Accounting for short samples and heterogeneous experience in rating crop insurance

Julia I. Borman
North Carolina State University, Raleigh, North Carolina, USA
Barry K. Goodwin
Agricultural and Resource Economics, North Carolina State University, Raleigh, North Carolina, USA
Keith H. Coble
Agricultural Economics, Mississippi State University, Mississippi State, Mississippi, USA
Thomas O. Knight
Agricultural and Applied Economics, Texas Tech University, Lubbock, Texas, USA, and
Rod Rejesus
Agricultural and Resource Economics, North Carolina State University, Raleigh, North Carolina, USA

Abstract

Purpose – The purpose of this paper is to be an academic inquiry into rating issues confronted by the US Federal Crop Insurance program stemming from changes in participation rates as well as the weighting of data to reflect longer-run weather patterns.

Design/methodology/approach – The authors investigate two specific approaches that differ from those adopted by the Risk Management Agency, building upon standard maximum likelihood and Bayesian estimation techniques that consider parametric densities for the loss-cost ratio.

Findings – Both approaches indicate that incorporating weights into the priors for Bayesian estimation can inform the distribution.

Originality/value – In most cases, the authors' results indicate that including weighting into priors for Bayesian estimation implied lower premium rates than found using standard methods.

Keywords Crop insurance, Loss distribution, Crops, Insurance

Paper type Research paper

Introduction

Crop insurance is one avenue available to agricultural producers to protect themselves against natural hazards. The US Department of Agriculture (USDA) Risk Management

This article was presented in the Finance Sections at the 2012 annual meeting of the Agricultural and Applied Economics Association (AAEA) in Seattle, Washington, August 12-14, 2012. The articles in these sessions are not subjected to the journal's standard refereeing process. Sessions organized by Nicholas Paulson, University of Illinois. Papers reviewed for publication by Nicholas Paulson and C.G. Turvey.

Agricultural Finance Review
Vol. 73 No. 1, 2013
pp. 88-101

0002-1466
DOI 10.1108/00021461311321339

Agency (RMA), which is in charge of administering the US Federal Crop Insurance program has, until recently, developed rates primarily based on a simple average of historical loss-cost ratios (i.e. the ratio of indemnity payments to total liability). When loss events occur independently, an insurer is able to use fewer years of loss experience to accurately determine the expected future loss-costs. In the case of agriculture, losses tend to be spatially correlated (Glauber, 2004). This creates the need to employ long time-series of yields or losses to accurately model yield risk. Although the RMA uses several decades of data to compensate for this information problem, using a simple average has the potential to improperly weight the significance of any single year's experience. For instance, in the summer of 2012, the US experienced one of the most extensive and extreme droughts since the 1950s. Such an event would be given the same weight as any other year's loss.

The US Federal Crop Insurance program has become the largest of the agricultural commodity programs. For 2012, the program carries a total liability just over \$116 billion and insures 281 million acres. The program typically spends approximately \$7 billion annually in premium subsidy payments and provides subsidies to private companies to administer the program. Currently, a new Farm Bill that remains under consideration with Congress may expand the current insurance coverage to include "shallow-losses." Both farmers adopting crop insurance plans and private insurance companies issuing plans receive significant taxpayer subsidies. Smith (2011) estimated that private insurers received \$1.44 for every dollar farmers have received in crop insurance subsidies. Private insurance companies are also provided with an advantageous taxpayer-supported reinsurance agreement. Important legislative changes, including the 1994 Crop Insurance Reform Act and the 2000 Agricultural Risk Protection Act, brought about significant changes in the program including increased subsidies and a substantial rise in participation. Figure 1 shows total acreage enrolled in the program in recent years and documents the significant increases in participation that followed program changes.

The Risk Management Agency recently commissioned two studies to evaluate the overall approach used to establish premium rates and the terms of coverage

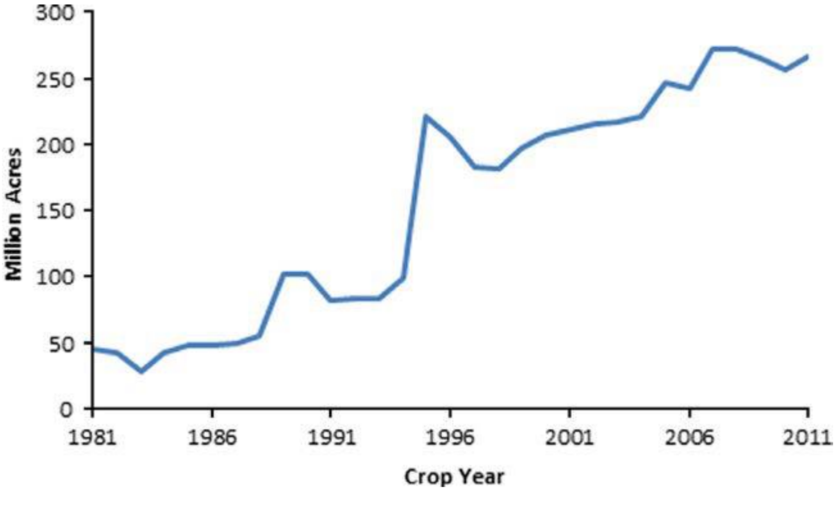

Figure 1.
Total acreage insured in the federal crop program

in individual plans of insurance[1]. The fundamental approach used in establishing base premium rates for yield coverage, which is described in detail in these studies, has largely involved a simple arithmetic average of historical annual loss-cost ratios. A number of caveats apply to this basic approach, including smoothing to accommodate catastrophic losses and various adjustments for other risks in the program such as prevented planting and differences in crop practices. Annual loss-cost data from 1975 through the present are typically used to calculate an unweighted average loss-cost which is then used to represent the base premium rate. Following recommendations raised in these rate reviews, the RMA adopted significant changes in this basic rating approach to recognize the problems associated with changes in the program, technological innovations, and the weighting of weather events to recognize the relatively short sample of annual data available for rating.

The changes made to address program changes over time are documented in the accompanying paper by Coble *et al.* In this paper, we discuss alternative approaches to empirically addressing the significant changes that have occurred in the program as well as the weighting of data to reflect longer-run weather patterns. The proper approach to recognizing these issues in rating remains an important research topic. Many different approaches are conceivable and there is little existing evidence to suggest the superiority of any single approach. In this analysis, we investigate two specific approaches that differ from those adopted by the RMA and instead build upon standard maximum likelihood and Bayesian estimation approaches that consider parametric densities for the loss-cost ratio that incorporate changes in participation and weighting of weather events in calculating premium rates. Our analysis is intended to stimulate academic inquiry into these rating issues and is not meant to suggest the superiority of any single approach.

Short samples, dependencies, and structural changes

The federal crop insurance program has some rather unique issues that are not commonly encountered in most commercial lines of property and casualty insurance. As noted, individual loss events tend to be highly correlated in the spatial dimension. This gives rise to the "sib-pairs" problem that is often encountered in case-control genetic association statistics, where dependencies among individual observations suggests that the number of "effective observations" that consists in a sample may be far less than the actual number of observations. In case-control genetic association studies, characteristics measured across multiple members taken from a common group (e.g. a family) are recognized to be non-independent. These related members are often termed "sib-pairs" (or sibling-pairs). For example, for a sample taken across N individuals, the effective number of observations may be expressed as λN , where $0 < \lambda < 1$. As the correlation among individuals approaches zero, λ approaches 1. However, as this degree of correlation increases, λ decreases. The effective number of observations is defined as the equivalent number of independent observations that lead to the same variance for the variable of interest. The central question as it pertains to the changing level of participation in the crop insurance program is whether 200 million insured acres gives ten-times as much information about risks as do 20 million acres.

A common measure of the number of effective observations in a correlated sample can be derived by considering the number of independent groups and then the number of individuals within the correlated groups. For example, consider a case of 100 observations made up of ten groups of ten. Across the groups, observations are independent. However, within the groups, observations are correlated with a Pearson (linear) correlation coefficient of ρ. The effective number of observations for this sample of 100 observations will be less than 100 if the correlation is greater than zero. In the case of m equally sized independent groups of sib-pairs, the effective number of observations is given by $N^e = N/(1+(m-1)\rho)$[2]. Figure 2 shows the relationship between actual and effective numbers of observations for correlated groups. Clearly, as the degree of dependence rises, the number of effective observations falls.

Accompanying this issue is the fact that many underlying structural factors that may be relevant to risks have changed over time. Such factors include a number of program changes beyond changing participation. Perhaps most relevant is the fact that crop insurance offerings have expanded to include a number of innovative plans. In recent years, over 70 percent of the total liability in the program has been for revenue coverage. It is also the case that production agriculture has realized a number of technological changes, including the advent of biotech crops which many believe to have lower risk than conventional varieties.

The changing structure of crop insurance and agriculture suggests that one should base measures of risk on the most recently available data, which are more representative of the contingencies being rated for. On the other hand, the dependencies of loss events and the systemic nature of weather suggests that data in the cross-section are likely to be highly correlated and therefore offer much less information that the total number of (annual and cross sectional) observations might imply. This suggests that one needs more information (i.e. a longer time-series) to adequately measure risk. Finally, a related sample size problem pertains to the fact that many of the relevant risks in agriculture may involve events that are only rarely observed. Such events may trigger catastrophic losses. Examples would include a $1-$ in $-$ 100 year drought such as the type experienced this summer. The Coble *et al.* reports noted that a much more extensive time-series of weather variables is available.

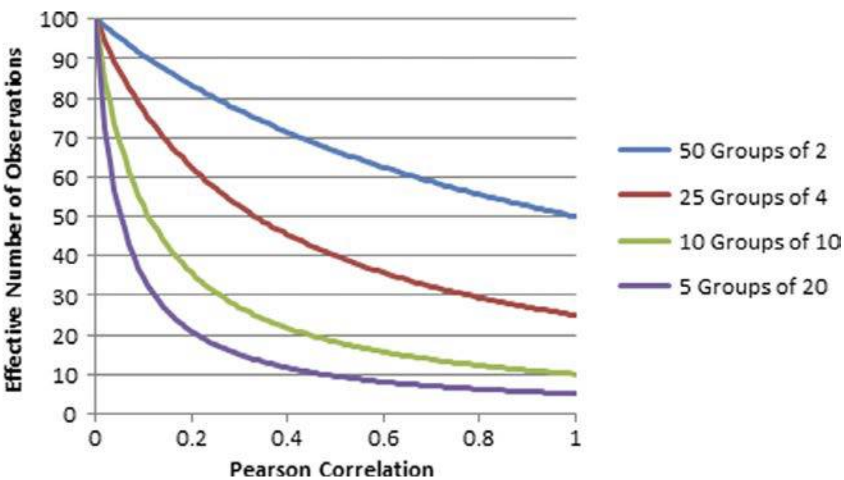

Figure 2.
Effective number of observations in dependent samples

They suggested a nonparametric smoothing approach that would properly address this issue. Here, we explore an alternative approach that uses weather experience to form empirical frequency priors in Bayesian estimates of a parametric loss-cost density.

Empirically rating with short samples and structural change

A variety of different approaches exist to deriving insurance premium rates that account for the aforementioned issues associated with the limited amount of data and changes that have characterized the underlying risk over time. RMA adopted a nonparametric approach to rating that involved using variable bin-width histograms that accounted for the relative frequency of weather events over a longer period of time. In addition, as is described in the companion piece to this paper, the loss-costs were adjusted for a structural break that corresponded to significant legislative changes in 1994. Here, we describe alternative approaches to accommodating these issues. The problem is inherently of a Bayesian flavor. Prior information drawn from a much longer set of weather data can be used to suggest alternative weights for individual years of loss-cost experience. Weights based upon the acreage insured may suggest that later experience data is more relevant to expected losses than the earlier history based upon experience with far fewer acres. We adopt a heuristic, approximate Bayesian approach that empirically derives informative priors from weighted maximum likelihood estimation of a loss-cost distribution. Weather weights are derived from the frequency of observed events over the much longer set of weather data (dating to 1895). The resulting posterior distributions for the parameters of interest incorporate the information contained in the weighted estimates[3].

Our basic approach is as follows. We first develop a weight based on historical weather or acreage and apply this weight to a maximum likelihood estimated density for the loss-cost ratio over the 1981-2010 period. A truncated normal defined over the [0,1] interval is used to represent the loss-cost ratio. RMA has adopted a truncated normal density in various aspects of its rating and the specification is quite flexible in representing skewness while maintaining consistency in restricting loss-costs to lie between zero and one. The shape parameter estimates and standard errors derived from maximizing this weighted likelihood function are then used to define empirical priors. Specifically, we adopted normal priors with the mean and standard errors set using the weighted maximum likelihood estimates[4]. Two versions of the priors are considered - one using the standard error from the weighted ML estimates as the standard deviation for the prior and a second which uses one-half of the standard error – representing tighter priors. In this paper, we focus on six important corn producing counties from 2010 in Iowa-Kossuth, Lyon, O'Brien, Plymouth, Sioux, and Story. The distribution of loss-costs for yield protection coverage is shown in Figure 3[5]. Note that the yield risk for the selected counties tends to be relatively low, with many instances of zero losses. This results in the data being concentrated at zero with a substantially skewed right tail.

Development of likelihood weights

Although the most recently available data may be more representative of the contingencies being rated for, the dependencies of loss events and the systemic nature of weather suggests that data in the cross-section are likely to be highly correlated.

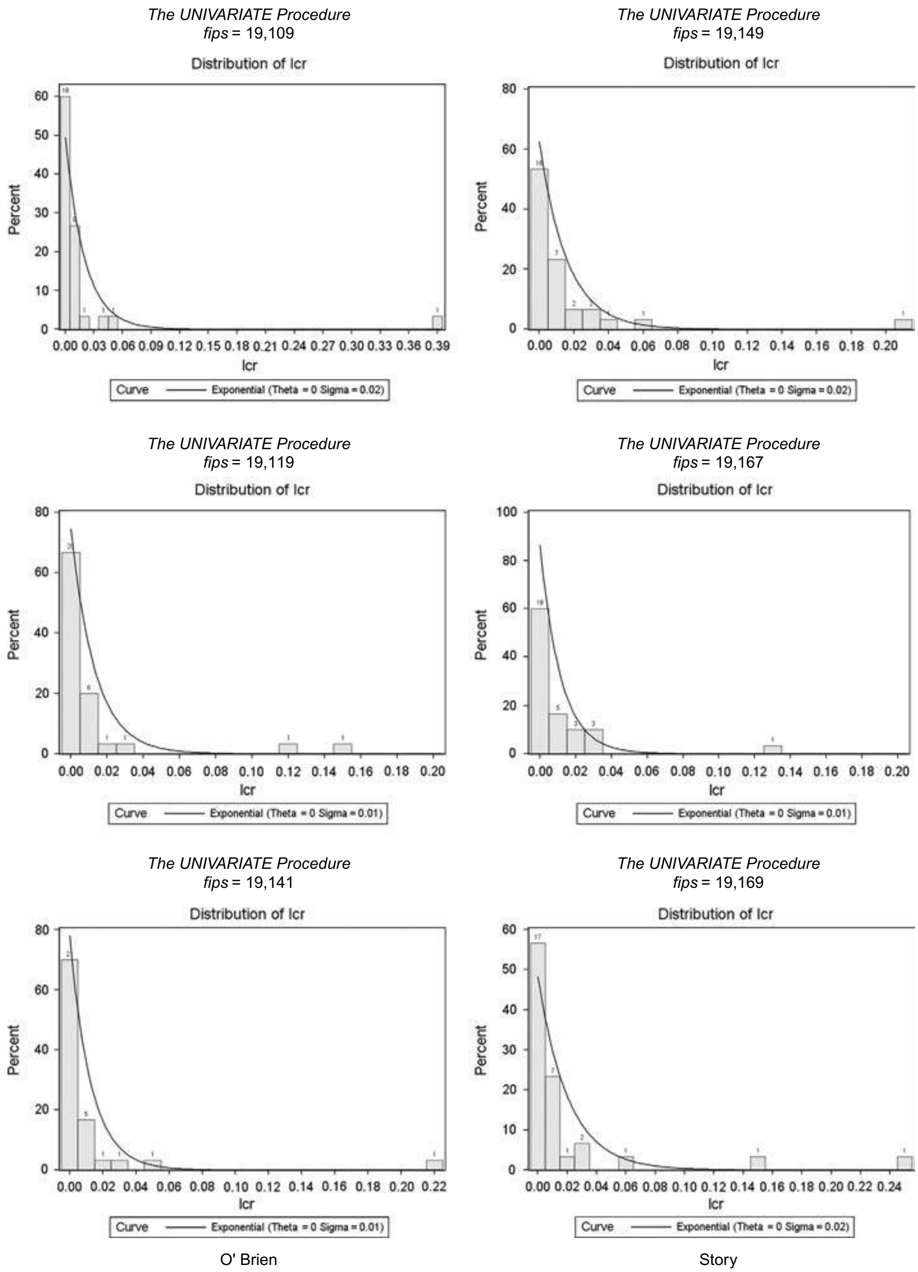

Figure 3.
Distribution of loss-costs for yield protection coverage

We would thus expect less information that the total number of (annual and cross sectional) observations might imply. This suggests that one needs more information (i.e. a longer time-series) to adequately measure risk. To incorporate this longer time-series of data, we develop two types of weights-one based on historical weather information and one on insured acreage.

We first fit a parametric density to the weather variable thought to be most relevant to yield risks – the Palmer Z drought index. The Palmer Z index is a short-term drought index available from the National Oceanic and Atmospheric Administration (NOAA)[6]. In this application, we use the index values for July only since it is a key month for corn yields. Experimentation with alternative parametric distributions suggested that a normal density was appropriate for representing the distribution of Palmer's Z index. The density of the index was then used to estimate the relative frequency with which each observation of the Palmer Z index in the sample of crop insurance experience data (1981-2010) occurred in the longer series of weather data. Letting p_i be the probability of each observation i, and $p = \sum_{t-1}^{T} pt$, be the sum of all observations at the county level, we can construct the weather weight as:

$$w_w = \frac{p_i}{p}$$

Acreage weights are constructed as a share of total acreage insured in the county over the years 1981-2010. Letting at be the acreage insured in the county at year t of each observation i, and $a = \sum_{t-1}^{T} at$, be the sum of all observations across all years at the county level, we can construct the acreage weight:

$$w_a = \frac{a_t}{a}$$

In this way, years with more insured acreage are given more weight in the likelihood function.

Estimation and results

The loss-cost ratio is the ratio of indemnity payments to liability, and thus is bounded on the unit interval. Most approaches to rating consider the mean of the loss-cost ratio as an estimate of the actuarially-fair premium rate. RMA has adopted a truncated normal density in various aspects of its rating and the specification is quite flexible in representing skewness while maintaining consistency in restricting loss-costs to lie between zero and one. The log likelihood function of interest is given by:

$$LLF(LCR; \mu, \sigma, 0, 1) = \delta w \times log\left(\frac{(1/\sigma)\phi((LCR - \mu)/\sigma)}{\Phi((1-\mu)/\sigma) - \Phi(-\mu/\sigma)}\right)$$

where $\delta = 0$ for the unweighted specification, and $\delta = 1$ for a weighted specification (either by weather or acreage). Maximizing the likelihood function yields ML estimates of the parameter vector $\theta = (\mu, \sigma)$, for each county of interest. These estimations form the priors for our Bayesian estimation using a random walk Metropolis algorithm.

The final step in our process involves using the metropolis algorithm in a Markov chain Monte Carlo (MCMC) estimation context to derive estimates of parameters of the posterior densities. The logarithm of the posterior density is as follows:

$$log(p(\theta|y)) = log(\pi(\theta)) + \sum_{k=1}^{n} log(f(y_k|\theta))$$

where θ are the priors determined by the weighted maximum likelihood estimations detailed above and $log(\pi(\theta))$ is the sum of the log of the prior densities. We again assume that the density of the loss-cost ratio ($f(y;\theta)$) follows a truncated normal distribution and that each observation in the data set is independent. The algorithm runs 10,000 iterations to obtain the posterior. We discard the first 1,000 iterations to account for starting values of the chain (e.g. the "burn-in"). The procedure cumulatively adds the log likelihood for each observation. The unweighted maximum likelihood estimates and posterior estimates for weighted specifications are presented in Table I. Figure 4 shows trace plots and posterior densities for selected counties. The plots are representative of those obtained for all cases considered and indicate proper convergence and satisfactory mixing, supporting the validity of the Bayesian estimates.

Maximum likelihood and Bayesian posterior estimates of the parameters (μ and σ) of the truncated normal loss-cost densities are presented in Table I. Three versions of the density are estimated for each county. The first uses standard maximum likelihood estimation techniques. The estimates include unweighted versions as well as estimates derived from weighted likelihood functions, with weights derived from historical acreage and weather variables. The estimated mean parameters of the truncated normal distributions are all negative, suggesting a positive skewed distribution. The standard deviations of the distributions are all relatively small, indicating relatively low loss-costs and thus low risks.

In general, weighting for the frequency of observed weather events or for relative changes in acreage tends to shift the maximum likelihood estimates. The biggest differences are naturally observed for the fully weighted maximum likelihood versions. In the case of the Bayesian posterior mean values of the parameters, the densities are a combination of the unweighted and weighted estimates. Specifically, the posterior estimates represent a "shrinkage" type estimate from the priors. The extent of shrinkage is determined by the variance of the priors, with a larger variance reflecting less confidence in the prior and therefore placing greater weight on the unweighted sample.

The most straightforward comparison of the implications for rating from the alternative sets of estimates can be derived from a consideration of the implied distributions of loss-costs. In particular, the mean loss-cost represents an estimate of the premium rate and various quantiles of the distribution can be taken to imply probable maximum loss (PML) values, which represent the loss-cost ratio that one would expect to exceed over a particular number of years[7]. Table II presents summary statistics for the implied loss-costs. In general, the differences in mean loss-costs across the alternative approaches to estimation are relatively modest.

	Unweighted		Weather weighted		Acres weighted	
	μ	σ	μ	σ	μ	σ
Maximum likelihood						
Kossuth	−0.1996	0.0749	−0.1805	0.0664	−0.0555	0.0276
	(0.4273)	(0.0527)	(0.0701)	(0.0085)	(0.2604)	(0.0384)
Lyon	−0.0488	0.0333	−0.0487	0.0312	−0.0430	0.0224
	(0.1326)	(0.0232)	(0.0234)	(0.0040)	(0.2025)	(0.0308)
O'Brien	−0.1212	0.0373	−0.0923	0.0386	−0.0393	0.0201
	(0.2459)	(0.0263)	(0.0373)	(0.0049)	(0.1838)	(0.0276)
Plymouth	−0.0552	0.0390	−0.0490	0.0351	−0.0406	0.0287
	(0.1526)	(0.0271)	(0.0249)	(0.0045)	(0.2201)	(0.0392)
Sioux	−0.0204	0.0213	−0.0198	0.0214	−0.0174	0.0174
	(0.0706)	(0.0147)	(0.0128)	(0.0027)	(0.1146)	(0.0236)
Story	−0.0798	0.0570	−0.0751	0.0482	−0.0398	0.0208
	(0.2221)	(0.0397)	(0.0362)	(0.0061)	(0.1882)	(0.0286)
Posterior means and standard deviations (prior 1: $\hat{\sigma}$)						
Kossuth			−0.1855	0.0674	−0.1893	0.0674
			(0.0428)	(0.0055)	(0.0659)	(0.0080)
Lyon			−0.0506	0.0317	−0.0624	0.0334
			(0.0149)	(0.0026)	(0.0289)	(0.0045)
O'Brien			−0.0956	0.0391	−0.1049	0.0402
			(0.0230)	(0.0031)	(0.0395)	(0.0050)
Plymouth			−0.0509	0.0357	−0.0629	0.0377
			(0.0152)	(0.0029)	(0.0307)	(0.0052)
Sioux			−0.0207	0.0217	−0.0260	0.0227
			(0.0080)	(0.0017)	(0.0149)	(0.0030)
Story			−0.0777	0.0489	−0.0663	0.0463
			(0.0226)	(0.0038)	(0.0279)	(0.0048)
Posterior means and standard deviations (prior 2: $\hat{\sigma}/2$)						
Kossuth			−0.1797	0.0666	−0.1492	0.0620
			(0.0285)	(0.0035)	(0.0476)	(0.0063)
Lyon			−0.0488	0.0313	−0.0560	0.0324
			(0.0095)	(0.0017)	(0.0244)	(0.0040)
O'Brien			−0.0919	0.0387	−0.0874	0.0378
			(0.0148)	(0.0020)	(0.0297)	(0.0042)
Plymouth			−0.0490	0.0352	−0.0591	0.0370
			(0.0100)	(0.0019)	(0.0270)	(0.0047)
Sioux			−0.0198	0.0215	−0.0245	0.0225
			(0.0053)	(0.0011)	(0.0136)	(0.0027)
Story			−0.0749	0.0484	−0.0434	0.0405
			(0.0146)	(0.0026)	(0.0190)	(0.0033)

Note: Standard errors/posterior standard deviations in parentheses

Table I.
Parametric loss-cost distribution estimates

In most cases, weighting by the frequency of observed weather events tends to result in lower premium rates. However, for Kossuth county, with the exception of the weather weighted maximum likelihood estimates, the results all indicate a higher premium rate than the unweighted specification. A comparison of the MCMC generated mean loss-cost ratios to the 2011 yield protection premium rates (evaluated at a rate yield equal to the county reference yield) for 70 percent coverage or above indicates our

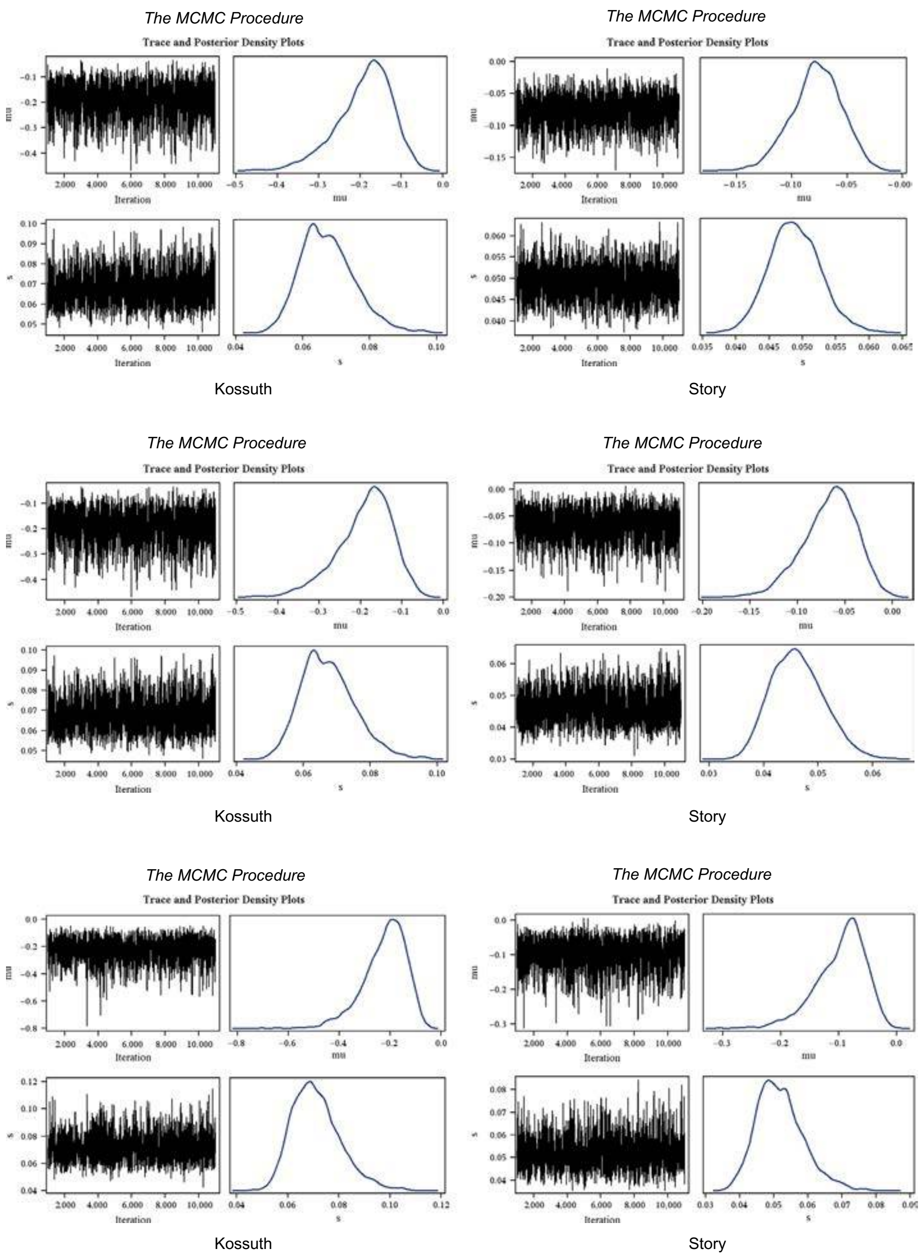

Notes: These plots serve as illustration of the MCMC process; the MCMC algorithm runs 10,000 iterations to obtain the posterior distribution; the trace plots provide strong evidence of proper convergence and satisfactory mixing

Figure 4.
Selected MCMC plots (under priors 1)

	Means						
	Maximum likelihood			*Priors 1: $\hat{\sigma}$*		*Priors 1: $\hat{\sigma}$*	
	Unweighted	*Weather*	*Acres*	*Weather*	*Acres*	*Weather*	*Acres*
Kossuth	0.0180	0.0102	0.0231	0.0199	0.0194	0.0216	0.0236
Lyon	0.0134	0.0086	0.0149	0.0129	0.0128	0.0135	0.0134
O'Brien	0.0129	0.0077	0.0083	0.0124	0.0124	0.0126	0.0128
Plymouth	0.0156	0.0128	0.0176	0.0159	0.0152	0.0152	0.0156
Sioux	0.0116	0.0090	0.0114	0.0113	0.0114	0.0115	0.0113
Story	0.0211	0.0080	0.0257	0.0207	0.0203	0.0207	0.0204
	1 in 10 PML (90th percentile)						
	Maximum likelihood			*Priors 1: $\hat{\sigma}$*		*Priors 1: $\hat{\sigma}$*	
	Unweighted	*Weather*	*Acres*	*Weather*	*Acres*	*Weather*	*Acres*
Kossuth	0.0415	0.0226	0.0539	0.0476	0.0424	0.0495	0.0489
Lyon	0.0300	0.0189	0.0331	0.0286	0.0284	0.0292	0.0300
O'Brien	0.0290	0.0176	0.0194	0.0281	0.0268	0.0287	0.0290
Plymouth	0.0347	0.0286	0.0387	0.0355	0.0336	0.0346	0.0342
Sioux	0.0253	0.0195	0.0249	0.0248	0.0250	0.0252	0.0248
Story	0.0470	0.0178	0.0562	0.0455	0.0455	0.0455	0.0458
	1 in 20 PML (95th percentile)						
	Maximum likelihood			*Priors 1: $\hat{\sigma}$*		*Priors 2: $\hat{\sigma}/2$*	
	Unweighted	*Weather*	*Acres*	*Weather*	*Acres*	*Weather*	*Acres*
Kossuth	0.0551	0.0279	0.0648	0.0613	0.0524	0.0623	0.0604
Lyon	0.0371	0.0243	0.0417	0.0373	0.0369	0.0373	0.0376
O'Brien	0.0361	0.0219	0.0205	0.0351	0.0352	0.0350	0.0370
Plymouth	0.0437	0.0355	0.0485	0.0446	0.0429	0.0428	0.0436
Sioux	0.0310	0.0244	0.0306	0.0308	0.0310	0.0312	0.0305
Story	0.0591	0.0219	0.0700	0.0561	0.0574	0.0560	0.0556

Table II.
Estimated loss-cost means and PML

estimates are all below actual rates. This suggests that extremely dry weather events, such as in a drought, may be receiving too much weight in a simple average. Likewise, weighting by acreage results in lower rates in most cases, reflecting the fact that those years with greater participation have tended to have more positive loss experience. The PML values are quite similar across the alternative estimates, with the unweighted and posterior estimates yielding very similar PML values.

Figures 5 and 6 show the loss-cost density functions estimated using the alternative approaches. In general, weighting by weather frequency results in the biggest differences in the implied densities. This is also reflected in the mean values implied by the distributions. The densities obtained from the posterior estimates are quite similar to the unweighted densities in most cases.

Conclusion

In this paper we present a new method for incorporating weather and participation into the parametric distribution for the loss-cost ratio. Past RMA methods have used simple averages in the construction of county level base rates, which may result in improperly weighting weather experience or compensate for structural changes which have changed participation rates. Our method expands upon previous maximum likelihood and Bayesian methods, and indicates that incorporating weights into the priors for Bayesian estimation can inform the distribution.

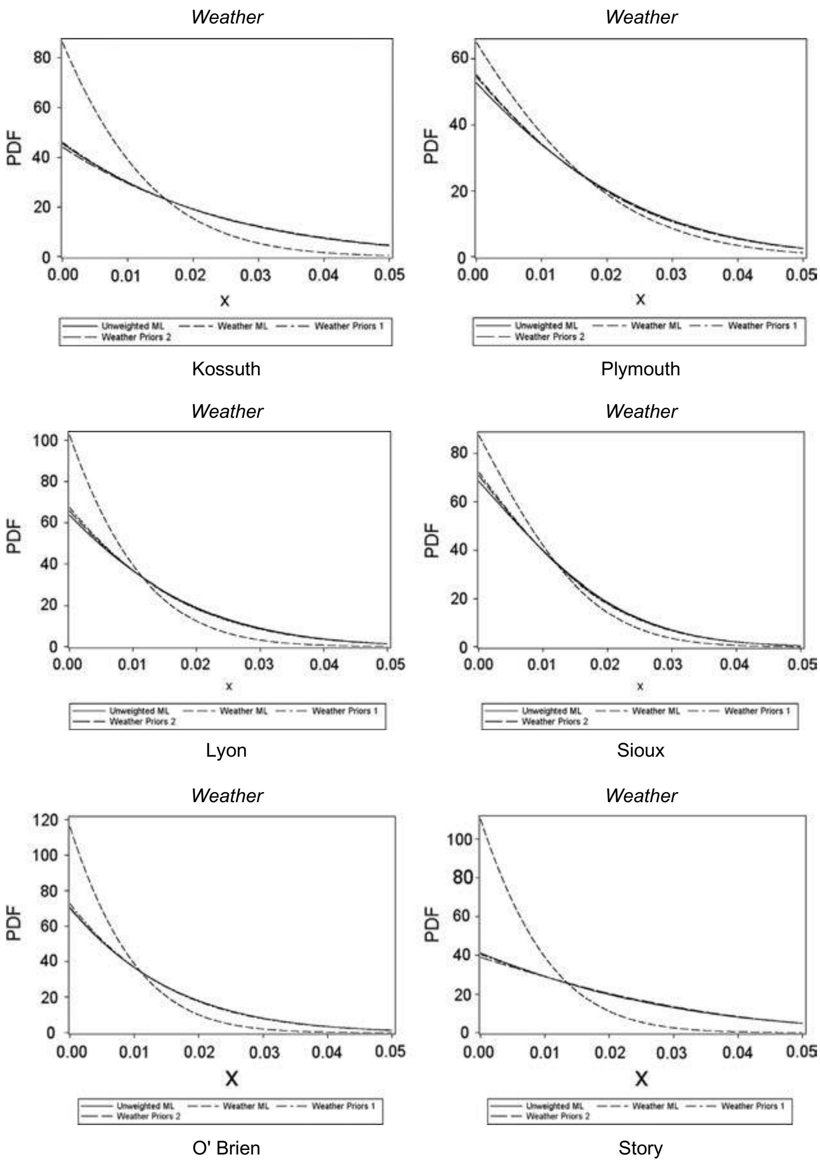

Notes: In each plot we present the probability distribution function yielded from the unweighted maximum likelihood specification, the specification weighted by weather and the two weather weighted MCMC results; for all counties the weather weighted ML specification is more negatively skewed than its unweighted counterpart, implying lower premium rates; for four counties (Plymouth, Lyon, Sioux and O'Brien), the MCMC procedure predicts the distribution to be between the two ML estimates, indicating a shrinkage type estimator resulting from this two stage process

Figure 5.
Weather PDFs

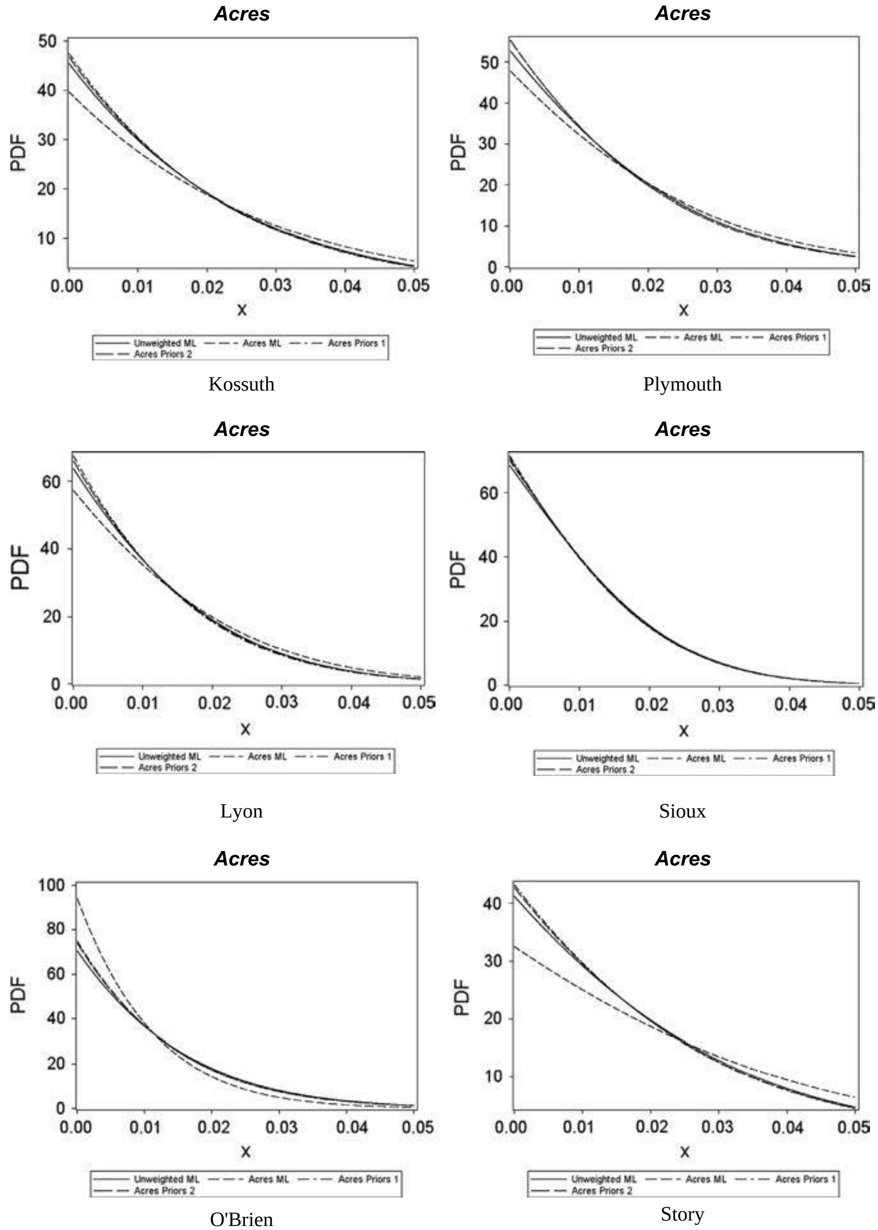

Figure 6.
Acreage PDFs

Notes: In each plot we present the probability distribution function yielded from the unweighted maximum likelihood specification, the specification weighted by acreage and the two acreage weighted MCMC results; in most cases, the MCMC weighting indicates lower rates; however, unlike in the case of weather weighting, the result of weighting the ML estimates does not skew the loss-cost distribution in a consistent way

Notes

1. See Coble *et al.* (2010, 2011) for details on each study.
2. See, for example, Yang *et al.* (2011) for a discussion of the determination of effective sample sizes among correlated groups of sibs.
3. Our approach is similar in spirit to the density smoothing methods of Whittle (1958) and the weighted likelihood bootstrap method of Newton and Raftery (1994)
4. The distribution of the standard deviation was restricted to be positive.
5. It is important to note that coverage has steadily shifted away from yield to revenue coverage. Thus, the total acreage from which the loss-cost ratios are drawn has declined in recent years after rises significantly prior to the mid-1990s. As is described in the Coble *et al.* (2010) review, RMA actually converts revenue coverage to equivalent yield coverage. However, the converted loss-costs are not publicly available and thus we rely on yield-only APH coverage to illustrate our method.
6. In cases where a county lies in multiple divisions, we adopt the same procedures as Coble *et al.* (2010) for determining which division counties are assigned to.
7. The PML values are equivalent to value-at-risk (VaR) values.

References

Coble, K.H., Knight, T.O., Goodwin, B.K., Rejesus, R.M. and Miller, M.F. (2010), "A comprehensive review of the RMA APH and COMBO rating methodology", Final Report. Technical Report, March 15.

Coble, K.H., Rejesus, R.M., Goodwin, B.K., Boyles, R., Knight, T.O. and Duffield, G. (2011), "Methodology analysis for weighting historical experience technical report methodology analysis for weighting of historical experience by contact: methodology analysis for weighting historical experience", Technical Report.

Glauber, J.W. (2004), "Crop insurance reconsidered", *American Journal of Agricultural Economics*, Vol. 86 No. 5, pp. 1179-1195.

Newton, M.A. and Raftery, A.E. (1994), "Approximate Bayesian inference with the weighted likelihood bootstrap", *Journal of the Royal Statistical Society. Series B (Methodological)*, Vol. 56 No. 1, pp. 3-48.

Smith, V.H. (2011), *Premium Payments: Why Crop Insurance Costs too Much*, American Enterprise, Institute, Washington, DC, pp. 1-19.

Whittle, P. (1958), "On the smoothing of probability density functions", *Journal of the Royal Statistical Society. Series B (Methodological)*, Vol. 20 No. 2, pp. 334-343.

Yang, Y., Remmers, E.F., Ogunwole, C., Kastner, D., Gregersen, P. and Li, W. (2011), "Effective sample size: quick estimation of the effect of related samples in genetic case-control association analyses", *Computational Biology and Chemistry*, Vol. 35 No. 1, pp. 40-49, available at: www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3119257\&tool=pmcentrez\&rendertype=abstract

Corresponding author
Barry K. Goodwin can be contacted at: barry_goodwin@ncsu.edu

This article has been cited by:

1. Martin Odening, Zhiwei Shen. 2014. Challenges of insuring weather risk in agriculture. *Agricultural Finance Review* **74**:2, 188-199. [Abstract] [Full Text] [PDF]