



**AgEcon** SEARCH  
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

*The World's Largest Open Access Agricultural & Applied Economics Digital Library*

**This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.**

**Help ensure our sustainability.**

Give to AgEcon Search

AgEcon Search  
<http://ageconsearch.umn.edu>  
[aesearch@umn.edu](mailto:aesearch@umn.edu)

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

# When Less is More: Historical Yield Data and Rating Area Crop Insurance Products <sup>1</sup>

Yong Liu and Alan P. Ker

Department of Food, Agricultural and Resource Economics

University of Guelph

## ABSTRACT

The Federal Crop Insurance Program -- operated by the United States Department of Agriculture's Risk Management Agency (RMA) -- offers various types of insurance, covers a multitude of crops, carries significant liability, and is the cornerstone of domestic farm policy. Currently, RMA uses county yield data from the 1950s onwards to set guarantees and estimate premium rates for their area yield and revenue insurance products but trims yield data prior to 1991 in rating their newer shallow loss products. The past 70 years reflect very significant innovations in both seed and farm management technologies; innovations that have likely moved mass all around the support of the yield distribution. Although the RMA rating methodology corrects for time-varying movements in the first two moments, it is unclear whether using the entire yield series remains appropriate. We use distributional tests and an out-of-sample retain-cede rating game to answer if RMA should or should not historically trim yields in estimating their premium rates. Despite small sample sizes and the need to estimate tail probabilities, the historical data appears to be sufficiently different such that trimming is justified. While we caution against extrapolation of our results, they do give cause for consideration in other empirical analyses using historical yield data.

*Some key words:* yield data, trimming, premium rates

Institute for the Advanced Study of Food and Agricultural Policy  
Working Paper 2019.5

---

<sup>1</sup>Contact info: Alan P. Ker, Professor Department of Food, Agricultural and Resource Economics, OAC Research Chair in Agricultural Risk and Policy, and Director Institute for the Advanced Study of Food and Agricultural Policy, University of Guelph, [aker@uoguelph.ca](mailto:aker@uoguelph.ca). Yong Liu is a Ph.D. student, Department of Food, Agricultural and Resource Economics, University of Guelph, [yliu@uoguelph.ca](mailto:yliu@uoguelph.ca)

## Introduction

The Risk Management Agency (RMA) uses historical county-level yield data to set guarantees, estimate premium rates, and calculate indemnities for their area programs. Moreover, these county level rates are used in estimating farm level rates. Many in the literature have used this data to consider a variety of issues related to rating crop insurance contracts (Miranda and Glauber, 1997; Goodwin and Ker, 1998; Goodwin and Hungerford, 2015; Yvette Zhang, 2017). In many cases, and certainly with respect to RMA rating methodology, yield data are detrended and adjusted for possible heteroscedasticity and then assumed to be independent and identically distributed. Zhu, Goodwin, and Ghosh (2011) denote this the “two-stage method”. For most major crop-region combinations, annual NASS county yield data exist from the 1950s onwards and reflect very significant innovations in both seed and farm management technologies, which have likely moved mass all around the support of the yield distribution. This begs the question to what extent, if any, do yield losses in the 1950s and 1960s inform much about yield losses in 2019? That is, despite correcting for movements in the first two moments of the yield data generating process (dgp), does the identically distributed assumption hold? The issue of trimming is exacerbated by the need to estimate tail probabilities. Interestingly, and in direct contrast to their area programs, RMA uses yield data starting in 1991 to rate the newer area-based shallow loss products. This highlights an underlying, but empirically undocumented notion, that the yield dgp has significantly changed over the past half-century, *possibly* rendering the more historical data useless or even harmful in estimating premium rates.

Changes in seed and farm management technologies and their effects on yields have been well documented in the agronomy literature. Notable examples include the introduction of biotech seeds and precision farming. Many have shown that corn, soybean and wheat yields in the United States have more than doubled from 1950 to mid-1990s (Reilly and Fuglie, 1998; Fernandez-Cornejo, 2004; Duvick, 2005; Egli, 2008; Fernandez-Cornejo et al., 2014; Assefa et al., 2017; Egli, 2017). They tend to suggest that roughly half of the yield gain is attributed to genetic seed improvements while the other half is attributed to improved agronomic practices. Although the agronomy literature has focused on changes in average yields, some have also documented increasing volatility in yields (Naylor, Falcon, and Zavaleta, 1997; Kucharik and Ramankutty, 2005; Challinor et al., 2014; Leng, 2017). Conversely, there has been a relatively large body of work on the changes in yield volatility by agricultural economists, primarily driven by issues related to crop insurance (Harri et al., 2011; Claassen and Just, 2011; Yvette Zhang, 2017). With respect to changes in the higher moments ( $> 2$ ) of the yield distribution, there has been markedly less work. Zhu, Goodwin, and Ghosh (2011), using

NASS county level yield data for corn, soybean and cotton, find changes in higher moments through time. Tack, Harri, and Coble (2012), using county level cotton data in Arkansas, Mississippi, and Texas, found that the third moment, or skewness, was changing with time for Mississippi and Texas. Note that changes in higher moments indicate the common approach of correcting for changes in the first two moments is not sufficient for the identically distributed assumption in most of the literature as well as the RMA rating methodology for area-based programs. However, given the need to estimate tail probabilities, the above results (which are very region-crop specific) do not necessarily suggest historically trimming yield data will lead to more accurate premium rates; the loss function for each is over very different subsets of the density space.

The objective of this manuscript is to answer the question whether RMA should or should not trim their historical yield data in estimating crop insurance rates. Using county-level NASS yield data for corn, soybean, and winter wheat we first, for completeness, consider nonparametric distributional tests to assess if the adjusted yield data may result from different dgps. Second, we use an out-of-sample retain-cede rating game -- commonly employed in the literature -- to compare premium rates from the full versus historically trimmed yield data. Specifically, we focus our trimming at 1991 to reflect the distinction between RMA's area based insurance programs and their shallow loss programs.

The remainder of this manuscript proceeds as follows. The next section details the NASS yield data, the RMA detrending methodology, and the RMA heteroscedasticity treatment. The third section presents the statistical results from testing the identically distributed assumption. The fourth presents the economic results using an out-of-sample retain-cede rating game. The final section summarizes our findings.

## **NASS Yield Data, Detrending Methodology, and Heteroscedasticity Treatment**

NASS provides 49 categories of field crops including beans, cotton, corn, grain, hay, peanuts, mint, rice, soybeans, and wheat. The data generally date back to the 1950s. We use county level yield data for corn, soybean, and winter wheat for the period 1951–2017 (67 years). Our corn and soybean analysis focuses on states that account for the majority of national corn and soybean production. We removed counties with one or more missing yield observations as well as any state that does not have 25 or more counties. We also removed all states that reported more than ten percent of their acreage as irrigated in the 2012 Census of Agriculture. After doing so, we are left with seven states for corn: Illinois (IL), Indiana (IN), Iowa (IA), Minnesota (MN), Ohio (OH), South Dakota (SD), and Wisconsin (WI). These states accounted for 57.8 percent of harvested acreage and

61.8 percent of national production in 2017. All corn states except South Dakota met the inclusion criteria for soybean. These six states accounted for 50.5 and 53.9 percent of national harvested acreage and production, respectively, in 2017. For winter wheat, we considered the top 15 states that had less than ten percent of their acreage irrigated in 2012 Census of Agriculture, only two of which met the inclusion criteria: Kansas (KS) and Michigan (MI). These two states accounted for 29.2 percent and 28.9 percent of national harvested acreage and production, respectively, in 2017. In total, our data comprises 414 corn, 373 soybean, and 64 winter wheat counties.

Premium rates are estimated using a two-step process in which a trend is first estimated and then residuals are adjusted for possible heteroscedasticity. A two-step process is by far the most common in the literature as noted by Zhu, Goodwin, and Ghosh (2011). RMA estimates the temporal process of yields, denoted  $y_t = (y_1, \dots, y_T)$ , for each crop-county combination using a robust two-knot linear spline:

$$(1) \quad y_t = \theta_1 + \theta_2 t + \delta_1 d_1(t - k_1) + \delta_2 d_2(t - k_2) + \varepsilon_t$$

with  $d_1 = 1$  if  $t \geq k_1$  and  $d_2 = 1$  if  $t \geq k_2$  for knots  $k_1, k_2 \in (1 + \bar{k}, \dots, T - \bar{k})$  and  $k_2 - k_1 \geq \underline{k}$ . The  $\underline{k}, \bar{k} \geq 10$  are *a priori* imposed bounds which prevent the knots from locating too close together ( $\underline{k}$ ) or too close to either endpoint ( $\bar{k}$ ). Knot locations  $k_i$  are selected using a grid search (least-squares criterion). The model is run with zero, one, and two knots and then the number of knots used is selected using AIC.<sup>2</sup> Given the number of knots, two robustness procedures are performed; the spline is iterated to convergence with Huber weights and then twice through a bisquare function. Specifically, let  $\tilde{\varepsilon}_t$  be the estimated residuals from the robust spline with the chosen number of knots and  $\tilde{\eta}_t = \tilde{\varepsilon}_t / \sqrt{T^{-1} \sum \tilde{\varepsilon}_t^2}$ . The Huber function assigns weight one to observations if  $|\tilde{\eta}_t| < c$  and weight  $c/|\tilde{\eta}_t|$  otherwise with a default  $c = 1.345$ . Similarly, the bisquare function weights observations  $(1 - (\tilde{\eta}_t/c)^2)^2$  if  $|\tilde{\eta}_t| < c$  and zero otherwise with default  $c = 4.685$ .<sup>3</sup>

Denote the residuals from the above detrending process as  $\hat{\varepsilon}_t$  and the fitted values as  $\hat{g}(t) = \hat{y}_t$ . The heteroscedasticity adjustment via the Harri et al. (2011) estimates:

$$(2) \quad \ln(\hat{\varepsilon}_t^2) = \alpha + \gamma \ln(\hat{y}_t) + v_t.$$

---

<sup>2</sup>We do not impose the spatial and temporal priors on knots used by the RMA.

<sup>3</sup>Given any results are dependent on the choice of detrending method, we considered three alternative methodologies to ensure robustness of our results: (i) linear model estimated by  $L_2$ ; (ii) linear model estimated by  $L_1$ ; and (iii) nonparametric local lines using out-of-sample cross validation for the smoothing parameter. Our results are robust to any of these alternative detrending methodologies.

Note, constant and proportional variance in the underlying yield data correspond to  $\gamma = 0$  and  $\gamma = 2$ , respectively. Yields are adjusted based on a one-step ahead forecast ( $\hat{y}_{T+1}$ ) and the heteroscedasticity coefficient ( $\hat{\gamma}$ ):<sup>4</sup>

$$(3) \quad \hat{y}_t^* = \hat{y}_{T+1} + \hat{\varepsilon}_t \left( \frac{\hat{y}_{T+1}}{\hat{y}_t} \right)^{\frac{\hat{\gamma}}{2}}$$

The adjusted yields are then used to generate the empirical premium rate for period  $T + 1$

$$(4) \quad \pi_{T+1}^* = \frac{1}{T} \sum_{t=1}^T \max \{0, \lambda \hat{y}_{T+1}^* - \hat{y}_t^*\}$$

where  $\lambda$  is the coverage level such that  $\lambda \hat{y}_{T+1}^*$  is the yield guarantee.

### Testing the Identically Distributed Assumption

When testing for structural change, generally a Chow-type test is used; the sample is split into different sub-populations and residuals from regression equations within the sub-populations are combined with the residuals from a regression equation spanning the two samples to form a Wald type test statistic. The Bai-Perron test is a sup-type test of the Chow test in that it does not assume the breakpoint is known or the number of breakpoints. The Wilcoxon rank sum test is like a Chow test and primarily has power against changes in location. Overall, these tests have power only against changes in the conditional mean function (first moment) and thus, unsurprisingly, resulted in very little rejections on the adjusted yields across the crop-county combinations.<sup>5</sup> We are interested in structural changes in the higher moments of the dgp, beyond the conditional mean or variance. A common choice is Kolmogorov-Smirnov (KS) test which considers the maximum difference between two empirical distribution functions and thus has power against differences in all moments. Note, the test is nonparametric in that the test statistic is a function of the two empirical distribution functions. Also, the KS test has been shown to have relatively low power in comparison to Chow or Bai-Perron tests as these tests have an infinitely smaller space of alternatives (Wilcox, 1997). Moreover, the difference between two empirical distribution functions is most pronounced for differences in the location, followed by differences in scale, and then higher moments in sequential order. Recall we will only be testing differences in the higher moments and thus the power of the KS test is further weakened in that the two samples we are comparing have near identical first two moments. The KS test statistic is denoted  $D_{n,m}$  and defined as:

---

<sup>4</sup>RMA uses a two-step ahead forecast because of data availability/timing issues. We choose a one-step ahead forecast for our analysis simply to gain an additional degree of freedom given we are truncating an already short time series.

<sup>5</sup>Given we have corrected for the time-varying changes in the first moment in our detrending process, any rejections reflect the inappropriateness of the underlying functional form in the detrending process.

$$(5) \quad D_{n,m} = \sup_x |F_{1,n}(x) - F_{2,m}(x)|,$$

where  $F_{1,n}$  and  $F_{2,m}$  are the empirical distribution functions of the first and the second samples respectively. Specifically, the entire yield series is detrended and corrected for heteroscedasticity and then split pre and post 1991 corresponding to the different RMA rating procedures; recall the area-based programs use all the yield data whereas the newer supplemental loss programs only use yields from 1991 onwards. The null of the KS test is rejected at level  $\alpha$  when

$$(6) \quad D_{n,m} > c(\alpha) \sqrt{\frac{n+m}{nm}},$$

where  $c(\alpha)$  is calculated from the Kolmogorov distribution.

We also consider a second test forwarded by Li (1996) and further developed by Li, Maasoumi, and Racine (2009) (denoted LMR test). This LMR test is similar to a Cramer-von Mises test in that rather than based on the supremum difference, it is based on the integrated squared difference. Specifically, the LMR test smooths the data using kernel methods and calculates the integrated squared difference. Moreover, Li, Maasoumi, and Racine (2009) find power is increased if one bootstraps the null, using randomization methods, rather than use an asymptotic expansion for the distribution of the test statistic. Specifically, the entire yield series is detrended, corrected for heteroscedasticity, and then divided into two subsets pre and post 1991. The test statistic is defined as:

$$(7) \quad LMR = \int_{-\infty}^{\infty} \left( \hat{f}_1(x) - \hat{f}_2(x) \right)^2 dx,$$

where  $\hat{f}_1$  and  $\hat{f}_2$  are kernel estimates based on the two subsets of data. Li, Maasoumi, and Racine (2009) suggest using least squares cross validation for bandwidth selection. Moreover, the kernel estimates under the bootstrap samples to recover the distribution of  $LMR$  under the null use the same two bandwidths in each bootstrap. In our application, 500 bootstrap samples were used to construct the null.

The test results are presented in table 1. As expected, the LMR test has significantly more power than the KS test given the LMR test statistic is calculated over the entire support and a randomization method is used for creating the null. Second, the results reject that the data pre and post 1991 come from the same distribution in many of the crop-state combinations despite the small number of observations. In corn, 30% of the counties reject at the 5% significance level

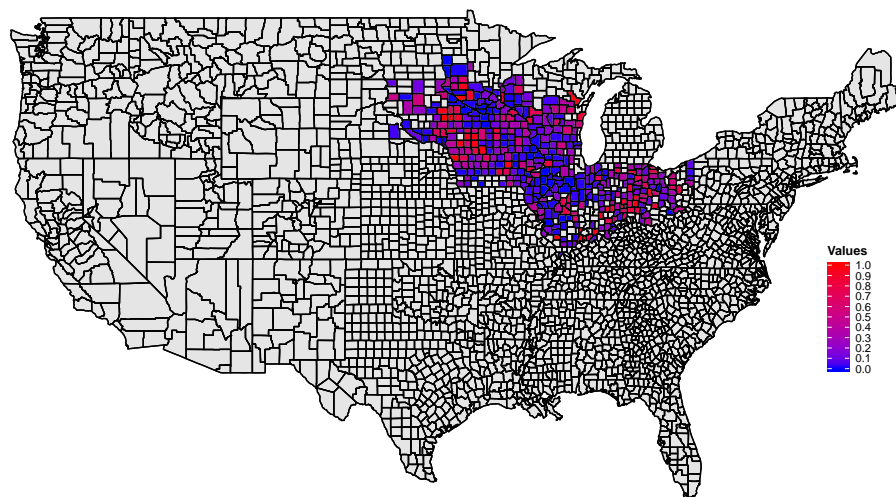
while 41% reject at the 10% significance level. The results are similar across the seven states. Soybean exhibits somewhat less significance as compared to corn; 14% of the counties reject at the 5% significance level while 24% reject at the 10% significance level. Winter wheat exhibits very little statistical significance, just 6% of the counties reject at the 5% significance level while 14% reject at the 10% significance level, barely above the size of the test. Interestingly, these results correspond to the level of research expenditures in the three crops over the past half-century. Corn has seen the most innovation while wheat has seen very little. The LMR test results ( $p$ -value) for corn, soybean, and winter wheat are graphically illustrated by county-crop combination in Figure 1. There does appear to be geographical clustering. For example, with respect corn the majority of the central counties -- the high production counties -- reject the null of identically distributed. With respect to soybean, the clustering rejections are in the more eastern counties. With respect to winter wheat, the clustering in Michigan is to the west while in Kansas it is to the southwest.

TABLE 1. Identically Distributed Test Results

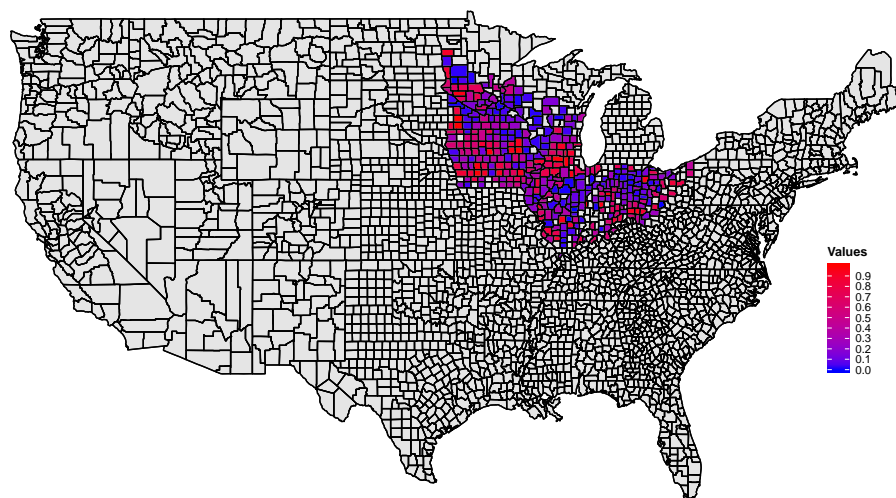
Crop-State	Number of Counties	Number of Rejections			
		KS Test		LMR Test	
		$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.05$	$\alpha = 0.10$
<i>Corn</i>					
Illinois	73	1	12	38	46
Indiana	60	2	4	13	20
Iowa	91	4	18	22	34
Minnesota	57	11	21	25	27
Ohio	58	1	5	5	12
Wisconsin	48	4	10	11	20
South Dakota	27	0	2	9	11
<i>Soybean</i>					
Illinois	82	1	3	14	18
Indiana	59	1	1	6	12
Iowa	93	0	0	5	15
Minnesota	55	3	6	11	17
Ohio	51	0	2	9	14
Wisconsin	33	0	0	9	14
<i>Winter Wheat</i>					
Kansas	35	0	0	1	3
Michigan	29	0	1	3	6

## Trimming and Estimating Crop Insurance Rates

Results from the previous section call into question the identically distributed assumption from a statistical perspective, but provide little information regarding economic importance. As previously

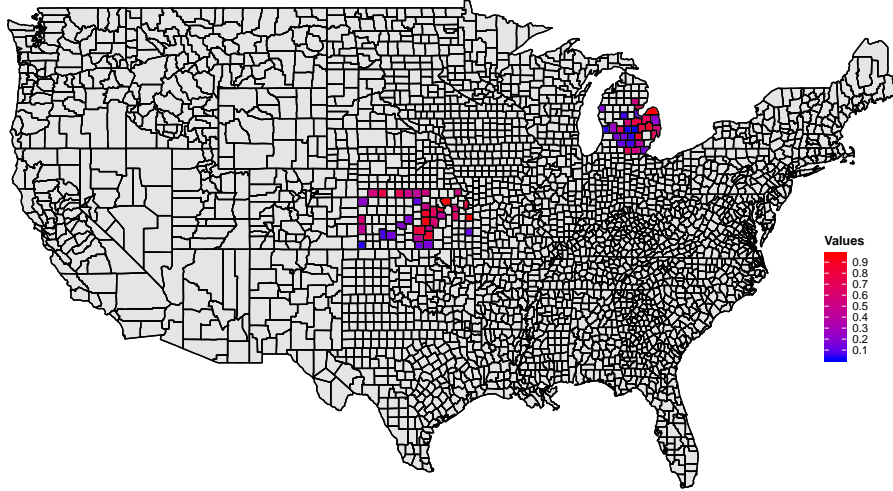


(A) Corn



(B) Soybean

FIGURE 1. Maps of  $p$ -values: LMR Test



(c) Winter Wheat

FIGURE 1. Maps of  $p$ -values: LMR Test, Continued

mentioned, the loss functions for estimating a distribution versus a premium rate are over different subsets of the density space. We consider the effect of trimming in rating crop insurance contracts by using an out-of-sample retain-cede rating game consistent with the literature. Specifically, the game allows two players using different methodologies to estimate premium rates and adversely select against one another. The game was first proposed by Ker and McGowan (2000) and has since been employed by Racine and Ker (2006), Harri et al. (2011), Annan et al. (2013), Tack and Ubilava (2015), Yvette Zhang (2017), and Shen, Odening, and Okhrin (2018) to justify alternative rating methodologies. The game was modified with an additional test of rating efficiency in Ker, Tolhurst, and Liu (2016). Park, Brorsen, and Harri (2018) utilized both tests in proposing an alternative rating methodology which exploits spatial closeness. The game was inspired by the retain-cede decision of private insurers in regards to the crop insurance contracts they sell. Some salient features of the U.S. crop insurance program are relevant to the game. First, RMA rather than private insurers set the rates for all policies. Second, the private insurer must sell all policies in a state that it operates in (even if it deems the policy to be under-priced). Third, the private insurer shares, asymmetrically, in the underwriting gains and losses of all policies it sells. Fourth, there

is a mechanism in which private insurers can significantly reduce their exposure on policies they deem unwanted.<sup>6</sup> Given these salient features, private insurers determine which policies to retain and which to cede. That is, private insurers retain policies they believe over-priced and expect an underwriting gain and cede policies they believe under-priced and expect an underwriting loss. As a result, private insurers necessarily develop their own rates in attempts to strategically adverse select against RMA and recover excess rents. Mimicking this allows one to hypothetically compare two sets of premium rates: one based on the full yield series and one based on the trimmed yield series. This is in contrast to the past literature that employs the retain-cede game to evaluate alternative rating methodologies using the same data.

Specifically, we assume RMA uses the full historical yield data from 1951-1997 on a county-crop basis to estimate the RMA premium rates for 1998. Conversely, the private insurer estimates their rates using a 25 year trimmed data set; that is yields from 1973-1997.<sup>7</sup> Both the RMA and the private insurer use the RMA rating methodology outlined above and as such the only difference in the two sets of rates is the result of trimming the historical yield data. Based on the two sets of rates, the private insurer identifies which contracts to retain and which to cede. The underwriting gains or losses for the set of retained and ceded contracts are calculated using the actual yields in 1998. This process is repeated for 20 years and the loss ratios (defined as the ratio of total underwriting losses to total premiums) for both the retained and ceded sets of contracts are calculated. We conduct the game for each crop-county combination at the 90% coverage level where the very large majority of area-based contracts are purchased.

As in the above cited literature, we undertake two hypothesis tests. The first tests whether the loss ratio from the retained contracts is less than the loss ratio from retaining contracts randomly (choosing which contracts to retain randomly is equivalent to the private insurer being indifferent between the two sets of competing rates). Like Li, Maasoumi, and Racine (2009), randomization methods are used to recover the  $p$ -value. Game 1 mimics the current reality of the US crop insurance program. However, the private insurer has an advantage because they react to the RMA premium rates. As such, whichever of the two competing rates the private insurer uses has an inherent competitive advantage in game 1. This advantage is nullified in game 2 by contrasting the changes in loss ratios under both sets of the competing rates (see Ker, Tolhurst, and Liu, 2016, for details). The number of contracts considered is the number of counties multiplied by 20 years; 8,280 contracts for corn, 7,460 contracts for soybean, and 1,280 contracts for winter wheat. The

---

<sup>6</sup>Specific details are outlined in the USDA-RMA Standard Reinsurance Agreement with approved private insurers.

<sup>7</sup>The 25 year trimming roughly corresponds to the 1991 cutoff. We also did trimming at 30 years and our results were qualitatively identical and quantitatively very similar.

results, which include percent retained by the private insurer, the government or ceded contracts loss ratio, the insurer or retained loss ratio,  $p$ -value of game 1, and  $p$ -value of game 2, are presented in Table 2.

TABLE 2. Out-Of-Sample Retain-Cede Rating Game

Crop-State	Number of Counties	Retained by Private (%)	Loss Ratio Government	Loss Ratio Private	Game 1 $p$ -value	Game 2 $p$ -value
<i>Corn</i>						
Illinois	73	66.1	0.581	0.519	0.0207	0.1316
Indiana	60	72.5	0.676	0.560	0.0013	0.0577
Iowa	91	59.1	0.411	0.270	0.0013	0.0577
Minnesota	57	75.9	0.228	0.195	0.0000	0.0577
Ohio	58	68.1	0.713	0.585	0.0013	0.2517
Wisconsin	48	77.0	0.543	0.378	0.0013	0.7483
South Dakota	27	81.3	0.810	0.545	0.0000	0.0002
<i>Soybean</i>						
Illinois	82	51.4	0.874	0.467	0.0059	0.2517
Indiana	59	70.5	1.019	0.538	0.0002	0.0207
Iowa	93	52.4	0.840	0.531	0.0059	0.1316
Minnesota	55	76.2	0.787	0.491	0.0013	0.0059
Ohio	51	76.1	0.734	0.640	0.0059	0.1316
Wisconsin	33	85.5	1.054	0.738	0.0013	0.2517
<i>Winter Wheat</i>						
Kansas	35	32.9	1.525	0.814	0.0059	0.0577
Michigan	29	49.0	0.346	0.349	0.2517	0.1316

Under a 25-year trimming decision rule, we find the private insurer's loss ratio is less than the RMA loss ratio for all 14 of the 15 state-crop combinations (Michigan wheat is higher only in the third decimal place). For corn, the private insurer loss ratio ranges from 66% to 89% of the RMA loss ratio. For soybean, the ratio ranges from 53% to 87%. Given only two states for winter wheat, the ratio is 53% for Kansas and 101% for Michigan. With respect to the first game,  $p$ -value 1 is significant in all state-crop combinations but Michigan wheat, suggesting that economically and statistically significant rents can be recovered by private insurers by trimming the yield data. With respect to the second game,  $p$ -value 2 is significant at the 10% level for seven of the 15 state-crop combinations, suggesting that trimming leads to statistically significantly more accurate premium rates. In no cases did **not** trimming lead to statistically significantly more accurate premium rates. Specifically, with respect to corn, four of the seven states were significant. With respect to soybean, two of the six states were significant. Finally, with respect to winter wheat, one of the two states

were significant. In summary, the out-of-sample retain-cede rating game provides strong evidence for trimming, consistent with the results from the last section.

## Conclusions

Historical yield data has been utilized in many empirical applications in literature, most notably in applications related to crop insurance. In general and with good cause, all available historical yield data used. Most applications account for time-varying lower moments but assume time-constant upper moments. In addition, RMA does the same in estimating the premium rates of their area products. However, there has been significant innovations in farm management and seed technologies in the past half century such that mass has likely moved all around the yield distribution. Not surprisingly, a few papers have found changes in the upper moments of the yield dgps thus questioning the standard approach of correcting the first two moments only. Our distributional test results find strong evidence of the inappropriateness of the identically distributed assumption for corn and soybean and markedly less so for winter wheat. These results are surprisingly strong in that the sample sizes are relatively small and thus the power of the tests against economically relevant alternatives are weakened. Our out-of-sample retain-cede rating game, which represents a different loss function over only a subset of the density space, is consistent with our distributional tests. That is, trimming does not increase estimation error in the rating process and is shown in approximately half of the crop-state combinations considered to decrease estimation error. This result is quite noteworthy suggesting that despite small sample sizes and the need to estimate tail probabilities, the historical data appears to be sufficiently different such that trimming is justified. Finally, our results across crops are fairly consistent with the research expenditures across crops in that we find the biggest efficiency gains from trimming in corn which has experienced the most innovation. While we caution against extrapolation, our results should give cause for consideration when using historical yield data in other applications as well.

## References

- Annan, F., J. Tack, A. Harri, and K. Coble. 2013. "Spatial Pattern of Yield Distributions: Implications for Crop Insurance." *American Journal of Agricultural Economics*, pp. 253–268.
- Assefa, Y., P. Prasad, P. Carter, M. Hinds, G. Bhalla, R. Schon, M. Jeschke, S. Paszkiewicz, and I.A. Ciampitti. 2017. "A New Insight into Corn Yield: Trends from 1987 through 2015." *Crop Science*, pp. 1–10.
- Challinor, A.J., J. Watson, D. Lobell, S. Howden, D. Smith, and N. Chhetri. 2014. "A meta-analysis of crop yield under climate change and adaptation." *Nature Climate Change* 4:287–291.
- Claassen, R., and R.E. Just. 2011. "Heterogeneity and distributional form of farm-level yields." *American Journal of Agricultural Economics* 93:144–160.
- Duvick, D.N. 2005. "The contribution of breeding to yield advances in maize (*Zea mays* L.)." *Advances in Agronomy* 86:83–145.
- Egli, D. 2008. "Comparison of corn and soybean yields in the United States: historical trends and future prospects." *Agronomy Journal* 100:1–79.
- Egli, D.B. 2017. *Seed biology and yield of grain crops*. CABI.
- Fernandez-Cornejo, J. 2004. "The seed industry in US agriculture: An exploration of data and information on crop seed markets, regulation, industry structure, and research and development." Working paper, United States Department of Agriculture, Economic Research Service.
- Fernandez-Cornejo, J., S. Wechsler, M. Livingston, and L. Mitchell. 2014. "Genetically engineered crops in the United States.", pp. .
- Goodwin, B.K., and A. Hungerford. 2015. "Copula-based models of systemic risk in US agriculture: implications for crop insurance and reinsurance contracts." *American Journal of Agricultural Economics* 97:879–896.
- Goodwin, B.K., and A.P. Ker. 1998. "Nonparametric estimation of crop yield distributions: implications for rating group-risk crop insurance contracts." *American Journal of Agricultural Economics* 80:139–153.
- Harri, A., K.H. Coble, A.P. Ker, and B.J. Goodwin. 2011. "Relaxing heteroscedasticity assumptions in area-yield crop insurance rating." *American Journal of Agricultural Economics*, pp. aar009.
- Ker, A.P., and P. McGowan. 2000. "Weather-based adverse selection and the US crop insurance program: The private insurance company perspective." *Journal of Agricultural and Resource Economics*, pp. 386–410.
- Ker, A.P., T.N. Tolhurst, and Y. Liu. 2016. "Bayesian estimation of possibly similar yield densities: implications for rating crop insurance contracts." *American Journal of Agricultural Economics*

- 98:360–382.
- Kucharik, C.J., and N. Ramankutty. 2005. “Trends and variability in US corn yields over the twentieth century.” *Earth Interactions* 9:1–29.
- Leng, G. 2017. “Recent changes in county-level corn yield variability in the United States from observations and crop models.” *Science of The Total Environment* 607:683–690.
- Li, Q. 1996. “Nonparametric testing of closeness between two unknown distribution functions.” *Econometric Reviews* 15:261–274.
- Li, Q., E. Maasoumi, and J.S. Racine. 2009. “A nonparametric test for equality of distributions with mixed categorical and continuous data.” *Journal of Econometrics* 148:186–200.
- Miranda, M.J., and J.W. Glauber. 1997. “Systemic risk, reinsurance, and the failure of crop insurance markets.” *American Journal of Agricultural Economics* 79:206–215.
- Naylor, R., W. Falcon, and E. Zavaleta. 1997. “Variability and growth in grain yields, 1950-94: does the record point to greater instability?” *Population and Development Review*, pp. 41–58.
- Park, E., B.W. Brorsen, and A. Harri. 2018. “Using Bayesian Kriging for Spatial Smoothing in Crop Insurance Rating.” *American Journal of Agricultural Economics* 101:330–351.
- Racine, J., and A.P. Ker. 2006. “Rating crop insurance policies with efficient nonparametric estimators that admit mixed data types.” *Journal of Agricultural and Resource Economics* 31(1):27–39.
- Reilly, J.M., and K.O. Fuglie. 1998. “Future yield growth in field crops: what evidence exists?” *Soil and Tillage Research* 47:275–290.
- Shen, Z., M. Odening, and O. Okhrin. 2018. “Adaptive local parametric estimation of crop yields: implications for crop insurance rate making.” *European Review of Agricultural Economics* 45:173–203.
- Tack, J., A. Harri, and K. Coble. 2012. “More than mean effects: modeling the effect of climate on the higher order moments of crop yields.” *American Journal of Agricultural Economics* 94:1037–1054.
- Tack, J.B., and D. Ubilava. 2015. “Climate and agricultural risk: measuring the effect of ENSO on US crop insurance.” *Agricultural Economics* 46:245–257.
- Wilcox, R.R. 1997. “Some practical reasons for reconsidering the Kolmogorov-Smirnov test.” *British Journal of Mathematical and Statistical Psychology* 50:9–20.
- Yvette Zhang, Y. 2017. “A Density-Ratio Model of Crop Yield Distributions.” *American Journal of Agricultural Economics* 99:1327–1343.
- Zhu, Y., B.K. Goodwin, and S.K. Ghosh. 2011. “Modeling yield risk under technological change: Dynamic yield distributions and the US crop insurance program.” *Journal of Agricultural and*

*Resource Economics*, pp. 192–210.