proceed directly to Section 4 of this paper.

A word of caution

To be usable, the *InsuranceSim* spreadsheet needs many rows of formulas, causing it to be a large spreadsheet to provide electronically. Thus, the version stored on the website has only a few rows in the simulation page showing – as is, it is still almost 3.5MB. Consequently, before results will have any meaning, you need to click on the “Copy formulas” button. This will copy the needed formulas automatically and will take several minutes. Once formulas have been copied, you can make changes to the inputs and see how they impact results. The file with the copied formulas is quite large (~18MB) and thus it is recommended you save it with a different name. If storage space is a concern, you would not have to save this “large” spreadsheet because you can duplicate it anytime you want. Regardless, of whether or not you save the larger file, it is recommended that you save the original file to your computer before copying formulas so that you do not have to download that file more than once.

Be certain that you type numbers in only blue cells of the spreadsheet. Typing numbers in cells that contain formulas can cause errors that are most difficult to track down, often resulting in you having to start over by once again downloading the spreadsheet from the website.

Section 1 – Background

In this paper the words “crop insurance” refer to multiple peril government subsidized crop insurance, not to private, specific peril, insurance such as the traditional hail and fire insurance. Use of subsidized crop insurance is becoming ever more common in the Great Plains and elsewhere. Availability of new types of insurance and changes within existing types complicate the crop insurance purchase decision. In some areas and for some crops, two similar products are available, CRC (crop revenue insurance) and RA (revenue assurance), but at different premium rates, further complicating the decision.

Crop insurance is fraught with rules, details, and technicalities. However, *InsuranceSim* accommodates only the level of detail thought sufficient for guiding meaningful crop insurance decisions. Also, although *InsuranceSim* allows consideration of several insurance products, it does not cover every product a manager might be interested in, nor is there any guarantee that the products it does cover are available in your state or county. Some products are sufficiently similar to noted products in *InsuranceSim*, that they can easily be accounted for, though premiums might be different. For example, IP (income protection insurance) can be represented by selecting RA (revenue assurance) with enterprise units. When it comes to the final decision, always contact your local insurance agent to be sure you know a product's constraints and limitations.

It should be remembered that most commercially available insurance (e.g., private hail, car insurance, life insurance) has a cost associated with transferring risk to others. That cost, however, is not equal to the premiums paid – since the premiums paid are partially offset by changes in your expected outcomes over time. For example, purchasers of hail insurance expect to collect indemnities over time that would at least partially offset losses in those years. Thus, a portion of premiums is simply "transferring money from good years to bad years." For many types of commercially available insurance, as a point of reference, the cost associated with risk transfer is approximately 30% of the premiums paid. Essentially, this 30% covers insurance company profits, insurance agent commissions, loss adjustment expenses, and the costs of running the insurance business. Then, from the policy holder's standpoint, the other 70% is a transfer of funds from good years to bad years, and, from the insurance company's perspective, the other 70% is a transfer of funds from those who did not have loss claims to those who did.

The crop insurance considered in *InsuranceSim* is typically much different than the "typical" insurance just discussed. First, a large part of the insurance company's "cost of doing crop insurance business" is covered by a direct subsidy by the government to the insurance company, which has the effect of reducing farmer paid premiums. This alone would make the cost of risk transfer for the farmer effectively 0 (assuming competition between insurance companies), which means that we might expect that premiums collected the insurance company would nearly all be paid out as indemnities, at least on average. But, the premiums collected by the insurance company are also subsidized by the government. For example, currently (2003), the government pays 38% of the premium cost on policies with 85% coverage, and 64% on policies with 60% coverage (see Table 1 for more detail on subsidy percentages). On average, across all coverage levels of buy-up products (those beyond CAT – where the government covers 100%) in Kansas for 2002, the government paid 57% of the premiums. What this means is that, if premiums were actuarially fair, where actuarially fair in this case means sufficient to cover indemnities but not to cover insurance company operating expense and profit (since that is covered in a different subsidy), then a buyer of crop insurance might be expected to earn around a 75% return on his investment in insurance premiums (1/0.57 - 1). Of course, none of this is to say that premiums are "actuarially fair" for a given farm. For example, some farms may have expected yields in excess of what the crop insurance company projects and expects, while yields for other farms might actually be below what the insurance company projects and expects. This routinely makes the crop insurance purchase decision a decision involving both profit maximization as well as risk assessment.

Section 2 – A Brief Description of Relevant Terms

This paper offers only a limited description of crop insurance terms. For additional detail it is suggested you go to the Barnaby webpage, www.agecon.ksu.edu/risk.

Crop insurance types or products often are described as being either "yield" policies or "revenue" policies. The simplest distinction between the two product classes is that changes in market prices after sign-up may impact indemnities paid on revenue policies but do not on yield policies. Essentially, at sign-up, both yield and revenue products begin with a bu/acre coverage, or yield to insure, multiplied by a price, which becomes the initial dollars of insurance coverage. Then, depending on the product, changes in market price after sign-up can cause changes in one or more of: frequency of a loss, dollars of coverage, revenue to count against coverage, and indemnities paid.

Over the years, different terminology has evolved for yield products than for revenue products, but, from a computer simulation standpoint, having different terminology can obscure an understanding of the underlying mathematics. Consequently, we borrow heavily from the terminology associated with revenue products, but use it to represent both yield and revenue products.

What follows is a brief description of several relevant insurance and/or *InsuranceSim* terms and a number of the products that can be simulated in *InsuranceSim*. Additional descriptions are found directly in the step-by-step depiction of *InsuranceSim*, which is delineated in Section 4 of this paper.

APH, **percent of yield coverage**, and **bushels of coverage**

APH (actual production history) is the term used by insurance companies to approximate the "expected" crop yield (as in bu/acre) for an insured unit. It typically is based on historical production reported by the farmer, adjusted county average yields for missing historical yields, along with rules associated with discarding or truncating yield outliers. The **APH** number is clearly reported on the insurance policy. When the **APH** value is multiplied by the **percent of yield coverage** selected by the farmer, the result is the **bushels of coverage** associated with a policy. The difference between **APH** and the **bushels of coverage** is the "deductible" associated with the policy, dimensioned as bu/acre.

Planting price, **percent of price coverage (only if MPCI)**, **revenue guarantee**, **harvest price**, and **revenue to count**

The **planting price** is a futures-based price for the revenue products and a USDA-determined price for the yield products. When the **planting price** is multiplied by the **percent of price coverage (only if MPCI)** and by the bushels of coverage, the result is the initial **revenue guarantee** offered by the policy, dimensioned as $/acre. From the farmer's standpoint, the **planting price** is predetermined prior to the insurance sign-up deadline. We use the term **harvest price** to represent the price ($/bu), when multiplied by the actual crop yield, that becomes the **revenue to count**, which is dimensioned as $/acre. For yield products, the **harvest price** is always equal to the **planting price** (adjusted by the price coverage if MPCI). For revenue products, **planting price** and **harvest price** typically are different.

For reference, Table 1 shows the yield and price coverages currently (2003) available on CAT- and MPCI-type subsidized insurance products, along with the percent subsidy. The subsidy row of the table applies also to subsidized revenue products such as CRC and RA.

Table 1. Yield and price coverages and subsidy percentages

	CAT	Buy-up coverage							
Yield Coverage %	50	50	55	60	65	70	75	80	85
Price Coverage %	55	100	91-100	84-100	77-100	72-100	67-100	63-100	59-100
Subsidy %	100	67	64	64	59	59	55	48	38

Adapted from the 2003 Crop Insurance Handbook available on the Risk Management Agency's website

Insured unit
Though there may be others, there are three main classes of unit structure associated with crop insurance. Insuring a farm's production in an **enterprise unit** means that all insurable acreage of the crop in question in the county in which the insured has a crop share is insured together as one unit. Essentially, all of that farm's crop production, irrigated and non-irrigated, continuous and summerfallow, etc. is combined together for the purpose of calculating indemnities. Using a **basic unit** means that all insurable acreage of the crop in the county in which the insured has a 100% crop share (cash rented and owned, irrigated and non-irrigated), is combined together for indemnity determination. With **optional units** (if eligible), crops in different legal sections of land are considered insured separately, as are irrigated and non-irrigated production in the same section (depending on certain farming practices). For the most part, *InsuranceSim* simulates crop yields associated with **optional units** and with an **enterprise unit** – though **basic units** could be accommodated simply by adjusting within year yield variability downwards, or simply by considering them to be **optional units** (which would cause little distortion).

CAT (catastrophic insurance), a yield product
With **CAT**, 50% of APH is covered at 55% of the MPCI price. The USDA pays the premium except that the policy holder pays a $100 administration fee. Only basic units apply to **CAT**. CAT-type coverage is also available on other policies provided by the USDA, for example, IP or GRP (group risk plan). Though IP can be simulated in *InsuranceSim*, IP-CAT cannot. Also, *InsuranceSim* does not consider GRP (either CAT or buy-up versions). Thus, in *InsuranceSim*, any reference to **CAT** means the traditional MPCI-based CAT (sometimes referred to as APH-based).

MPCI (multi-peril crop insurance), a yield product
With **MPCI**, the farmer selects a percent of coverage (50% to 75%, and in some places up to 85%, in 5% increments) to be applied against the APH. **MPCI** pays out based on the difference between the bushels of coverage and the bushels of production, multiplied by the planting price (and price %). The planting price is determined by the USDA and typically is announced near sign-up time. **MPCI** is available with basic units or with optional units, but not generally with enterprise units.

CRC (crop revenue coverage), a revenue product

CRC pays out based on the difference between the revenue to count and the revenue guarantee. The revenue to count is the yield times the harvest price. The revenue guarantee is the bushels of coverage multiplied by the higher of the planting price and the harvest price. The harvest price is capped at the planting price plus $2.00 (wheat), or $1.50 (corn and milo), or $3.00 (soybeans), and is cupped at the planting price less these values. Planting and harvest prices are determined from publicly traded commodity futures prices of harvest time contracts observed near planting time and near harvest time (i.e., the average daily futures closing price for particular time periods, contracts, and exchanges). For example, for Kansas (and certain other states), the wheat planting and harvest prices are based on the August 15 through September 14 average daily closing price (the planting price) and the June 1 through June 30 average daily closing price (the harvest price) of the July Kansas City Board of Trade (KCBOT) wheat futures contract. Corn and milo planting and harvest prices are based on the February average (planting) and the October average (harvest) of the December Chicago Board of Trade (CBOT) corn futures contract (milo based on 95% of corn). Soybean planting and harvest prices are based on the February average (planting) and the October average (harvest) of the November CBOT soybean futures contract. All three unit structures are available with **CRC**.

RA-HPO (revenue assurance with harvest price option), a revenue product

This contract is nearly identical to the CRC contract. However, it has these exceptions. First, harvest prices are not cupped and capped as they are with CRC. Second, for Kansas, the harvest price is the July 1 through July 14 average of the July KCBOT wheat futures contract rather than the June average. Third, the harvest price for corn and milo are based on the November average of the December CBOT corn futures contract rather than the October average. Based on historical price patterns, there is little reason to think of **RA-HPO** and CRC as different types of insurance. Thus, when both are available, the insured is expected to select the one with the lowest premium. All three unit structures are available with **RA-HPO**.

RA (revenue assurance), a revenue product

RA is identical to RA-HPO except that the harvest price is set equal to the planting price. This means that the revenue guarantee does not rise with prices that happen to rise into harvest. All three unit structures are available with **RA**.

Section 3 – Fundamental Assumptions of *InsuranceSim*

InsuranceSim relies on a number of fundamental assumptions that are either conceptual or statistical. This section briefly discusses those assumptions. Do not worry if you do not fully understand the descriptions provided. Enough understanding to make you a successful user of *InsuranceSim* will easily emerge when you read the step-by-step process described in Section 4 of this paper.

Efficient futures markets

InsuranceSim assumes efficient futures markets for the revenue contracts. The efficient market hypothesis suggests that the planting price of a futures contract is the best guess at that time of where that contract will be trading at harvest time. It means that the odds of prices rising into harvest are merely 50:50. This means that the average (over many years) of the futures price at planting time should be close to the many-year average of the harvest time futures price. The idea of efficient grain futures markets has been well supported across a wide body of research for many years. Further, it is the cornerstone on which revenue insurance rating mechanisms have been built. From a futures hedging standpoint, efficient markets mean that the expected futures hedging or options trading profit will be zero over many years (actually slightly negative to account for commissions and slippage). However, assuming efficient markets in *InsuranceSim* does not preclude an analysis of risk associated with different insurance and price risk management strategies.

Random simulations

The insurance decision is about uncertainty and risk, and there essentially are two ways to consider risk. One way would be to look at what actually happened in the past over many years. For example, one might ask questions like, In the last 20 years, what was the lowest observed price? Or, what was the lowest observed yield for a unit on my farm? Follow-up questions would relate to questions like, What would my farm's loss have been had it purchased MPCI insurance in that bad year? What if it had purchased RA insurance? Following these procedures, the analyst would declare that "I believe the past is a reasonable caricature of the future, thus, what would have been the best thing to do in the past (by some risk measure), likely will also be the best thing to do in the future."

A potential problem with looking at actual price and yield data from the past (along with posited risk management strategies), is that 20, or even 30 years, is a fairly short series from a statistical standpoint, and thus could lead to inappropriate inferences about the future. An alternative approach is to create ("simulate") thousands of "possible" outcomes for the future, with alternative risk management strategies assessed across these many possible outcomes. Choosing a many-year simulation method over one using only known historical data is based on the belief that it will be a more reliable caricature of the future. For example, the percent of specific bad outcomes calculated from the simulations is believed to be a more reliable prediction of such measures of interest for the future. After all, a short historical series might be plagued with an uncharacteristically high (or low) percentage of especially bad (or good) outcomes – because uncharacteristic weather had occurred during more historical years than one might expect. Regardless, many-year simulations typically do depend on some measures from the past. In particular, they often depend on measures of expected values and variability, such as the mean (average) and standard deviation, to guide the mathematics of the simulations. The hope is that such measures are somewhat stable over many years.

Simulation outcomes in *InsuranceSim* depend heavily on random draws from a statistically normal distribution. A normal distribution is one that has the ususal bell-shaped curve in a graph of outcome frequency on the y-axis and outcome value on the x-axis. That is, most of the observations will occur near the mean (average) of the distribution, with fewer and fewer outcomes occurring at values that depart from the mean in either the positive or negative direction. Simulating a random normal distribution depends on a user-supplied mean and a user-supplied standard deviation (a measure of variability among the outcomes). Means tend to be somewhat well-known by users because they represent the expected long run average. Standard deviations tend to be less well-known by users, and thus require more information from others.

Because standard deviations tend to be scale dependent (series with higher means tend to have higher standard deviations), *InsuranceSim* typically requires user-input measures that are more proportional. This makes *InsuranceSim* more usable across a wide range of users, for example, for those who expect lower dryland corn yields because they are in western Kansas where rainfall is limiting, but also for those who expect higher dryland corn yields because they are in the Corn Belt where rainfall is higher and more dependable. That is, rather than ask a user to input a standard deviation, *InsuranceSim* asks the user to input a CV (coefficient of variation), which is simply the standard deviation divided by the mean. In short, measures like CV tend to be more generalizable across broad classes of *InsuranceSim* users.

Harvest price simulations
In *InsuranceSim*, many possible outcomes of harvest price are simulated given the known planting price and a measure of the variability of prediction errors, where a prediction error is the difference between the price at harvest (actual price) and the planting-time price (predicted price). Excessively high and excessively low prices, as determined by the user, are truncated. See Appendix A for a more comprehensive description of how harvest prices are simulated.

Basis simulations
In *InsuranceSim*, a basis value (cash price less harvest futures price) is added to each simulated harvest price to provide an associated cash price outcome. Cash price outcomes are used to determine crop sales (thus profitability) and LDP (loan deficiency payment) outcomes against a local loan price. They also come into play for assessing changes in risk associated with hedging strategies. But, harvest basis is never known with certainty at planting time, when the insurance decision is being made. Consequently, *InsuranceSim* simulates a randomly normal basis series from a user-supplied expected value and a user-supplied measure of basis variability.

Crop yield simulations
InsuranceSim works with only one crop at a time. Thus, words like "whole crop" or "farm" refer to the same thing. But, given a crop, a "farm" typically is made up of several fields or groups of fields, which typically constitute optional units in an insurance setting. Thus, there are at least two measures of crop yield variability of interest here. One is the variability in a farm's yield across years and the other is the variability in yield across fields or units, but within the same year. The idea is that, at planting time, the producer does not know what his farm will yield come harvest. Though, he probably does know his farm's long run expected yield, and he probably does have some idea about how variable that farm yield across years has been and thus is expected to be in the future. Further, at planting time, the producer would not know what his unit yields will be that year. For example, he does not know which fields and how many might get hailed out that year – or which fields might end up "catching" yield-enhancing rains at critical times.

InsuranceSim starts by simulating farm yields from a normal distribution, relying on user-supplied information about the mean and the variability of that distribution. That is, the user provides an expected yield and a CV measure thought appropriate for across-years farm yield variability. Then, given a single farm yield outcome from the simulation, unit yields are simulated from a normal distribution using that farm yield outcome as the mean and a standard deviation based on a CV value associated with that particular farm yield outcome (more mathematical detail is provided in Appendix B). The simulated farm yield is only used to guide the simulation of unit yields and is not used directly in the economic analysis structure? Rather, the relevant farm yield for the economic analysis is calculated from simulated unit yields as weighted by acres for each unit. As with harvest price simulations, *InsuranceSim* truncates all yield outcomes above or below user-determined thresholds.

Section 4 – Using the *InsuranceSim* Spreadsheet

A large number of steps (inputs) are required to set up the spreadsheet for your use. However, once set up, only a handful are required to consider scenarios involving other insurance products or other price risk management stategies.

Step 1. Enter the crop type number (wheat =1, corn =2, milo =3, soybeans =4)
Currently, only wheat, corn, milo, and soybeans are handled by *InsuranceSim*. Depending on demand, future versions may accommodate other crops. Alternatively, with careful study of this paper and the spreadsheet, it is fairly easy for a user to tailor a spreadsheet for other crops.

Step 2. Select a crop insurance type from the list provided in the spreadsheet
Keep in mind that not all insurance types are handled by *InsuranceSim*. Also, the types listed are really just for convenience. Other policies not listed can often be handled in the spreadsheet. For example, IP (income protection) insurance (available in some places) is little more than RA (revenue insurance) with enterprise units. Thus, it can easily be handled.

Step 3. Select a unit structure
Select a 1 for optional units or a 2 for enterprise units. *InsuranceSim* does not explicitly consider basic units separate from optional units. Thus, if your farm's crop insurance is generally based on basic units, you should enter a 1, which is the same as for optional units. Most Great Plains users are expected to enter a 1 in this category – because enterprise units typically are not thought to adequately protect against location-specific risk such as hail or localized drought. However, the user might learn valuable information by also considering enterprise units. For example, a user might find that an enterprise unit policy at a higher coverage level would afford roughly the same risk protection as an optional unit policy at lower coverage. Then, a simple comparison of premiums across the competing policies would complete his analysis.

Step 4. Enter the percent yield coverage for the insurance
This is the percent that, when multiplied by the APH, determines the initial yield coverage provided by the insurance. Put another way, 100 minus this value is the "percent deductible" associated with the policy. Typically, most insurance products are available in the coverage range of 50% to 75% (in 5% increments), and occasionally at 80% or 85%. However, the CAT contract is for 50%. **For a scenario with no insurance, enter a 0 for percent coverage.**

Step 5. Enter the percent price coverage for the insurance
This spreadsheet entry is only used by *InsuranceSim* when MPCI scenarios are being considered (i.e., when the crop insurance type is set to 2). As such, it is the percent of the maximum MPCI price (provided by the USDA and entered in Step 21) that is covered by the insurance policy. See Table 1 listed earlier for recommended percentages. Percent price coverages for CAT (55%)

and the revenue products (100%) are hardwired in the spreadsheet and thus do not depend on this spreadsheet entry.

Step 6. Enter a $/acre premium for your farm if you don't want to enter premium by unit
If you want to enter a single $/acre premium applied to your whole farm, enter it here. If not, enter a 0 here. If you enter a 0 here, then you should enter unit-specific premiums as described in Step 11 below.

Step 7. Enter a 1 for the units you wish to represent your farm
Part of a farm's risk is associated with within-year across-unit yields, and the more accurately the spreadsheet can reflect those units, the more accurately will the farm risk assessment be. Consequently, enter information for as many units as you have on your farm even if you have to "make a best guess." Again, as already noted, do not worry that you may not know the exact premium for each unit However, *InsuranceSim* can handle up to only 10 units at one time. If your farm has more than 10 units, enter only those units thought most representative of your farm. The analysis will not be significantly distorted by doing that. **You must enter a 1 for those units you want to consider and a 0 for those you do not want to consider**. Any information you provide or that is calculated for units assigned a 0 will be ignored in the analysis. That means if you want to exclude a unit that you had previously considered, you need only insert a 0 in that column. You do not have to remove other information you had previously entered for that unit.

Step 8. Enter the acres for each unit
The acres for each unit are needed to get at a whole farm crop yield and to assess risk at the farm level rather than at the unit level. After all, one unit may have an especially low APH but it may represent only a few acres for your farm.

Step 9. Enter the expected yield in bu/acre for each unit
For many cases the expected yield may not be different across units. On the other hand, having different units could be especially important for farms that have different practices that must be assigned the same crop insurance type and coverage, for example, a western Kansas farmer who has both continuous wheat and summerfallow wheat – which likely have different expected yields. Another possibility is considering both irrigated and non-irrigated crops. After unit-specific yields are entered, the spreadsheet calculates an acres-weighted expected farm yield based on what you entered for unit acres and unit expected yields. You should check to be sure that this number agrees with what you believe is a reasonable estimate for your farm. This is NOT a yield goal such as you might enter in fertilizer recommendations. Rather, it should be your best guess of what your farm will do yield wise, **on average**, which means while accounting for crop failures related to hail, drought, or whatever.

Step 10. Enter the APH in bu/acre for each unit
Unit-specific APH numbers typically are provided by the insurance agent writing the policy. If that information is unavailable at the time of spreadsheet use, put in values close to the expected yields you entered in those categories.

Step 11. Enter the unit-specific premium rate in $/acre
If you choose to enter unit-specific premium rates, here is where they are entered. But, if you have any value greater than 0 in the cell described in Step 6, then the unit-specific premiums you enter here will be ignored. The reason you have the opportunity to enter premiums as a single per-acre farm-wide value is because often you do not know exactly the amount of the premium for each unit at the time of insurance determination – and especially the unit-specific premiums for competing policies you might want to consider. In general, having accurate unit-specific premiums is not nearly as important as having good estimates for such things as expected yields. That is because much of the analysis summary is based on assessment around "fair" premiums, which is what they would be in a completely actuarially fair scenario – that is, the total simulated indemnities divided by the total acres. Of course, the more you are using *InsuranceSim* to select between two competing nearly identical insurance products (like CRC and RA-HPO), the more important it is that you plug in accurate premiums for each policy's scenario.

A cautionary word about unit information in Steps 7-11
Entering unit information is very important because many insurance policies pay indemnities based on unit risk, not farm-wide risk. Thus, even when you want to consider enterprise-unit-type policies, you should still enter the unit-specific information. Otherwise, you will not be able to effectively compare competing policies where some are based on optional units and others on enterprise units. Thus, the rule should be as follows. **Across all crop insurance-risk management strategies you wish to consider, keep this unit-information section the same – except for premiums.** For example, when you put a 2 in the unit structure cell, indicating you want to examine the enterprise unit structure, do not change anything in the unit section except the premiums.

Step 12. Enter the planting time futures price
This is the futures price on which revenue insurance products explicitly depend. That is, in Kansas, it means the following prices (other areas of the country may or may not be different):

- wheat – the August 15 through September 14 average daily closing price of the harvest time July KCBOT wheat futures price
- corn – the February average daily closing price of the December CBOT corn futures price
- milo – use the same value as for corn (**do NOT enter 95% of corn futures** because that 95% is hardwired in the spreadsheet)
- soybeans – the February average daily closing price of the November CBOT soybean futures price

If you are using *InsuranceSim* before you know the planting price, enter your best guess assuming efficient markets. With efficient futures, your best guess today is merely today's price of the relevant futures contract. Using corn as an example, here is the planting price you would enter based on when you might be using *InsuranceSim*:

- prior to February 1 – today's December CBOT corn futures price
- between February 1 and February 28 – a "February" daily average of December CBOT corn futures prices, where each daily price already known is used in the average, and where today's December futures price is assigned to each and every unknown trading day.

after February 28 – the average of February's daily price closes for December corn futures

Having a value entered in the cell for planting price is important even when considering insurance products that do not explicitly depend on this price (for example, MPCI). That is because simulated harvest time cash prices depend on this value, as do simulated price risk management strategies.

Step 13. Enter maximum and minimum allowable futures prices
These values must be entered to keep the simulation from including price outcomes that are clearly unbelievable. That is, simulated harvest futures price outcomes are truncated at these values. For example, if you selected $6.00 as a maximum, any random draw above $6.00 would be assigned the value of $6.00. In general, values of around "planting price plus $2.00-$2.50" for wheat, corn, and milo, and "planting price plus perhaps $2.50-$3.00" for soybeans, should be adequate for assigned maximums. Reasonable minimums would be "planting price less perhaps $1.50-$2.00" for wheat, corn, and milo, and perhaps "planting price less $2.00-$2.50" for soybeans.

Step 14. Enter the simulation standard deviation of the percentage error
This value determines the variability about simulated harvest futures prices. Smaller values will decrease the variability and larger values increase the variability. For example, if this value is set to 0, then all harvest price simulations will equal the planting price. Based on historical price relationships, we suggest the following crop-specific values: wheat 0.153 (15.3%), corn and milo 0.175 (17.5%), and soybeans 0.147 (14.7%).

Step 15. Enter the expected harvest time basis in $/bu
Expected basis is defined as the expectation of cash price less nearby futures at harvest time, where harvest time is defined as the same period used to determine insurance indemnities associated with revenue products. For Kansas, the relevant contracts are July KCBOT wheat, December CBOT corn, and November CBOT soybeans. Expected basis information can readily be obtained from local grain elevators, university extension sources, or a farm's own information on historical cash and futures prices. Grain elevators typically refer to basis in terms like "30 under" or "10 over," which would be entered as -$0.30 (i.e., a negative 30 cents) and +$0.10.

Step 16. Enter the expected value for the ratio: standard deviation of basis / futures price
InsuranceSim users probably will not have a reliable expectation for this value. Based on research involving 20 years and multiple Kansas locations, we believe that a value between 0.02 and 0.04 (2% to 4%) would be a reliable number that is not particularly crop specific nor location specific. Reducing this ratio will reduce the variability of the simulated basis series; increasing it will increase that variability. Of course, as is true with any simulation, the user should examine the simulated series to see if the values appear reasonable. To aid the user's assessment, the spreadsheet conveniently reports minimums, maximums, averages, and standard deviations for most of the simulated columns of data (at the top of each column).

Step 17. Enter the maximum and minimum allowable crop yields
The spreadsheet truncates all random draws outside the bounds established by these values.

Typically, the minimum should be set at 0 because there often are a number of reasons that a unit might end up with a 0 yield (e.g., hail, extreme drought). The maximum should be set to the highest yield you believe possible for the crop being simulated. Often, a value that is around twice the expected yield would be a reasonable number to enter for the maximum.

Step 18. Enter the across-years CV of farm yield
We suggest a CV (coefficient of variation, which is the standard deviation divided by the mean) of around 0.20 to 0.30 (20% to 30%) for dryland crops and around 0.10 to 0.15 (10% to 15%) for irrigated crops.

Step 19. Enter the across-units CV of crop yield within a year
Based on limited research, we suggest a value that is around 3/4 of the value entered for the across-years CV in Step 18. For example, a value of around 0.18 (18%) is thought reasonable had you entered 0.24 (24%) in Step 18. Because the research behind this particular suggested number is so limited, it is important that you examine the simulated unit yields associated with a particular simulated farm yield to see if you think they are reasonable. If you believe the unit yield variability is too high then you should lower this number (and vice versa).

Step 20. Enter your local government loan price for the crop being examined
The loan price determines simulated LDP payments (loan deficiency payments), which are an important element of any risk simulation for farms that participate in government farm programs. Local loan prices are readily available from local FSA offices or from the FSA's national website. If you want to simulate risk without LDP payments, set the loan price to 0.

Step 21. Enter the Maximum MPCI price
This is the planting-time price associated with the traditional MPCI crop insurance whenever the maximum MPCI price is selected. Always enter the *maximum* available MPCI price. If you are purchasing MPCI insurance at less than the maximum price (see Table 1), that is handled by entering a value less than 100% in the percent price coverage cell described in Step 5. The MPCI price for the crop in question can readily be obtained from your local insurance agent or from the website of USDA's Risk Management Agency (RMA). Since the MPCI price is only used when simulating MPCI insurance scenarios, if you do not wish to consider MPCI, then this number would be irrelevant.

Step 22. Enter the percent of expected crop yield you are hedging with short futures
InsuranceSim considers only a limited hedging strategy. That is, some percent of expected yield is assumed to be hedged at planting time using the same futures contract and price used by revenue insurance products. The hedge is assumed to be lifted at harvest time at the simulated harvest futures price. Other, more selective, hedging strategies are not considered by *InsuranceSim*. Because the assumption of efficient markets is used by *InsuranceSim*, the expected (i.e., the average across all simulated observations) profit associated with hedging would only coincidentally be something different than the loss associated with futures commissions. In fact, each simulated outcome is adjusted by the small amount needed to make this identically true. Nonetheless, consideration of scenarios involving hedging provides worthwhile information about how insurance and hedging are expected to interact to impact risk. If you do not want to consider futures hedging, set this value to 0.

Step 23. Enter the futures brokerage commission and slippage charge in $/bu
Trading futures involves brokerage commissions and slippage allowances, often considered to be around $50 to $100 per contract, which is $0.01 to $0.02 per bu with 5,000 bu contracts.

Step 24. Enter the percent of expected crop yield you are forward contracting
Forward contracts are assumed to take place at the planting futures price plus expected basis less a risk premium. That is, grain elevators routinely require a risk premium to take on basis risk (entered in Step 25). If you do not want to consider forward contracting, set this value to 0.

Step 25. Enter the contract fee in $/bu
This is the risk premium that grain elevators often assess against expected basis to account for basis risk. A value of $0.04/bu typically is considered reasonable based on published research. No other contract default fee is considered. For example, if you contract more bushels than you deliver, it is assumed that undelivered bushels can be purchased from the elevator at the market price in order to make the contract good. Essentially, the contract fee is just like futures commissions – the cost of transferring risk to others.

Step 26. Enter the percent of expected crop yield you are covering with PUT options
InsuranceSim allows for limited consideration of futures options strategies. Namely, only at-the-money options are considered and only at the planting price. Like futures hedging, options are considered unwound (let expire or sell for a profit) at harvest time based on the simulated harvest futures price. Most importantly, the cost for options is considered to be the fair premium cost determined by an efficient markets assumption. Brokerage commissions and slippage charges are considered to be the same as those input for futures hedging. As with futures hedging, because of the assumption of efficient markets, the expected profit associated with hedging equals the loss associated with the cost of commissions and slippage. Nonetheless, consideration of scenarios involving options provides worthwhile information about how insurance, hedging, and options are expected to interact to impact risk. If you do not want to consider futures options, set this value to 0.

Step 27. Enter the percent of expected crop yield you are covering with CALL options
This user input is fundamentally identical to that associated with put options in Step 25, only it deals with call options.

A word about hedging, options, and contracting
In *InsuranceSim*, all pricing strategies are additive. Thus, if you want to consider only hedging or only forward contracting, be sure to set the other one to 0. Also, if you enter the same percentage for puts and calls, risk will not change but expected profit will drop by double the commission cost. To simulate a minimum price contract, enter the desired percentage in the forward contracting cell and enter the same percentage in the call options cell. In short, with careful thought, a number of different price risk strategies can be considered by using the appropriate combination of possibilities in *InsuranceSim*. But, more exotic strategies involving such things as spreads and windows cannot be considered. Nor can strategies that depend on futures markets being inefficient. Of course, if you believe that the markets are inefficient, then pricing strategies could easily be devised that allow profits without even considering actual crop production (i.e., speculative profits).

Section 5 – Analysis Summary

For each simulation outcome in *InsuranceSim*, the spreadsheet calculates each of the values needed to determine the farm's revenue associated with crop sales, LDP payments, short futures hedging, forward contracting, put and call options trading, and insurance indemnities. All values are dimensioned as $/acre. Next, the average (across all outcomes) of only crop sales revenue is considered to be the economic "cost" associated with this farm's production of this crop. Then, that "cost" is subtracted from each simulation outcome to provide a measure of "profit" associated with that outcome. Given efficient markets, this approach to calculating "profit" would be approximately consistent with economic theory that says economic profits (after including a charge for the farmer's equity and for his unpaid labor) are 0 in the long run.

It should be noted that the way *InsuranceSim* will be used, the particular planting price associated with the current year will be entered, not some long run expected planting price. This means that the "cost" benchmark just described is not exactly appropriate to get at true economic profit (because farms are actually more profitable in high planting price years and less profitable in low planting price years). Thus, the reported profit should be thought of as a profit "relative to this year." Regardless, the most relevant results will focus on comparing relative risk across different simulation scenarios. And, this can be correctly done without accounting for the fact that reported profit is not quite true economic profit.

InsuranceSim calculates several measures of risk. The first measure is the standard deviation of profit across all simulated outcomes. The standard deviation is the square root of the average squared deviation of an outcome from the across-all-outcomes mean. For a statistically normal distribution, the standard deviation is an important measure of variability. That is, approximately 2/3 of the outcomes are expected to lie within an interval made up of the mean plus or minus 1 standard deviation. Thus, for a profit distribution that had an average of $0/acre profit and a standard deviation of $40/acre, we would expect about 1/6 of the observations to have outcomes with losses greater than $40 (i.e., profits < -$40) and 1/6 of the observations to have outcomes with profits greater than $40. Then, the risk associated with different *InsuranceSim* scenarios can be compared by comparing the different reported standard deviations. Scenarios with larger standard deviations are considered more risky because the variability of profit is greater.

The standard deviation measure is sometimes criticized as a measure of risk because it considers both good variability (profitable outcomes) and bad variability (unprofitable outcomes), when the farmer is presumed to be interested in only the bad outcomes from a risk standpoint. This could be especially important for comparing crop insurance strategies, which intrinsically mitigate bad outcomes by bringing in indemnity payments. One alternative to the standard deviation measure that has been proposed is a measure referred to as the one half standard deviation (1/2STD). This measure is calculated by first squaring only each of the negative deviations from the average profit, and assigning a 0 to each of the positive deviations. Then the square root of the average of this series is computed. As with standard deviation, higher values of 1/2STD imply greater risk.

A third measure of risk considered by *InsuranceSim* is simply the worst possible outcome (e.g., the largest loss) among all simulated outcomes. Obviously, scenarios with worse worst outcomes (greater largest losses) are considered more risky.

The "worst outcome" measure of risk is sometimes criticized because it depends on a single roll of the dice, so to speak, rather than on the distribution of possible outcomes. Consequently, *InsuranceSim* computes the probability of incurring a loss beyond certain thresholds. *InsuranceSim* currently considers five such thresholds, where the thresholds are based on 10% of the average "crop sales only" revenue measure, 20%, 30%, 40%, and 50%. If desired, the user can change these percentage measures to better suit his needs (on the simulation page of the spreadsheet). The probability of loss is presented in terms of "1 year in X years." Additionally, *InsuranceSim* reports the expected loss conditional upon the loss being larger than the stated threshold. Thus, both the frequency and magnitude of loss is considered using this approach.

For users who want to use *InsuranceSim* to compare multiple scenarios involving different combinations of insurance and price risk management strategies, the best way to proceed is as follows. Start with a scenario of interest. Then copy the information in the Analysis Summary section of the spreadsheet somewhere else (copying values only, not formulas), for reference later. Then repeat the process for an alternative scenario, and so on. Of course, you may want to also copy the whole set of simulation parameters you had entered – to ensure that you will know later what situation had generated a particular summary. Alternatively, you may just want to print out the whole user input section along with the Analysis Summary for each scenario.

In using *InsuranceSim*, the number of changes you make from one scenario to another is based on your needs. However, with complex scenarios, such as those involving both insurance and price risk management strategies, you may want to change only 1 item at a time (for example, the percent of crop hedged) before you print out the results. Otherwise, you may not know what caused the change in risk or profit across scenarios.

Appendix A – Simulating Harvest Prices

Given that planting price is a reasonable guess of eventual harvest price (at least if we observe the two prices over many years), and because they would only be coincidentally equal, we can write:

[Eq. A1] $H = P + E$,

where H represents harvest price, P is planting price, and E represents the prediction error (the planting price is a prediction of the harvest price), known only after harvest: $E = H - P$. If H and P were normally distributed variables, then E would also be a normally distributed variable, centered on 0. That is, on average, we would expect the error (E) to be 0, leading to an expected harvest price prediction equal to the observed planting price P. Then, conditional on some known planting price, say $3.00, we could simulate many reasonable and possible outcomes of H, simply by adding to P a series of randomly generated errors that have an average of 0 and the appropriate variability (i.e., standard deviation).

Unfortunately, crop prices probably are not quite normally distributed, in part because we cannot observe negative prices (or even near-0 prices for that matter). That is, starting with a planting price of $2.00, we might end up with a harvest price of $4.50 (a $2.50 positive move), but certainly would not end up with a harvest price of -$0.50 (a $2.50 move in the opposite direction). Consequently, crop prices tend to be distributed more log normally (where the natural log of price is thought to be normally distributed). Such distributions would have more outcomes greater than $1 above the mean than the number of outcomes less than $1 below the mean.

One easy-to-understand way to move towards a more log normal framework is to characterize prediction error as a percentage error rather than as an actual error:

[Eq. A2] $\%E = (H - P) / H.$

For example, if the planting price for some year was $3.00, and the harvest price turned out to be $3.60, then the %E associated with that year would be (3.6 - 3)/3.6 = 0.167 expressed as a decimal, which is 16.7% expressed as a percentage.

Working in a %E framework, a normally distributed %E variable, centered on 0, could be created, where each outcome is used along with some constant value of P to determine an outcome of harvest price H. The calculation can be determined by solving Eq. A2 for H:

[Eq. A3] $H = P / (1 - \%E).$

For example, if the planting price were $3.00, and a random draw on %E happened to be 0.10 (i.e., 10%), the simulated H outcome is 3 / (1 - 0.10) = $3.33. In a simulation framework, where *E*[.] denotes the statistical expectations operator, P is a constant (the known planting price), and %E is a random normal variable, Eq. A3 becomes

[Eq. A4] $E[H] = P * E[1 / (1- \%E)].$

One problem associated with Eq. A4 is that, in an expectations framework,

[Eq. A5] $E[1 / (1- \%E)]$ is NOT $= 1 / (1 - E[\%E]).$

That means that $E[1 / (1- \%E)]$ does not become 1 when %E is 0. And, through Eq. A4, it means that, over many outcomes, the average for H will NOT equal the constant value of P, a feature desired from the standpoint of our assumption of efficient markets.

The "problem" just noted can be addressed by considering a Taylor series expansion of $E[1 / (1-\%E)]$ around the point $E[\%E] = 0$. The ending result is:

[Eq. A6] $E[H]$ approx. $= P * (1 + var[\%E]),$

where *var*[.] denotes statistical variance. Since a variance is always positive, Eq. A6 makes it clear that using a formula like Eq. A3 to simulate outcomes of H would result in an *E*[H] value that will be greater than P. To "fix" this problem we can simply divide each H outcome by (1 +

var[%E]). This results in the simulation formula:

[Eq. A7] $H = P / \{(1 - \%E)*(1 + var\%E)\}$,

where P is the known planting price, %E is a random variable simulated from a normal distribution with a mean of 0 and a known variance referred to as var%E, and H is the simulated harvest price. Now, over many outcomes, the average simulated H value will approach P. Then, we can use historical price data to guide us towards an appropriate measure of var%E.

A second "problem" associated with the price simulations is that, even after correcting for the problem already noted, it is still possible to simulate certain price outcomes that are simply unbelievable. That is, they are obviously "too high" or "too low." In previous work, this paper's authors have redrawn such outlier prices from an appropriately specified triangular distribution. But, we do not believe the marginal benefits associated with that technique merit the substantially increased complexity it would bring to the *InsuranceSim* spreadsheet. Thus, based on a user's assessment of the highest and lowest possible price that might reasonably be expected, *InsuranceSim* simply truncates price outcomes that fall outside of those bounds. As long as the bounds are not particularly confining, this approach should not materially distort the results.

Appendix B – Simulating Unit Yields

How should the across-units within-year CV parameter of the unit yield simulation be determined at each farm yield outcome? Call that measure $auCV_i$ to indicate that it will vary for each simulated farm yield outcome i. We begin by asking the user to input an across-units expected CV measure (called auCV). Based on very limited research, we suggest a value that is around 3/4 of the across-years CV (called ayCV) that the user had input. For example, had the user supplied an ayCV measure of 0.24 (i.e., 24%), we would suggest an auCV measure of around 0.18 (i.e., 18%).

The problem is that, while the user-supplied auCV measure probably is appropriate for simulated farm yield outcomes near the user-supplied expected yield (called EY), it probably becomes less and less appropriate as simulated farm yield outcomes depart from that EY value. Consequently, and again based on only limited research, we arbitrarily do the following. For a farm yield outcome (Y_i) equal to EY we let $auCV_i = auCV$, which is the value the user had supplied for that measure. Then, for yield outcomes that are ever higher than EY, we scale the $auCV_i$ downwards until a minimum value of $auCV_i = 2/3(auCV)$ is reached when $Y_i = maxY$, where maxY is the user-supplied maximum yield considered possible. On the downside, for Y_i outcomes < EY, we scale the $auCV_i$ value upwards to a maximum value of 2(auCV) that is used when $Y_i = minY$, where minY is the user-supplied minimum yield considered possible. The formula determining the across-units yield simulation CV is:

[Eq. B1] $auCV_i = auCV + [(EY - Y_i) / EY] * auCV *$ (if $Y_i <= EY$ then 2, else 2/3).

The $auCV_i$ measure obtained from Eq. B1 is multiplied by the farm yield outcome Y_i to become the simulation standard deviation for a unit yield. However, one further adjustment is needed.

The user is allowed to input a different expected yield for each unit, for example EY1, EY2, etc., as well as different acres for each unit. Thus, the farm yield EY measure used earlier is actually calculated from an acres-weighted measure of unit-specific expected yields (the EY1, EY2, etc.). Since different units are allowed to have different expected yields across-years, simulated unit yields within each year must be adjusted upwards or downwards according to whether the EY associated with that unit is higher or lower than the acres-weighted expected farm-yield EY. We handle this by simply multiplying each simulated yield from some unit (say unit 1) by the proportion that that unit's EY is of the acres-weighted expected farm yield. Thus, for unit 1, we would multiply all simulated yields by the factor EY1/EY.