

On Technological Change in Crop Yields

Tor Tolhurst & Alan Ker*

ABSTRACT. Technological changes in agriculture tend to alter the mass associated with a segment or subpopulation of the yield distribution as opposed to shifting the entire distribution upwards. We propose modeling crop yields using mixtures with embedded trend functions to account for potentially different rates of technological change in different subpopulations of the yield distribution. By doing so we can test some interesting and previously untested hypotheses about the data generating process of yields. For example: (1) is the rate of technological change equivalent across subpopulations; and (2) are the probabilities of subpopulations constant over time? Our results -- technological change is not equivalent across subpopulations and probabilities have not changed significantly over time -- have implications for modeling yields. While we consider the impacts for rating crop insurance contracts, accurate modeling of technological change is relevant to issues such as food sustainability, economic development, feeding a rapidly growing world population, biofuels markets and policy, and climate change.

July 2013 Working Paper Series - 13-02

Institute for the Advanced Study of Food and Agricultural Policy

Department of Food, Agricultural and Resource Economics

OAC

University of Guelph

*Tor Tolhurst, former M.Sc. student and current Research Associate, Department of Food Agricultural and Resource Economics, University of Guelph (ttolhurs@uoguelph.ca). Alan Ker, Professor and Chair, Department of Food Agricultural and Resource Economics, University of Guelph (aker@uoguelph.ca). The authors would like to thank the Ontario Ministry of Agricultural and Food, the Ontario Ministry of Rural Affairs, and the Institute for the Advanced Study of Food and Agricultural Policy (Department of Food, Agricultural and Resource Economics, OAC, University of Guelph) for their generous financial support.

Introduction

Crop yields are agriculture’s principle unit of productivity measurement and, as a result of numerous changes in technology, agriculture has experienced dramatic and widespread yield increases over the past 75 years. These advances impact food sustainability, economic growth, world hunger, energy markets, and our ability to mitigate or enhance potential climate change effects. The rate of technological change has been exclusively measured at the mean implying technological developments result in a location or location-scale shift of the yield distribution. However, evidence in the crop science literature indicates technological developments alter the mass associated with a segment or subpopulation of the yield distribution (e.g. Barry *et al.* 2000; Dunwell 2000; Ellis *et al.* 2000; Badu-Apraku, Menkir and Lum 2007; De Bruin and Pederson 2008; Gosala, Wania and Kang 2009; Edgerton *et al.* 2012). For example, *triple-stacked* seeds were developed to increase resilience to a variety of pests as well as high winds thereby reducing mass in the lower tail (Edgerton *et al.* 2012). In contrast, *racehorse* seeds were developed to increase mass in the upper tail under relatively optimal growing conditions (Lauer and Hicks 2005). These developments suggest that the rate of technological change may vary across subpopulations and the probability of those subpopulations may change as well.

We propose modeling crop yields using a mixture of normals to account for the different subpopulations or components of the yield distribution. This is not new as others have first estimated a trend and then using the residuals estimated a mixture of two normals (e.g. Ker 1996; Goodwin, Roberts and Coble 2000; Woodard and Sherrick 2011). However, the mixture model is more flexible than previously employed in that it can accommodate different rates of technological change within different components and as such can be used to model yields without limiting technological change to location or location-scale shifts of the yield distribution.

To illustrate and provide some intuition for the proposed model we present in Figure 1

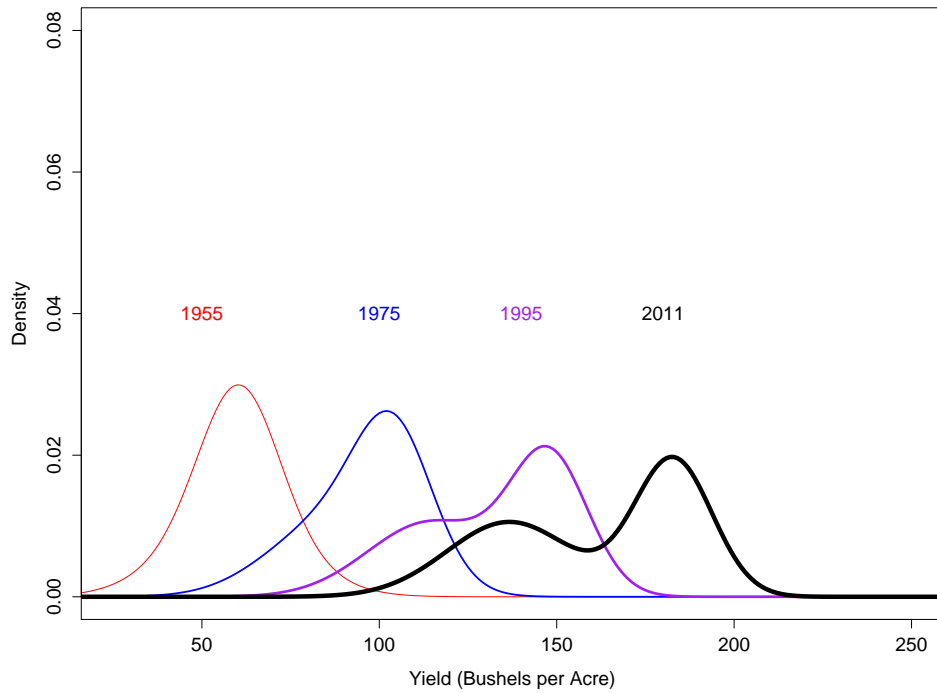



Figure 1: County-level corn yields in Adams County, Illinois and estimated densities with the proposed two-component trend.





Figure 2: County-level corn yields in Adams County, Illinois and estimated densities following the traditional estimation procedure.

the estimate for corn yields from Adams County, Illinois assuming a mixture of two normals with linear trends.¹ The rate of technological change in the *upper* component appears greater than the rate of technological change in the *lower* component, which implies technological change is increasing yields in the *upper* component faster than yields in the *lower* component. The bottom panel of Figure 1 illustrates the accompanying estimated yield densities at four different time periods. The shape of the estimated yield densities changes noticeably over time. In 1950 the estimated yield density appears relatively normal, while in 1970 the estimated density displays negative skewness. As time increases and the effect of differing rates of technological change become more prevalent, the estimated densities become increasingly bimodal and the overall variance increases (giving rise to the presence of heteroscedasticity). In contrast, Figure 2 illustrates the estimated trend and mixture densities for the same yield series by first estimating a single trend and then estimating the mixture from the heteroscedasticity-corrected residuals as commonly done in the literature. The estimated densities are, somewhat surprisingly, quite different given both assume a normal mixture and use the same data. Indeed, the crop insurance premium rates derived from these estimated densities are also markedly different: at the 75% and 90% coverage levels the single trend model results in rates 5.43 and 1.89 times larger than the rates from the trend mixture model respectively. These are non-trivial differences for USDA’s Risk Management Agency’s (RMA) area-yield programs which carried \$3.7 billion in liability in 2012 as well as the proposed shallow loss programs that are likely to be part of the next farm bill.²

Using a mixture model with embedded trend functions in each mixture enables us to test some interesting hypotheses about the data generating process of yields that have been previously untestable. First, we test if the rate of technological change is equivalent across subpopulations. Second, we test if the probabilities of the subpopulations are constant over the sample period. Third, given the size of RMAs area-yield programs and the likelihood

¹We discuss the number of mixtures and functional form issues in the empirical methods section.

²Both the area-yield programs and proposed shallow-loss program use county-level detrended heteroscedasticity-corrected yield data to estimate premium rates. Adams county had over 18,000 acres of corn insured under area-yield type insurance programs in 2012.

of those growing with the potential introduction of shallow loss programs, we test if economically and statistically significant rents can be recovered by adversely selecting against premium rates derived from the one trend mixture model using premium rates derived from the two trend mixture model. To do so, we use county level NASS data from Illinois, Indiana, and Iowa for corn and soybeans and Kansas, Nebraska, and Texas data for wheat.

The manuscript proceeds as follows. The next section discusses modelling crop yields with a brief survey of key contributions in the literature. Then we present the data, empirical model, and discuss our adjustments to the traditional Expectation-Maximization (EM) algorithm. The final two sections present the results and concluding remarks.

Modeling Crop Yields

Most often the approach to estimating conditional yield densities is to: (i) estimate a trend using the yield data; (ii) test the residuals from (i) for heteroscedasticity and adjust if necessary; and (iii) estimate a parametric or nonparametric conditional yield density given (i) and (ii). Relative to estimating technological change and issues of heteroscedasticity, estimating yield densities has received far greater attention in the literature.

A wide variety of density estimation approaches have been proposed. In 1958, Botts and Boles first suggested the use of “normal curve theory” to determine crop insurance premium rates. Day (1965) argued crop yield densities displayed non-normal attributes such as significant skewness. In response, Gallagher (1987) suggested the gamma distribution while Nelson and Preckel (1989) suggested the beta distribution. Goodwin and Ker (1998) proposed non-parametric kernel density methods while Just and Weninger (1999) argued deviations from normality were the result of inconsistencies in methods and data. A semi-parametric approach was forwarded by Ker and Coble (2003). Later parameteric specifications included the logistic (Atwood, Shaik and Watts 2003) and Weibull distributions (Sherrick *et al.* 2004). Inverse sine transformation methods were used by Moss and Shonkwiler (1993), Ramirez (1997), Ramirez, Misra and Field (2003), Ramirez and McDonald (2006). Normal mixtures

have been used by Ker (1996), Goodwin, Roberts, and Coble (2000), and Woodard and Sherrick (2011).

With the exception of Harri *et al.* (2011) and Just and Weninger (1999) heteroscedasticity has received surprisingly little attention in the literature considering the magnitude of its effect on crop insurance rates. Deterministic and stochastic approaches have been considered in estimating technological change in yield data. Deterministic approaches have dominated the literature and include a simple linear trend, two-knot linear spline (Skees and Reed 1986), and polynomial trend (Just and Weninger 1999). Stochastic approaches include the Kalman filter (Kaylen and Koroma 1991) and ARIMA(p, d, q) (Goodwin and Ker 1998). More recently, Ozaki and Silva (2009) and Claassen and Just (2011) incorporated spatial information into their temporal model.

We propose modeling yields as a mixture of normals with embedded trend functions in each mixture. That is:

$$y_t \sim \sum_{j=1}^J \lambda_j N(h_j(t), \sigma_j^2) \quad (1)$$

where the unknown parameters λ_j, σ_j^2 and functions $h_j(t)$ are estimated with a maximum likelihood approach using the heuristic EM algorithm for the j components of the mixture. The proposed model offers advantages in all three aspects of modeling yields. First, the normal mixture can approximate any continuous distribution and by default the distributional structures associated with yields. Second, embedding possibly unique trend functions within each mixture does not restrict the effect of technological developments to a location or location-scale shift of the yield distribution. Third, while the prevalence of heteroscedasticity in yield data is often corrected for, its presence has not yet been well explained in the crop science or agronomic literature. Note, that with the proposed model differing rates of technological change can lead to heteroscedasticity even with homoscedastic component variances.

Data and Empirical Methods

As with many empirical applications using yield data there is a trade-off between length of the time series and disaggregation keeping in mind what conditional yield distribution is sought. Ideally we would use farm-level yield data -- particularly considering technological adoption decisions are made at the producer level -- however the data is not sufficiently long to estimate the mixture model with any economically relevant degree of statistical significance. Therefore we use county-level National Agricultural Statistics Service (NASS) yield data for our analyses. While this averages the farm-level yield data and thus mixes the adopted technologies across farmers within a county, a county is a sufficiently small region having *relatively* similar weather patterns such that the distributional structure of county yields should not be markedly different than farm yields. A side benefit of using county level yield data is that our results are particularly relevant for USDA's area-yield crop insurance programs and the shallow loss programs proposed in the new farm bill.

We estimate our proposed model using county-level corn, soybean, and wheat yields. Corn and soybean data are from Illinois, Indiana and Iowa for the period 1955 to 2011. These states are major producers of both corn and soybeans. In 2011 they accounted for 15.6%, 7.2% and 17.3% of national corn production and 13.7%, 7.8% and 15.4% of national soybean production, respectively (NASS, 2013). For wheat we use Kansas, Nebraska and Texas with a slightly shorter time series due to the limited availability of data; we use 1968 - 2011 for Texas and 1956 - 2011 for Kansas and Nebraska.³ These three states were also major producers, accounting for 24.2%, 4.3% and 8.6% of national wheat production in 2011, respectively. In total, we have 754 crop-combinations after excluding any county without a full yield history.

We consider a mixture of two normals where the mean of each normal is replaced by a

³Data for Kansas are available for a longer period but were truncated at 1956 to align with the Nebraska data.

linear trend representing technological change. That is:

$$y_t \sim (1 - \lambda)N(\alpha_l + \beta_l t, \sigma_l^2) + \lambda N(\alpha_u + \beta_u t, \sigma_u^2) \quad (2)$$

where the unknown parameters λ , α_l , β_l , σ_l^2 , α_u , β_u , and σ_u^2 are estimated using maximum likelihood but optimized with an EM algorithm. Mixture models are commonly estimated using the EM algorithm because convergence issues arise with direct optimization of the likelihood.⁴ We modified the traditional EM algorithm (Dempster, Laird and Rubin, 1977) to embed the trend functions. This required replacing the weighted means estimate in the traditional EM algorithm with a weighted least squares estimate where the diagonal of the weighting matrix for the weighted least squares is the weighting vector and off diagonal elements are zero.⁵

While more than two components could be used we assume two for the following reasons. First and foremost, two components are sufficiently flexible to accommodate the variety of distributional structures that are associated with yield data (namely symmetric, skewed, long-tailed and bimodal) and estimated yield densities using a mixture of two normals are nearly identical to estimates from nonparametric kernel methods. Second, most climatic variables which are relevant to crop growth are spatially correlated within a county. As a result, producer yields tend to be correlated and central limit theorems do not apply. However, if one conditioned yields on climate, the resulting *conditional* yield density would be normally distributed. Thus partitioning climate into two subpopulations -- poor weather years and non-poor weather years -- would suggest a mixture of two normals for the unconditional

⁴Karlis and Xekalaki (2003), citing Bohning (1999) and McLachlan and Peel (2000), summarize the advantages and disadvantages of the EM algorithm.


⁵The main limitation of the EM algorithm is that it may converge on local maxima, particularly when the log-likelihood function is relatively flat or has multiple peaks. The problem of local maxima can be reduced by choosing multiple starting values. Starting values may be either chosen for the parameters or for the probability that a given realization belongs to a given component. We attempted three different approaches and found identical results in almost all cases. First we assigned a given yield realization probability zero to the *lower* component if it was greater than one standard deviation below the mean trend and one otherwise. Second we assigned a given yield realization probability zero to the *lower* component if it was below the mean trend and one otherwise. Third we choose starting values for the parameters λ , α_l , β_l , σ_l^2 , α_u , β_u , and σ_u^2 .

(with respect to weather) yield distribution at the county level.


We estimated a variety of functional forms for $h_j(t)$ such as linear, quadratic, logarithmic, and exponential. Most estimates exhibited a very linear structure (others overfit the data) and thus we assumed a linear trend for the functional form of $h_j(t)$. One caveat of note, in cases where there does not exist a very low yield in the first 5-10 years the linear trend for the *lower* component crossed the linear trend in the *upper* component. Considering least squares and linear specifications this is not surprising. In these cases we constrained the intercept from the *lower* component to be equal to the *upper* component to prevent the trend lines from crossing.

Estimation Results


Figure 3 presents the estimated temporal process for representative county-crop combinations. Roughly 90% of the cases are similar to the presented cases: the rate of technological change in both components is positive and is higher in the *upper* component. Interestingly, despite the yields of these three crops being quite different, the estimated trends look similar across crops and regions. The figures also illustrate how diverging component means create heteroscedasticity: as technological change in the *upper* component outpaces the *lower* component, dispersion increases (despite homoscedastic component variances). A more complete picture of the relationship between $\hat{\beta}_u$ and $\hat{\beta}_l$ is provided in Figure 4 which maps $\hat{\beta}_u$ against $\hat{\beta}_l$ for all county-crop combinations. The solid line represents equivalent rates of technological change between the two subpopulations and corresponds to the assumption of using a single trend. It is clear that the rate of technological change in the *upper* component has generally outpaced the rate in the *lower* component by a considerable margin for all three crops. Only a small number of cases have $\hat{\beta}_u < \hat{\beta}_l$ and fall below the solid line: 5.3% of corn, 8.7% of soybean and 24.8% of wheat counties. It is also readily apparent the rate of technological change in corn -- regardless of *upper* or *lower* component -- has significantly




(a) Jones, Iowa Corn




(b) Clinton, Iowa Corn




(c) Chickasaw, Iowa Soybean



(d) Decatur, Iowa Soybean



(e) Clay, Kansas Wheat



(f) Hansford, Texas Wheat

Figure 3: Representative two-component technological trend estimates.




Figure 4: Rate of technological change in *upper* year versus *lower* year yields across all states.

outpaced soybean and wheat.⁶ For corn, $\hat{\beta}_u$ is never below one and the majority of $\hat{\beta}_l$ exceed one. For soybean and wheat, $\hat{\beta}_u$ never exceeds one.

Table 1 reports summary statistics of the slope ratios broken down by crop and region. Also reported is the likelihood ratio test results from the hypothesis $\beta_l = \beta_u$. The results clearly suggest that the rate of technological change varies across the two components: 84.0% for corn, 82.3% for soybean and 64.0% for wheat reject the null hypothesis.⁷ These results suggest that technological developments -- of adopted technologies -- have not resulted in a simple location or location-scale shift in the yield distribution. This finding, consistent with the crop science literature, calls to question the use of a single trend to model technological change in yields.

Adoption of new technologies could alter the probabilities or mass of the mixture components. Ideally technological change would result in more resilient crops and production practices that would increase the mass in the *upper* component. Conversely, climate change may increase or decrease the mass in the *upper* component depending on geographic location. We test whether the probability of a draw from the *upper* component has changed over time by regressing the probability of membership in the *upper* mixture -- $\hat{\gamma}_t$ from the EM algorithm -- against time. If the probability of the *upper* component was increasing through time then the probability of a draw from the *upper* component would also be increasing through time. Table 2 summarizes the results and Figure 5 illustrates the expected membership over time corresponding to corn in Adams county, Illinois (Figure 1). Since the dispersion of expected memberships increases over the sample period we use robust standard errors for all of the *t*-tests. Only a small minority of county-crop combinations reject the null: 12.0% for corn, 5.4% for soybean and 3.0% for wheat. These results indicate the probability of a given component has not statistically changed over the sample period.⁸

⁶This is not surprising as significantly more research dollars have focused on corn productivity and the ability of seed providers to retain property rights.

⁷Not surprisingly the rejection rates are higher for corn and soybean: hybrid seeds have been developed for corn and soybeans but not wheat.

⁸This does not necessarily imply technological change or, say climate change, has had no effect -- it could also imply the *net* effect has been zero.

Table 1: Hypothesis Test One Results.

Crop	n	Summary Statistics for Ratio of $\hat{\beta}_R/\hat{\beta}_P$ Ratio					Rejection Rate of $\hat{\beta}_R = \hat{\beta}_P$
		Minimum	Mean	Median	Maximum	Std. Dev.	
Corn							
Illinois	97	0.56	1.35	1.34	2.52	0.23	81.4%
Indiana	79	0.70	1.53	1.38	8.48	0.89	84.8%
Iowa	99	0.57	1.69	1.32	12.24	1.25	85.9%
Soybean							
Illinois	97	-34.20	1.12	1.38	4.68	3.68	83.5%
Indiana	82	0.92	1.39	1.34	3.03	0.36	80.5%
Iowa	98	0.82	1.79	1.51	4.60	0.86	82.7%
Winter Wheat							
Kansas	88	-478.79	-5.16	1.54	7.57	51.54	61.3%
Nebraska	42	-0.10	1.25	1.27	2.83	0.54	38.1%
Texas	72	-400.83	-2.84	1.32	86.83	48.71	57.0%

Note: Rejection rate is the per cent of counties rejecting the null hypothesis evaluated at the 5% significance level. The extreme values (for example a maximum ratio of 86.83 for Texas wheat and a minimum ratio of -478.79 for Kansas wheat) are extreme because $\hat{\beta}_P \rightarrow 0$, which inflates the ratio. These values are extremely high when they approach zero from the righthand side and extremely low when they approach zero from the lefthand side. These values are apparent in Figure 4 near the vertical axis and there is nothing to indicate they are problematic.

Table 2: Hypothesis Test Two.

Crop-State	Number of Counties Rejecting Null		
	Positive	Negative	Total
Corn			
Illinois	0.0%	16.5%	16.5%
Indiana	0.0%	15.2%	15.2%
Iowa	0.0%	5.1%	5.1%
Soybean			
Illinois	0.0%	4.1%	4.1%
Indiana	2.4%	2.4%	4.9%
Iowa	1.0%	6.1%	7.1%
Winter Wheat			
Kansas	3.4%	5.7%	9.1%
Nebraska	0.0%	2.4%	2.4%
Texas	2.8%	0.0%	2.7%

Note: Statistical significance evaluated at the 5% significance level using least squares t -test with robust standard errors.




Figure 5: Expected membership ($\hat{\gamma}_t$) over time for corn in Adams County, Illinois.

Implications for Crop Insurance

As noted above, one of the benefits of using county-level yield data is that our results are particularly relevant for the area-yield programs such as GRP, GRIP, GRIPH which use detrended heteroscedasticity corrected county-level yield data to estimate premium rates. This program carried \$3.2 billion in total liability in 2012. In addition, if the shallow loss programs currently under the proposed farm bill become reality, detrended heteroscedasticity corrected county-level yield data will be used to estimate these rates as well.

We compare the crop insurance rates estimated from the mixture model with differing rates of technological change (denoted the two-trend model) to the mixture model with a single rate of technological change (denoted the one-trend model). To compare rates, we use an out-of-sample simulation using historical yield data to both estimate the models, derive the premium rates, and calculate losses. The simulation involves two agents: (1) the Risk Management Agency (RMA) which we assume derives rates from the one-trend model, and (2) a private insurance company, which we assume derives rates from the two-trend model. The simulated contract rating game has been commonly used in the literature to compare two rating methods (Ker and McGowan 2000; Ker and Coble 2003; Racine and Ker 2006; and Harri *et al.* 2011).

The simulation imitates the decision rules of the Standard Reinsurance Agreement, which effectively allows private insurers to retain or cede contracts of their choice *ex ante*, that is, averse select against the government.⁹ Premium rates are estimated out-of-sample based on a two-step-ahead forecast of the expected yield. Let $\hat{\pi}_{tk}^p$ be the estimated premium rate of the private insurance company for county k in year t based on yield data from 1955 to $t - 2$. Also let $\hat{\pi}_{tk}^g$ be RMAs estimated premium rate for county k in year t again based on yield data from 1955 to $t - 2$. The private insurance company retains policies with rates lower than the government rates ($\hat{\pi}_{tk}^p < \hat{\pi}_{tk}^g$) because *ex ante* it expects to earn a profit

⁹The SRA contains multiple funds and is more complicated but essentially a private insurance company can significantly reduce their exposure to unwanted policies.

on those policies. Loss ratios are calculated for the set of retained policies and the set of ceded policies using actual yield realizations. If loss ratios for the set of retained policies is statistically significantly lower than those of the ceded policies then we conclude that the two-trend model leads to more accurate premium rate estimates than the one-trend model. The simulation is performed on a by crop and state basis and p -values are calculated using randomization methods (1000 randomizations) as in Ker and McGowan (2000), Ker and Coble (2003), Racine and Ker (2006), and Harri *et al.* (2011).¹⁰ We ran the simulations for 20 years for all nine crop-state combinations.

Table 3 summarizes the results of the out-of-sample simulation for all state-crop combinations at the 75% and 90% coverage level. Overall the private company loss ratio (using the two-trend method) is lower in 77.8% of state-crop combinations across both coverage levels lending good support to the two-trend versus one-trend model. Of these, seven state-crop combinations are statistically significant at the 5% significance level and eight at the 10% significance level. Interestingly, six of the seven *statistically* significant cases are for wheat where claims are common and thus the power of the test is high. Conversely, the results are not *statistically* significant for corn and soybean where claims are rare and the power of the test is low. Notably, in all cases where the government loss ratio is lower the results are not statistically significant. The two-trend model also appears to perform relatively better at higher coverage levels where the percentage of claims is higher. Summarizing, the out-of-sample simulation shows good (not strong) support for the two-trend versus one-trend model.¹¹

¹⁰ p -values close to 0.00 represent statistical significance that the two-trend model is superior to the one-trend model whereas p -values close to 1.00 represent statistical significance that the one-trend model is superior to the two-trend model.

¹¹ It is worth noting that this simulation compares the ability of the two models to estimate the lower tail of the conditional yield density **only** and by default has relatively low power compared to the likelihood ratio tests in the previous section which showed very convincing support for the two-trend versus one-trend model.

Table 3: Out-of-sample simulation results for corn, soybeans and wheat.

Set	Coverage	Retained by Private	Psuedo Loss Ratio		<i>p</i> -value	% of Policies to Payout
	Level		Private	Government		
Corn						
Illinois	75%	85.9%	0.092	0.026	0.787	3.0%
	90%	87.8%	0.287	0.465	0.012	18.6%
Indiana	75%	81.1%	0.164	0.134	0.604	3.7%
	90%	82.4%	0.395	0.434	0.373	19.8%
Iowa	75%	83.9%	0.357	0.409	0.361	6.6%
	90%	86.8%	0.406	0.444	0.291	12.2%
Soybean						
Illinois	75%	77.7%	0.366	0.374	0.431	3.0%
	90%	67.4%	0.540	0.474	0.695	12.9%
Indiana	75%	86.9%	0.257	0.258	0.479	3.4%
	90%	78.4%	0.611	0.746	0.201	19.6%
Iowa	75%	80.6%	0.809	0.467	0.757	6.1%
	90%	78.7%	0.751	0.911	0.076	16.1%
Wheat						
Kansas	75%	52.7%	1.414	2.396	0.001	21.6%
	90%	48.6%	1.291	1.833	0.000	40.2%
Nebraska	75%	76.5%	0.771	2.470	0.001	11.2%
	90%	70.2%	0.932	1.692	0.001	35.9%
Texas	75%	69.4%	1.166	2.356	0.000	23.2%
	90%	64.2%	1.193	1.876	0.000	44.2%

Note: Out-of-sample simulation for 20 years with equally-weighted counties. Results are similiar at different coverage levels, simulation lengths and with acreage-weighted counties and are available from the authors upon request.

Conclusions

We propose modeling yields as a function of normal mixtures with embedded trend functions in each mixture. The normal mixture is beneficial in that it can accommodate the variety of distributional structures associated with yield data. Embedding possibly unique trend functions within each mixture does not restrict the effect of technological developments to a location or location-scale shift of the yield distribution and thus is more consistent with the agronomic and crop science literature. Using county level NASS data for corn, soybean, and wheat, we found that the large majority of the county-crop combinations rejected the null of identical trend functions and find that the rate of technological development is statistically higher in *upper* versus *lower* components. This finding also provides a reasonable explanation for the prevalence of heteroscedasticity in yield data. The two-trend versus one-trend model is also supported by the out-of-sample rating simulation which suggested the two-trend approach leads to more accurate premium rate estimates. This is of particular interest given area-yield type insurance programs carried \$3.7 billion in liability in 2012 and will likely grow if shallow loss programs are introduced. A caveat worth noting is that higher rates of technological change do not necessarily suggest that there have been greater developments for *upper* versus *lower* components but rather that *adopted* technologies have lead to more pronounced yield increases in *upper* versus *lower* components. Interestingly, this is consistent with the incentives provided by subsidized crop insurance; an insured producer will more readily adopt the technology that increases mass in the upper tail of the yield distribution versus a competing technology that decreases mass in the lower tail.

The proposed model presents an opportunity to address a number of other important questions in future research. For example, the proposed model is sufficiently flexible to consider how climate change will effect different aspects of the yield data generating process. Possible questions include: will climate change increase or decrease the probability of the *lower* component; will climate change adversely effect the variance within the *lower* com-

ponent; will climate change reduce the rate of technological change in the *lower* component but not the *upper* component thereby increasing the variance of yields; and how do different endowments and management strategies (especially soil quality and crop rotation) effect technological change in the *lower* versus *upper* component? More generally, food sustainability, economic growth, world hunger and population impacts, biofuels and accompanying policies, and the impacts of potential climate change are very much dependent on our understanding and modelling of technological change in crop yields as well as the underlying conditional yield distribution. The proposed mixture model with embedded trend functions provides an alternative and we suggest more probable model of technological change in crop yields.

References

- Atwood, J., S. Shaik, and M. Watts. 2003. "Are crop yields normally distributed? A reexamination." *American Journal of Agricultural Economics* 85 (4):888 – 901.
- Badu-Apraku, B., A. Menkir, and A.F. Lum. 2007. "Genetic Variability for Grain Yield and Its Components in an Early Tropical Yellow Maize Population Under *Striga hermonthica* Infestation." *Journal of Crop Improvement* 20(1-2):107–122.
- Barry, B.D., L.L. Darrah, D.L. Huckla, A.Q. Antonio, G.S. Smith, and M.H. O'Day. 2000. "Performance of Transgenic Corn Hybrids in Missouri for Insect Control and Yield." *Journal of Economic Entomology* 93(3):993–999.
- Bilmes, J.A. 1998. "A Gentle Tutorial of the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models." Working paper No. TR-97-021, International Computer Science Institute: Berkeley, CA.
- Bohning, D. 1999. *Computer Assisted Analysis of Mixtures & Applications in Meta-analysis, Disease Mapping & Others*. CRC Press, Boca Raton, FL.
- Botts, R.R., and J.N. Boles. 1958. "Use of normal-curve theory in crop insurance ratemaking." *American Journal of Agricultural Economics* 40(3):733–740.
- Claassen, R., and R.E. Just. 2011. "Heterogeneity and distributional form of farm-level yields." *American Journal of Agricultural Economics* 93 (1):144 – 160.
- Day, R.H. 1965. "Probability distributions of field crop yields." *Journal of Farm Economics* 47 (3):713 – 741.
- De Bruin, J.L., and P. Pedersen. 2008. "Yield Improvement and Stability for Soybean Cultivars with Resistance to *Heterodera Glycines Ichinohe*." *Agronomy Journal* 100:1354–1359.

- Dempster, A.P., N. Laird, and D. Rubin. 1977. "Maximum likelihood from incomplete data via the EM algorithm." *Journal of the Royal Statistical Society Series B* 39:1–38.
- Dunwell, J.M. 2000. "Transgenic approaches to crop improvement." *Journal of Experimental Botany* 51 (suppl. 1):487–496.
- Edgerton, M.D., J. Fridgen, J.R.A. Jr., J. Ahlgrim, M. Criswell, P. Dhungana, T. Gocken, Z. Li, S. Mariappan, C.D. Pilcher, A. Rosielle, and S.B. Stark. 2012. "Transgenic insect resistance traits increase corn yield and yield stability." *Nature Biotechnology* 30:493–496.
- Ellis, R., B. Forster, D. Robinson, L. Handley, D. Gordon, J. Russell, and W. Powell. 2000. "Wild barley: a source of genes for crop improvement in the 21st century?" *Journal of Experimental Botany* 51(342):9–17.
- Gallagher, P. 1987. "U.S. soybean yields: Estimation and forecasting with nonsymmetric disturbances." *American Journal of Agricultural Economics* 69 (4):796 – 803.
- 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 (1):139 – 153.
- Goodwin, B.K., M.C. Roberts, and K.H. Coble. 2000. "Measurement of price risk in revenue insurance: implications of distributional assumptions." *Journal of Agricultural and Resource Economics*, pp. 195–214.
- Gosala, S.S., S.H. Wania, and M.S. Kang. 2009. "Biotechnology and Drought Tolerance." *Journal of Crop Improvement* 23(1):19–54.
- Harri, A., K. Coble, A.P. Ker, and B.J. Goodwin. 2011. "Relaxing heteroscedasticity assumptions in area-yield crop insurance rating." *American Journal of Agricultural Economics* 93 (3):707 – 717.

- Just, R.E., and Q. Weninger. 1999. "Are crop yields normally distributed?" *American Journal of Agricultural Economics* 81 (2):287 – 304.
- Karlis, D., and E. Xekalaki. 2003. "Choosing initial values for the EM algorithm for finite mixtures." *Computational Statistics & Data Analysis* 41:577–590.
- Kaylen, M.S., and S.S. Koroma. 1991. "Trend, weather variables, and the distribution of US corn yields." *Review of Agricultural Economics* 13(2):249–258.
- Ker, A.P. 1996. "Using SNP Maximum Likelihood Techniques to Recover Conditional Densities: A New Approach to Recovering Premium Rates." Unpublished, Working paper, Department of Agricultural and Resource Economics, University of Arizona.
- Ker, A.P., and K. Coble. 2003. "Modeling conditional yield densities." *American Journal of Agricultural Economics* 85 (2):291 – 304.
- Ker, A.P., and P. McGowan. 2000. "Weather-based adverse selection and the U.S. crop insurance program: The private insurance company perspective." *Journal of Agricultural and Resource Economics* 25(2):386–410.
- Lauer, J., and D. Hicks. 2005. "Betting the Farm on Racehorse Hybrids." In *Paper presented at the Wisconsin Crop Management Conference (Jan 18 - 20)*.
- McLachlan, G., and P. D. 2000. *Finite Mixture Models*. Wiley, New York.
- Moss, C.B., and J.S. Shonkwiler. 1993. "Estimating yield distributions with a stochastic trend and non-normal errors." *American Journal of Agricultural Economics* 75 (5):1056 – 1062.
- Nelson, C.H. 1990. "The influence of distributional assumptions on the calculation of crop insurance premia." *North Central Journal of Agricultural Economics* 12 (1):71 – 78.
- Nelson, C.H., and P.V. Preckel. 1989. "The conditional beta distribution as a stochastic production function." *American Journal of Agricultural Economics* 71:370 – 378.

- Ozaki, V.A., and R.S. Silva. 2009. "Bayesian ratemaking procedure of crop insurance contracts with skewed distribution." *Journal of Applied Statistics* 36(4):443–452.
- 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*, pp. 27–39.
- Ramirez, O.A. 1997. "Parametric model for simulating heteroskedastic, correlated, nonnormal random variables: The case of Corn Belt corn, soybean and wheat yields." *American Journal of Agricultural Economics* 79 (1):191 – 205.
- Ramirez, O.A., and T. McDonald. 2006. "Ranking crop yield models: A comment." *American Journal of Agricultural Economics* 88 (4):1105 – 1110.
- Ramirez, O.A., S. Misra, and J. Field. 2003. "Crop-yield distributions revisited." *American Journal of Agricultural Economics* 85 (1):108 – 120.
- Sherrick, B.J., F.C. Zanini, G.D. Schnitkey, and S.H. Irwin. 2004. "Crop insurance valuation under alternative yield distributions." *American Journal of Agricultural Economics* 86 (2):406 – 419.
- Skees, J.R., and M.R. Reed. 1986. "Rate making for farm-level crop insurance: Implications for adverse selection." *American Journal of Agricultural Economics* 68 (3):653 – 659.
- Woodard, J.D., and B.J. Sherrick. 2011. "Estimation of mixture models using cross-validation optimization." *American Journal of Agricultural Economics* 93 (4):968 – 982.