

**Actuarial Review of “Methodology Analysis for Weighting Historical Experience”
for USDA’s Risk Management Agency**

Review completed by AgRisk Management LLC of Ames Iowa

Executive Summary

The rating methods used by RMA to set Yield Protection rates will be improved substantially if the new approach recommended made by Coble et al are adopted. RMA's current rating methods need improvement because equal weight is given to each year of loss experience. In the early years of the crop insurance program when participation rates were low, those farmers who chose to participate were more likely to make an insurance claim than those farmers who chose not to participate. This inflated losses in the early years above the level of losses that would have occurred if participation had been at today's levels. Thus using the inflated losses from the early years of the program on an equal basis with recent years increases premium rates above where they should be. Also, there is strong evidence that changes in farming technologies and practices over the last 20 years have made at least corn and soybean yields more tolerant of adverse growing conditions. For these two reasons, today's risk profile for corn and soybeans cannot be accurately measured by earlier periods of loss experience without some adjustment. An additional reason why giving equal weights to loss experience is not appropriate is that if RMA has 30 years of loss experience, then under current RMA rating methods each year receives a weight of 3.33 percent. But if it is known that a particular weather event in this 30-year history is actually a 1-in-100 year event or a 1-in-15 year event, then losses in that year should be given a weight of either 1 percent or 6.67 percent. Thus, RMA's current methods either give 70 percent too little weight or 100 percent too much weight to losses in that year. These fundamental flaws in current rating methods would be largely corrected if the recommendations made by Coble et al were followed. Five recommendations are made by Coble et al. My summary review of each is as follows:

Recommendation 1. RMA should use Climate Division Data for calculating crop-specific weather indexes. This recommendation consists of two parts. The first is that RMA should construct weather indexes. The second is that RMA should use a particular data set to construct the weather indexes. There are sound reasons why RMA should use weather indexes to help it set premium rates for yield insurance. Construction of weather indexes potentially allows for better estimates of the likelihood of future weather events because weather records cover a much longer time span than is covered by RMA's loss cost data. Use of a 100-year weather history can provide better insight into whether the probability of recurrence of a 1993-type weather event (lack of heat in parts of the Corn Belt and excess rain) is a 1-in-30 year event or a 1-in-100 year event. If it is a 1-in-100 year event then the loss cost for 1993 should be given 70 percent less weight in rate making than it now receives. Coble, et al do a good job justifying why the Climate Division Data is the most appropriate data set to use to construct the weather indexes. The only significant weakness of the Climate Division Data for use in rating the most important crops in the crop insurance program is the use of regression analysis prior to 1931 for allocating state average growing conditions to climate divisions. However, with the possible exception of geographically large and diverse states, such as Texas and Montana, there is a high correlation between state average weather and climate division weather, so the regression estimates should adequately capture significant departures from normal weather.

Recommendation 2. RMA should use fractional logit models to estimate weather indexes with Palmer Drought Severity Index (PDSI) and Cooling Degree Days (CDD) in the regressions. Because loss-cost data are by definition limited to between zero and one it makes sense to use an estimation technique that explicitly accounts for these limits. But, as the authors point out the results of the weather index regressions are only used to rank years in terms of severity of losses. It would be surprising if OLS rankings would differ in a meaningful way from rankings obtained from fractional logit regressions. If the computational burden of fractional logit is small relative to OLS, then I agree with the recommendation. If the computational burden is high, then I recommend using OLS instead. Use of PDSI and CDD to construct the weather indexes makes sense. PDSI captures lack of precipitation, excess precipitation, and excessive heat, all of which can lead to crop losses. CDD captures the amount of heat during the growing season. I make some suggestions in the main part of my review for improvement in how the weather indexes should be estimated. My judgment is that the recommended approach is sound, but RMA should verify the reliability and performance of the regression equations for each climate division before implementing the procedures. A simple plot of predicted loss cost rank against actual loss cost rank provides insight into this reliability.

Recommendation 3. RMA should place each year of loss cost experience into discrete probability categories that are defined by a long-term history of weather. Some method must be used to determine how likely a set of growing conditions that occurred in a particular year in RMA's loss experience history will occur again in the future. The proposed method of using probability bins is a robust method that when combined with Recommendations 4 and 5, should result in a much more reliable rating system. My only reservation with this approach is that it could lead to poor rate making if prediction errors from the weather index equations are large.

Recommendation 4. RMA should change its method of calculating catastrophic loads by adopting a 90th percentile load cap, by spreading the load to the climate division instead of the state, and by dampening the weight given to the most extreme weather years. Spreading catastrophic loads to the state level subsidizes farmers who live in climate regions within a state that are prone to large losses and penalizes farmers who live in climate regions that are not. This proposal is consistent with the sound proposal to use make climate division as the basis for categorizing weather as it impacts crop insurance losses. Coble et al do not clearly justify a 90th percentile cap versus an 80th percentile cap. Either would work with the rest of their proposal.

Recommendation 5. A discrete adjustment should be made to pre-1995 losses and a 20 year loss history should be used for base rates. There is ample justification for making a discrete adjustment to losses in the early period of the crop insurance program. However, Coble et al's justification for choosing 1995 as the year to make this adjustment rather than 1998 or 1999 is not adequately documented. Use of a 20-year history for base rates combined with a longer time period for a catastrophic load is a simple, straightforward change that would result in current premium rates that reflect both modern production practices and all available observations of high loss years.

Research Report

(1) Description of the methodology used by the expert reviewer.

The method used in this review was to carefully read the extensive analysis with an objective of determining whether the report's recommendations 1) make sense given the extensive literature that suggests that RMA's current rate making methods could be improved; 2) are internally consistent and could be implemented by RMA; and 3) were justified by data and documentation provided in the report. A more extensive review would have replicated key results and would have estimated the impacts of alternative approaches that may make sense. But a lack of results that were included in the report and the large amount of data involved in generating results made a more extensive analysis infeasible.

The rest of the research report is organized as follows. First an overview of why RMA should change its basic rate making methods is provided to dispel any doubts whether change is needed. Second, each of the specific recommendations contained in the report are reviewed and suggestions for improvement are made. And third, answers to the standard review questions are addressed.

The Basis for a New Rating Method

RMA currently uses all available loss cost data back to 1975 to determine its premium rates for yield insurance (Yield Protection and APH). Premium rates for Revenue Protection are based on Yield Protection rates so premium rates for much of the entire crop insurance program are based, at least in part, on losses that were paid out to insured farmers as long as 35 years ago.

There is nothing inherently wrong with using data from 35 years ago if that data reveals useful information about current yield risk. For example, we know significant indemnities were paid to corn farmers in some parts of Illinois in 1989 because of lower-than-average yields caused by hot and dry weather. The existence of hot and dry weather in parts of Illinois in that year that led to yield losses is part of the historical weather record that should be used to determine the likelihood that yield losses caused by hot and dry weather will occur in the future. Thus, today's premium rates should be based on the knowledge that hot and dry weather led to yield losses in 1989.

RMA currently gives each year's loss cost experience an equal weight. So the loss cost experience in Illinois in 1989 is given equal weight to the Illinois loss cost experience in 1999 and 2010. However, there may be a number of reasons why the actual loss cost (indemnity divided by liability) from 1989 should be given a different weight than loss costs in either 1999 or in 2010.

Change in Insurance Pool over Time

The 1988 Corn Belt drought drove a large number of farmers to buy crop insurance in 1989. About 30 million acres of corn were insured with crop insurance, which was about 41 percent of planted acres. By 1993 this number had decreased to 22 million acres or about 30 percent of planted acres. In response to the large increase in premium subsidies that were made available beginning in 1998, a much larger proportion of corn acreage is insured. In 2010 about 83 percent of corn is insured with crop insurance.

Academics, who have studied who participated in crop insurance in the 1980s and 1990s (Just, Calvin, and Smith (1999); Quiggin, Karagiannis, and Goodwin (1993); and Goodwin (1993)), a time period when participation was low conclude that those who were most likely to receive indemnities are those who were most likely to buy insurance. This means that average indemnities paid out to farmers who bought insurance during this period are higher than the average indemnities that would have been paid out had a greater proportion of farmers bought crop insurance. This means that if the exact same growing conditions from 1989 or 2010 were to have repeated themselves in 2010, and the exact same management practices and production practices that were used in 1989 were used in 2010 than in 1989, then average per-acre crop insurance payments would be significantly lower in 2010 because a much higher proportion of crop acres would have been insured. This demonstrates that the loss cost experience from 1989 overstates what the losses would be if 1989 weather would occur again today. One way of adjusting the 1989 loss cost experience to reflect this over-estimate is to reduce the weight that is given to this experience when RMA calculates the average lost cost from in its historical record. To adjust only for the change in adverse selection it would be ideal if the amount of discount would reflect the extent to which a repeat of 1989 growing conditions and production practices but with today's participation rates would lower indemnities.

Different Yield Volatility

The increase in crop insurance participation rates since the 1980s and 1990s means that the crop insurance pool today is less prone to adverse selection than it used to be. This necessarily implies that loss cost experience from the 1980s and 1990s should be given less weight than loss experience under the same growing conditions from the 2000s.

But suppose that farmers in the 1980s and 1990s had purchased crop insurance with the same enthusiasm as today's farmers buy it. Would it make then make sense to given the earlier loss experience the same weight as more recent loss experience? One might think not because corn yields today are so much higher than corn yields in the past. On average, corn yields have increased by about 20 bushels per acre each decade. But the level of yields does not determine whether a loss cost from a lower-yielding period should be given the same weight as the more recent period. What matters is whether the yield risk—as measured by the coefficient of variation (standard deviation of yield divided by expected yield)—has changed over time. If yield risk is decreasing over time, then a return of poor growing conditions from the 1980s in 2012 would lead to lower loss cost ratios than those that actually occurred in the 1980s even if adverse selection were

not an issue. The reason for this decrease is that loss costs measure the percentage of liability that is paid out in indemnities. If the percentage yield loss from, for example, a return of 1988-drought conditions, would today result in a lower percentage yield loss, then loss costs today would be lower than in 1989.

Two recent studies that I was part of give insight into whether yield volatility has changed over time. Yu and Babcock (2010) show that percentage yield declines for corn and soybeans at the county level due to hot and dry growing conditions have decreased since 1980. This result implies that loss cost data from the 1980s and 1990s from a drought of a given severity overstate the losses that would occur today if the same drought were to occur in 2012. More generally, Yu and Babcock (2011) conclude that after controlling for growing conditions, the coefficient of variation of corn yields at the county level has declined over time for most major corn growing areas. Together these results imply that loss cost data from the 1980s and 1990s should be given less weight (after controlling for weather and changes in adverse selection) than more recent loss cost data.

That today's corn and soybean crop yields are less susceptible to poor growing conditions is no surprise to farmers and farm technology providers. The advent of biotech corn makes corn plants more resilient. More efficient weed control with herbicide tolerant crops lowers soil compaction and preserves soil moisture. These new technologies have been rolled out on a continuing basis since the mid-to-late 1990s. Over this time period, a larger and larger proportion of the nation's corn and soybean crops have been planted to biotech crops. The corn biotech endorsement that provides farmers a discount on their crop insurance is recognition that today's crops are less risky than crops from the 1980s.

Low or High Probability of Recurrence of Growing Conditions

Relative to trend yields, corn yields in Iowa in 1993 dropped by a greater amount (38 percent) than in any year since at least 1950. This yield drop was caused by a combination of lack of summer heat (purportedly helped by the eruption of Mount Pinatubo in June of 1991) and excessive rainfall. The 1993 year's losses are given equal weight to losses from every other year in RMA's loss cost history. This is an appropriate method if the chance of a recurrence of 1993 is the same as the chance of any other year happening again. How can one tell?

If the only weather history available is that which corresponds to RMA's loss history, then RMA is justified in giving each year's loss cost experience equal weight (subject to the proviso of no change in the insurance pool and no change in yield volatility over time). But weather records go back to before the beginning of the 20th century. Using the entire history of weather records can give far more accurate estimates of the probability distribution of weather than the limited history that corresponds to RMA loss cost history. For example, if RMA has 35 years of lost cost data for Iowa, then it implicitly assigns a 1/35 probability (0.0286) that 1993 will occur next year. But suppose RMA has 115 years of weather records available and over this record there has never been a year with less heat and more rainfall than 1993. The extended weather history

then assigns a probability of 1/115 (0.0087) to 1993. By choosing not to avail itself of the longer weather record RMA would assign more than three times too much weight to the loss cost experience of 1993.

Use of a longer weather record can also work to lower weights. The Corn Belt has not had a major drought since 1988. This does not necessarily mean that a 1988-style drought will only occur only one year in 35. If over a 115 year time series, there have been six 1988-style droughts, then this suggests that there is a 5.7 percent chance that 1988 will occur again, rather than RMA's assigned weight of 2.86 percent. In this hypothetical case 1988 loss costs should be given double the weight that they currently receive.

The foregoing discussion points out that there are sound reasons for RMA to modify its current rate making methods. In brief, these are that 1) adverse selection has decreased; 2) yield risk at least for corn and soybeans has decreased; and 3) a long time period of weather records is available to provide additional insight into the likely recurrence of a particular year's growing conditions. RMA should be commended for commissioning the Coble et al analysis which develops five recommendations and an implementation plan to give RMA a method for modifying the weights given to the crop insurance program's historical experience. Each of the five recommendations are reviewed in the order that they are presented.

It is a maintained assumption in the Coble et al report that RMA will continue to use its loss cost history as the basis for the yield insurance rating system. This recommendation will not be reviewed here.

Recommendation 1. RMA should use Climate Division Data for calculating crop-specific weather indexes. This recommendation consists of two parts. The first is that RMA should construct crop-specific weather indexes. The second is that RMA should use a particular weather data set to construct the indexes.

A weather index is simply a translation of measurements of weather variables into an index of growing conditions for a crop. Because different crops in the same region respond differently to weather and because the same crop planted in different regions may respond differently to weather, the indexes should be both crop and region specific. For example, in Illinois, a critical period for corn yields is the middle to end of July when pollination occurs. Ample soil moisture and slightly lower-than-average temperatures with cooler-than average night temperatures are ideal conditions for pollination. Soybeans in Illinois can withstand high summer heat but need August rainfall. Average to slightly above average July temperatures are ideal in Minnesota for corn. Thus the translation of growing season weather conditions into an index of crop conditions should be allowed to vary across crops and regions.

The justification for constructing a weather index is that if a longer time period of weather records are going to be used to more accurately estimate the probability distribution of growing conditions that will affect future crop yields and associated loss

cost ratios, then estimates of how weather affects yields or loss costs must be obtained. Otherwise, one cannot accurately characterize the probability that a particular year contained in RMA's loss cost history will occur again in the future.

Coble et al argue convincingly that the Climate Division Data is the best data to use to characterize growing conditions back to the 1890s. The data are already aggregated into spatial units (essentially crop reporting districts) that make them ideally suited for analysis of county aggregate loss cost data. The data are updated on a monthly basis by NOAA. The data are available back to 1895. And the data contain precipitation and temperature data that are key to understanding crop insurance losses. The only potentially major weakness that Coble et al find in the data is the use of regression analysis prior to 1931 for allocating state average growing conditions to climate divisions. However, with the possible exception of geographically large and diverse states, such as Texas and Montana, there is a high correlation between state average weather and climate division weather, so the regression estimates should adequately capture significant departures from normal weather that lead to crop losses.

Recommendation 2. RMA should use fractional logit models to estimate weather indexes with Palmer Drought Severity Index (PDSI) and Cooling Degree Days (CDD) in the regressions.

Coble et al construct a weather index by regressing observed loss cost ratios on observed weather variables. Both the loss cost ratios and the weather data are aggregated to the climate division. The predicted values from the regression equations are then used as the weather index for each climate division. The weather index takes on values both within the period for which RMA has loss cost data and during the period before loss cost data are available. Thus the index can be used to estimate what loss costs would have been had crop insurance been available and purchased in earlier periods.

One issue that could arise in estimating the weather indexes is that Coble et al use the actual observed loss costs in their regressions. But, Recommendation 5 of the report is to make adjustments to actual loss cost ratios in the early period of the loss cost historical record when there is evidence that there has been a structural shift in loss costs since 1995. That is, before 1995 a given set of growing conditions that lead to positive loss costs would lead to a positive, but lower, lost cost in the current period. For consistency, it is important that if loss costs are going to be used to construct the weather indexes, then adjusted loss costs in the regressions should be used. As discussed in the implementation report (Figure 4.2, page 49) use of adjusted loss costs is recommended. This recommendation should be followed.

Because loss-cost data are by definition limited to between zero and one it seems to make sense to use an estimation technique that explicitly accounts for these limits, both in terms of prediction and in terms of being able to estimate regression equations with loss cost values equal to zero. The fractional logit technique is one such technique. (Tobit regressions are another.) As the authors point out the results of the weather index regressions are only used as a ranking device that ranks weather conditions according to

how they impact loss costs. It would be surprising if OLS rankings would differ in a meaningful way from rankings obtained from fractional logit regressions. If the computational burden of fractional logit is small relative to OLS, then use of the fractional logit procedures is warranted. However, if the computational burden is high, then use of OLS would likely give just about the same results.

Use of the PDSI and CDD in the weather index regressions makes sense. PDSI captures lack of precipitation, excess precipitation, inadequate soil moisture and excessive heat, all of which can lead to crop losses. CDD captures the amount of heat during the growing season. Coble et al improve on some of their earlier work by including separate regression coefficients for both positive values of PDSI and negative values of PDSI. As shown in Table 4.4 for Indian soybeans, increases in PDSI when PDSI is negative (drier and hotter than normal conditions) has a marginal effect on loss costs of -0.8383 whereas increases in PDSI when it is positive have a marginal effect on loss costs of 0.2246. This means that when conditions are hot and dry, reductions in hot and dry conditions (increases in the PDSI) decrease loss costs. When conditions are wet and cool, then increases in PDSI increase loss costs. This illustrates the importance of allowing for differential marginal response to changes in the PDSI variable depending on whether it is negative or positive.

I am surprised that the same logic was not applied to CDD. It appears that only one marginal response to changes in CDD was allowed. But for many crops, increases in CDD are beneficial up to some point at which further increases cause losses. For example, Yu and Babcock (2011b) estimate that in the Northern Corn Belt, corn yields increase modestly with increases in CDD when CDD is low. But when CDD is high, increases in CDD cause sharp yield losses.

Including differential marginal responses of loss costs to changes in PDSI but not to changes in CDD will likely cause the models to fit better with PDSI rather than CDD. I would recommend that if this procedure is implemented again, that differential responses to CDD be included. This is especially critical (as shown below) for states in which lack of heat leads to large crop insurance losses.

Coble et al recognize that there is a continuum of possible combinations of variables that could be used in the regressions. Each of these variables can be specified in terms of monthly or bi-monthly amounts. They can be both included at the same time or regressions can be run separately. It is reasonable to limit the possible combinations of these weather variables to include. The method that is proposed for selecting among the possible combinations of weather variables is to make out-of-sample predictions of loss cost for each combination of weather variables and then to choose the combination that minimizes prediction errors. Although not explicitly explained, it seems that the combination of weather variables that minimizes the sum of prediction errors summing across all climate divisions in a state is the combination selected. The loss cost response (the regression parameters) to the weather variables included in a combination can vary across climate divisions under each trial combination.

The choice of basing the model selection on out-of-sample performance rather than in-sample performance is the right approach. Basing model selection only on in-sample performance will tend to result in model over-fitting in that models with too many variables and parameters will tend to be selected. This suggests that in-sample model selection criteria will tend to lead to spurious results out-of-sample relative to using out-of-sample selection criteria. Simpler models will tend to be selected by out-of-sample criterion and then will tend to lead to less spurious results.

But this advantage only manifests itself if the out-of-sample data used to judge the performance of models represents the range of data that the models are likely to encounter when they are used to predict loss costs in the historical period. For example, suppose that the out-of-sample data on which model performance is based has no extreme weather events that lead to large loss costs. A model that successfully predicts loss costs due to extreme weather events will have difficulty competing with a model that does well at predicting low or moderate loss costs if the out-of-sample data only includes observations of low and moderate loss costs. Thus it is important to include a wide range of loss costs in the data used to judge out-of-sample performance.

I could not find any discussion in the report about how the data used for out-of sample prediction were selected so there is no way to judge if the procedures followed by the analysts are subject to this criticism. If cross-validation was used by fitting the models many times using subsets of the entire sample and predicting the rest of the sample, then the out-of-sample data used for model selection contains all the data that is available. But if only a few years of the sample were dropped, and the models were used to predict just these years, then the model selection criteria method used is flawed. Additional details of what approach was used would have been beneficial.

The benefit of using out-of-sample data to judge model performance is that more parsimonious models will tend to be selected. An alternative approach would be to use agronomic knowledge of crop development to select a parsimonious model. For example, it is well-known that heat during pollination in July negatively impacts corn yields and that rain in August helps soybeans. Why not use this prior information to select which variables to include in the analysis?

Table 4.2 in the implementation report is labeled as hypothetical. It is not clear if this means that these are made up numbers. For now, assume that they are not and that hypothetical means “representative.” If these numbers are actual estimates, then they reinforce the point made above about the lack of a differential response to CDD and raise questions about the ability of the selected models to predict loss cost in Central Iowa. Table 1 below reproduces the Table 4.2 results but ranks the actual loss costs from high to low. The predicted loss cost for each year is also shown.

Table 1. Actual and predicted loss cost for Central Iowa (from table 4.2 of Coble et al)

Year	Actual		Predicted	
	Loss Cost	Rank	Loss Cost	Rank
1988	0.13574	1	0.052011	1
1993	0.124284	2	0.007346	23
1983	0.039556	3	0.029771	3
2009	0.014345	4	0.006317	27
2008	0.01166	5	0.006276	28
1989	0.011598	6	0.007655	20
1998	0.009496	7	0.0471	2
1991	0.009129	8	0.013833	9
1980	0.0085	9	0.018607	5
1990	0.008043	10	0.012985	11
1984	0.006547	11	0.009911	15
2007	0.006309	12	0.021054	4
1986	0.005422	13	0.00751	21
1995	0.004531	14	0.017045	7
2001	0.004379	15	0.010698	14
1985	0.004289	16	0.008525	18
2002	0.004131	17	0.009899	16
2004	0.002627	18	0.007787	19
1982	0.00209	19	0.009397	17
1996	0.001729	20	0.005732	29
2003	0.001688	21	0.01218	12
1981	0.001656	22	0.006697	26
1997	0.001591	23	0.00672	25
1992	0.001455	24	0.003945	30
2006	0.001017	25	0.011144	13
1994	0.000968	26	0.007347	22
2005	0.000671	27	0.015349	8
1987	0.000637	28	0.013779	10
1999	0.000574	29	0.006773	24
2000	0.00022	30	0.017802	6

Droughts led to large Central Iowa losses for corn in 1988 and 1983 and some losses in 1989. Losses due to lack of heat, late planting, excessive moisture and possibly large hail losses led to large losses in 1993 and significant losses in 2009 and 2008. Notice that the (reported) regression predictions do a good job at picking out the 1988 and 1983 losses, correctly categorizing them as the number 1 and number 3 loss years. But the regressions do an abysmal job at picking out the number 2 loss year in 1993—predicting that that 1993 had a rank of 23 (the 8th lowest loss cost year)—and the number 4 and 5 loss years

in 2008 and 2009—categorizing them as the 4th and 3rd lowest loss cost years in the sample. A possible explanation for this abysmal performance is given in Table 4.5 of the implementation report in which it is reported that Total CDD and the sum of June and July CDD were the best set of explanatory variables in Central Iowa. As noted previously, if a differential response to CDD is not allowed, then decreases in CDD will always decrease loss cost because increases in CDD are associated with excessive heat and large loss costs. If a differential response to CDD had been allowed for low levels of CDD, then it is likely that the model could have more correctly predicted the large losses of 1993 and the significant losses in 2008 and 2009.

Figure 1 below plots predicted loss costs from Table 1 above against the actual loss costs. The pattern in Figure 1 of a few years with large loss costs and many years with very low loss costs is likely typical of most major corn and soybean producing climate divisions. When it comes to getting rates right, it would seem critical to be able to predict the frequency of large loss years.




Figure 1. Plot of actual vs predicted loss costs for Central Iowa

In Central Iowa there have only been three large loss years. Two of the loss large loss years are correctly identified by the model. One is not and the model incorrectly categorizes 1998 as a large loss year. What this means is that when looking back in time, the model would not categorize any repeat of 1993 conditions as leading to large losses, thereby understating the frequency with which we should expect to see losses as high as experience in 1993.¹ Furthermore, the model would categorize years with 1998-like

¹ Without access to the weather data it is not possible to determine if in fact, the historical record suggests that there have been repeats of 1993-like growing conditions during the 1895 to 1975 period.

conditions as being high loss years, when in fact, 1998 had a loss cost ratio of less than 0.01.

The difficulty that the model has in Central Iowa (assuming that hypothetical means representative and not made up) is perhaps better revealed in Figure 2 which plots the predicted rank of loss cost against the actual rank. As can be seen the weather index has some prediction ability in Central Iowa as shown by an upward-sloping best fit line, but the prediction ability is quite modest with a slope equal to about 0.2 and an R-Squared value of 0.04. A perfect ability to rank loss cost years would reveal an intercept of 0.0 and a slope of 1.0 on the trend line. A complete lack of fit would generate a slope of 0.0. That Coble et al would try to minimize prediction error is somewhat puzzling because I could not see where the loss cost predictions are ever used in their analysis. What are used are the rankings of loss cost.

It would be useful to see if minimizing errors in ranking would lead to significantly different model selection criteria such as proposed by Rosset, Perlich, and Zadrozny (2005). I expect it would. Because the purpose of the weather regressions is to be able to rank loss costs in the historical weather period, why not use a regression criterion function that gets the fitted line between predicted and actual rank as close as possible to 45 degrees?




Figure 2. Plot of predicted vs actual loss cost rank for Central Iowa

If a ranking accuracy criteria function is not feasible, might a criteria that places large penalties on large under- and over-prediction of lost cost be more appropriate for a model that is being used to identify which years in the historical record would lead to large loss costs and which years would lead to low loss costs? That is, using a performance function that includes a problem-specific loss function rather than simply minimizing the sum of squared errors might be more appropriate. Figure 3 shows that there were really only three large squared prediction errors made with this model in Central Iowa. But two of the prediction errors could have real consequences in selecting which years are high loss years.




Figure 3. Squared prediction errors for Central Iowa

If a “customized” loss function approach is used for model selection it is important to carefully determine if it is more important to avoid under-prediction of loss or if it is more important to avoid over-prediction of loss or if both are equally important. If they are equally important than a super-weighting of large misses (either positive or negative) could lead to better model selection.

But perhaps it is only important for the ranking regressions to find the high loss years. Figure 4 below shows that the ranking procedure does a good job at picking out the high loss years in Central Illinois. The Figure 4 data were taken from Table 4.13 in the implementation report (which is not labeled as being hypothetical). The three highest years are ranked correctly. The fourth highest year is ranked seventh, and the fifth highest loss cost year is ranked correctly. After these large losses are accounted for the model’s ability to predict rank breaks down substantially with a slope of only about 0.15 as revealed by Figure 5.




Figure 4. Plot of predicted versus actual rank of loss in Central Illinois.




Figure 5. Plot of predicted versus actual rank for Central Illinois, eliminating the top five loss years

One last suggestion about model fitting is that it is quite likely that cross-climate division (within a state) restrictions on parameter estimates for a crop would not be rejected. Imposing such restrictions might provide more robust results by allowing greater variations in weather observations across climate divisions to be able to “inform” parameter estimates of the weather indexes. This could allow better estimates to be made. This type of restriction is already implemented by Coble et al for climate divisions for which the chosen model specification is not statistically significant.² So an alternative to using separate climate division models when they are statistically significant and a state model for divisions that are not, is that one could simply test if cross division restrictions hold and then simply estimate one model for the multiple divisions for which the restrictions are not rejected.

Recommendation 3. RMA should place each year of loss cost experience into discrete probability categories that are defined by a long-term history of weather.

The first step in reviewing this recommendation is to discuss what it is that Coble et al are actually proposing. It is clear that the first step in their proposal is to predict loss costs from 1895 to 2009 for each climate division. This will yield 116 predicted loss costs. These lost costs are then sorted low to high. Then 15 bins are created such that an equal number of predicted loss costs fall in each bin. Given that 116 divided by 15 is 7.73, most bins will have 8 years and some will have 7 years. Then, the years for which there are actual loss costs are placed in the 15 bins.

At this point it is not clear if the predicted loss costs are used to determine where a year is placed or if the actual loss costs are used to determine which bin a year goes in. It would seem that the predicted loss costs must be used because otherwise there is no guarantee that the upper limit of the 15th bin is greater than the maximum loss cost or that the lower limit of the 1st bin is less than or equal to the lowest loss cost. In addition, on page 28 of the “technical report” the predicted loss cost for 1988 is used to placed 1988 in bin 10 in their example. So for the purposes of this review, I assume that it is predicted loss costs that determine in which bin a year goes.

Note that using predicted loss cost at the climate division level instead of actual loss costs to determine in which bin a year is placed means that prediction error will cause some high loss cost years to be placed in lower bins than if they were categorized according to actual loss costs. Additional mixing of years will take place because the years are placed in bins according to climate division predicted loss costs and not county loss costs. Every county within a climate division will have a given year of loss costs being placed in the same bin despite likely large differences in lost costs across counties. Table 4.3 in Coble et al shows that mixing at the county level does occur with 1991 being placed in the 10th

² With regards to statistical significance, because the best model fit is one that has minimum prediction error out of sample, and that statistical significance is a within sample measure, it is not clear how Coble et al tested for statistical significance. Furthermore, it could well be that there are statistically significant climate division models that were not chosen.

bin, which is the second highest loss cost bin despite 1991 being in the lower half of the loss cost years. It must be that the predicted loss cost for 1991 at the climate division level categorizes it as a high loss year. This example reinforces the point made above that accuracy in predicting the rank of loss costs for each year in the historical record with a regression equation is more valuable to this exercise than minimizing out-of-sample prediction error.

The next step is to take a county's loss cost data and place it in each bin. If Recommendation 5 is adopted and 20 years of loss cost data are used to determine base rates, then these 20 years of loss cost data would be placed in one of the 15 bins according to their predicted loss costs. If a bin is not populated with at least one lost cost, then 14 bins are created by defining bin widths such that most bins have 8 years and some have 9 years. The 20 years are then allocated to these bins. The number of bins is continually reduced until all bins have at least one year of lost cost data with a minimum number of bins being set at five. Not addressed is whether the boundaries of the bins have a unique solution. I would expect not.

So what could go wrong with this procedure? First of all, let's assume that we have accurate ranking regressions. After each year has been placed in a bin and it is determined that each bin contains at least one year of data, then the "weather weighted" average loss cost for the county is calculated (ignoring for now the catastrophic cap on loss cost). Then the average loss cost for the county that will serve as the base rate for the county equals the average of each bin's average loss cost. Suppose that with 20 years of data we have five bins. Each bin is then given a 20 percent weight. The weight given to each year's lost cost equals 20 percent divided by the number of years that have been placed in the bin. For bins that have only one year, that single year's lost cost is given a 20 percent weight in determining the base rate. At the extreme case where there are four bins with only one year of data and one bin with 16 years of data, then four years of loss costs determine 80 percent of the base premium rate. The other 16 loss cost years determine 20 percent of the base load.

Could such an extreme situation happen? Coble et al calculate that there is an eight percent chance that a bin could contain a single observation with 20 year of data (page 20 of their report). What this implies is that with 20 years of data eight percent of the counties will be giving assign 20 percent weight to at least one year of data. This makes sense because the bin definitions are based on the full 116 years of predicted lost costs. It could be that an entire sample of 20 years would fall almost entirely into a single bin. The question then becomes, is this cause for concern? After all, this is exactly the type of situation that this new rating method is supposed to account for.

Suppose that the last 20 years of growing conditions have been extremely favorable for crop yields in a climate division relative to what the full 116 year history would indicate. That is, with bins being defined by the full 116 year history, 16 years of the last 20 are placed into bin 1. The other four years are placed in bins 2 through 5 which are associated with higher loss costs. That is, based on the historical record, each of the bins, by definition, contains 20 percent loss costs. Thus it is entirely appropriate to place a 20

percent weight on each of the loss cost observations that fall into bins 2 through 5 and only a 1.25 percent weight on each loss cost observation in bin 1. This is the motivation behind this binning procedure. It can accurately translate historical weather patterns directly into appropriate weights to account for variations from “normal” weather in the recent loss cost history.

However, consider what can happen with the apparent magnitude of prediction error that is illustrated in Table 1. Suppose that 1993 in Iowa is categorized as a low loss year, as in Table 1 and that the looking back over the last 116 years reveals that 1993-type conditions manifest themselves every 10 years. An accurate characterization of 1993 would place it alone in the highest bin. But suppose the high concentration of low loss years in the last 20 years means that 12 of the last 20 years are placed in bin 1. Being a predicted low loss year (7th lowest out of 20) 1993 is placed with 12 other years that have a low predicted lost cost. Thus with five total bins, 1993 would receive a weight of 0.01667 (0.2/12) instead of 0.10.

This hypothetical example illustrates that the proposed method can lead to poor results if the weather index model cannot accurately translate weather conditions into loss cost rankings. Thus I recommend that RMA proceed with caution in terms of implementing this bin-based rating procedure until the accuracy of ranking regressions (weather indexes) are demonstrated to a satisfactory level. A plot of predicted rank versus actual rank for the selected model will provide insight into the accuracy of the weather index. Such a plot also suggests that selecting a model generates a rank goodness of fit line that has a slope as close as possible to 1.0 might do a better job as a weather index than simply minimizing sum of squared prediction error.

Recommendation 4. RMA should change its method of calculating catastrophic loads by adopting a 90th percentile load cap, by spreading the load to the climate division instead of the state, and by dampening the weight given to the most extreme weather years.

The current procedure of spreading catastrophic loads from counties to the entire state subsidizes farmers who live in regions within a state that are prone to large losses and penalizes farmers who live in climate regions that are not. If all climate regions are equally prone to large losses, this procedure makes sense. But there is large heterogeneity between climate regions within states as homogeneous as Iowa, let alone Texas. Thus the current procedure serves to subsidize producers who farm in catastrophic-prone regions of a state and taxes those who do not.

RMA now has up to 35 years of loss cost data on which to base catastrophic loads. It is likely (although I do not have the data to prove it) that included in this history are a reasonable number of high loss years in almost all climate regions of the country. Thus there is a much stronger basis today to move from state catastrophic loading to a more disaggregate load area. The proposal to spread the catastrophic load to counties within climate divisions is consistent with the proposal to use the climate division as the basis for categorizing weather as it impacts crop insurance losses.

Given that the recommendation to move the catastrophic load to climate divisions is adopted, moving to a 90th percentile cap from an 80th percentile cap on catastrophic load would seem to not make much of a difference. There is large spatial correlation in growing conditions within a climate division. Thus it is likely that all or practically all counties within a climate division will suffer catastrophic losses in the same year. Thus all will be subject to the cap in the same year. The catastrophic load will essentially average out the losses in excess of the cap between all the counties. Because all counties in major production regions will likely be subject to a 90th percentile cap as well as an 80th percentile cap, the movement from an 80th percentile cap to a 90th percentile cap will have little impact on the ultimate premium rate. Rather, the proposal to move to the 90th percentile cap would serve to increase the base rate and decrease the catastrophic load.

I could not find a justification by Coble et al for moving to a 90th percentile cap. If RMA chooses to not move to the climate division as the basis for rate making, then moving to a 90th percentile cap will reduce the subsidy that low risk climate divisions are giving to high risk climate divisions, which would improve the efficiency of the program. But the efficiency gains of such a move are likely much smaller if the climate division is adopted as the basis for rate making.

One aspect of the moving to a 90th percentile cap that is not discussed in the proposal arises if RMA adopts the proposed procedure for using probability categories for weather weighting. As described above, the combination of prediction errors combined with a highly skewed loss-cost history could result in a high loss year being given too much weight or not enough weight. If the high loss year is given too much weight because of prediction error, then staying with the 80th percentile cap will decrease the impact of the error on the base rate because more of the loss will be removed from the base rate calculation. On the other hand if too little weight is given to a high loss year because of prediction error, then moving to a 90th percentile cap will keep a bit more of the loss in the base rate calculation. But because in this situation, the high loss cost is being divided by a relatively large number of years in the bin, adding a bit more loss to the bin will only increase the base rate by a small amount. Thus, on average, consideration of prediction error might lead one to consider staying at the 80th percentile cap, particularly if the climate division will be serving as the basis for rate making.

The proposal to base catastrophic loads on all loss cost experience is a good one. After all, a longer history of loss costs provides more information about extreme events than a shorter history. But if the long history of loss costs is used, it is important that Step 2 (adjusting loss costs due to a change in the insurance pool) in the flow chart on page 49 of the implementation report be applied to the loss cost years that are not used for base ratemaking. This is the procedure that Coble et al recommend on page 58 of the implementation report and it needs to be adopted.

The proposal to lower the weight given to the most extreme events is a good use of the full history of weather data. After all, it makes no sense to assume that each excess loss year has an equal probability of occurring again in the future. This proposal should be

adopted if RMA has confidence that the prediction errors of the ranking regressions are reasonable.

Recommendation 5. A discrete adjustment should be made to pre-1995 losses and a 20 year loss history should be used for base rates.

As explained at the beginning of this research report, there is ample justification for adjusting lost costs in the early period of the crop insurance program. The crop insurance pool today is less prone to adverse selection and (at least for corn and soybeans) farming practices and technologies have led to a decrease in yield risk as measured by the coefficient of variation of yields. Because adoption of technologies happens over a period of years and because new technologies are continually introduced, a reduction in yield risk should occur gradually over time. Using a discrete change on a certain date to capture the aggregate impacts of a gradual change in yield risk between two time periods will work, but a procedure that estimates how yield risk changes over time on an annual basis may be more appropriate. If there was a discrete change in program rules or some other discrete change that affects aggregate risk, then a model that measures the impacts on aggregate risk (loss cost) at the time that the change was made is appropriate.

Coble et al recommend that a discrete adjustment be made to loss cost data generated before 1995. Justification for making this adjustment in 1995 rather than in some other year could come from two sources. The first is a statistical test that demonstrates that a break in loss cost (accounting for weather variability) is best explained by making the break in 1995 rather than in alternative years. The second justification could be that there was a discrete break in the rules governing the program or that there was a discrete break in the insurance pool that occurred in 1995.

Coble et al seem to offer both a statistical justification for choosing 1995 as well as a program-rules-change justification. On page 19 of the implementation report, the authors state “We have observed that there is a significant discontinuity in the data for many crops that occurs around 1995.” This statement could be based on a statistical test. But no statistical test was provided that identifies 1995 as being the best year in which to make the break. Rather the authors go on to argue that the increase in acres insured beginning in 1995 is evidence of the kind of discrete change in the program that justifies a break-point in 1995.

A large increase in the proportion of a crop’s acres that are insured with crop insurance will decrease adverse selection and average loss costs. Legislative changes in 1994 caused an upsurge in program participation. However, the only policies that RMA uses to calculate loss costs are buy-up policies, not catastrophic policies. And there was no surge in buy-up acres in 1995. As shown in Figure 6 there really is no year that looks like it constitutes a discrete difference in program participation. In fact, 1994 looks like it had a larger change in buy-up acres than 1995, both in terms of absolute and percentage changes. This can be verified by looking at the percent change in acres insured (Figure 7) as well as the absolute change in acres insured (Figure 8).

What these figures illustrate is that there was a gradual change in total acres insured in buy-up coverage. The increases really began to take off in 1999, which was the first year that buyers of revenue insurance policies knew that dramatically higher premium subsidies were available to them before they made their purchase decision.

Figure 9 shows that corn farmers perhaps began to increase their use of crop insurance in a big way in 1997 rather than 1999. But there is no clear break in program participation in 1995. Arguably the most significant change in the crop insurance program was made in 2000 with passage of ARPA. And the most significant ARPA rule was dramatically increased premium subsidies. But farmers were notified of higher premium subsidies before the 1999 sales closing date, so the impacts of higher premium subsidies on the insurance pool likely began in 1999.




Figure 6. Acres insured with crop insurance buy-up policies.

Source: RMA Summary of Business reports)




Figure 7. Year-over-year change in buy-up acres since 1989




Figure 8. Percent change in buy-up acres since 1989




Figure 9. Change in corn acres insured with buy-up coverage since 1994

The above discussion points out that more justification would have been helpful for explaining why 1995 was the discrete point taken in the empirical analysis and not 1999 or some other year. If 1999 had been used as the discrete change year rather than 1995, then it is likely that an even greater adjustment would be made to older loss cost data because the average reduction in losses due to lower adverse selection and lower risk crops (at least for corn and soybeans) would be greater after 1998 than after 1994.

After reflecting on this recommendation for quite some time, I see wisdom in the joint proposal to move to a 20-year basis for rate-making and a discrete adjustment to older loss costs. Moving to a rolling average of 20 years for making base rates will, in a few years, remove from the loss cost experience the period in which adverse selection clearly played an important role in program experience. Making a discrete adjustment (whether it be pre-1995 or pre-1999) will allow continued use of the early data for base rate determination for a short time, while allowing the early data to be used on a continuing basis for calculating catastrophic loads. If RMA chooses to not move to a 20-year rolling average for base rate making, then I would argue for a loss adjustment approach that would estimate how loss costs have changed over time (holding weather constant) for the different crops and regions. The joint proposal to make a discrete adjustment to early-period loss costs and to move to a 20-year average history is a reasonable approach that is easy to implement. My only recommendation is to better justify 1995 as the break-point using statistical tests. If statistical tests cannot be used, then see if moving to a 1999 break point is better supported by the data.

References

- Goodwin, B. K., 1993, An Empirical Analysis of the Demand for Multiple Peril Crop Insurance, American Journal of Agricultural Economics, 75(2): 425-434.
- Just, R. E., L. Calvin, and J. Quiggin, 1999, Adverse Selection in Crop Insurance: Actuarial and Asymmetric Information Incentives, American Journal of Agricultural Economics, 81(4): 834-849.
- Quiggin, J., G. Karagiannis, and J. Stanton, 1993, Crop Insurance and Crop Production: An Empirical Study of Moral Hazard and Adverse Selection, Australian Journal of Agricultural Economics, 37(1): 95-113.
- Rosset, S., C. Perlich, and B. Zadrozny. "Ranking-Based Evaluation of Regression Models." doi.ieeecomputersociety.org/10.1109/ICDM.2005.12
- Yu, T. and B.A. Babcock. "Are U.S. Corn and Soybeans Becoming More Drought Tolerant?" American Journal of Agricultural Economics, 92(2010):1310-1323.
- Yu, T. and B.A. Babcock. (2011a) "Weather Effects on Trend, Variance, and Skewness of Corn Yields." Under review.
- Yu, T. and B.A. Babcock. (2011b) "Estimating Non-linear Weather Impacts on Corn Yield--A Bayesian Approach." Working paper 11-WP 522. Center for Agricultural and Rural Development, Iowa State University. April 2011.

Standard Review Questions

(1) Protection of producers' interests.

(A) Does the policy provide meaningful coverage that is of use to producers, and provide it in a cost-efficient manner?

Answer: Yes.

Rational: The proposal would provide much needed reform to RMA rate making that would improve program efficiency and would offer producer premium rates that better reflect modern production practices and the actual risk that producers bring to the insurance pool

(B) Is the policy clearly written such that producers will be able to understand the coverage that they are being offered? Does the policy language permit actuaries to form a clear understanding of the payment contingencies for which they will set rates? Is it likely that an excessive number of disputes or legal actions will arise from misunderstandings over policy language?

Answer: Not applicable.

(C) Is the mechanism for determining liability (i.e., the amount of coverage) clearly stated and supported by an example?

Answer: Not applicable.

(C) Is the mechanism for determining the amount of premium clearly stated and supported by an example?

Answer: Yes.

Rational: The worked example in the implementation report helped resolve many issues I had with understanding exactly how the proposed method for rate making would be implemented. It is now clear.

(D) Are the mechanisms for calculating indemnities clearly stated and supported by an example?

Answer: Not applicable.

(E) In the case of price or revenue policies, are the mechanisms for establishing price clearly stated

Answer: Not applicable.

- (G) Is adequate, credible, and reliable data available for establishing expected market prices for insured commodities? Is it likely that the data will continue to be available? Is the data vulnerable to tampering if the proposed policy is approved? Is the data likely to be available when needed? Is the proposed system for publishing prices feasible?

Answer: Not applicable.

- (H) Does the policy avoid providing coverage in excess of the expected value of the insured crop?

Answer: Not applicable because the proposal only deals with premium.

- (I) Does the policy contain indemnity or other provisions that cannot be objectively verified by loss adjusters, underwriters, or auditors?

Answer: Not applicable because the proposal only deals with premium.

- (J) Is the policy likely to treat all similarly-situated producers the same?

Answer: Yes

Rational: Adoption of these proposals will lead to improvement in the actuarial fairness of premiums. Midwest corn and soybean farmers have been paying higher-than-actuarially fair premiums for crop insurance. This proposal will help to rectify this situation. It seems that many spring wheat producers have been paying premiums that are lower-than-actuarially fair. This proposal will rectify this situation.

- (K) Will insureds be able to comply with all requirements of the policy?

Answer: Not applicable.

- (L) Does the policy create vulnerabilities to waste, fraud, or abuse?

Answer: No. The proposal does not impact coverage, only premium.

- (M) Is the product likely to adversely affect the agricultural economy of the crop that is proposed for coverage, or of other crops or areas?

Answer: No. Crop acreage will not materially change if this proposal is adopted because it only affects premium.

(2) Actuarial soundness

- (A) Is adequate, credible, and reliable rate-making data available? Is it likely that the data will continue to be available? Is the data vulnerable to tampering if the proposed policy is approved?

Answer: Yes.

Rational: The best most complete data that is available anywhere in the world regarding crop insurance experience and weather were used to develop the proposed procedures.

- (B) Are the explicit and implicit assumptions used in the rating process reasonable?

Answer: For the most part, yes

Rational: It is not clear if the authors of the report expected that their report would be peer-reviewed. If they had then perhaps they would have been more explicit in providing details about the procedures they used. For example, I still have no idea what methods they used to select out-of-sample data for model selection. I expect they did some form of cross validation but they do not say. In addition, I expect that they might have done some analysis regarding their selection of 1995 as a break point, but again, they do not say. For the most part I was able to infer what they did. It would have been helpful though if the authors had provided predicted and actual loss cost for more climate districts. I recommend that RMA ask for this data so that RMA can judge whether the regression equations used to predict loss cost are accurate enough to warrant immediate implementation of the proposal.

- (C) Are the technical analyses (e.g., stochastic and other simulations) technically correct? Do they provide credible, relevant results?

Answer: As far as I could tell.

Rational: The large amount of data and analysis that went into these two reports made it impractical for me to replicate key aspects of the report. What I read and what was reported led me to conclude that the analysis was done correctly. I make some suggestions about how they might want to alter their objective function for model selection, and how to improve use of CDD, but what they did seems correct.

- (D) Is the data used for the analyses appropriate, reliable, and the best available?

Answer: Yes

Rational: The best most complete data that is available anywhere in the world regarding crop insurance experience and weather were used to develop the proposed procedures.

- (F) Does the actuary certifying the submission's rates provide adequate and accurate support for the certification?

Answer: Not applicable.

- (F) Does experience from prior years and relevant crops and areas support the validity of the proposed rates?

Answer: Yes

Rational: Then last 10 to 15 years of experience with the crop insurance program shows that corn and soybean farmers are paying too much for their crop insurance. The results of this analysis are consistent with this experience.

- (G) Is the product likely to be sold in a sufficient number such that actuarial projections would be credible?

Answer: Not applicable.

- (H) Does the submission increase or shift risk to another FCIC-reinsured policy?

Answer: No. All rates for YP and RP are based on YP rates.

- (I) Are the proposed premium rates likely to cover anticipated losses and a reasonable reserve?

Answer: More than reasonable reserve.

Rational: Keeping the 0.88 load seems overkill given the long history we have in actual loss costs and the long (116 year) history of weather records. I do not see why the 0.88 number was not at least reduced given that we are adding so much more weather data to rate making.

- (H) Is the actuarial method appropriate for the proposed policy?

Answer: Yes. This proposal will improve RMA rate-making procedures. The proposed procedures are more appropriate than current procedures.

(3) Administrative burden

Does the policy place an unreasonable administrative burden on the insureds, AIPs or the Federal crop insurance program?

Answer: No.

Rational: The proposed changes will result in more work for RMA staff. But this is a small price to pay for better rates so this “burden” is reasonable and justified on a cost-benefit basis.

(4) Marketability

Is the submitter’s determination of marketability reasonable and supported by the documentation?

Answer: Not applicable.

(6) Other review areas

(A) Does this policy provide coverage that, in whole or in part, is generally available from the private sector?

Answer: Not applicable.

(B) Does the policy propose to insure a peril that is not authorized by the Act?

Answer: Not applicable.

(C) To the extent of the reviewer’s knowledge, does the policy comply with all requirements of the Act and the public policy goals of the FCIC?

Answer: Yes. This proposal will improve the actuarial soundness of the crop insurance program.

Biosketch of Reviewer

Bruce Babcock

Bruce Babcock is a senior partner of AgRisk Management, a professor of economics, the Director of the Biobased Research Center, and the Cargill Endowed Chair in Energy Economics at Iowa State University. Dr. Babcock received his Ph.D. in Agricultural and Resource Economics from the University of California at Berkeley in 1987. Dr. Babcock received his B.S. and M.S. degrees from the University of California at David in 1980 and 1981 respectively.

Dr. Babcock has conducted economic and statistical research on risk management and crop insurance issues since 1982. He has published widely on risk management issues, pioneering new uses of yield distributions to model economic behavior. His research has led directly to the development of new rating procedures for crop insurance products. Rates for Revenue Assurance, Group Risk Income Protection, Livestock Gross Margin, and a number of private insurance products were estimated using new rating procedures developed by Dr. Babcock and his colleagues at Ag Risk Management.

