Journal of Agricultural and Resource Economics 40(2):306–324 ISSN 1068-5502

Accounting for Weather Probabilities in Crop Insurance Rating

Roderick M. Rejesus, Keith H. Coble, Mary Frances Miller, Ryan Boyles, Barry K. Goodwin, and Thomas O. Knight

This article develops a procedure for weighting historical loss cost experience based on longer time-series weather information. Using a fractional logit model and out-of-sample competitions, weather variables are selected to construct an index that allows proper assessment of the relative probability of weather events that drive production losses and to construct proper "weather weights" that are used in averaging historical loss cost data. A variable-width binning approach with equal probabilities is determined as the best approach for classifying each year in the shorter historical loss cost data used for rating. When the weather-weighting approach described above is applied, we find that the weather-weighted average loss costs at the national level are different from the average loss costs without weather weighting for all crops examined.

Key words: crop insurance, premium rating, weather weighting

Introduction

The U.S. crop insurance program uses historical loss experience as the foundation of its premium rating system. In particular, the Risk Management Agency (RMA) formerly used equally weighted, adjusted, historical, average loss cost data (i.e., the ratio of indemnity payments to total liability) for a crop in a county as the backbone of the premium rating procedure. The former system used county loss cost experience data back to 1975 (where available) and gave each year's experience equal weight. The simple average of the equally weighted historical loss cost data then served as the main basis for crop insurance premium rates, consistent with the fundamental principle of insurance ratemaking in which the rate represents an estimate of the expected value of future losses.

The weighting of historical loss data (say, for thirty-eight years from 1975 to 2012) is an important issue because it is directly related to the question of whether this series truly represents the "longer term" weather experience that must be captured to accurately estimate premium rates (Borman et al., 2013). In many lines of insurance, thirty-eight years of loss history would be considered a very long time-series to use in rate making (Vergara et al., 2008). However, thirty-eight years may be a relatively short series for accurately reflecting probabilities of weather events that are the dominant factor in determining crop losses. For example, when loss cost data is simply averaged

---

Roderick M. Rejesus is an associate professor in the Department of Agricultural and Resource Economics at North Carolina State University. Keith H. Coble is the W.L. Giles Distinguished Professor in the Department of Agricultural Economics at Mississippi State University. Mary Frances Miller is Senior Consulting Actuary at Select Actuarial Services in Nashville, Tennessee. Ryan Boyles is the State Climatologist and Director at the State Climate Office of North Carolina and an extension associate professor in the Department of Marine, Earth, and Atmospheric Sciences at North Carolina State University. Barry K. Goodwin is the William Neal Reynolds Professor in the Departments of Economics and Agricultural and Resource Economics at North Carolina State University. Thomas O. Knight is the Paul Whitefield Horn Professor and Emabeth Thompson Professor of Risk Management in the Department of Agricultural and Applied Economics at Texas Tech University.
The authors would like to thank Larry Makus and two anonymous reviewers for their helpful comments on earlier versions of this manuscript. Support for this research provided in part by USDA Hatch Project No. NC2442 (for Rejesus).
Review coordinated by Larry Makus.

to calculate county base rates, the recent 2012 drought year is given 1/38 weight, even though the long-term frequency of the weather events that drove these losses may be greater or less than 1/38. It could be that the 2012 drought was a one-in-twenty-year event rather than a one-in-thirty-eight-year event. If so, a larger weight than 1/38 would be appropriate for that year. Alternatively, it could be that a 2012-level drought only occurs once in fifty years in a longer weather time series and should be given less weight than 1/38. The intent of weather weighting loss cost data is to bring additional information from a longer series of weather variables to more properly weight the loss cost data used to calculate average county rates.

This study develops a methodology for weighting the historical loss cost experience used to calculate crop insurance premium rates (specifically, the rates for the Yield Protection product, which also underlies the rates for Revenue Protection and Revenue Protection with Harvest Price Exclusion). A detailed investigation is performed to develop an optimal methodology for weighting or otherwise adjusting RMA's historical loss cost data in order to maximize its statistical validity for developing premium rates. Multiple weighting approaches are evaluated based on statistical validity, feasibility, sustainability, and a balance of improvement versus complexity. In particular, we explicitly consider the Palmer Drought Index and other weather variables in the development of the weighting methodology. The RMA has now largely adopted the technique proposed in this study.[1] Because of the complexity of the issue and variation in climate response across crops, soils, and geographical location, we believe that this is a fruitful area for further research.

Data Issues and Conceptual Considerations in Weather Weighting Historical Loss Data

Historical Loss Cost Data

The objective of setting insurance premium rates is to provide an estimate of the expected value of future losses (or costs). In the U.S. crop insurance program, rates are typically set separately for each crop because different crops are subject to different perils and, consequently, varying loss costs. In addition, one of the most important components of the U.S. crop insurance rating system is the use of aggregate county-level loss cost data to first estimate a county base rate.[2] County-level data are initially used in crop insurance rating because it is rare for an insured individual to have a sufficiently long time series of historical loss data to be able to directly calculate an estimate of individual expected losses. Therefore, crop insurance rating in the United States aggregates experience data of individual farmers by county.

The starting point for constructing historical county-level experience data is to collect observed indemnity and liability data from individual insureds within a county for a particular crop. The individual experience data are then "normalized" by removing and/or adjusting for several coverage options so that the resulting experience data are comparable (i.e., normalized to a common base). For example, prevented planting is not considered to be a production loss, since the crop was not planted (i.e., typically due to weather constraints). Hence, prevented planting indemnities and liabilities are typically excluded when constructing the historical county-level data used for estimating county base rates (called the StatPlan database).[3] Revenue insurance experience is also used to construct

---

[1] Although the methodology developed in this study has been adopted by RMA, we recognize that there may be other conceivable approaches to weather weighting loss cost data for rating purposes (see Borman et al., 2013, for example, which proposed a Bayesian approach). Using nonparametric kernel densities or fitting parametric distributions using a longer weather data series are also viable alternative approaches that should be investigated in the future. Our intent is for this study to serve as a starting point for further academic inquiry on this topic without implying the superiority of any single approach.

[2] The county base rate is then adjusted based on different factors (such as choice of coverage level, crop type, insured's yield relative to the county average, etc.) to provide an individualized rate for each insured. See Coble et al. (2010) for more details of the full RMA rating procedure.

[3] A detailed discussion of the individual adjustments made to indemnity and liability data to construct the StatPlan database can be found in Coble et al. (2010).

the county-level data by recomputing indemnity and liability data for these revenue policies as if they were for a yield insurance policy.

Once the indemnities and liabilities are normalized, the historical loss cost can be calculated and the county base rate estimated. More formally, an average of the equally weighted yearly county-level loss costs is used to calculate a county base rate for which the individual premium rates are derived:

$$(1) \qquad County\ Base\ Rate_i = E(LC_{ij}) = \sum_{j=1}^{n} f(LC_{ij}) \times LC_{ij} = \frac{1}{n}\sum_{j=1}^{n} LC_{ij},$$

where LC_{ij} is the loss cost ratio from the StatPlan database for county i and year $j = 1, \ldots, n$, and $f(LC_{ij})$ is the calculated probability based on probability density function (pdf) of LC_i over n discrete time periods. A catastrophic loading procedure is also used in the estimation of county base rates to reduce the influence of outliers in the historical loss experience. Catastrophic loading is intended to remove anomalous experience from the county data while preserving normal loss experience. In general, losses deemed catastrophic are spread across all counties for a crop in a state. The previous catastrophic loading procedure censors the county loss experience at the 80th percentile of the historical county experience. All indemnities above the censoring point are aggregated to the state level. Once the catastrophic portion has been removed from the county experience, the uncensored observations below the 80th percentile and the censored values of the censored observations remain; these values are used in equation (1).

Using equal weights in the averaging process in equation (1) implicitly assumes a uniform pdf in which observed loss cost values are equally likely to occur. But given that weather is the major factor that drives crop losses, the equal weights assumption may be problematic since weather distributions may not follow a uniform distribution. Conceptually, probability information from a long time series of weather data may help augment the inherently small sample in the StatPlan database and better estimate the weights when averaging loss cost data.

Another related issue is the need to capture a full range of *LC* values, which typically comes from a long time series of loss cost data. However, as mentioned above, the StatPlan database typically starts from 1975; data from a time series of this length may not be able to provide a complete range of LC values. Hence, it may be desirable from this perspective to have the longest possible time series of loss data for rating purposes.

But it should be noted that there is an inherent tension between the need to use long time series data to properly capture the probability and magnitudes of infrequent but catastrophic weather events versus the concern that, with a long time series, loss data from the more distant past may not be truly representative of the current production and insurance environment such that the measured losses at that time would not be comparable to more recent losses. For example, contract specifics will have changed over time and, thus, modified expected losses. Also, data quality and credibility has improved over time. Finally, agricultural production technology has evolved rapidly such that the effect of a particular weather event on insurance losses may be different from what would have occurred in the past. With a proper weather-weighting procedure, it is possible to use a shorter series of loss data to assure improved representativeness but also use a long series of weather data to more properly weight infrequent catastrophic weather events that cause crop loss.

Weather/Climate Data

In developing a system to weight loss experience data using longer weather or climate data, one has to consider the following issues: (1) the weather or climate data that will be used for weighting (e.g., the length of the data, the degree of coverage and/or level of aggregation, the relationship of such weather to losses, and the availability of weather variables) and (2) the development of a procedure to properly weight each year in the short loss data (e.g., categorizing each loss data year and creating weights for each year in a manner that is consistent with other parts of the rating process).

Weather Data Considerations

There is an abundance of weather data available in the United States that can be used for weather-based weighting of loss experience data. However, there are several issues to consider in choosing the weather data to be used. First, the length of the different climate data series that are available must be considered. In the context of weighting insurance data, the longest series of historical weather data available would help ensure that different weather outcomes, especially the rare extreme weather events that cause losses, would be adequately represented in the longer data series. Information about the probabilities of different weather events will be better captured using a very long climate data series. Various approaches have been taken to address these issues (The rapidly growing literature on this issue is nicely summarized in Tack, Harri, and Coble, 2012; Vergara et al., 2008; Tack, 2013).

However, the need for a long data series must be balanced with a second issue—the degree of coverage and level of aggregation. For example, there may be weather data that are available for 200 years, but these data sets may only contain data for a particular part of the country and/or only at the national level. Crop insurance covers a large portion of the United States and so weather data covering most or all states are needed. In addition, there is significant heterogeneity of the weather events that drive losses at the county level for a particular year. There is value in having data at a lower level of aggregation (i.e., county level or 5×5-mile grids) rather than at the national level only. However, weather/climate data at lower levels of aggregation may have used data interpolation methods in constructing the data, especially at the subcounty level where many locations lack weather stations.

Another factor to consider in choosing the weather or climate data to use in weighting loss experience is the availability of different weather variables. A longer series of climate data may be available for some basic variables like temperature or precipitation, but variables like drought indexes may not be available for this longer period of time. Climate data at lower levels of aggregation and with wider coverage may only be available for certain weather variables and may be absent for others. Hence, the availability of different weather variables in a particular climate data set is also an important consideration to allow for flexibility in determining the weather variables that can help to explain losses.

Finally, in choosing climate data for weather weighting crop insurance loss cost data, the source and future availability of the data are also important considerations. The source of the climate data has to be reliable and must have a good reputation in terms of reporting weather/climate data. Moreover, there should be a reasonable expectation that the weather/climate data will continue to be available in the future to support updating of weather weights as more data become available.

Weather Data Choice

Several datasets partially meet the four weather data considerations above. First, the National Oceanic and Atmospheric Administration (NOAA) Climate Prediction Center (CPC) Unified Precipitation Analysis—an interpolation of the available point-based precipitation gauge data collected by both NOAA and the U.S. Geological Survey (USGS)—meets the second and fifth considerations listed above but only provides information on precipitation and has data only since 1948. Important information on temperature and drought are not provided, and these data do not allow for characterization of the relative frequency of known extreme drought events in the 1920s and 1930s nor hurricane or flooding events prior to 1948.

A national analysis of Palmer Drought Severity Index developed by Dai, Trenberth, and Qian (2004) meets some of the considerations listed above but is not updated regularly and provides drought severity information only every 250 kilometers, which is insufficient to explain local loss experience.

Another group of data that partially meet the considerations listed above is atmospheric model simulations, including the National Centers for Environmental Prediction (NCEP) reanalysis and the

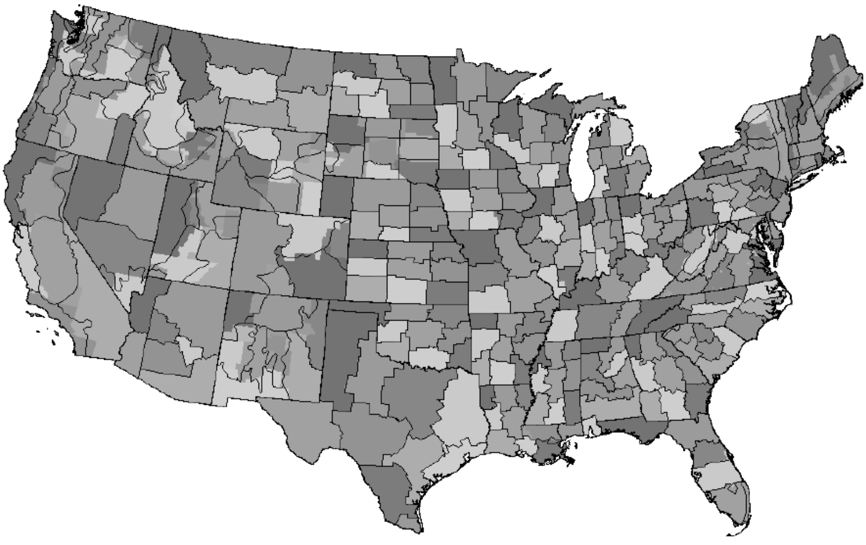

Figure 1. Climate Division Boundaries and County Assignment within Climate Divisions

North American Regional Reanalysis (NARR). These products meet the second, third, and fourth criteria, but NCEP reanalysis only provides information since 1948 and NARR only since 1979.

The data collection that best meets all weather data considerations listed above is the National Climatic Data Center's (NCDC) Time Bias Corrected Divisional Temperature-Precipitation-Drought Index data, also called the Climate Division Data. Climate Division Data provide monthly, serially complete information on temperature, precipitation, relative severity of dry and wet periods using drought indexes, and degree-day metrics of heat and cold accumulation since 1895 for the continental United States, grouped into 344 divisions. Updates are operationally provided each month by the NOAA National Climatic Data Center. Guttman and Quayle (1996) provide a nice description of the history and current status of this Climate Division Data. More technical details on the data and adjustment methods are provided in National Climatic Data Center (1994) and Karl et al. (1986).[4]

Climate Division Data are produced using more than 5,000 National Weather Service Cooperative observer gauge reports. Climate division boundaries in this data set group stations of similar climate into regions that follow state political borders. In most cases, the climate division boundaries also follow county boundaries. However, in regions with more complex geography (including some states with complex topography and/or shorelines), climate division boundaries follow river basins within each state. While climate divisions were originally designed in 1912, boundaries were adjusted in the 1940s to align with crop reporting districts or drainage basins. In some instances, climate divisions cross or split counties. The Climate division boundaries and consequent assignment of counties to particular climate divisions are shown in figure 1. This allocation is based on relative area, geography, and other factors.

There are limitations to using the Climate Division Data. Climate division boundaries are not always delineated for climate homogeneity. Especially in the mountainous terrain of the western United States, the boundaries follow drainage basins, and not all locations within those boundaries

[4] We also considered the issue of climate change. In general, our reading of the climate change literature generally forecasts extreme changes and places us at something of an inflection point of this change (Intergovernmental Panel on Climate Change, 2007; Beach et al., 2009). That is, our climate is not dramatically different from fifty years ago, but it is likely to be significantly different in another fifty years. Since crop insurance typically only forecasts two years ahead, we are likely just beginning to see these effects. Complicating the issue is technological change, which also affects yield risk. One could make some climate change adjustments for recent decades, but our reading of the literature suggests that the adjustments at this point would be minor.

are likely to have very similar climate characteristics, as climate changes quickly with changes in elevation. Another weakness is that the station network used for each division's calculations is not constant. Stations move, cease operation, and new ones are introduced throughout the history of the observing network. This introduces some error to any divisional calculations. Another weakness is the accuracy of division-level data prior to 1931, when regression equations are used to estimate division-level data from the statewide average data that were standard during that period.[5]

Despite these weaknesses, the Climate Division Data provide the best operationally available climate information for crop loss analysis. They provide serially complete national coverage (with no missing data) at a geographic scale sufficient to characterize local climate extremes with a period of record sufficient to identify the relative frequency of climate events that may largely drive loss experience. In 2014, NOAA's National Center for Environmental Information released an updated version of the Climate Division Data that address several of these limitations by providing a time-consistent, homogenized gridded product that accounts for topographic complexity and will provide national coverage with county-level precision for the critical weather variables of interest since 1895.[6]

Merged Loss Cost and Weather Data

The development of the weather-weighting procedure starts by merging the climate data set with the StatPlan loss experience data. Note that the Climate Division Data are observed at the climate division level as described above, while the RMA StatPlan data are reported at the county level.[7] Therefore, all counties within a particular climate division have the same weather data and the loss data also must be aggregated to the climate division level.[8] This is done by summing the indemnities and liabilities of all counties within a climate division and then calculating loss cost ratios based on these summed amounts. The climate division data can then be used to generate a weather index to classify loss years, while the county data can be used to average loss costs and calculate a county base rate.

Empirical Approach to Weather Weighting Historical Loss Cost Data

Weather Index Development

A critical component in the development of a weather-weighting approach is the choice of the weather variables used to determine the relative weights assigned to each year of loss data. One can use a single weather variable or a combination of different weather variables. Based on Wilhelmi, Hubbard, and Wilhite (2002) and the expert opinion of the climatologist on our study team, we chose to examine a parsimonious set of weather variables—the Palmer Drought Severity Index (PDSI) and Cooling Degree Days (CDD). The PDSI is a particularly good weather variable to examine because it subsumes effects of both precipitation and temperature and provides a locally relative scale ranging from very wet to very dry conditions. Wilhelmi, Hubbard, and Wilhite (2002) show

[5] As suggested by one referee, we conducted the same analysis described in this study using only weather data from 1931 to 2009. Overall, there were only minor differences in the average weather-weighted loss cost estimates from only using the 1931–2009 data vis-à-vis using the full 1895–2009 weather data (as presented in table 5 below). However, as explained above, we would like to have the longest series of historical weather data possible to ensure that different weather outcomes (especially the rare extreme events that cause losses) would be adequately represented in the data. We also believe that the pre-1931 data is still of sufficient quality to justify its inclusion in the weather-weighting process. Results of the alternate analysis using the 1931–2009 data are available from the authors upon request.

[6] Details about the updated data are available at http://www.ncdc.noaa.gov/monitoring-references/maps/us-climate-divisions.php

[7] The county loss data used in this study are typically aggregated for all types/practices (with the exception of wheat, for which the data are separated to identify winter and spring wheat). This type of aggregation is consistent with the county-level data used in calculating the base county rate (see Coble et al., 2010, p. 38).

[8] The analysis drops experience prior to 1980, as it was often missing and considered less reliable.

that much loss experience is associated with drought conditions, but the PDSI also allows for very wet (flood) conditions that may also be associated with loss. CDD allows for examining heat units for a particular time period that affects crop growth. CDD is equivalent to Growing Degree Days (GDD) at base 65°F and allows exploration of loss experience that may be associated with extended cold or heat that would not be captured in the PDSI.

We created two variables to represent positive PDSI and negative PDSI values. Positive PDSI values represent wet spells (i.e., larger positive numbers indicate more moisture) and negative PDSI values represent drought conditions (i.e., larger negative numbers represent more severe drought conditions). In addition, the positive and negative PDSIs are limited to the May-June and July-August periods (i.e., average May-June and average July-August PDSIs are used in the study). In summary, four PDSI measures are examined in the development of our weather index—May-June PDSI for positive values (mj_pdsi_p), May-June PDSI for negative values (mj_pdsi_n), July-August PDSI for positive values (ja_pdsi_p), and July-August PDSI for negative values (ja_pdsi_n). The CDD variables used to develop the weather index are total season CDD (from May to September) ($total\_cdd$) and June-July total CDD (jaj_cdd). Crop growth is frequently adversely affected by heat units in the June-July periods.[9]

Based on these six weather variables, we created an index by estimating a fractional logit regression model (at the climate division level) in which the dependent variable is the climate division adjusted loss cost ratio and the independent variables are the six weather variables discussed above (see Papke and Wooldridge, 1996). The full model specification[10] for the fractional regression is expressed as

$$\begin{aligned} LC_{ij}^{CD} = \; & \beta_0 + \beta_1 \times (mj\_pdsi\_p)_{ij} + \beta_2 \times (mj\_pdsi\_n)_{ij} + \beta_3 \times (ja\_pdsi\_p)_{ij} + \\ & \beta_4 \times (ja\_pdsi\_n)_{ij} + \beta_5 \times (total\_cdd)_{ij} + \beta_6 \times (jaj\_cdd)_{ij} + \varepsilon_{ij}, \end{aligned} \tag{2}$$

where LC_{ij}^{CD} is the loss cost ratio for climate division (CD) i in year $j = 1, \ldots, n$, the independent variables are as defined above, and ε_{ij} is the error term.

Fractional logit regression is used to account for the proportional nature of the data and censoring of loss costs at 0 and 1 (since $0 \le LC_{ij}^{CD} \le 1$). A fractional logit regression assumes that the conditional mean of the "fractional" dependent variable (LC_{ij}^{CD}) follows a logistic functional form and ensures that predicted values from the model do not fall below 0 or above 1.[11] This model is consistently estimated using quasi-maximum likelihood estimation techniques (Papke and Wooldridge, 1996; Ramalho, Ramalho, and Murteira, 2011). Examples of fractional regression results for corn (in Illinois) and soybeans (in Indiana) are presented in table 1.

Based on our investigation of the degree of censoring of the data at 0, we believe that using the fractional logit is appropriate in this application. The degree of 0 censoring in the data ranges from 6–11% for corn and soybeans to about 30% for barley and potatoes. On the other hand, the degree

[9] These six weather variables apply to all crops except winter and spring wheat. For winter wheat, the following variables are used: Sept/Oct average PDSI (positive and negative), April/May average PDSI (positive and negative), September to May total season CDD, and March to April total CDD. For spring wheat, the following variables are used: April/May average PDSI (positive and negative), June/July PDSI (positive and negative), April to August total season CDD, and May to June total CDD. Durum wheat type has been aggregated with spring wheat. The variable choices made here were aimed at relative simplicity to facilitate RMA implementation of the weighting procedure developed and for ease of understanding for interested stakeholders.

[10] As discussed below, different "combinations" of the weather variables in this full specification is tested to determine the variables that best predict losses in a particular climate division.

[11] An ordinary least squares (OLS) regression can also be used to estimate the index. The disadvantage of OLS is that predictions are not constrained to lie on the [0,1] interval. Nevertheless, one can argue that the predicted loss costs here are only used as a "tool" to rank the years in terms of having "good" vs. "bad" weather (i.e., one could interpret negative values as indicating good weather years). The magnitudes of the predictions are not used per se. Using OLS to estimate the model did not result in significantly different classifications of the loss years (relative to the fractional logit model). However, we recommend using the fractional logit given the degree of censoring in the data and the intuitive concept of limiting predicted loss costs between 0 and 1.

Table 1. Example Fractional Logit Regression Results

Results for Corn: Climate Division 5, Illinois		Results for Soybeans: Climate Division 1, Indiana	
Explanatory Weather Variables	**Parameter Estimate**	**Explanatory Weather Variables**	**Parameter Estimate**
Total Season CDD	0.010	July-August Negative PDSI	−0.838
	(0.018)		(1.424)
June-July total CDD	0.005	July-August Positive PDSI	0.225
	(0.033)		(1.397)
Intercept	−17.635	Intercept	−4.945
	(15.792)		(3.281)
Goodness-of-fit measures:		Goodness-of-fit measures:	
Pearson test (p-value)	< 0.001	Pearson test (p-value)	< 0.001
R-squared	0.796	R-Squared	0.563
Predicted-Actual Correlation	0.894	Predicted-Actual Correlation	0.7529

Notes: The dependent variable in these regressions is the adjusted loss cost ratios aggregated up to the climate division level. Figures in parentheses are the standard errors. Predicted-Actual Correlation is the correlation between the actual loss cost ratio values and the predicted loss cost ratio values. All fractional logit results for all "state-climate division-crop" combinations are available from the authors upon request.

of censoring at 1 is below 1% for most crops (with the the exception of apples, which has censoring at 1 of about 1.1%).

An out-of-sample competition for each state is conducted to determine which combination among the six initial weather variables best predicts losses (i.e., in this case, which combination best predicts adjusted loss cost out-of-sample). A minimum mean square error (MSE) criterion is used to evaluate the model with best out-of-sample predictions:

$$MSE = \left(\frac{1}{n} \sum_{i=1}^{n} e_i^2 \right), \tag{3}$$

where e_i is the difference between the actual adjusted loss cost ratio and the corresponding predicted loss cost ratio (out-of-sample) and n is the number of climate divisions.

A lower MSE means that there is a smaller discrepancy between the actual and predicted adjusted loss cost ratios, and the combinations of weather variables that produce the lowest MSE values are preferable. We run independent regressions for each climate division within the state (i.e., climate divisions do not cross state lines) but undertake the out-of-sample competition to find the best combination of weather variables for the entire state. This implies that each regression model is estimated independently but a common specification, in terms of the weather variables included in the regression model, is applied for all climate divisions within a state for each individual crop. In other words, for a crop in a state, the same weather variables are used in the loss cost regression though parameters on weather variables may differ across climate divisions.

To facilitate the out-of-sample competition for each state, we limit the number of weather variable combinations to be considered to seven: (1) May-June PDSI positive and May-June PDSI negative; (2) July-August PDSI positive and July-August PDSI negative; (3) total season CDD and June-July total CDD; (4) May-June PDSI positive, May-June PDSI negative, July-August PDSI positive, and July-August PDSI negative; (5) May-June PDSI positive, May-June PDSI negative, total season CDD, and June-July total CDD; (6) July-August PDSI positive, July-August PDSI negative, total season CDD, and June-July total CDD; and (7) May-June PDSI positive, May-June PDSI negative, July-August PDSI positive, July-August PDSI negative, total season CDD, and June-July total CDD. Limiting the combinations to these seven choices and estimating the model for each crop, covering all states, allows for less of a computational burden. It is also supported by previous research identifying critical periods (Dale and Daniels, 1995; Payero et al., 2006).

Table 2. Hypothetical Example of Bin Classification: Soybeans in Mississippi (State=28), Climate Division 1, Bolivar County (1980–2009)

Year	Bin Classification	Predicted Loss Cost Ratios	Actual Adjusted Loss Cost Ratios
1980	4	0.116	0.387
1981	8	0.151	0.140
1982	2	0.098	0.154
1983	5	0.118	0.338
1984	4	0.112	0.108
1985	8	0.154	0.098
1986	9	0.206	0.399
1987	1	0.096	0.140
1988	10	0.271	0.225
1989	8	0.153	0.124
1990	4	0.109	0.236
1991	6	0.130	0.108
1992	5	0.117	0.002
1993	2	0.098	0.066
1994	4	0.112	0.020
1995	1	0.097	0.093
1996	1	0.088	0.031
1997	5	0.122	0.033
1998	10	0.241	0.209
1999	5	0.121	0.131
2000	8	0.154	0.181
2001	5	0.122	0.075
2002	4	0.116	0.048
2003	5	0.117	0.020
2004	3	0.107	0.031
2005	7	0.151	0.024
2006	9	0.185	0.154
2007	8	0.156	0.101
2008	6	0.128	0.126
2009	4	0.115	0.069

Notes: This climate division had a sufficient number of observations to run a credible fractional regression model and calculate a predicted loss cost (weather index). Optimal number of bins for this division (without empty bins) = 10. The Actual Adjusted Loss Cost Ratio is the actual loss cost ratios calculated from the RMA database. The Predicted Loss Cost Ratios (also called the Weather Index) are the predicted values from the fractional regression models used to classify the years into equal probability (variable-width) bins.

Once the optimal combination of weather variables is chosen for a particular crop and state, this combination of weather variables is used to produce a weather index for all of the climate divisions within the state producing the crop. Essentially, the predicted loss cost ratio values of the chosen regression model specification are used as the weather index for each year of weather data. Table 2 presents an example of predicted loss cost values vis-à-vis actual loss costs for soybeans in climate division 1 in Mississippi. Figure 2 presents a plot of predicted versus actual values for corn in Illinois (climate division 4).

Using predicted values (i.e., predicted loss costs in this case) makes it possible to "backcast" a weather index for each year in which weather data are available, even when no available loss experience data exists for the precrop insurance years. This predicted value is conditional on modern technology and production practices rather than technologies or practices used in previous decades since we use more recent data to estimate the parameters of the regression models. This approach is consistent with the objective of the weather-weighting exercise we are conducting.

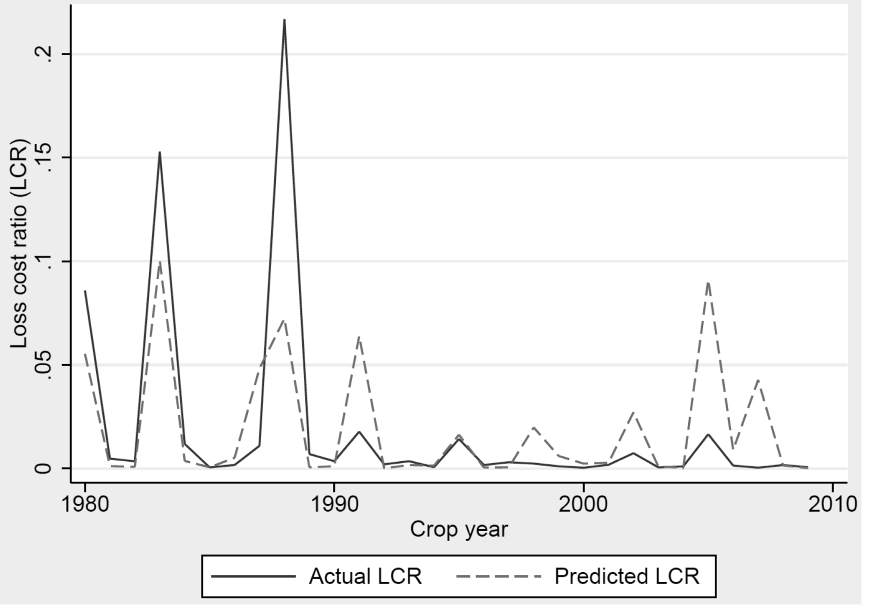

Figure 2. Actual Loss Cost Ratio versus Predicted Loss Cost Ratio for Climate Division 4 in Illinois: Corn

With these predicted values, which we also call the weather indexes for each year, the relative probability of an extreme weather event (or an extreme loss event) can be assessed over a 115-year time span (1895–2009). For example, the weather index for 1988 can be compared to other years from 1895 onward to determine the relative probability of this weather event (i.e., the 1988 drought) occurring in the larger sample.[12]

A concern with using the predicted values is that there may be cases when even the "best" combination of weather variables does not produce a statistically significant model that explains losses over time. For example, in some climate divisions the Pearson chi-square test of overall model fit for the preferred model specification is not statistically significant and the correlation of the predicted values with the actual loss costs is actually negative. This means that the weather variables we considered do not have enough power to explain the pattern of losses observed over time and that there is no significant positive correlation between the model predictions and actual loss costs. We flag these cases, and the weighting methods based on the weather index developed are not applied in these flagged counties. In the case of corn and soybeans, for example, approximately 24% of insured counties were flagged.[13]

Approaches to Loss Year Classification and Weight Assignment

Using the predicted loss cost values from the regression model, each year needs to be classified and assigned a weight that represents its likelihood as indicated by the longer weather series. There are

[12] As a concrete example, the weight for the 1988 severe loss year with simple averaging would have been 3.33% (or 1/30) for Iowa (climate division 5). But using the approach described above, the 115-year predicted loss cost data would have given a weight of 1.73% for the 1988 loss year. This suggests that when one considers a longer time span the likelihood of the 1988 loss year is actually lower than the likelihood implied from a simple averaging procedure.

[13] The proportion of counties flagged varies across crops. The lowest proportion is for corn at 24.11% and the highest proportion is for apples at 59.02%. For cotton, rice, and spring wheat, the proportion of counties flagged ranged from 35–45%. In general, the best results are found for spring-planted crops in rain-fed production regions (i.e., corn and soybeans). For crops like rice, which is primarily irrigated, the weather variables provide little in terms of explaining rice losses. The results for apples and potatoes are similar. More detailed documentation of the flagged counties is available from the authors upon request.

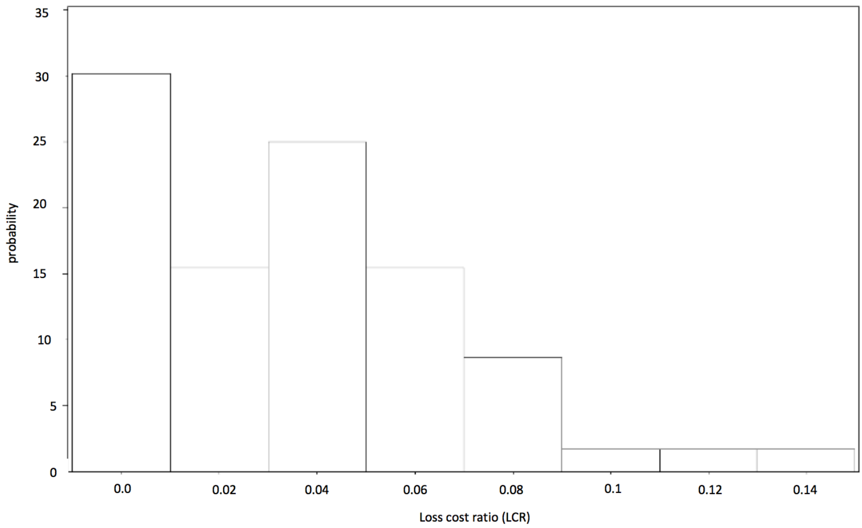

Figure 3. Example Histogram with Equal Bin Widths and Variable Probabilities for Each Bin

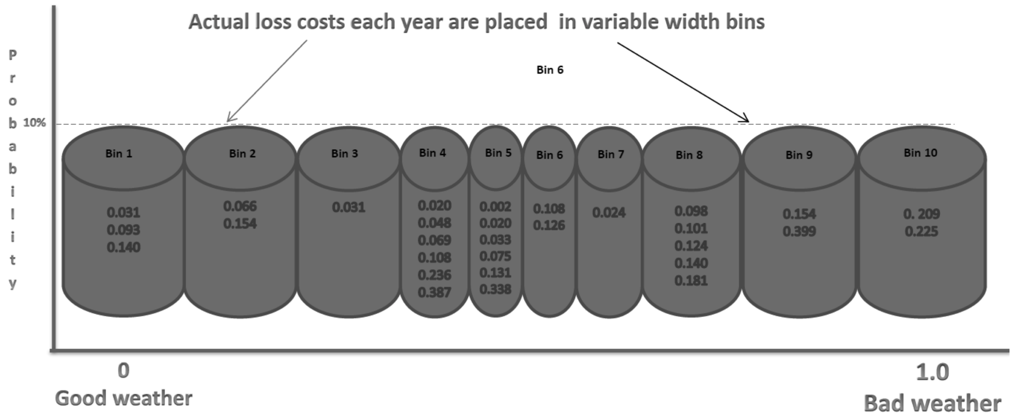

Figure 4. Bolivar County, Mississippi Example of Variable-Width Bins with Equal Probability for Each Bin

Notes: Actual adjusted loss costs for Bolivar County are reflected in the bins above (consistent with table 2)

a number of ways to classify a year and assign a weight. One approach is to generate a histogram with equal bin widths and variable probabilities (or frequencies) (see Coble et al., 2010, p. 85 and figure 3).[14] The bins or groupings with equal widths can then be used to classify each year of the loss experience (i.e., for which bin does the loss year belong to given the actual experience) and the probability associated with the bin assigned to the year will serve as the weather weight. An alternative to the histogram approach is to develop variable bin (or grouping) widths with equal probabilities associated with each bin (see figure 4). The bins or groupings will again be used to classify each year, but since these are variable-width bins with equal probability, there is no need to have differential weights for each actual year of experience.

[14] Alternative methods such as generating kernel densities or fitting parametric distributions can also be used instead of histograms (Borman et al., 2013). However, one should recognize that these more complex procedures may have implications for implementation. One has to weigh the relative benefits of more complex approaches against the efficiency and ease of more simple approaches (like using a histogram).

In both of these loss classification procedures, one has to evaluate the number of bins to be used and make sure that all bins are represented in the shorter loss data used in averaging the loss costs (i.e., the empty bin problem). If not, the weighted average may not fully reflect the available historical experience. In addition, the complexity of the procedure and the ease of implementation should also be a considered in choosing the approach to classify and assign weights to the actual loss years.

With these considerations in mind, we believe that the variable bin width approach is better than a standard histogram approach because it mitigates the "empty bin" issue described above. That is, the likelihood of having empty bins for the years with loss data (1980–2009) can be reduced under this approach as compared to a histogram approach with equal bin widths and variable probabilities. Moreover, the variable bin width with equal probability approach is a fairly straightforward method compared to using nonparametric kernel densities or parametric distributions. This "simplicity" facilitates the practical implementation of this procedure for multiple crops and for nationwide coverage.

Description of Variable Bin Width Approach to Weather Weighting

The variable bin width approach to weather weighting is implemented by first determining the number of bins or percentiles and assigning the predicted loss costs to the appropriate bin or percentile cut-off. For example, assuming that we are interested in ten bins, we would like to find the predicted loss costs in the long history of weather data that correspond to the 10th, 20th, 30th, 40th, 50th, 60th, 70th, 80th, 90th percentiles, in addition to the minimum and maximum values. In this case, we have variable-width bins, since the ranges of the loss cost values used to delineate the bins are not equal across bins, but the probability of falling into each bin is always equal to 10%. If the predicted values are normally distributed, the tails (at both ends of the distribution) tend to have wider bin ranges since only a few observations fall in these areas, but the middle bins tend to have smaller widths because a large number of observations fall in these middle bins.

Once the variable-width bins are delineated, the predicted loss cost value for each year (from 1895 onward) can be classified and assigned to the bin in which it falls. Using the above example, if the bin width for the 10th bin (from the 90th percentile to the maximum) is, say, from 0.09 to 0.15 and the predicted loss cost for year 1988 (a high-loss year) is 0.13, then year 1988 is in the 10th bin. Each year is similarly classified using predictions from the fractional logit regression models. Since the probability of each bin is equal in this approach, there is no need to assign a specific differential weight to each bin.

As mentioned above, one issue that needs to be addressed is the number of bins to assume and the possible existence of empty bins during years with loss cost data (1980–2009). Once the years from 1895 onward are classified based on the weather index, the RMA's actual adjusted loss cost data from 1980 to 2009 are used to calculate the average loss cost for a county. Hence, it is possible that years from 1980 to 2009 do not contain a dispersion of data such that each bin has one or more loss costs (i.e., not all bins are represented in the 1980–2009 period). For example, it may be that no year in the 1980–2009 period is classified as falling into bin 9. This will have adverse implications for the calculation of the average loss costs if not all bins are represented in the 1980–2009 period (i.e., not all bin probabilities are represented).

In particular, a range of observed weather history is not being captured in the weighting of loss costs. Therefore, to address the issue of empty bins and determine the appropriate number of bins, the approach we pursue is to first look at fifteen bins and then move down one bin at a time (i.e., from fifteen to two bins) to establish the largest number of bins for which there are no cases of empty bins in the years with loss data (from 1980 to 2009). This is done for each climate division, so the number of bins may vary for each climate division within a state.

A hypothetical example of bin classification results for soybeans in Mississippi is presented in table 2. In this example, there are ten bins, which assures that there are no empty bins from

1980 to 2009. All bin classifications are represented in the 1980–2009 data. The actual adjusted loss cost ratios for Bolivar County, Mississippi, and the bin classification process for Misissippi climate division 1 (which includes Bolivar County) is also depicted in figure 4. In table 2, the model insignificance flag is equal to 0, which means that the model fit results for this climate division are significant and at least ten observations are used in the estimation.

Loss Cost Averaging Procedure

After each year is classified into a particular bin at the climate division level (for all 115 years), the classified data for each year and the insignificance flags (based on regression model) are then merged back with the original county-level loss data. Since the regressions and year classifications based on the weather indexes are done at the climate division level, all counties within a particular climate division will have the same year classification and insignificance flags.

The average loss costs are next calculated using the 1980–2009 data for which there are actual loss cost values available in the StatPlan data. We first calculate the aggregate loss cost for each county, which is the current procedure used for computing the county base rate. Then we do a "weather-weighting" average of loss costs for each county. This weather weighting is done by first taking the average loss cost within each of the defined bins and then taking the "average of the average loss costs" across the bins. For example, if there are nine bins within a county, we first calculate a simple average of the loss costs within each of these nine bins (i.e., one average loss cost for each bin that results in nine "average" observations). Then we take the average of the nine average loss costs for the nine bins (i.e., "average of the average loss costs"). Since the bins are constructed to have equal probabilities, there is no need for taking a "weighted average of the average loss costs."

In the approach described above, a "recency weighting" procedure can also be applied when taking the average loss cost within a bin. That is, more recent years of data can be given more weight relative to older years within each bin. Alternatively, the procedure above can be easily implemented with less than the thirty years in the 1980–2009 data series, using, for example, data from 1990–2009 only, if the older loss data's loss environment is sufficiently different than the more recent time period. The procedure above also allows for consistency with the current catastrophic loading procedure. In this case, we also calculate the unweighted and weather-weighted average loss costs where the adjusted loss cost data are censored at the 80th percentile. A similar calculation can be done where the censoring is done at the 90th percentile (since there was a recent recommendation to increase the censoring for catastrophic loading to this level).

Results and Discussion

An example case in which county-level loss costs are merged with the bin classification data can be seen in table 3 for corn in Dewitt County, Illinois. The unweighted and weather-weighted average loss costs at the county level can be calculated using the type of data presented in table 3. The bin classification column allows us to conduct the weather-weighting procedure described above. If the insignificance flag for model fit is equal to 1 in any county, we do not recommend using weather weighting for the county and we do not report a weather-weighted average in these cases. While hundreds of models were estimated, two examples are reported in table 1 to illustrate the fractional logit results.

Table 4 presents examples of unweighted and weather-weighted average loss costs for several counties in Iowa. We calculate four loss costs averages (i.e., four weighting types) per county in table 4: (1) average loss cost calculated with no weather weighting or censoring, (2) average loss cost calculated with weather weighting and no censoring, (3) average loss cost calculated with no weather weighting but with censoring at the 80th percentile, and (4) average loss cost calculated with weather weighting and censoring at the 80th percentile. In the example counties in table 4,

Table 3. Hypothetical Example County-Level Data Used to Calculate Weather-Weighted Average Loss Costs for De Witt County (County=39), Climate Division 4, Illinois (State=17): Corn

Year	Actual Adjusted Loss Costs	Bin Classification
1980	0.123	10
1981	0.008	3
1982	0.004	2
1983	0.128	11
1984	0.008	5
1985	0.000	2
1986	0.000	5
1987	0.000	9
1988	0.132	10
1989	0.001	2
1990	0.003	3
1991	0.001	10
1992	0.001	1
1993	0.001	3
1994	0.000	3
1995	0.018	8
1996	0.000	2
1997	0.000	2
1998	0.001	8
1999	0.000	6
2000	0.000	4
2001	0.001	4
2002	0.012	9
2003	0.000	3
2004	0.001	1
2005	0.003	10
2006	0.001	7
2007	0.002	9
2008	0.001	3
2009	0.003	1

Notes: The climate division in this example had a sufficient number of observations to run a credible fractional regression model and calculate a predicted loss cost (weather index). Actual Adjusted Loss Costs are the loss cost ratios for the county where RMA adjustments were used to make the resulting data comparable. Optimal number of bins for this division (without empty bins) = 11.

Table 4. Hypothetical Example of Unweighted and Weather-Weighted Loss Costs at the County Level for Boone (County=15), Dallas (County=49), and Grundy (County=75) Counties within Climate Division 5 in Iowa (State=19): Corn

County	Avg. Loss Costs w/o Weather Weighting (No Censoring)	Avg. Loss Costs w/ Weather Weighting (No Censoring)	Avg. Loss Costs w/o Weather Weighting (80th Percentile Censoring)	Avg. Loss Costs w/ Weather Weighting (80th Percentile Censoring)
Boone	0.0096	0.0076	0.0028	0.0027
Dallas	0.0100	0.0097	0.0059	0.0058
Grundy	0.0092	0.0051	0.0013	0.0010

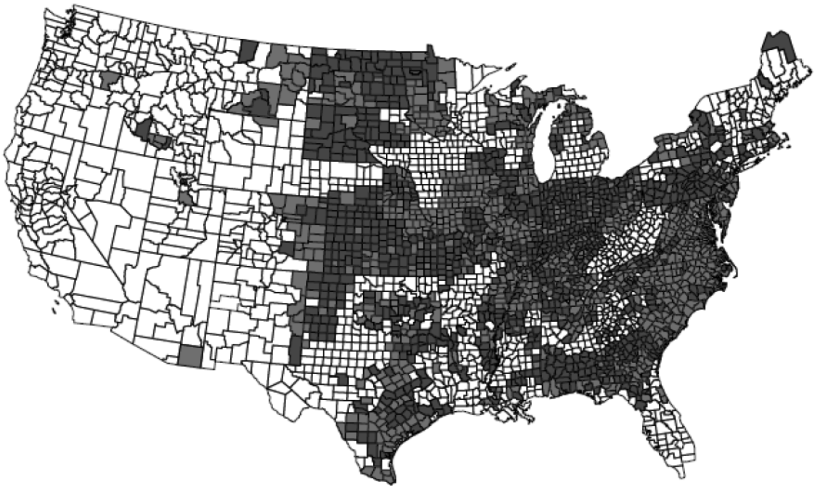

Figure 5. Map of the Difference between the Unweighted Average Loss Cost and the Weather-Weighted Loss Costs for Corn

Notes: Negative difference (e.g., weather weighted < unweighted) is in dark gray and positive difference (e.g., weather weighted > unweighted) is in light gray.

the weather-weighted average loss cost tends to be smaller than the unweighted average loss cost. However, this is not a pattern observed in every county-crop combination. There are cases where the weather-weighted average loss costs are higher than the unweighted average loss costs. These results may simply be a result of the particular sample used, and no additional implications should be drawn from them.

Table 5 presents the national average of the calculated unweighted and weighted loss costs for all crops examined. This is the liability-weighted average across counties (i.e., the liability-weighted average (not simple average) of the average county-level loss costs based on the 2009 liability of each county). For apples, barley, cotton, potatoes, rice, and spring/winter wheat, the weather-weighted average loss costs (at the national level) tend to be smaller than the unweighted loss costs. However, for corn, cotton, sorghum, and soybeans the weather-weighted average loss costs (at the national level) tend to be larger. Ultimately, we believe that the effect of weather weighting on premium rates is an empirical issue. The resulting premium rate outcomes for each crop in a county greatly depend on how the embodied weather distribution in the shorter loss sample compares to the weather distribution when incorporating information from the longer weather series. For the county-crop combinations in which weather weights lower premium rates, some low probability-high loss events were receiving too much weight in the previous approach where a simple average is taken (and vice-versa for county-crop combinations that resulted in higher premium rates).

A map showing the pattern of the differences between unweighted and weighted average loss costs for corn is presented in figure 5. Around 51% of counties have weather-weighted average loss costs lower than the unweighted loss costs. The effect of using climate division data shows up in some cases, as results tend to follow climate division lines. In general, these results suggest that weather-weighting produces rates that are different than the unweighted alternatives, and these weather-weighted rates should be more accurate because they better integrate weather information into the calculation.

As mentioned in the previous section, it is also possible to apply the proposed weighting procedure using less than thirty years of data. For example, twenty years of the most recent experience data (in our case from 1990 to 2009) can be used so that the more recent production environment is given more importance because this information is likely more credible than that

Table 5. Liability Weighted National Average (across Counties) of Unweighted and Weather-Weighted Average Loss Costs for Apples, Barley, Corn, Cotton, Potatoes, Rice, Sorghum, Soybeans, Spring Wheat, and Winter Wheat

Crop	No. of Counties	Unweighted Loss Cost Ratios (No Censoring)	Weather-Weighted Loss Cost Ratios (No Censoring)	Weather-Weighted Loss Cost Ratio as a Percentage of Unweighted Loss Cost Ratio (No Censoring)	Unweighted Loss Cost Ratios (Censoring at 80th Percentile)	Weather-Weighted Loss Costs Ratios (Censoring at 80th Percentile)	Weather-Weighted Loss Cost Ratio as a Percentage of Unweighted Loss Cost Ratio (Censoring at 80th Percentile)
Apples	140	0.183	0.176	96%	0.151	0.146	97%
Barley	646	0.103	0.095	92%	0.072	0.068	94%
Corn	1,930	0.050	0.052	104%	0.029	0.029	100%
Cotton	437	0.143	0.146	102%	0.110	0.111	101%
Potatoes	128	0.083	0.081	97%	0.066	0.065	98%
Rice	84	0.026	0.025	96%	0.015	0.015	100%
Sorghum	750	0.121	0.132	109%	0.089	0.092	103%
Soybeans	1,523	0.054	0.054	100%	0.038	0.038	100%
Spring Wheat	244	0.122	0.117	96%	0.089	0.087	98%
Winter Wheat	951	0.098	0.085	87%	0.072	0.065	90%

Notes: These are the national average loss costs across all counties (i.e., liability weighted average) in which the insignificance flags are not equal to 1. All weighted and unweighted loss costs for each county are available from the authors upon request.

Table 6. Changes in National Average Loss Cost Ratios when Experience Data Used for Setting County Base Rates Are from Twenty Years of Experience Instead of Thirty Years of Experience

Crop	Unweighted Loss Costs (No Censoring)	Weather-Weighted Loss Costs (No Censoring)	Unweighted Loss Costs (Censoring at 80th Percentile)	Weather-Weighted Loss Costs (Censoring at 80th Percentile)
Apples	106%	106%	107%	107%
Barley	80%	85%	89%	92%
Corn	82%	88%	88%	90%
Cotton	106%	97%	109%	103%
Potatoes	97%	98%	100%	100%
Rice	82%	90%	97%	98%
Sorghum	101%	102%	101%	102%
Soybeans	84%	87%	84%	87%
Spring Wheat	86%	96%	90%	95%
Winter Wheat	104%	105%	109%	107%

Notes: Figures reported in columns 2–5 above reflect the national average loss cost ratio based on twenty years of experience data as a percentage of the national average loss cost ratio based on thirty years of experience data (for the specific weighting and censoring scenarios examined).

from earlier years. Table 6 reports the changes in the estimated average loss costs when using twenty years instead of thirty years of data in the proposed weighting procedures. Results for apples, cotton, and winter wheat indicate that limiting the loss cost histories to twenty years instead of thirty would lead to higher rates. Conversely, barley, corn, soybeans, and spring wheat are observed to have lower rates. These are based on averages, and significant variation exists within each crop.

Conclusions and Implications

This article develops a procedure for weighting the historical loss cost experience used in calculating crop insurance premium rates. The idea is to use longer time-series information about weather variables to augment the shorter historical county loss cost data used for crop insurance rating, thereby improving statistical validity of premium rates. In developing the weighting methodology, the following factors were explicitly considered: statistical validity, feasibility, sustainability, and a balance of improvement versus complexity. The 2012 drought in the Midwest proved a useful illustration of the need for such an approach. Prior to 2012, much of the affected region had not experienced a drought since 1988. Not surprisingly, a twenty-three-year period between droughts made assessing the probability weights for this region difficult.

Our evaluation suggests that the National Climatic Data Center's Time Bias Corrected Divisional Temperature-Precipitation-Drought Index data, also called the Climate Division Data, is the most appropriate data set to use when weather weighting historical loss cost data. Fractional logit models can be used to relate loss cost experience to weather variables like the Palmer Drought Severity Index (PDSI) and Cooling Degree Days (CDD). Out-of-sample forecasting competition can then be used to select the specific weather variables the best explains weather experience for each particular climate division. This process creates a weather index from 1895 to the present, which allows the likelihood of the loss experienced in each year to be better assessed using a longer series of weather data. A variable-width, equal probability "binning" approach can then be implemented on the historical, county-level loss experience data to more properly calculate the expected (or average) loss cost used in estimating county base rates.

Results of this study show that a weather-weighting approach is indeed feasible within the context of U.S. crop insurance rating, and the approach developed in this study provides a way to capture longer term weather experience to augment shorter historical loss cost data used in estimating premium rates. Given that previous studies have provided evidence that asymmetric information problems are prevalent in the U.S. crop insurance program due to the inability to estimate premium rates commensurate with the actual level of risk (Goodwin and Smith, 1995), the weather-weighting approach developed here may be viewed as another step toward more accurately estimating premiums and hopefully reducing asymmetric information problems such as adverse selection and moral hazard. In this case, "hidden" weather probability information, which is not captured in the shorter historical loss cost data used for calculating base rates, can now be used to better assess county-level weather risk and improve the estimation of county base rates, which in turn can reduce asymmetric information problems in the long run.

In addition, the weather-weighting procedure proposed in this study can also be viewed as a viable method for addressing the effects of climate change on agricultural losses. First, this approach adds structure to the weighting of historical experience; thus it would be feasible to impose a climate change forecast such as temperature rise on rates. With continuous updating of the weather weights as additional climate information comes in, the proposed weighting procedure could consistently account for systematic changes in weather probabilities due to climate change.

Second, this approach lends itself to shortening the loss experience time series to twenty years. If there are trends in climatic conditions, then lags are reduced but not eliminated. Crop insurance rates are typically a two-step-ahead forecast, which means actuarial soundness would require adjusting the last twenty years of experience but would not require adjustments further into the future. Interestingly, Coble et al. (2013) find that loss cost ratios have not increased over time but rather decreased. We suspect technology gains have overwhelmed climate effects in recent years, as suggested by Yu and Babcock (2010).

[Received March 2014; final revision received April 2015.]

References

Beach, R. H., C. Zhen, A. Thomson, R. M. Rejesus, P. Sinha, A. W. Lentz, D. V. Vedenov, and B. A. McCarl. "Climate Change Impacts on Crop Insurance." RTI Project Number 0211911, U.S. Department of Agriculture, Risk Management Agency, Kansas City, MO, 2009.

Borman, J. I., B. K. Goodwin, K. H. Coble, T. O. Knight, and R. M. Rejesus. "Accounting for Short Samples and Heterogeneous Experience in Rating Crop Insurance." *Agricultural Finance Review* 73(2013):88–101. doi: 10.1108/00021461311321339.

Coble, K. H., T. O. Knight, B. K. Goodwin, M. F. Miller, and R. M. Rejesus. "A Comprehensive Review of the RMA APH and COMBO Rating Methodology: Final Report." Report prepared for the USDA Risk Management Agency by Sumaria Systems, Inc., U.S. Department of Agriculture, Risk Management Agency, Washington, DC, 2010.

Coble, K. H., T. O. Knight, M. F. Miller, B. K. Goodwin, R. M. Rejesus, and R. Boyles. "Estimating Structural Change in U.S. Crop Insurance Experience." *Agricultural Finance Review* 73(2013):74–87. doi: 10.1108/00021461311321320.

Dai, A., K. E. Trenberth, and T. Qian. "A Global Dataset of Palmer Drought Severity Index for 1870–2002: Relationship with Soil Moisture and Effects of Surface Warming." *Journal of Hydrometeorology* 5(2004):1117–1130. doi: 10.1175/JHM-386.1.

Dale, R. F., and J. A. Daniels. "A Weather-Soil Variable for Estimating Soil Moisture Stress and Corn Yield Probabilities." *Agronomy Journal* 87(1995):1115–1121. doi: 10.2134/agronj1995.00021962008700060013x.

Goodwin, B. K., and V. H. Smith. *The Economics of Crop Insurance and Disaster Aid*. Washington, D.C: American Enterprise Institute, 1995. Available online at http://www.aei.org/wp-content/uploads/2014/07/-the-economics-of-crop-insurance-and-disaster-aid_1/0212473810.pdf .

Guttman, N. B., and R. G. Quayle. "A Historical Perspective of U.S. Climate Divisions." *Bulletin of the American Meteorological Society* 77(1996):293–303. doi: 10.1175/1520-0477(1996)077<0293:AHPOUC>2.0.CO;2.

Intergovernmental Panel on Climate Change. *Climate Change 2007: Impacts, Adaptation and Vulnerability*. Cambridge, UK: Cambridge University Press, 2007.

Karl, T. R., C. N. Williams, P. J. Young, and W. M. Wendland. "A Model to Estimate the Time of Observation Bias Associated with Monthly Mean Maximum, Minimum and Mean Temperatures for the United States." *Journal of Climate and Applied Meteorology* 25(1986):145–160. doi: 10.1175/1520-0450(1986)025<0145:AMTETT>2.0.CO;2.

National Climatic Data Center. "Time Bias Corrected Divisional Temperature-Precipitation-Drought Index." Documentation for dataset TD–9640, DBMB, NCDC, NOAA, Asheville, NC, 1994.

Papke, L. E., and J. M. Wooldridge. "Econometric Methods for Fractional Response Variables with an Application to 401(k) Plan Participation Rates." *Journal of Applied Econometrics* 11(1996):619–632. doi: 10.1002/(SICI)1099-1255(199611)11:6<619::AID-JAE418>3.0.CO;2-1.

Payero, J. O., S. R. Melvin, S. Irmak, and D. Tarkalson. "Yield Response of Corn to Deficit Irrigation in a Semiarid Climate." *Agricultural Water Management* 84(2006):101–112. doi: 10.1016/j.agwat.2006.01.009.

Ramalho, E. A., J. J. S. Ramalho, and J. M. R. Murteira. "Alternative Estimating and Testing Empirical Strategies for Fractional Regression Models." *Journal of Economic Surveys* 25(2011):19–68. doi: 10.1111/j.1467-6419.2009.00602.x.

Tack, J. "A Nested Test for Common Yield Distributions with Applications to U.S. Corn." *Journal of Agricultural and Resource Economics* 38(2013):64–77.

Tack, J., A. Harri, and K. H. Coble. "More than Mean Effects: Modeling the Effect of Climate on the Higher Order Moments of Crop Yields." *American Journal of Agricultural Economics* 94(2012):1037–1054. doi: 10.1093/ajae/aas071.

Vergara, O., G. Zuba, T. Doggett, and J. Seaquist. "Modeling the Potential Impact of Catastrophic Weather on Crop Insurance Industry Portfolio Losses." *American Journal of Agricultural Economics* 90(2008):1256–1262. doi: 10.1111/j.1467-8276.2008.01213.x.

Wilhelmi, O. V., K. G. Hubbard, and D. A. Wilhite. "Spatial Representation of Agroclimatology in a Study of Agricultural Drought." *International Journal of Climatology* 22(2002):1399–1414. doi: 10.1002/joc.796.

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. doi: 10.1093/ajae/aaq074.