Received 7 May 2013
Revised 19 December 2013
Accepted 18 February 2014

Crop yield distributions: fit, efficiency, and performance

Bruce J. Sherrick
Department of Agricultural and Consumer Economics and TIAA-Center for Farmland Research, University of Illinois, Urbana, Illinois, USA
Christopher A. Lanoue
Nautilytics, LLC, Boston, Massachusetts, USA
Joshua Woodard
Dyson School of Applied Economics and Managesment, Cornell University, Ithaca, New York, USA, and
Gary D. Schnitkey and Nicholas D. Paulson
Department of Agricultural and Consumer Economics, University of Illinois, Urbana, Illinois, USA

Abstract

Purpose – The purpose of this paper is to contribute to the empirical evidence about crop yield distributions that are often used in practical models evaluating crop yield risk and insurance. Additionally, a simulation approach is used to compare the performance of alternative specifications when the underlying form is not known, to identify implications for the choice of parameterization of yield distributions in modeling contexts.

Design/methodology/approach – Using a unique high-quality farm-level corn yield data set, commonly used parametric, semi-parametric, and non-parametric distributions are examined against widely used in-sample goodness-of-fit (GOF) measures. Then, a simulation framework is used to assess the out-of-sample characteristics by using known distributions to generate samples that are assessed in an insurance valuation context under alternative specifications of the yield distribution.

Findings – Bias and efficiency trade-offs are identified for both in- and out-of-sample contexts, including a simple insurance rating application. Use of GOF measures in small samples can lead to inappropriate selection of candidate distributions that perform poorly in straightforward economic applications. The β distribution consistently overstates rates even when fitted to data generated from a β distribution, while the Weibull consistently understates rates; though small sample features slightly favor Weibull. The TCMN and kernel density estimators are least biased in-sample, but can perform very badly out-of-sample due to overfitting issues. The TCMN performs reasonably well across sample sizes and initial conditions.

Practical implications – Economic applications should consider the consequence of bias vs efficiency in the selection of characterizations of yield risk. Parsimonious specifications often outperform more complex characterizations of yield distributions in small sample settings, and in cases where more demanding uses of extreme-event probabilities are required.

Originality/value – The study helps provide guidance on the selection of distributions used to characterize yield risk and provides an extensive empirical demonstration of yield risk measures across a high-quality set of actual farm experiences. The out-of-sample examination provides evidence of the impact of sample size, underlying variability, and region of the probability measure used on the performance of candidate distributions.

Keywords Crop insurance, β distribution, Burr XII distribution, Mixture of normals, Weibull distribution, Yield risk

Paper type Research paper

Agricultural Finance Review
Vol. 74 No. 3, 2014
pp. 348-363

0002-1466
DOI 10.1108/AFR-05-2013-0021

Considerable debate remains about the most appropriate methods for representing uncertain yields in many contexts. Crop yield risk and related insurance evaluation applications provide an obvious and direct set of applications, with numerous models to evaluate the likelihood of yields, identify coverage differentials, select policies and coverage levels, and increasingly in ratings and development of private product pricing applications. Likewise, successful usage of other risk management tools by producers hinges on an accurate understanding of yield risk and its impact on revenue risk. Policy responses to farmers' risks, and evaluations of alternative policy designs, however, often begin with historically calibrated cases, and empirical experiences that are then fitted to a preferred candidate distribution and generalized to more complete depiction of the possible outcomes. Even applications such as the performance of so-called "Olympic averages" of yields and prices depend in a crucial way on the underlying uncertainty that generates the sequences of observations used in the construction. In each case, being able to appropriately represent yield risk has critically important implications.

When faced with the need to represent crop yield distributions, the most common response is to collect as representative a set of data as possible, and then use informed judgment and stylized requirements of the application to select a functional form to represent the data generating process and its resulting distributional form. There is an extensive literature on goodness-of-fit (GOF) measures and informational measures relating the data to the fitted distribution, and often simple in-sample measures suffice for selection. However, some applications of the fitted distributions are arguably more concerned with economic consequences associated with specific types of errors. For example, most insurance applications involve a region of insurance indemnification, such as the lower tail, and are not as concerned with other outcomes. In such a case, the focus is often on a bias in implied insurance rates generated from incomplete regions of alternative distributions. Sometimes the focus is on accuracy of the mean (e.g. aggregate yield forecasting) or on extreme tails (capital requirements, values-at-risk, and crop insurance applications) or variability around a level (for hedging and other insurance applications, option pricing, etc.), or on overall fitting performance (simulation of policy impacts, forecasting, etc.).

A slightly more subtle issue involves the treatment of first-stage fitting errors, and evaluation of unknown distributional candidates in small samples. The true data generating process is rarely known, and as a result, the flexibility of the candidate distributions can influence small sample fitting results, even if not as similar in large samples. For example, a β distribution with flexible limits may fit a small sample of data generated from a lognormal process better than a fitted lognormal. Out-of-sample techniques and other validation methods do reasonably well to help limit this type of impact and identify "best" overall performance given more flexible fitting objectives. However, it is less obvious how to compare the performance of alternative candidates in more complex economic applications such as insurance evaluation when the underlying "true" distribution is not known, and where the small sample performance in the economic application may not be consistent. To examine this latter issue, one could use known data generating processes and simply examine the performance of the candidate distributions across sampling features to help understand the likely performance of the candidates in various other conditions typical of cases where yield models are often used.

This study contributes to the empirical evidence about yield distributions and then uses two of the more plausible candidates to test their performance in small sample applications of the type often confronting those modeling yield distributions. With a high-quality farm-level corn yield data set from the Illinois Farm Bureau Farm

Management Association (FBFM) from 1972 to 2008, alternative in-sample GOF measures and associated insurance rates are examined across a set of flexible two- and three-parameter distributions, as well as a semi-parametric and non-parametric distribution[1]. Then, to investigate the efficiency and rate sampling properties of alternative distributions, simulations are conducted in which different sized samples are drawn from known parametric distributions, and are fitted to candidate distributions; the fitted distributions are then used to estimate insurance rates under each candidate distribution and compared to the known true values. This process is repeated in order to derive the rate sampling distribution for each candidate distribution. Given that the simulation begins with known starting "true" distributions, it thereby essentially allows for a type of "out-of-sample" evaluation in which we assess the underlying rate distribution directly – a related but more specific objective in insurance contexts. Not surprisingly, the results show that more flexible parametric and semi-parametric forms fit to data better (in terms of in-sample bias and precision/efficiency) due to flexibility from more parameters or less structure, but perform worse out-of-sample relative to more parsimonious forms, such as the Weibull, β, and Burr XII. An insurance pricing context is used to identify a meaningful economic loss function to compare performance across distributions under a common and important application.

Background

A long debated question among agricultural economists is the choice of the "best" distribution for modeling crop yields, yet, no single family of distributions or method of selection is widely accepted for any particular rating application. Among the parametric family of distributions, many studies have examined normality, typically rejecting its use because of the negative skewness and excess kurtosis generally observed in crop yield data (Day, 1965; Ramirez, 1997; Atwood *et al.*, 2003; Ramirez *et al.*, 2003). In contrast, Just and Weninger (1999) argued that the rejection of the normal distribution in many previous studies is potentially at fault due to methodological problems in typical yield distribution analyses. Other works (Nelson and Preckel, 1989; Nelson, 1990; Hennessy *et al.*, 1997) use a β distribution to model crop yields. The β distribution is arguably the most highly examined parametric form in empirical crop yield modeling literature. The β distribution is flexible enough to take on varying forms of skewness and kurtosis, and is being bounded between zero and a maximum value. It has also been argued to fit particularly well in lower tail regions and be desirable due to its simplicity in use. Still other works attempt to examine alternative parameterizations of crop yield distributions. Gallagher (1986) and Pope and Ziemer (1984) use the γ distribution, while Sherrick *et al.* (2004) find the Weibull distribution to be a good candidate for Illinois farm yields. Both the γ and Weibull distributions are potentially desirable due to their ability to exhibit negative skewness within a parsimonious specification. Claassen and Just (2011) examine a large number of relatively short series from RMA of insurance purchaser's data across extensive sets of crop/county sets to assess intra-county heterogeneity and find a version the reverse lognormal to do well in depicting yields. The inverse Gaussian distribution is commonly used to model first-passage time in insurance applications and as a result is also included as a flexible alternative.

Recently, non- and semi-parametric distributional forms have received more attention in the empirical literature (Goodwin and Ker, 1998; Turvey and Zhao, 1999; Ker and Goodwin, 2000; Norwood *et al.*, 2004; Wang and Zhang, 2002; Ker and Coble,

2003) because of their increased flexibility, particularly in modeling thicker distribution tails. This increased flexibility allows the distributional form to cover a broader set of skewness and kurtosis values, although forecasting may suffer from efficiency problems due to their tendency to over-fit sample data. Accumulating evidence exists which finds that in particular cases semi-parametric methods can outperform parametric methods in similar applications in terms of out-of-sample efficiency (see e.g. Norwood *et al.*, 2004).

Data

This study utilizes a high quality, extensive farm-level corn yield data set from the Illinois FBFM data set from 1972 to 2008. FBFM is a cooperative educational-service program that assists farmers with farm and tax management decisions by providing a comprehensive system of firm specific financial statements and standardized business reports. FBFM has roughly 2,500 commercial scale grain farms in its record keeping association, providing dependable and extensive yield histories on a set of common standards for measuring and reporting financial and production data. This data set is unique in the USA for its long panel of certified yield data, and is representative of a wide cross-section of Illinois commercial grain farms. Criteria for including farms samples included requirements to have: first, at least 20 years of yields; second, farm size >80 acres; and third, less than two consecutive years without production data. These screens result in 2,088 remaining corn farms, of which 768 farms have more than 30 years of data.

The yield data are detrended to control for increases in productivity through time. This study adopts a linear time trend as has been found to be a reasonable specification for corn yields in this region (Zanini, 2001; Ozaki *et al.*, 2008a). Four different data aggregation levels are examined and results are compared for data detrended at the state, district, county, and individual farms. The choice of aggregation when detrending is highly debated. Previous work (see e.g. Atwood *et al.*, 2003) suggests that farm level detrending may result in inefficient trend estimates, while the state level trends may exhibit unacceptably high degrees of bias. The overall implications of the study were not sensitive to the level of aggregation when detrending, except that farm-level results are considerably more variable, and deemed unreliable. Hereafter, results are reported for the district level detrending procedure with others available for more detailed comparisons. Summary statistics for the detrended yield data are provided in Table I. The production intensity in the regions identified by FBFM as East, Central, and Northwest generally correspond to the heart of the Corn Belt areas with the greatest corn production.

In-sample GOF analyses

GOF tests across the candidate distributions are considered in two related ways. First, each parametric distribution is compared and ranked using common GOF tests individually, and in composite form. Then, to examine the performance in a straightforward economic application, implied insurance rates are compared across the fitted parametric and non-parametric distributions and to the directly calculable empirical rates from the sample data.

The parametric distributions in this study (and number of parameters) include the β (four), γ (two), inverse Gaussian (two), normal (two), Burr XII (four), and Weibull (two). Three most common GOF tests performed are the Kolmogorov-Smirnov (K-S), Anderson-Darling (A-D), and χ^2-tests. Each of the GOF tests differs slightly in their focus.

Table I.
Sample characteristics from FBFM corn farms

	Sample summary statistics										
Region[a]	Trend[b]	Mean[c]	SD	CV	Skew.	Kurt.	Max	Min	Farm count	Yield count	% Sample acreage
NW	2.07	173.81	23.97	0.14	−0.83	1.35	252.34	44.27	395	10,959	19.50
NE	1.86	172.69	24.46	0.14	−0.51	0.36	265.10	46.03	176	4,662	8.30
West	2.12	175.58	28.17	0.16	−0.60	0.60	275.27	48.37	94	2,647	5.06
Central	2.03	181.82	26.23	0.14	−0.75	0.93	281.36	59.10	519	14,292	25.75
East	1.84	171.49	28.37	0.17	−0.86	0.81	253.68	57.58	296	7,849	13.90
WSW	1.87	177.59	24.45	0.14	−0.58	0.43	257.84	62.08	151	4,106	6.93
ESE	1.61	155.54	26.73	0.17	−0.37	−0.09	239.09	44.24	253	6,682	12.17
SW	1.71	136.81	26.02	0.19	−0.30	−0.11	228.57	29.00	119	3,426	3.41
SE	1.68	142.55	24.77	0.17	−0.33	−0.01	235.20	37.76	85	2,277	4.97
State total	1.87	169.55	25.95	0.20	−0.66	0.69	277.79	30.10	2,088	56,900	100

Notes: [a]FBFM farms grouped by Crop Reporting District. [b]Trend rates based on NASS Yield data. [c]Summary statistics for FBFM yield in bu./acre based on simple farm weighted values

The K-S test is a simple and well-known distribution-free measure of the congruence between the fitted and empirical distributions, essentially measuring the vertical distances between the empirical distribution function (EDF) and the fitted distribution. It is regarded as being most influenced by the center of the distribution where data are naturally more densely sampled. The A-D is calculated as:

$$A_n^2 = n \int_{-\infty}^{+\infty} [F_n(x) - F(x)]^2 \omega(x) f(x) dx,$$

where n is the number of observations, $f(x)$ the fitted density function, $F(x)$ the fitted cumulative distribution function, and $F_n(x) = N_\chi/n$, where N_x is the number of X_i's less than x (Stephens, 1974). The weight function used: $\omega(x) = n/F(x)[1–F(x)]$ results in more weight given to the departures in the tails of the distribution. Finally, the χ^2-test places equal weight on the center and tails of the distribution, and is calculated as:

$$\chi^2 = \sum_{i=1}^{K} \frac{(O_i - E_i)^2}{E_i},$$

where O_i is the number of observations in the ith bin, $E_i = N \times (F(Y_u)–F(Y_l))$ is the number of expected data points in the ith bin, $F(Y_u)$ is the CDF of the upper limit of the ith bin of the distribution, and $F(Y_l)$ is the CDF of the lower limit of the ith bin of the distribution (Snedecor and Cochran, 1989).

Each of the candidate distributions is fit to each of the individual farm-level yield series using maximum likelihood estimation. Then, the three GOF test statistics are used to rank the performance of each of the six distributions. The GOF rankings for the parametric distributional forms are summarized and tabulated in aggregate and within each of the nine Illinois National Agricultural Statistics Service Crop Reporting Districts (CRD) to group areas where production risks are likely similar. The results are also summarized and tabulated by various farm characteristics, including the mean and standard deviation of farm yield, the total number of yield observations, and farm size as measured by total acreage.

GOF results

The results for the parametric GOF tests are tabulated in Table II[2]. The first column of Table II identifies the GOF test followed by the distribution examined. The body of the table contains the weighted average of the test ranks constructed by the fraction of times each rank was achieved. For example, under the A-D test in the Northwest district, the Burr XII is ranked in order from one to six in 45, 37, 14, 3, 1, and 0 percent of the cases, resulting in a combined ranking score $0.45 \times 1 + 0.37 \times 2 + 0.14 \times 3 + 0.03 \times 4 + 0.01 \times 5 + 0.00 \times 6 = 1.78$. The table is shaded in blocks of each test across six distributions by district, and ordered within each test type panel by the overall rank.

Under the A-D test which is often favored when there is particular concern about fitting in the extreme regions or both tails, the Burr XII distribution has the lowest combined rank score in all districts. The Burr XII distribution also performs best across all 1,956 farms, with a combined ranking score of 1.88. The Weibull and β distributions are the next best fitting parametric forms, having similar combined ranking scores across all regions and overall as well. The result patterns are similar under the KS, and differ only slightly under the χ^2-test where the β has a slightly better average rank than the Weibull. An argument can be made for the highly flexible, four- parameter Burr XII to be the best fitting in-sample across all farms, though it is also expected to be more prone to over-fitting, especially as sample size declines. The normal distribution performs relatively better in districts where the skewness is lower (S and SE districts), compared to districts where the average skewness is highest (East district), but is still not likely to be viewed as the "best" under typical metrics or applications. The γ and inverse Gaussian distributions are in fifth and sixth place across every district and GOF test conducted.

		District									
Test	Distribution	NW	NE	West	Central	East	WSW	ESE	SW	SE	Total
K-S	Burr XII	1.98	2.18	2.11	1.92	1.92	2.19	2.29	2.35	2.38	2.07
	Weibull	2.49	2.51	2.61	2.49	2.35	2.64	2.67	2.68	3.05	2.54
	β	2.53	2.51	2.43	2.60	2.31	2.69	2.85	2.69	2.41	2.56
	Normal	3.47	3.29	3.44	3.44	3.61	3.21	3.02	3.04	3.01	3.35
	γ	4.53	4.51	4.40	4.55	4.80	4.26	4.17	4.24	4.14	4.47
	InvGauss	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00
A-D	Burr XII	1.78	1.90	1.92	1.72	1.89	1.96	2.10	2.04	2.17	1.88
	Weibull	2.45	2.66	2.64	2.47	2.34	2.59	2.84	2.77	2.88	2.56
	β	2.67	2.50	2.42	2.68	2.54	2.93	2.90	2.71	2.28	2.66
	Normal	3.50	3.40	3.45	3.50	3.61	3.18	2.95	3.05	3.20	3.38
	γ	4.67	4.64	4.63	4.71	4.75	4.47	4.33	4.48	4.53	4.61
	InvGauss	5.92	5.90	5.94	5.92	5.88	5.87	5.89	5.95	5.95	5.91
χ^2	Burr XI I	2.49	2.37	2.81	2.51	2.43	2.66	2.54	2.63	2.72	2.53
	β	2.69	2.62	2.48	2.50	2.39	2.78	2.82	2.91	2.45	2.61
	Weibull	2.67	2.60	2.92	2.73	2.76	2.94	2.70	2.71	2.88	2.74
	Normal	3.26	3.40	3.13	3.25	3.31	3.09	3.19	3.09	3.20	3.24
	γ	3.88	4.01	3.67	4.02	4.11	3.53	3.75	3.66	3.75	3.89
	InvGauss	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00
Weighted	Burr XII	1.85	1.97	2.06	1.80	1.81	2.06	2.22	2.30	2.45	1.96
	Weibull	2.45	2.51	2.66	2.45	2.39	2.56	2.67	2.66	2.96	2.52
	β	2.72	2.61	2.50	2.66	2.48	2.82	2.87	2.77	2.36	2.67
	Normal	3.46	3.36	3.44	3.47	3.58	3.19	3.03	3.01	2.93	3.35
	γ	4.52	4.55	4.34	4.61	4.75	4.38	4.22	4.27	4.30	4.50
	InvGauss	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00	6.00

Table II.
Goodness-of-fit rankings: Illinois districts, FBFM corn farms

Empirical rating analysis
It could be the case that a particular distribution fits best in-sample, but that the region of the distribution of greatest concern in an economic application results in different rankings. For example, it could be the case that a very flexible distribution could fit in the tails better than a distribution that accurately characterizes the mean, but has less ability to relate mass in the tails. This argument has been advanced in favor of the β distribution in certain insurance contexts as its fit in the lower tail may be suitable, and its fit in the upper regions inconsequential. To compare the results in a context with economic importance, the fitted distributions are used to construct simple insurance rates (conditional expected values in specified regions) and compared to the empirical rates across the sample. To first provide a visual understanding of the process, four representative corn farms are taken from Dekalb, Marion, and McLean counties, as typical representations of the Northern, Southern, and Central parts of Illinois, respectively. The four highest ranking distributions from the previous section are investigated (Burr XII, Weibull, β, and normal), along with one semi-parametric (two-component mixture-of-normals (TCMN)), and one non-parametric distribution (Gaussian kernel density estimator) to extend the analysis to popular flexible forms.

The Gaussian kernel density estimator places a kernel (or bump) at each yield realization, and then the sum of the densities of the kernels forms the shape of the non-parametric curve. The Gaussian kernel density estimator probability density function (PDF) is:

$$f_h(x) = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{x - xi}{h}\right)$$

where $K(x - x_i/h) = 1/\sqrt{2\pi}e^{-(x-x_i)^2/2h^2}$ is the Gaussian kernel function, and Silverman's rule of thumb is used to set the value for the smoothing parameter, h (Silverman, 1986).

To provide a visual illustration, empirical CDF plots of four Illinois farms are provided in Figure 1. In examining the two taken from McLean, note that both have a notable irregular step in the left portion of the EDF. The kernel density estimator and the mixture-of-normals distribution are more flexible and tend to associate more mass near this "spike," and also appear to fit well throughout the center and right tail of the distribution. With respect to the parametric distributions, the Weibull and Burr XII distributions virtually overlap and tend to reflect the left tail better than the normal and β distributions. In general, when a relatively extreme value appears in the left tail, the Weibull and Burr XII distributions exhibit more flexibility to associate mass with that point than either the normal or β distributions. The normal distribution performs the worst, most likely due to its inability to capture skewness, but when the data are almost symmetric, such as in DeKalb-316, the normal performs comparably to the Burr XII and Weibull distributions. However, the flexibility that is exhibited by the mixture-of-normals and kernel density estimator can result in over-fitting in sample and thus may limit their out-of-sample forecasting ability. On the other hand, the Weibull and Burr XII distributions are better able to reflect out-of-sample variation, and thus may be more desirable in out-of-sample contexts such as that of crop yield forecasting, but might do a worse job in applications that have more significant consequences of

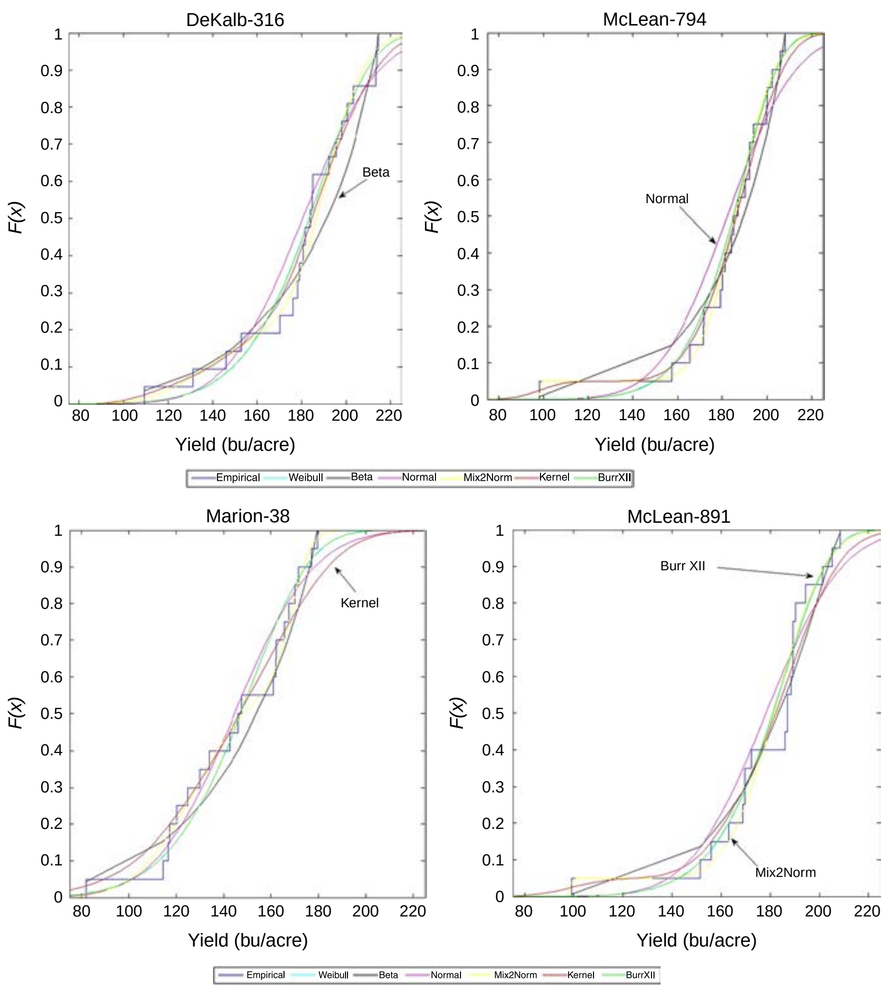

Figure 1.
Corn yield distributions

mischaracterizing extreme events. The context in which the fitted distributions are used thus likely matters to the selection of the "best" distribution, based on the economic costs of mis-calibration to different regions of the probability measure.

To illustrate the economic implications of the choice of distributional form, an insurance application is used in which bias and efficiency can be examined at various probability partitions. Specifically, all distributions are first fit to each farm's data; each are then used to calculate farm-level yield insurance rate estimates at coverage levels of 85, 75, and 65 percent. The expected yield insurance rate is expressed for each farm and candidate distribution as:

$$InsRate_{i,d} = \int_{0}^{k_i} Max\{0, k_i - Y_i\} f_{d,i}(Y_i|\theta_{i,d}) \mathrm{d}Y_i$$

where $k_i = Cover \times E(Y_i)$, $E(Y_i)$ and Y_i are the expected and realized yields for farm i, *Cover* the coverage level, and $f_{d,\,i}(Y_i \mid \theta_i)$ the probability density function for distribution d defined by parameter set $\theta_{i,\,d}$. The rate estimates are then compared to the average empirical rates from the underlying farm-level yields to evaluate in-sample bias and rate efficiency across all farms. The bias for each distribution is calculated as the difference between the implied insurance rate and the empirical or "burn" rate. A measure of in-sample efficiency is constructed using the root mean squared error (RMSE) across farms: $RMSE_d = \sqrt{\frac{1}{m}\sum_i \left(InsRate_{i,\,d} - BurnRate_i\right)^2}$ where m is the number of farms in the sample being evaluated. Taken together, these provide a useful empirical indication of important performance issues in an insurance rating application.

Rating analysis results

The results for the in-sample rating analysis are provided in Table III. Table III is organized with sections for average, bias, and efficiency tabulated for each of the distributions examined, and then the body of the table provides results for three possible partitions of the fitted probability measures corresponding to 85, 75, and 65 percent coverage insurance. Individual results are tabulated for farms in each of the nine CRDs and the overall state results. The bias and efficiency (RMSE) values are reported as percentages of their associated empirical burn rate. The shaded areas within the bias and efficiency statistics indicate the distributional form that contains the lowest absolute value under each district and coverage level. For example, in the case of the NW district at the 85 percent coverage level, the absolute value of the bias for the β distribution is highlighted as it is closest to zero when compared to the other distributions. The bolded values within the bias section indicate distributional rate estimates that are larger than the empirical/burn rates (positive rate bias). For example, all bias values under the kernel density estimator are bolded as they are all greater than the empirical rate.

The average of the empirical rates ranges from a high of 3.84 bushels/acre in the East district, to a low of 2.34 bushels/acre in the West Southwest district at an 85 percent coverage level. The averages of the empirical rates for each FBFM farm across all districts are 3.00 bushels/acre, 1.23 bushels/acre, and 0.43 bushels/acre for the three coverage levels. Focussing on the bias test statistic, the β distribution performs best in 15 of the 27 coverage level-district pairs, and also at a 75 and 65 percent coverage level for the state as a whole. The TCMN and Weibull fit the next best in terms of bias with six each of the 27 combinations, respectively. The normal distribution performs least well in all districts. Interestingly, the rate estimates from both the β and kernel density estimator distributions consistently overstate the empirical rate. The rate estimates from the Weibull, Burr XII, and TCMN distributions, on the other hand, consistently understate the empirical rates 86.7, 96.7, and 100.0 percent of the time, respectively. A surprising characteristic of the differences between the empirical rates and the fitted rates from the distributional forms is the fact that the kernel density estimator actually performs the worst in terms of bias relative to the empirical burn rate. In contrast and unsurprisingly, the kernel density estimator is generally the most efficient distribution in-sample as measured by the RMSE statistic. The kernel density estimator is the most efficient in 16 of the 27 coverage level/district combinations, while the TCMN is the most efficient in the remaining 11.

Table III. Insurance rate comparisons across fitted distributions and empirical

Statistic	District Distribution	NW 85%	NW 75%	NW 65%	NE 85%	NE 75%	NE 65%	West 85%	West 75%	West 65%	Central 85%	Central 75%	Central 65%	East 85%	East 75%	East 65%	WSW 85%	WSW 75%	WSW 65%	ESE 85%	ESE 75%	ESE 65%	SW 85%	SW 75%	SW 65%	SE 85%	SE 75%	SE 65%	Total 85%	Total 75%	Total 65%
Average[a]	Empirical[b]	2.44	1.01	0.38	2.66	1.05	0.36	3.39	1.50	0.51	2.84	1.13	0.37	3.84	1.73	0.68	2.34	0.81	0.21	3.33	1.27	0.42	3.70	1.56	0.56	3.14	1.23	0.42	3.00	1.23	0.43
	Weibull	1.97	0.63	0.18	2.30	0.78	0.24	2.98	1.09	0.35	2.30	0.74	0.21	2.75	0.99	0.31	2.13	0.69	0.20	3.21	1.29	0.46	3.57	1.55	0.60	3.02	1.22	0.44	2.53	0.90	0.29
	Normal	1.84	0.48	0.11	2.12	0.62	0.16	2.80	0.86	0.23	2.13	0.55	0.12	2.91	0.94	0.26	1.85	0.47	0.10	2.97	1.05	0.33	3.36	1.33	0.47	2.82	1.02	0.32	2.39	0.73	0.20
	Burr XII	1.86	0.58	0.16	2.17	0.73	0.22	2.82	1.00	0.31	2.17	0.68	0.19	2.70	0.96	0.30	2.02	0.64	0.18	3.06	1.21	0.42	3.39	1.45	0.55	2.88	1.14	0.40	2.41	0.84	0.26
	β	2.60	0.99	0.36	2.83	1.14	0.44	3.52	1.42	0.51	2.93	1.12	0.40	3.87	1.72	0.71	2.67	1.00	0.35	3.77	1.70	0.71	3.88	1.78	0.75	3.37	1.48	0.60	3.18	1.31	0.50
	TCMN	2.11	0.73	0.22	2.56	0.95	0.32	3.20	1.25	0.39	2.67	0.96	0.29	3.71	1.58	0.57	2.13	0.67	0.16	3.26	1.19	0.38	3.62	1.47	0.51	3.07	1.12	0.35	2.83	1.06	0.34
	Kernel	2.82	1.18	0.46	3.10	1.27	0.47	3.94	1.74	0.67	3.24	1.32	0.47	4.28	2.00	0.83	2.77	1.01	0.31	3.98	1.68	0.63	4.33	1.99	0.81	3.75	1.61	0.62	3.46	1.48	0.56
Bias (%)[c]	Weibull	−19%	−38%	−54%	−13%	−26%	−34%	−12%	−27%	−31%	−19%	−34%	−44%	−28.4%	−43.0%	−54.5%	−9%	−15%	−8%	−3%	*1%*	*10%*	−4%	0%	*7%*	−4%	−1%	*4%*	−15.7%	−26.5%	−33.3%
	Normal	−25%	−53%	−71%	−20%	−41%	−54%	−17%	−42%	−55%	−25%	−51%	−67%	−24.4%	−45.9%	−61.7%	−21%	−42%	−51%	−11%	−17%	−20%	−9%	−15%	−17%	−10%	−18%	−23%	−20.3%	−40.6%	−53.5%
	Burr XII	−24%	−43%	−58%	−18%	−31%	−39%	−17%	−33%	−38%	−24%	−40%	−50%	−29.7%	−44.5%	−55.9%	−13%	−22%	−17%	−8%	−5%	*1%*	−8%	−7%	−3%	−8%	−8%	−5%	−19.7%	−31.4%	−39.0%
	β	*6%*	−2%	−7%	*6%*	*8%*	*22%*	*4%*	−5%	*2%*	*3%*	−1%	*7%*	*0.8%*	−0.7%	*3.8%*	*14%*	*23%*	*63%*	*13%*	*33%*	*71%*	*5%*	*14%*	*33%*	*8%*	*20%*	*41%*	*5.8%*	*6.6%*	*17.0%*
	TCMN	−14%	−28%	−43%	−4%	−10%	−12%	−5%	−17%	−24%	−6%	−15%	−21%	−3.6%	−8.5%	−15.4%	−9%	−17%	−26%	−2%	−7%	−10%	−2%	−5%	−9%	−2%	−9%	−17%	−5.8%	−13.3%	−20.9%
	Kernel	*15%*	*16%*	*20%*	*16%*	*20%*	*32%*	*16%*	*17%*	*32%*	*14%*	*17%*	*26%*	*11.5%*	*15.6%*	*22.4%*	*18%*	*25%*	*47%*	*20%*	*32%*	*50%*	*17%*	*28%*	*44%*	*20%*	*31%*	*47%*	*15.4%*	*20.4%*	*30.6%*
RMSE (%)[d]	Weibull	42%	81%	137%	37%	73%	127%	32%	57%	92%	38%	71%	125%	39.6%	67.9%	105.5%	38%	80%	180%	27%	53%	101%	21%	43%	87%	26%	51%	104%	36.0%	68.0%	118.7%
	Normal	32%	77%	135%	28%	67%	121%	24%	56%	93%	31%	70%	124%	28.8%	60.4%	100.2%	30%	74%	169%	18%	42%	89%	15%	35%	74%	18%	43%	94%	27.3%	62.6%	113.7%
	Burr XII	40%	80%	137%	35%	72%	126%	30%	56%	92%	37%	70%	123%	39.4%	67.7%	105.5%	36%	78%	176%	24%	49%	97%	19%	41%	83%	23%	50%	103%	34.6%	66.9%	117.5%
	β	30%	63%	119%	27%	57%	122%	19%	32%	63%	25%	55%	111%	20.3%	41.6%	77.0%	35%	74%	189%	28%	67%	148%	18%	42%	91%	22%	52%	110%	25.3%	54.1%	109.8%
	TCMN	36%	64%	110%	14%	27%	49%	18%	35%	65%	21%	39%	73%	14.7%	25.6%	46.9%	28%	42%	155%	11%	23%	50%	9%	19%	43%	12%	26%	58%	20.5%	36.7%	71.3%
	Kernel	18%	24%	35%	19%	27%	50%	18%	21%	46%	16%	22%	38%	13.7%	20.0%	31.3%	21%	32%	72%	21%	38%	72%	19%	32%	56%	22%	37%	64%	18.1%	27.1%	47.1%

Notes: The highlighted values represent the distributional form with the lowest absolute bias or efficiency value by district and coverage level. The italic values represent bias values for which the fitted rate is greater than the empirical /burn rate. [a]Expected value of fitted distributional rates over all farms by district and coverage level. [b]Empirical /burn rate [max(0,Yield Guarantee-Estimated Rate)] for all farms by district and coverage level. [c]Bias (percent) is expected value of [empirical /burn rate-fitted distributional rate] for all farms divided by empirical /burn rate by district and coverage level. [d]RMSE (percent) is the expected value of the RMSE for all farms divided by empirical /burn rate by district and coverage level

The RMSE values of the Weibull and Burr XII distributions are highest among the distributions tested with percentages relative to the empirical rate of 33.4 and 31.6 percent, respectively, at an 85 percent coverage level. For all the distributions in this section, the average RMSE across every district increases as the coverage level decreases. For instance, the RMSE values for the β are 25.0 percent at the 85 percent coverage level, 53.6 percent at the 75 percent coverage level, and 114.5 percent at the 65 percent coverage level.

Overall, the non- and semi-parametric distributional forms are found to perform well in terms of in-sample efficiency. Nevertheless, in both the GOF examination and the empirical insurance application, the β, Burr XII, and Weibull distributions substantially outperform the other parametric distributional forms (normal, γ, inverse Gaussian).

Simulation evaluation of distributional performance

Next, specific types of out-of-sample fitting errors are examined. Starting with known data generating processes, samples of differing sizes are drawn, and then four candidate distributions – the Weibull, β, empirical, and TCMN – are fit to the samples and evaluated using a standard insurance rating framework where the implied insurance rates are compared to the underlying true rate. This process is repeated 5,000 times for each candidate generating distribution, and the resulting rates implied by the fitted distributions then compared to the known true rate. This exercise allows an out-of-sample assessment of the bias and efficiency of these distributions in terms of rates generated relative to known underlying distributions as an approach to demonstrate the likely performance of the distributions in other applications where the underlying true data generating process cannot be known with certainty. Importantly, if a particular form generally results in "best estimates" of rates regardless of the actual data generating process, then its use may be justified regardless of the underlying data generating process. On the other hand, if a given distribution results in large mistakes, even if it is correctly identified as the DGP, then it may not be the preferred candidate for modeling risk in the economically important application.

Both Weibull and β are used as the data generating distributions over various sample sizes (ten, 15, 20, and 30) and over ranges of yield statistics consistent with the samples from the Illinois farms in the Central region. Specifically, the Weibull and β distributions are constructed to represent pseudo-farms defined with mean yields of 160 and 180 bushels/acre, and three standard deviation levels of 20, 30, and 40 bushels/acre to provide a meaningful context to evaluate performance. The Weibull distribution parameters are established for each case using method-of-moments approximations of Garcia (1981). For the β distribution, the starting lower limit is bound at zero and the upper limit is set equal to three standard deviations above the mean[3]. Next, a method-of-moments approach is used to approximate the shape and scale parameters of the Weibull distribution. After drawing samples from the defined distributions, the parameters of the fitted Weibull, fitted β, and TCMN distributions are estimated using maximum likelihood estimation, while the fitted insurance rates from the empirical distribution are calculated in the same manner as the empirical burn rates in the previous section.

Simulation results

Summary results for a case viewed to be typical in much of Illinois with an expected corn yield of 160 and yield standard deviation of 20 bushels/acre are reported in Table IV[4]. The table is organized with four stacked vertical panels containing each of

		Data generating process					
		Weibull Coverage level			β Coverage level		
		85%	75%	65%	85%	75%	65%
Statistic	"True"	1.5873	0.4309	0.0952	1.4242	0.3021	0.0436
n = *10*	Distribution						
Average	Empirical	1.60	0.44	0.10	1.42	0.30	0.04
	β	2.80	1.39	0.70	2.81	1.39	0.70
	Weibull	1.61	0.49	0.13	1.67	0.51	0.14
	TCMN	1.30	0.32	0.07	1.19	0.23	0.04
Bias (%)	Empirical	*0.7%*	*2.4%*	*2.4%*	−0.2%	−0.7%	*1.6%*
	β	*77%*	*222%*	*632%*	*97%*	*359%*	*1503%*
	Weibull	*1%*	*14%*	*39%*	*18%*	*69%*	*212%*
	TCMN	−18%	−25%	−24%	−16%	−25%	−11%
RMSE (%)	Empirical	121%	224%	438%	115%	233%	592%
	β	215%	562%	1707%	242%	815%	3741%
	Weibull	90%	139%	228%	98%	198%	498%
	TCMN	117%	185%	307%	111%	181%	383%
n = *15*							
Average	Empirical	1.58	0.43	0.10	1.42	0.30	0.04
	β	2.48	1.07	0.47	2.36	0.98	0.41
	Weibull	1.60	0.47	0.12	1.69	0.50	0.13
	TCMN	1.38	0.33	0.07	1.28	0.25	0.04
Bias (%)	Empirical	−0.4%	*0.9%*	*5.8%*	*0.0%*	−0.8%	−4.7%
	β	*56%*	*149%*	*393%*	*66%*	*224%*	*840%*
	Weibull	*1%*	*9%*	*26%*	*19%*	*67%*	*195%*
	TCMN	−13%	−24%	−25%	−10%	−18%	−5%
RMSE (%)	Empirical	97%	180%	368%	92%	184%	430%
	β	155%	378%	1051%	158%	488%	1998%
	Weibull	72%	105%	160%	80%	161%	390%
	TCMN	96%	151%	252%	91%	152%	291%
n = *20*							
Average	Empirical	1.58	0.42	0.09	1.43	0.30	0.05
	β	2.12	0.80	0.30	2.05	0.74	0.27
	Weibull	1.60	0.46	0.11	1.70	0.50	0.13
	TCMN	1.44	0.34	0.07	1.33	0.27	0.04
Bias (%)	Empirical	−0.7%	−2.1%	−5.3%	*0.5%*	*0.4%*	*7.2%*
	β	*33%*	*86%*	*214%*	*44%*	*146%*	*513%*
	Weibull	*1%*	*7%*	*20%*	*19%*	*65%*	*187%*
	TCMN	−9%	−21%	−24%	−6%	−12%	*1%*
RMSE (%)	Empirical	83%	150%	293%	83%	167%	414%
	β	112%	249%	624%	119%	342%	1304%
	Weibull	62%	89%	131%	72%	144%	347%
	TCMN	84%	132%	222%	82%	142%	286%
n = *30*							
Average	Empirical	1.59	0.44	0.10	1.42	0.30	0.04
	β	1.87	0.60	0.18	1.76	0.53	0.15
	Weibull	1.59	0.45	0.11	1.70	0.49	0.12
	TCMN	1.50	0.37	0.08	1.36	0.27	0.04
Bias (%)	Empirical	*0.5%*	*1.5%*	*4.1%*	−0.1%	−0.3%	−2.4%
	β	*18%*	*40%*	*93%*	*23%*	*75%*	*239%*
	Weibull	0%	*4%*	*12%*	*19%*	*63%*	*175%*
	TCMN	−5%	−13%	−16%	−5%	−10%	*1%*

(*continued*)

Table IV. Bias and Efficiency Performance of known Distributions

Table IV.

		Data generating process					
		Weibull			β		
		Coverage level			Coverage level		
		85%	75%	65%	85%	75%	65%
RMSE (%)	Empirical	67%	125%	252%	66%	133%	320%
	β	80%	159%	360%	78%	195%	641%
	Weibull	50%	71%	101%	59%	119%	281%
	TCMN	69%	115%	188%	66%	117%	237%

Notes: Shaded entries have lowest absolute bias of efficiency value by district and coverage level. The italic values represent bias values for which the fitted rate is greater than the empirical/burn rate. All values in table from distributions with $\mu = 160$; $\sigma = 20$

the four sample sizes examined. The left hand set of results correspond to an underlying Weibull distribution used to generate the samples and the right hand side with an underlying β distribution. Coverage levels of 85, 75, and 65 percent are tabulated under each case. For each sample size, the average rate (average) in bushels/acre, the bias, and the RMSE across all coverage levels for each of the four distributions are tabulated. The bias and efficiency (RMSE) are presented in terms of percentages relative to the known theoretical rate.

The true rates from the Weibull distributional form are slightly larger than the true rates resulting from the β at all coverage levels given common sample statistics. However, the true rates from the distributional forms do converge as the standard deviation is increased.

Remarkably, the fitted β consistently overstates rates, and by an even greater amount when the "true" data generating process starts as a β than when the data are actually sampled from a Weibull. The bias is rather large, ranging from 17.3 to 65.6 percent in small samples ($n = 10$ or 15) at the 85 percent coverage level, with an increasing bias as the coverage level is reduced. At low coverage levels, the bias is typically several hundred percent, though the actual numeric value is often small. These findings hold across sample size. In contrast, the Weibull and TCMN perform quite well in terms of bias regardless of the mean – a robust result to changes in the standard deviation, sample size drawn, and across both true underlying distributions. An implication is that if the underlying data generating process is not known, the β may not be a good candidate relative to the Weibull for parameterizing the ratings distribution, even if its in-sample GOF performance is superior.

The bias of the empirical distribution is lowest among the distributions tested across all categories, except for six in which either the Weibull or TCMN is best; though the empirical rarely displays the lowest RMSE. At an 85 percent coverage level, the rate bias of the Weibull and TCMN distributions is typically < 3 percent, while the rate bias of the empirical distribution is < 1 percent when the standard deviation of the underlying is 40 bushels/acre and $n = 15$; and at a 65 percent coverage level the bias is always < 10 percent for the Weibull and *mixture-of*-normals and < 1.5 percent for the empirical. In general, the bias decreases as the sample size increases, and increases at lower coverage levels.

In terms of rate efficiency, the RMSE of the fitted β rates is, on an average, 34.8 percent greater than the RMSE of the Weibull fitted rates at a coverage level of 85 percent; 114.0 percent greater at a 75 percent coverage level; and 424.2 percent greater at 65 percent. In comparison to the TCMN, the Weibull typically performs better in terms of out-of-sample fitting, although in a few cases they perform similarly. The efficiency of the fitted distribution appears to be lower at lower coverage levels, and naturally also increases as the sample size increases. At a sample size of 30, the fitted rates of the Weibull are more efficient than the fitted rates of the TCMN by 7.5, 13.4, and 18.1 percent at coverage levels of 85, 75, and 65 percent, respectively.

Overall, the parsimonious two-parameter Weibull tends to consistently outperform the TCMN, empirical, and β in terms of out-of-sample efficiency, and is also comparable in performance to the mixture-of-normals in terms of bias, but is slightly more biased than the empirical distribution in fitting to samples used in a context to establish insurance rates.

Conclusions

Issues surrounding the choice of distribution for modeling yields, as well as the manner in which one should go about evaluating and comparing them, remains a contentious issue. This study analyzes the issue in a direct empirical manner using standard GOF and crop insurance rating perspectives with a comprehensive data set from the Illinois FBFM covering corn yields from 1972 to 2008. With three standard GOF tests, this study examines the in-sample fitting performance of six commonly cited in the crop yield field parametric distributional forms. This study develops an obvious extension to examine in-sample rate bias and efficiency of several alternative parametric and non-parametric distributions. Finally, this study uses a simulation approach to compare the out-of-sample bias and efficiency of the β, Weibull, empirical, and TCMN distributions.

The results from the first section show that the Burr XII, Weibull, β, and normal distributions perform better than the γ and inverse Gaussian distributions at representing yield samples across virtually all farm conditions. While the results from the second section demonstrate that the TCMN and kernel density estimators are the least biased in-sample, the results from the simulation analysis suggest that the more parsimonious Weibull distribution outperforms both the β and the TCMN on the basis of out-of-sample efficiency, particularly in small samples. The results of the simulations illustrate the bias-efficiency tradeoff when evaluating distributions with different levels of parameterization, and also add insight to the in-sample vs out-of-sample question as it relates to crop insurance rating and distribution selection. Most striking is that the use of simple GOF measures in small samples can lead to the inappropriate selection of candidate distributions that perform poorly in straightforward economic applications. In particular, this study finds that the fitted β consistently overstates rates even when fitted to data generated from an underlying β distribution. The small sample performance favors the Weibull distribution in this application, though other distributions not tested could potentially perform even better.

This study employs a single series approach (i.e. single farm) in evaluating GOF and generated insurance rates. Yet, insurers typically group large numbers of like farms together when making rates. Thus, further research is needed in order to assess the sampling distribution questions addressed herein in a more comprehensive framework comparable to that typically experienced by insurers. Also, the out-of-sample simulation analysis is based on simulated pseudo-data from known parametric distributions, and thus the out-of-sample results found here may not always carry over to cases representing actual data for any particular application (e.g. if the data have larger tails than the fitted β and Weibull distributions used here). Thus, further work is needed to evaluate out-of-sample rate performance of other starting distributions, such as fitting a kernel density estimator to actual yields and then drawing yields from the fitted kernel density estimator.

Notes

1. Most studies have examined county level data which are more stable than farm-level data and tend to have longer samples available from publicly collected and maintained sources. The data in this study represent the largest consistently collected farm level data that we are aware exists.
2. In some cases, the β distribution did not converge regardless of the choice of upper limit, and as a result, some samples were excluded from further analyses; the remaining sample was not qualitatively affected. The number of FBFM farms in each region and the total FBFM farms in Illinois for this application after eliminating cases with no convergence of the β are: Northwest – 379, Northeast – 154, West – 88, Central – 492, East – 280, West Southwest – 140, East Southeast – 236, Southwest – 111, Southeast – 76, and state total – 1,956 of the 2,088 original farms.
3. Other limits are also tested and the results for insurance applications were found to be insensitive.
4. Results from other cases identified earlier are qualitatively similar and are available from authors upon request.

References

Atwood, J., Shaik, S. and Watts, M. (2003), "Are crop yields normally distributed? A reexamination", *American Journal of Agricultural Economics*, Vol. 85 No. 4, pp. 888-901.

Claassen, R. and Just, R.E. (2011), "Heterogeneity and distributional form of farm-level yields", *American Journal of Agricultural Economics*, Vol. 93 No. 1, pp. 144-160.

Day, R.H. (1965), "Probability distributions of field crop yields", *Journal of Farm Economics*, Vol. 47 No. 3, pp. 713-741.

Gallagher, P. (1986), "US corn yield capacity and probability: estimation and forecasting with nonsymmetric disturbances", *North Central Journal of Agricultural Economics*, Vol. 8 No. 1, pp. 109-122.

Garcia, O. (1981), "Simplied method-of-moments estimation for the weibull distribution", *New Zealand Journal of Forestry Science*, Vol. 11 No. 3, pp. 304-306.

Goodwin, B.K. and Ker, A. (1998), "Nonparametric estimation of crop yield distributions: implications for rating group-risk crop insurance contracts", *American Journal of Agricultural Economics*, Vol. 80 No. 1, pp. 139-153.

Hennessy, D.A., Babcock, B. and Hayes, D. (1997), "Budgetary and producer welfare effects of revenue insurance", *American Journal of Agricultural Economics*, Vol. 79 No. 3, pp. 1024-1034.

Just, R.E. and Weninger, Q. (1999), "Are crop yields normally distributed?", *American Journal of Agricultural Economics*, Vol. 81 No. 2, pp. 287-304.

Ker, A.P. and Coble, K. (2003), "Modeling conditional yield densities", *American Journal of Agricultural Economics*, Vol. 85 No. 2, pp. 291-304.

Ker, A.P. and Goodwin, B. (2000), "Nonparametric estimation of crop insurance rates revisited", *American Journal of Agricultural Economics*, Vol. 82 No. 2, pp. 463-478.

Nelson, C.H. (1990), "The influence of distributional assumptions on the calculation of crop insurance Premia", *North Central Journal of Agricultural Economics*, Vol. 12 No. 1, pp. 71-78.

Nelson, C.H. and Preckel, P. (1989), "The conditional *beta* distribution as a stochastic production function", *American Journal of Agricultural Economics*, Vol. 71 No. 2, pp. 370-378.

Norwood, B., Roberts, M. and Lusk, J. (2004), "Ranking crop yield models using out-of-sample likelihood functions", *American Journal of Agricultural Economics*, Vol. 86 No. 4, pp. 1032-1043.

Ozaki, V.A., Goodwin, B.K. and Shirota, R. (2008a), "Parametric and nonparametric statistical modeling of crop yield: implications for pricing crop insurance contracts", *Applied Economics*, Vol. 40 No. 9, pp. 1151-1164.

Pope, R.D. and Ziemer, R. (1984), "Stochastic efficiency, normality, and sampling errors in agricultural risk analysis", *American Journal of Agricultural Economics*, Vol. 66 No. 1, pp. 31-40.

Ramirez, O.A. (1997), "Estimation and use of a multivariate parametric model for simulating heteroskedastic, correlated, nonnormal random variables: the case of corn belt corn, soybean, and wheat yields", *American Journal of Agricultural Economics*, Vol. 79 No. 1, pp. 191-205.

Ramirez, O.A., Misra, S. and Field, J. (2003), "Crop-yield distributions revisited", *American Journal of Agricultural Economics*, Vol. 85 No. 1, pp. 108-120.

Sherrick, B.J., Zanini, F., Schnitkey, G. and Irwin, S. (2004), "Crop insurance valuation under alternative yield distributions", *American Journal of Agricultural Economics*, Vol. 86 No. 2, pp. 406-419.

Silverman, B.W. (1986), "Density estimation for statistical and data analysis", Monographs on Statistics and Applied Probability, Chapman and Hall, London.

Snedecor, G.W. and Cochran, W.G. (1989), *Statistical Methods*, Iowa State University Press, Ames, IA.

Stephens, M.A. (1974), "EDF statistics for goodness of fit and some comparisons", *Journal of the American Statistical Association*, Vol. 69 No. 347, pp. 730-737.

Turvey, C.G. and Zhao, J. (1999), *Parametric and Non-Parametric Crop Yield Distributions and their Effects on All-Risk Crop Insurance Premiums*, University of Guelph, Guelph.

Wang, H. and Zhang, H. (2002), "Model-based clustering for cross-sectional time series data", *Journal of Agricultural, Biological, and Environmental Statistics*, Vol. 7 No. 1, pp. 107-127.

Zanini, F.D. (2001), "Estimating corn and soybean farm-yield distributions in Illinois", PhD thesis, University of Illinois, Urbana-Champaign, IL.

Further reading

Ozaki, V.A., Ghosh, S.K., Goodwin, B.K. and Shirota, R. (2008b), "Spatio-temporal modeling of agricultural yield data with an application to pricing crop insurance contracts", *American Journal of Agricultural Economics*, Vol. 90 No. 4, pp. 951-961.

About the authors

Bruce J. Sherrick is the Fruin Professor of the Farmland Economics and the Director of the TIAA-Center for Farmland Research. Bruce J. Sherrick is the corresponding author and can be contacted at: sherrick@illinois.edu

Christopher A. Lanoue is the Chief Executive Officer at the Nautilytics, LLC.

Dr Joshua Woodard is an Assistant Professor at the Cornell University in the Dyson School of Applied Economics and Management.

Gary D. Schnitkey is a Professor of Farm Management.

Nicholas D. Paulson is an Associate Professor at the University of Illinois, Department of Agricultural and Consumer Economics.