Journal of Agricultural and Resource Economics 25(1):195–214

Measurement of Price Risk in Revenue Insurance: Implications of Distributional Assumptions

Barry K. Goodwin, Matthew C. Roberts, and Keith H. Coble

A variety of crop revenue insurance programs have recently been introduced. A critical component of revenue insurance contracts is quantifying the risk associated with stochastic prices. Forward-looking, market-based measures of price risk which are often available in the form of options premia are preferable. Because such measures are not available for every crop, some current revenue insurance programs alternatively utilize historical price data to construct measures of price risk. This study evaluates the distributional implications of alternative methods for estimating price risk and deriving insurance premium rates. A variety of specification tests are employed to evaluate distributional assumptions. Conditional heteroskedasticity models are used to determine the extent to which price distributions may be characterized by nonconstant variances. In addition, these models are used to identify variables which may be used for conditioning distributions for rating purposes. Discrete mixtures of normals provide flexible parametric specifications capable of recognizing the skewness and kurtosis present in commodity prices.

Key words: mixture distributions, price risk, revenue insurance

Introduction

A variety of crop revenue insurance programs have recently been developed to supplement the standard Multiple Peril Crop Insurance that has existed in the U.S. since the 1930s. In general, these programs guarantee producer revenues by protecting against any revenue-diminishing combination of low prices and/or low crop yields. The revenue insurance contracts guarantee producers a minimum level of revenues. If, because of any combination of poor yields and/or low prices, revenues are beneath the guaranteed level, insured farmers receive an indemnity payment equal to the difference between realized and guaranteed revenues. Increased planting flexibility and recent farm program changes which included the elimination or reduction of direct price supports have led many to anticipate increased price risk and uncertainty. Such concerns have heightened interest in the revenue insurance products.

Goodwin is professor and Roberts is graduate research assistant, both in the Department of Agricultural and Resource Economics, North Carolina State University; Coble is assistant professor, Department of Agricultural Economics, Mississippi State University.

We are grateful to two anonymous reviewers for helpful comments and suggestions. This research was supported by the Economic Research Service of the USDA. The opinions and conclusions expressed herein are those of the authors and do not necessarily reflect the views of the Economic Research Service or the U.S. Department of Agriculture.

Three alternative crop revenue insurance products currently exist: Crop Revenue Coverage (CRC), Income Protection (IP), and Revenue Assurance (RA).[1] Conventional crop insurance programs have been hampered by actuarial problems that have led to significant losses. In particular, program outlays exceeded $8.9 billion between 1990 and 1997 [U.S. General Accounting Office (U.S. GAO)]. These high losses have been attributed to adverse selection and moral hazard issues. Adverse selection occurs when producers have more information about their risk than do insurers, such that premium rates are inaccurate. Moral hazard occurs when insuring producers alter their behavior in order to increase the likelihood of collecting indemnities.

Inaccurate premium rates and performance monitoring problems underlie the actuarial shortcomings of crop insurance programs. Conventional yield insurance programs need accurate measurements of an individual producer's distribution of expected yields in order to determine actuarially fair premium rates. In the case of revenue insurance, an additional critical component of the proper insurance premium is setting a rate that accurately reflects the price dimension of risk. A variety of methods for measuring price risk have been proposed. A report recently released by the General Accounting Office is critical of the actuarial methods underlying all three revenue insurance plans (U.S. GAO).

It is important to note that yields and prices are likely to be negatively correlated since low yields are typically accompanied by high prices. The extent of this correlation for an individual producer depends upon the degree of correlation between the producer's yields and an aggregate yield, such as the national average. This, in turn, depends upon the spatial correlations of yields over principal production regions. The three primary revenue insurance contracts have different approaches to addressing this yield-price correlation issue. CRC, the largest of the three main revenue insurance programs, simply treats yield and price risk as though they are independent. Standard Multiple Peril Crop Insurance (MPCI) rates are added to a premium component that represents the price side of revenue risk. Negative correlation implies that the risks associated with a revenue shortfall are probably less than those associated with price and yield shortfalls when the latter are considered in isolation. This is because low yields would typically be expected to increase price, thus offsetting a portion of the revenue shortfall. In this manner, CRC is sometimes said to be "conservatively" rated—i.e., the CRC rate is higher than a rate which recognized the negative yield-price correlation would be.

The emphasis of our analysis is on evaluating methods for rating the price side of risk in the largest revenue insurance program, CRC. As is the case with current CRC rating methods, we do not consider the issue of yield and price correlation, though we do discuss below how our analysis might be extended to consider such correlation. In contrast to the CRC plan, the RA and IP plans do attempt to account for yield and price correlation, though the adjustments made to account for correlation have been questioned (see U.S. GAO, pp. 63–64, 71).

---

[1] Additional forms of revenue insurance, including insurance which utilizes Schedule F tax return information as a basis for insurance and an areawide version that utilizes county average yields as a basis for insurance, are currently under development. Although the issues discussed in this study are pertinent to all three products, the specific provisions of the contract and examples are taken from CRC.

Within the academic community, the Risk Management Agency (RMA), and throughout the insurance industry, there has been a collaborative focus on the development of proper actuarial methods for rating revenue insurance contracts. Stokes, Nayda, and English review research on revenue insurance pricing methods and discuss options-based pricing methods for rating revenue insurance. Considerable disagreement exists regarding the proper approach for rating price risk. While it is widely recognized that forward-looking, market-based measures of price risk are to be preferred, it is also the case that such market-based mechanisms do not exist for several of the crops covered by revenue insurance. For example, appropriate options markets do not exist for soft white wheat, which is covered under new revenue insurance contracts. In addition, because rates are set several months before planting, options markets for many crops have very low volumes and thus are inappropriate for rating.

Recent discussions have addressed three alternative approaches to rating price risk. The current CRC program uses a historical series (1973–98) of futures prices, quoted at planting time (F_t) and harvest time (P_t) to derive a "forecast error" ($e_t = P_t - F_t$), which is then assumed to be normally distributed. The portion of premium associated with price risk is then calculated using standard results for a normal distribution. An approach which utilizes proportional errors (e_t/P_t) under the assumption of normality has been recommended as an alternative. This approach assumes that errors are proportionally larger as prices are higher, and is thus somewhat analogous to assuming a lognormal distribution for prices since lognormality suggests a proportional relationship between the variance and the mean of the observed data. A third approach to rating price risk utilizes existing options markets to derive market-based measures of price risk. As noted, this approach, while clearly preferable, is not appropriate for all revenue insurance contracts since the necessary options contracts do not exist for all crops currently insured. The revenue assurance (RA) version of revenue insurance utilizes corn and soybean options premia to rate revenue insurance contracts.

The assumption of lognormality has considerable precedent in the financial literature. Models of price variability and options price determination have typically assumed that prices are lognormally distributed. In particular, the Black-Scholes option valuation formula, which is based on the assumption of lognormally distributed prices, has gained widespread acceptance. However, relatively little attention has been given to evaluating the extent to which prices adhere to distributional assumptions and the potential implications of distributional misspecification. More recent research (see, e.g., Cornew, Town, and Crowson; Hudson, Leuthold, and Sarassoro; Hall, Brorsen, and Irwin; Hsieh) has documented leptokurtosis, skewness, and other distributional characteristics that may be inconsistent with normality and lognormality. Recognition of these problems has led to the development of a variety of approaches to easing distributional restrictions and providing modeling techniques that allow for nonnormal distributions.

The distribution of market prices also may be sensitive to market conditions. Distributional shifts may occur if market conditions change. If the variance of prices is time-dependent, and if this time dependence is not explicitly modeled, the distribution of prices observed over time may involve a mixture of different variances and thus may exhibit characteristics incompatible with normality.[2] The price series may also display other distributional characteristics such as skewness, kurtosis, and multiple modes.

---

[2] A mixture of two zero mean normal processes with different variances will typically imply a distribution that exhibits kurtosis.

Recent research has applied alternative techniques to derive price distributions that reflect characteristics not consistent with normality (see, e.g., Hall, Brorsen, and Irwin; Hsieh). In one approach, finite mixtures of known distributions are used to represent distributional characteristics that are not compatible with normality. This approach is often motivated by the assumption that, although a standard distribution is appropriate under a given set of market conditions, changing market conditions may result in different distributions. Thus, when the entire series of prices is observed, the underlying process describing the aggregate distribution is a mixture of several distributions. In other research, mixed-jump processes have been used to represent nonstandard distributions. Jump processes are appropriate in situations where random shocks shift the entire distribution. In both cases, the resulting distributions are capable of representing characteristics of a series that may not be consistent with normality or lognormality. For example, a simple mixture of two normals is capable of representing a standard, symmetric normal distribution as well as nonsymmetric distributions, skewness, bimodality, and leptokurtosis.

Figure 1 (panels A and B) illustrates implied volatilities for corn and wheat, respectively. In both cases, the volatilities appear to imply two general states of nature.[3] In the first and most common state (perhaps 75% of the time), volatilities are relatively stable at around 15%. In the second and less frequent state, volatilities are much higher. Of course, the patterns of volatility also reflect seasonality in variance. While the implications of such a cursory examination of weekly intraseason data for the annual price data required for revenue insurance products are unclear, the illustration provides at least anecdotal evidence consistent with a mixture of a low variance and a high variance state.

The objective of this analysis is to explore the distributional characteristics of corn and wheat prices, focusing on the measurement of price risk for determining premium rates for crop revenue insurance programs. We utilize time-dependent conditional heteroskedasticity models and mixture distribution models to evaluate price risk. The conditional heteroskedasticity models evaluate the role of time to maturity, contract quote and expiration dates, and annual fixed effects in modeling wheat and corn price variability. Implications for improving actuarial methods utilized in the revenue insurance programs are also considered.

The article proceeds in the following manner. We begin with a description of crop revenue insurance products available in the U.S. The econometric methods applied to the analysis of price risk are then developed. In the next section we analyze the time dependency of the variance of prices and provide a discussion of conditional heteroskedasticity models that relate price variation to a number of explanatory factors. Models of conditional corn and wheat price distributions obtained under alternative distributional assumptions, including normality, lognormality, and discrete mixtures, are also presented. The final section offers a brief review of the analysis and some concluding remarks.

---

[3] Of course, one could argue in favor of more than two states. As we point out below, identification of multiple states is an exercise generally constrained by the number of observations available for empirical work.

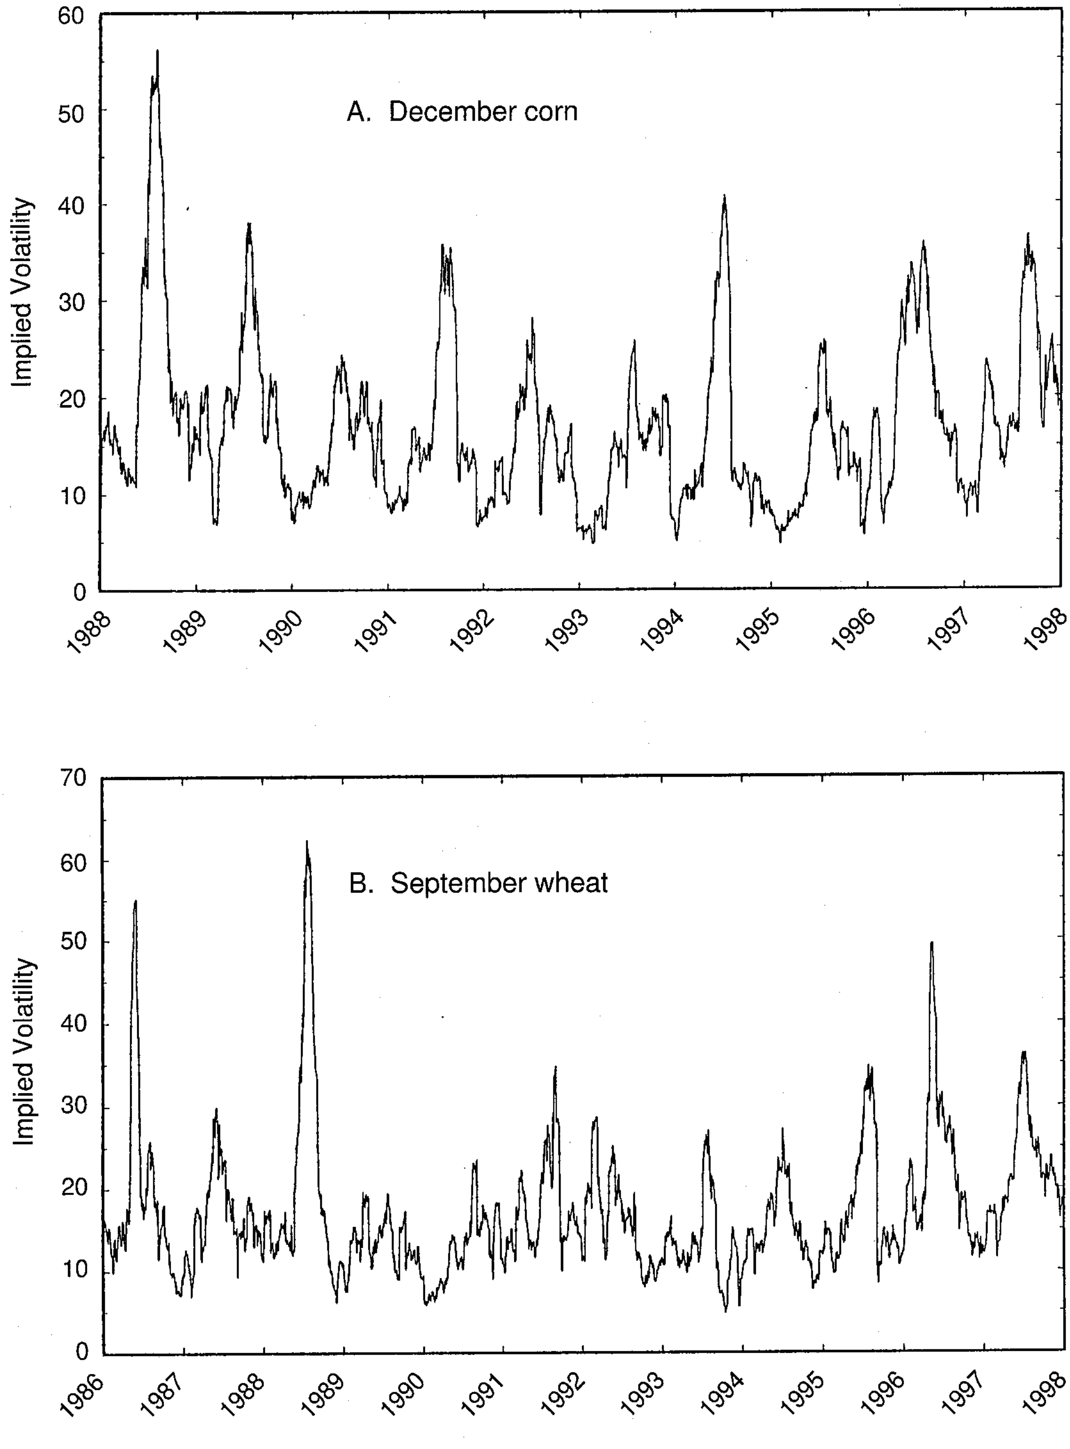

Figure 1. Implied volatilities for December corn and September wheat options contracts

Revenue Insurance Programs

Standard Multiple Peril Crop Insurance (MPCI) has been in existence in various forms since the 1930s. This insurance pays indemnities at a predetermined price whenever realized yields are less than guaranteed yields. A shortcoming of standard MPCI can be observed in the price (determined prior to planting) at which indemnities are paid. When yield losses are widespread, market prices are likely to be higher. Farmers receiving indemnities for lost yields may actually be reimbursed somewhat less (in bushel terms) than their guarantee since their indemnities likely reflect a price that is lower than the market price at harvest time. Revenue insurance had its beginnings with an optional rider that paid indemnities at harvest-time market prices. This, in conjunction with a put option contract, allowed producers to guarantee a minimum level of crop revenues. This coverage was extended to form the basis for individual Crop Revenue Coverage (CRC). CRC is currently available in major corn, soybeans, wheat, cotton, and grain sorghum growing regions. CRC has been quite successful, accounting for over 26% of corn crop insurance sales in 1997; the latest sales figures indicate that revenue insurance plans currently account for about 23% of the total acreage insured (RMA).

Income Protection (IP) was developed at Montana State University under a directive of the Federal Crop Insurance Reform Act to create a pilot insurance plan based on the actual costs of production. IP insurance is available for corn, soybeans, grain sorghum, cotton, and wheat in major growing regions. IP guarantees a minimum level of crop revenues, based on forecasted prices, individual farm yields, and area yields. If realized revenues fall beneath the revenue guarantee, producers receive an indemnity payment for the amount of the shortfall.

Revenue Assurance (RA) was developed by the Iowa Farm Bureau as a pilot program for corn and soybeans in Iowa. RA provides the option for "whole-farm" insurance in which producers insuring both corn and soybeans receive significant premium discounts. RA provides a guaranteed minimum level of revenue which is determined by individual farm yields and futures prices (adjusted for the local historical basis). If realized revenues are beneath the guarantee because of either low prices or low yields, or both, farmers receive an indemnity payment for the amount of the shortfall. A unique characteristic of the RA program is the utilization of market-based measures of price risks that are available in options markets. In contrast, the CRC and IP programs utilize historical futures prices to develop measures of price risks. RA actuarial procedures employ estimates of a beta distribution to model yield risks.

Econometric Methods

Revenue insurance contracts require a forecast of harvest-time prices, made conditional on information available prior to planting time. In addition, a measure of the uncertainty associated with the price forecast is needed to construct a premium rate reflecting the risk of adverse movements in prices. In all three revenue insurance plans, futures prices are used to forecast harvest-time prices. In the case of RA, options markets are used to gauge the uncertainty associated with prices. IP and CRC utilize historical price movements to evaluate price risks. The measurement of price risk in both the RA and CRC programs is heavily dependent upon assumptions regarding the parametric distributions underlying price movements. RA adopts standard Black-Scholes results to construct implied volatilities from observed options prices. As noted above, this approach assumes lognormally distributed prices (or, to be more precise, this model assumes a geometric Brownian motion process for prices, which is the continuous time equivalent of a lognormal distribution). In contrast, CRC assumes normally distributed prices in the construction of the price component of the revenue insurance premium. IP utilizes a nonparametric "empirical distribution" approach.[4]

In this analysis we employ two distinct approaches for evaluating price risk. In the first, the interest lies in determining if the variance of historical prices, which is used in rating revenue insurance products, is constant. Maximum-likelihood estimates of conditional heteroskedasticity models are used to evaluate the exogenous determinants of price variability. If an empirical analysis confirms that variances are constant or, alternatively, identifies factors which underlie a nonconstant variance, model estimates can provide conditional variance forecasts to be used in revenue insurance contract construction. In the second segment of the analysis, a set of annual price data is utilized to estimate price distributions and to evaluate insurance premia under alternative distributional assumptions.

In the conditional heteroskedasticity models, the variance of conditional prices (i.e., price differences) is assumed to be proportional to a function of several exogenous factors which are hypothesized to be related to price variability. In particular, it is assumed that the variance of prices for an individual contract i quoted at time t is given by the following:

$$\sigma_{it}^2 = \sigma^2 f(Z_{it}\gamma). \tag{1}$$

We assume that the conditional variance function $f(Z_{it}\gamma)$ is the square of a linear index function—i.e., $(Z_{it}\gamma)^2$. Such a model of multiplicative heteroskedasticity is widely applied in the literature. (For a detailed discussion of this model and its many variants, see Harvey.) Our specification ensures nonnegative variances for all observations. Under the assumption of normality, the following log-likelihood function is maximized to obtain estimates of γ and, if applicable, of parameters of a conditional mean equation ($\mu_{it} = X_{it}\beta$):[5]

$$\ln \mathrm{L} = -\frac{n}{2}\left[\ln(2\pi) + \ln(\sigma^2)\right] - \frac{1}{2}\sum_{i=1}^{n} \ln\left((Z_{it}\gamma)^2\right) - \frac{1}{2\sigma^2}\sum_{i=1}^{n}\frac{(y_{it} - \mu_{it})^2}{(Z_{it}\gamma)^2}. \tag{2}$$

---

[4] Nonparametric density estimation techniques offer complete flexibility in representing characteristics of a distribution. Such flexibility, however, does not come without a significant loss in efficiency. Thus, the nonparametric techniques may not be appropriate for the small samples which are commonly available for measuring price risk. In that probability density functions are commonly used as kernel functions in nonparametric density estimation, the nonparametric techniques are analogous to mixtures of a large number of components. For example, nonparametric estimation with Gaussian kernels is analogous to a mixture of n normals with equal variance terms (i.e., as determined by the kernel bandwidth).

[5] A conditional mean equation represents movements in expected prices, conditioned on observable data (typically expressed as $y_{it} = X_{it}\beta$). In our application, y_{it} represents the price difference, and no conditioning variables are added to the mean equation. As we explain below, this equation is modified to allow for first-order autocorrelation by replacing y_{it} with $y_{it} - \rho y_{it-1}$.

The second part of the analysis evaluates the distributional properties of the price data commonly used to rate revenue insurance. Finite mixture distributions represent a flexible, parametric approach to modeling probability distribution functions whose intrinsic characteristics are largely unknown. A k-component mixture density function is given by:

$$f(x) = \sum_{i=1}^{k} \left[\lambda_i f_i(x)\right], \tag{3}$$

where the probability weights (λ_i) satisfy the conditions that $\sum_{i=1}^{k} \lambda_i = 1$, and $\lambda_i > 0$ for all i. In our application, we consider only a mixture of two distributions, such that there is a single mixing parameter λ. Various densities are commonly applied in representing the underlying components of the mixture. The most common approach involves utilizing normal densities:

$$f_i(x) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{(x-\mu_i)^2/-2\sigma_i^2}. \tag{4}$$

Mixtures of normals nest a conventional normal distribution (obtained when $\mu_1 = \mu_2 = \ldots = \mu_k$, and $\sigma_1 = \sigma_2 = \ldots = \sigma_k$). Asymmetric and bimodal distributions may result when the μ_i's are not all equal. Kurtosis is implied when the μ_i's are not all identical.

Standard maximum-likelihood estimation techniques are commonly used to estimate mixture distributions. There are, however, particular characteristics of mixture problems that may complicate estimation. Nonlinear estimation techniques may have a tendency to concentrate component densities on individual points. In such a case, the σ_i associated with that point goes to zero and the likelihood function becomes numerically unstable. To prevent such instabilities, the λ and σ_i terms must be constrained to be positive. Estimation must also recognize that the mixing parameter λ must be constrained to lie in the interval (0, 1). Constrained maximum-likelihood estimation techniques are used in this study to estimate the components of the mixture. We constrain σ_i to be greater than 1E-9, and λ to lie in the closed interval $[0 + \varepsilon, 1 - \varepsilon]$ for ε = 1E-9.

The fact that the mixing parameter must be constrained and can lie on the boundary of the parameter space raises special concerns for hypothesis testing. In particular, test statistics may not have conventional distributions when the true parameter value is on the boundary of the parameter space. Likewise, under the null hypothesis that the mixing parameter λ is 0 (or, equivalently, 1), the parameters of the component distributions may not be fully identified. Problems in tests where a subset of parameters may be unidentified under the null hypothesis are common and can be addressed (see, e.g., the extensive literature underlying structural change tests with unknown break points, including the work of Andrews and of Hansen). In this application, the parameters characterizing the components of the distribution (i.e., σ_i and μ_i) are unidentified if $\lambda = 0$.[6] This precludes the application of standard hypothesis testing techniques for determining the number and nature of the component distributions. Bootstrapping techniques are commonly used as an alternative to construct empirical

[6] It should be noted that component parameters of individual distributions can be estimated when λ is constrained to be positive, even when the estimate of λ is very small.

distributions from which appropriate critical values and associated p-values can be obtained.

McLachlan, and Feng and McCulloch (1994, 1996) discuss bootstrapped likelihood-ratio tests for evaluating hypotheses in finite mixture models. These tests are particularly appropriate for determining the number of components to include in a mixture density. The tests also may be used to evaluate a particular parametric distributional specification. For example, an evaluation of a normal versus a lognormal distribution can be considered using a mixture of the form:

$$f(x) = \lambda\varphi(x) + (1 - \lambda)\phi(x), \tag{5}$$

where $\varphi(\cdot)$ represents the lognormal probability density function (pdf), and $\phi(\cdot)$ represents the normal pdf. The estimate of λ indicates whether the distribution is normal or, under the alternative, is a mixture of normal and lognormal densities. Because the likelihood-ratio test statistic [given by $-2(L_R - L_U)$, where L_R and L_U represent the restricted and unrestricted maximum log-likelihood function values, respectively] and parameters of the component distributions are defined even when λ is at the boundary, this approach provides a straightforward means for evaluating the number of components. McLachlan recommends a parametric bootstrap, whereby the data are simulated using estimates obtained under the null hypothesis. For each bootstrap replication, the alternative model is fit and the likelihood-ratio test statistic is constructed. The associated p-values, which can be used to evaluate the significance of the likelihood-ratio test statistic obtained from the estimation sample, can be calculated using the replicated test statistics. We follow this bootstrapping procedure to evaluate the number and nature of components in the price distributions.

The random variable x may also represent a conditional mean as in the standard linear regression problem. In this case, x may be replaced by $y - X\beta$ in equation (3), and the parameters of the conditional mean equation (β) may be estimated jointly with the parameters of the probability distribution (σ_i, μ_i, and λ). We follow this approach in our analysis. It should be noted that an additive intercept term is not identified when the mixture does not restrict estimates of μ_i. Estimates for the intercept can be recovered by imposing restrictions on the means of the component distributions. In particular, the implicit assumption of a zero mean for the errors provides identification.

Estimated Models and Results

The proper treatment of nominal prices observed over a long period of time is an important issue, especially in the second component of our analysis which uses data collected between 1899 and 1998. In particular, one must consider whether the prices should be deflated. Indeed, this issue has arisen in actuarial debates over the CRC program, where it was decided that nominal prices should be used. Of course, inappropriate deflation causes heteroskedasticity. Standard price deflators such as the CPI are not appropriate since they imply unreasonably high prices for distant periods. This is because agricultural prices have not followed the tendency of aggregate prices to rise over time.

Models utilizing logarithmic transformations of prices imply that the residuals (or price differences) are proportional to price levels, and thus that higher prices would be expected to correspond to larger price differences. Such an implication would suggest that the mean level of residuals in a logarithmic model (or differences in logarithmic prices) should be relatively stable over time. In contrast, models expressed in price levels suggest that the mean level of residuals (or price differences) does not depend upon the price level. To the extent that prices (or residuals) are being driven by movements in the overall price level, a plot of price differences should reveal increasing variability over time. Such plots (not presented here) were considered for both price levels and logarithmic transformations of the prices. We did not find evidence that price differences had trended upward in a manner consistent with aggregate price changes.[7] This was especially the case when logarithmic prices were considered. In light of these results and current rating practices used in the CRC program, both segments of our analysis employ nominal prices.

The Bridge database of daily settlement prices is used to construct monthly average futures prices for all contracts in all months over the period 1959–97. Expiration prices were the average in the month preceding the contract's expiration. This approach is analogous to the treatment of futures prices in constructing CRC premium rates. These data are used to estimate the conditional heteroskedasticity model [equation (2)] to determine if prices are characterized by nonconstant variances and to provide appropriate conditional forecasts of price variances. Since the pooled data set consists of many overlapping contracts, a complex form of moving-average error correlation is inherent in the price differentials. To allow for such correlation, we specify a first-order autoregressive correlation process among the monthly price differences. The correlation structure is restricted to prevent correlation corrections across alternative contracts.

Maximum-likelihood estimates and summary statistics for the conditional heteroskedasticity models ($Z_{it}\gamma$) are presented in tables 1 and 2 for corn and wheat, respectively. The models were expressed both in price levels (corresponding to a normal distribution), and in logarithms of prices (corresponding to a lognormal distribution). For corn, the default (omitted dummies) is a September contract quoted in the previous January. For wheat, the default is a July contract quoted in the previous January.

Though the magnitudes of the estimates differ, the results for the models expressed in levels and logarithms are quite similar. The results strongly confirm that the variance of the price differentials is not constant. They reveal that increased months to maturity decreases price volatility. This is consistent with the "Samuelson hypothesis" (Samuelson) which maintains that prices will reflect more information and thus be more volatile as contract expiration nears. In contrast to our findings, Hennessy and Wahl obtained results that were not consistent with the Samuelson hypothesis.

Our estimates also indicate that there are significant differences in price variability across alternative contracts. Contracts which expire in the months immediately preceding harvest (July for corn and May for wheat) appear to have the most volatile prices. Significant differences in the variability of prices over the growing season are also revealed in the estimates. Corn prices appear to be the most variable in June, July, and August—the most critical growing period. Likewise, wheat prices appear to be more variable in April. Wheat prices also appear to be quite variable in June and August,

[7] These plots are available from the authors on request. We should note that, although no trend in price differences was apparent, the variability of price differences did appear to be larger after 1970, though this pattern was not reflected in the logarithmic prices.

Table 1. Maximum-Likelihood Parameter Estimates and Summary Statistics for Conditional Price Heteroskedasticity Models: Corn

	Linear		Logarithmic	
Variables	Estimate	Std. Error	Estimate	Std. Error
ρ	0.9476	0.0028*	0.9427	0.0032*
Intercept	8.5877	0.3319*	0.0351	0.0014*
Months to Maturity	-0.0220	0.0021*	-0.0229	0.0021*
March Contract	0.0312	0.0296	0.0355	0.0325
May Contract	0.0320	0.0308	0.0018	0.0337
July Contract	0.1030	0.0326*	0.0422	0.0338
September Contract	—	—	—	—
December Contract	-0.0026	0.0292	0.0378	0.0326
February Quote	0.0261	0.0408	-0.0274	0.0429
March Quote	-0.0695	0.0400	-0.1222	0.0387*
April Quote	0.2981	0.0503*	0.1364	0.0478*
May Quote	0.1693	0.0570*	0.1626	0.0611*
June Quote	0.8233	0.0656*	0.7982	0.0686*
July Quote	0.9112	0.0872*	0.8698	0.0961
August Quote	1.2120	0.1024*	1.0467	0.1001*
September Quote	0.0903	0.0464	0.0820	0.0508
October Quote	0.1971	0.0463*	0.0564	0.0463
November Quote	0.1148	0.0469*	0.0341	0.0484
December Quote	0.0819	0.0443	0.0079	0.0482
R^2	0.9441		0.9381	
N	2,575		2,575	
Bera-Jarque Test (no annual effects)	1,384.52*		438.53*	
Bera-Jarque Test (annual effects)	47.65*		39.47*	
Test of Annual Effects	1,701.01*		815.71*	

Note: An asterisk (*) denotes statistical significance at the $\alpha = 0.05$ or smaller level.

perhaps reflecting harvest realizations or growing conditions for substitute spring wheats. The strong seasonality in the prices confirms the findings of other studies as well as conventional wisdom.

The parameter estimates allow a forecast of the variance, conditional on contract and month of quote. This is a forecast of the "average" variance for the particular month of quote and contract over the years of available data. This forecast could be used in conjunction with a price forecast to construct premium rates for the price component of revenue risk. This conditional variance may not be constant across years. Changes in other factors that affect price volatility (e.g., stocks, production, demand shocks, etc.) from year to year would result in annual differences in the conditional variances. Annual dummy variables were added to estimate an expanded model (not presented). Likelihood-ratio tests (tables 1 and 2) strongly confirm the significance of annual effects, implying that the conditional variances are not constant over the years of the analysis.

Table 2. Maximum-Likelihood Parameter Estimates and Summary Statistics for Conditional Price Heteroskedasticity Models: Wheat

	Linear		Logarithmic	
Variables	Estimate	Std. Error	Estimate	Std. Error
ρ	0.9065	0.0045*	0.9023	0.0049*
Intercept	16.8720	0.7339*	0.0393	0.0020*
Months to Maturity	-0.0117	0.0028*	-0.0112	0.0038*
March Contract	0.0090	0.0299	0.0119	0.0396
May Contract	0.0277	0.0286	0.0380	0.0380
July Contract	—	—	—	—
September Contract	-0.0033	0.0296	0.0444	0.0393
December Contract	-0.0002	0.0308	0.0155	0.0413
February Quote	-0.1745	0.0414*	-0.0574	0.0520
March Quote	-0.1045	0.0506*	0.0406	0.0756
April Quote	0.3641	0.0575*	0.3364	0.0659*
May Quote	-0.0890	0.0467	0.1251	0.0670
June Quote	0.2392	0.0582*	0.4219	0.0801*
July Quote	-0.1097	0.0505*	0.0584	0.0712
August Quote	0.5917	0.0618*	0.9330	0.0860*
September Quote	-0.0634	0.0434	0.0731	0.0585
October Quote	0.1148	0.0491*	0.2023	0.0649*
November Quote	-0.2031	0.0394*	-0.0438	0.0543
December Quote	0.0526	0.0421	0.1251	0.0512*
R^2	0.9103		0.9163	
N	2,080		2,080	
Bera-Jarque Test (no annual effects)	1,949.81*		683.57*	
Bera-Jarque Test (annual effects)	9.72*		12.99*	
Test of Annual Effects	1,881.61*		5,955.15*	

Note: An asterisk (*) denotes statistical significance at the $\alpha = 0.05$ or smaller level.

However, the expanded model cannot be used for forecasting variances out of sample since only information available at the time a forecast is made can be used to condition the forecast. The model which omits the annual dummy variables provides an "average" variance forecast which could be conditioned upon the months of contract and quote and used to forecast the variance, and thus rate price risk. Such models may offer advantages over current procedures which utilize only a single contract quoted at a single period of time by allowing use of a much larger sample (derived from using many contracts quoted over many different periods), thereby potentially improving the statistical efficiency of forecasts.

Bera-Jarque conditional moment (chi-square) tests of normality were also applied to evaluate the extent to which the models were consistent with normality and, in the case of the logarithmic models, lognormality. When the test is applied to the price-level models presented in tables 1 and 2, normality is strongly rejected in every case.[8] Likewise, the tests reject normality in each of the logarithmic versions of the model, suggesting that lognormality is also unsupported. When the test is applied to the models containing annual dummy variables, normality is still rejected, though at a much lower level of significance. This suggests that omission of fixed annual effects which are related to factors that influence variability from year to year results in a distribution that is much less consistent with normality than when such annual effects are accounted for. The residual nonnormality in the model without dummy variables to account for shifting annual variances may in part result from the implied mixture of (possibly normal or lognormal) distributions with different variances being associated with each year's distribution. Such a conclusion is somewhat tenuous, however, in light of the fact that the models containing annual effects still reject normality, albeit at a much lower level of significance. Alternatively, these results may suggest that the distribution of futures prices is inconsistent with either normality or lognormality, and thus that more flexible conditional heteroskedasticity models perhaps should be considered.[9]

In summary, the results from the first model show that futures price variability may be conditioned upon a number of explanatory factors, including months to maturity, month of contract, and month of price quote. These findings should be useful for constructing more accurate premium rates for the price-risk component of revenue insurance contracts. The proposed modeling approach allows a much larger sample to be used in constructing premium rates, potentially improving inferences and the accuracy of premium rates. Conditional moment tests reject normality, which may in part result from the mixing of time-varying variance distributions.

The second segment of the analysis utilizes a long series of annual observations on planting- and harvest-time futures prices. Corn and wheat futures were collected from selected issues of the Chicago Board of Trade's *Annual Report of the Trade and Commerce of Chicago* for the period covering 1899 to 1960. Data for subsequent years were taken from the Bridge financial database. Monthly observations for contracts expiring at harvest (September for corn and July for wheat) were constructed by taking the midpoint of the monthly high and low price quotes at planting times (January for corn and December for wheat).[10] The "harvest-time" price for each contract was that quoted in the month preceding the contract's expiration.

Maximum-likelihood techniques were employed to estimate alternative models of the annual price differentials in the second part of the analysis. A price relationship of the form $P_t = \alpha + \beta F_t$ was estimated, where P_t represents the harvest-time price, and F_t is the planting-time futures price.[11] In light of the prevailing assumption of lognormality for price distributions, five separate models differing in their distributional assumptions

---

[8] It should be noted that we do not report robust standard errors, and thus our estimated standard errors may be inconsistent if remaining residual heteroskedasticity is present.

[9] For example, Ramirez presented a flexible autoregressive conditional heteroskedasticity (ARCH) model that accounts for unimodal nonnormality. Such models may have promise in applications such as this one.

[10] This approach was necessitated by the available data—daily prices were not available before 1959. An evaluation of the difference in the monthly price constructed in this manner and a monthly average of daily closing prices revealed no significant difference. In particular, the average differential between the alternative monthly prices was nearly zero. Our use of these particular contracts was also necessitated by the availability of data.

[11] Similar results were obtained when the models were constrained according to an "efficient-markets" type of relationship, such that $\alpha = 0$, and $\beta = 1$. We estimate the parameters rather than constraining them in order to allow for any biases or premia which may exist in the relationship between prices. This is analogous to a linear mean forecast conditioned upon futures prices at planting.

were estimated. These included normality, lognormality, a mixture of two normals, a mixture of two lognormals, and a mixture of a lognormal and a normal. The mixture models permit testing of standard distributional assumptions using the bootstrapping procedures described above.

Tables 3 and 4 present the estimation results for the models of corn and wheat futures price relationships, respectively. Note that, with the exception of the variance estimate, maximum-likelihood estimates obtained under normality are equivalent to ordinary least squares (OLS). Estimates labeled as "OLS" in tables 3 and 4 are actually the equivalent maximum-likelihood estimates obtained under normality. Bera-Jarque normality tests are used to assess the extent to which the OLS residuals are consistent with normality and lognormality. As was the case above for the large sample of contracts, the tests reject normality and lognormality for both corn and wheat. These results question the validity of the assumptions of normality and lognormality used in the construction of revenue insurance premia. They suggest that alternative, flexible distributional specifications may be preferred. The OLS estimates for the normal and lognormal models have price coefficients which are slightly less than one. The mixture of normals model for corn has a price coefficient of 0.81—somewhat far from the expected value of one which would correspond to the futures price being an unbiased forecast of the future harvest-time price. The price coefficients for the other corn and wheat models are very similar, with values of about 0.90–0.96.

Recall that the mixing parameter λ characterizes the frequency of the alternative regimes. The estimated mixture of normals models points to an environment characterized by a mixture of a frequent (75–89% of the time) low-variance regime and a less frequent (11–25% of the time) high-variance regime. A similar pattern of variability is implied by the lognormal mixtures.[12] Such a finding is consistent with the pattern observed for options premia (figure 1). Our mixture approach is somewhat analogous to modeling heteroskedasticity in that two different variance estimates are used to characterize the aggregate distribution, though each distribution is permitted to have a unique mean. Analogously, our mixture models represent the nonnormal distribution that results when distributions with different variances are combined (i.e., *mixed*).

The bootstrapped testing approach described above was used to calculate the p-values associated with standard likelihood-ratio test statistics for the number of components (one versus two) to include in the mixture models. These are equivalent to testing H_0: $\lambda = 0$.[13] The evaluation of a mixture of normals versus a single normal results in a strong rejection, implying that a standard normal distribution is not suitable for either corn or wheat. In the tests of a mixture of lognormals versus a single lognormal, the test statistic has a value of 9.17 for corn (table 3) and 10.96 for wheat (table 4). The bootstrapped probability values indicate that these test statistics do not allow for rejecting H_0: $\lambda = 0$. Specifically, the corn and wheat test statistics have p-values of 0.29 and 0.25, respectively. This suggests that a single lognormal distribution is sufficient to model the price differentials when compared to a mixture of two lognormals. The final model includes a mixture of a normal and a lognormal distribution. Estimation of this mixture

---

[12] A reviewer has correctly noted that confidence intervals for σ_2 contain σ_1 in several cases. Though this is not a valid test of the significance of the differences in the two alternative variance parameter estimates, it does suggest that one should be cautious in concluding that the variance terms are different.

[13] In cases where the mixture consists of two identical distributions (e.g., two normals), this is also analogous to a test of $\lambda = 1$, since the estimates are equivalent under either null. The simulations used 300 replications.

Table 3. Maximum-Likelihood Parameter Estimates and Summary Statistics: Corn

Parameter	Normal OLS	Lognormal OLS	Mixture of Two Normals	Mixture of Two Lognormals	Mixture of Lognormal/ Normal
α	13.8443	0.4676			
	(7.1627)*	(0.1693)*			
β	0.9120	0.9063	0.8087	0.9158	0.9063
	(0.0438)*	(0.0353)*	(0.0293)*	(0.0334)*	(0.0349)*
σ	34.2225[a]	0.2063[a]			
	(2.5943)*	(0.0156)*			
λ			0.8928	0.9657	0.0000[b]
			(0.0522)*	(0.0282)*	—
μ_1			18.6397	0.4008	18.3780
			(4.2833)*	(0.1616)*	(1.0991)*
σ_1			18.5889	0.1739	39.7344
			(1.8038)*	(0.0164)*	(4.6821)*
μ_2			108.7290	1.0252	0.4676
			(27.6845)*	(0.2197)	(0.1693)*
σ_2			38.9203	0.1143	0.2063
			(18.2851)*	(0.1062)	(0.0156)*
Log-Likelihood Func.	-430.8066	13.8852	-408.4302	18.4701	13.8852
Test of Mixing			44.7563	9.1696	0.0000
p-Value			0.0000	0.2933	—
$\hat{P}$	239.4381	240.7328	228.5558	242.5003	240.3953
Pr $\{P < \hat{P}\}$	0.5015	0.5467	0.6246	0.5840	0.5436
Rate	5.6702	8.3928	5.1426	8.0266	8.2271
Bera-Jarque Test	13,698.38	1,156.53			

Notes: Numbers in parentheses are standard errors. An asterisk (*) denotes statistical significance at the $\alpha = 0.05$ or smaller level.

[a] Maximum-likelihood estimate of standard deviation.

[b] Value fixed by estimate on boundary of parameter space.

model produced estimates under which the density collapsed into a single lognormal distribution (i.e., λ was at its boundary value of 1E-9). The maximized log-likelihood function was nearly identical to that of the single lognormal. This obviates formal hypothesis testing, though the implication is clear—lognormality has strong support over normality. It should be noted that estimates of components of the mixture distributions are obtained even when the mixing parameter estimate lies on its boundary. Such estimates are difficult to interpret since they apply only to a very small fraction of the sample, as is implied by the restricted value of the mixing parameter.

It is also desirable to test the lognormal specification against the mixture of normals. The bootstrapping method was extended to consider a composite distribution comprised of a lognormal distribution and a bivariate mixture of normals. In each case, the

Table 4. Maximum-Likelihood Parameter Estimates and Summary Statistics: Wheat

Parameter	Normal OLS	Lognormal OLS	Mixture of Two Normals	Mixture of Two Lognormals	Mixture of Lognormal/ Normal
α	11.5468	0.2652			
	(6.5464)	(0.1207)*			
β	0.9371	0.9486	0.9648	0.9486	0.9486
	(0.0296)*	(0.0235)*	(0.0263)*	(0.0234)*	(0.0235)*
σ	31.0776[a]	0.1376[a]			
	(2.2665)*	(0.0100)*			
λ			0.7464	0.0208	0.0000[b]
			(0.0860)*	(0.0150)*	—
μ_1			7.8977	0.6628	-30.4253
			(4.1681)*	(0.1124)*	(0.0001)*
σ_1			14.1964	0.0612	2.2109
			(1.8793)*	(0.0375)*	(0.0001)*
μ_2			24.3102	0.1719	0.2652
			(14.9832)*	(0.1204)	(0.1207)*
σ_2			54.9053	0.1197	0.1376
			(10.2585)*	(0.0091)	(0.0034)*
Log-Likelihood Func.	-480.6846	15.9152	-459.7164	61.3998	55.9152
Test of Mixing			41.9503	10.9593	0.0000
p-Value			0.0000	0.2500	—
$\hat{P}$	330.4983	331.9241	330.2235	336.2555	332.4651
Pr $\{P < \hat{P}\}$	0.4934	0.5275	0.5553	0.5277	0.5276
Rate	3.7265	5.5187	2.9860	5.2874	5.4899
Bera-Jarque Test	7,895.48	1,658.61			

Notes: Numbers in parentheses are standard errors. An asterisk (*) denotes statistical significance at the $\alpha = 0.05$ or smaller level.
[a] Maximum-likelihood estimate of standard deviation.
[b] Value fixed by estimate on boundary of parameter space.

parameter defining the composite mixture between the lognormal and the discrete mixture of normals was on the boundary corresponding to an estimate of one. This indicates that the distribution again collapsed to a lognormal for both corn and wheat. Because the maximized log-likelihood function was again nearly identical to that of the single lognormal, formal hypothesis testing is again precluded. This does, however, indicate strong support for a single lognormal distribution when compared to a mixture of normals.

Prices were forecast for the last observation (1997) and insurance rates were based on a guarantee of 100% of this forecasted level. An insurance premium rate is given by expected loss over total liability. Expected loss is given by the product of the probability of a loss and the expected price given that a loss occurs. Numerical integration was used to estimate these probabilities and expected loss levels. With the exception of the mixture of normals case for corn, the predicted prices (given by $\hat{P}$ in tables 3 and 4) are very similar. As would be expected, rates based on lognormality are considerably higher than those based on normality. This reflects the positive skewness inherent in the lognormal distribution and larger variance estimates. In contrast, rates for the mixture of normals are somewhat lower than those under normality or lognormality, particularly in the case of corn. This lower rate in part reflects the lower forecasted price, which implies a lower price guarantee. The mixture of normals generates wheat premium rates that are somewhat smaller than those under normality. The premium rate estimates using the lognormal models and the models involving mixtures with lognormals are similar due to the similarity of the price distributions implied by the estimated models. The results suggest that rating procedures assuming normality may underestimate the price component of risk. Lognormality appears to provide more accurate estimates.

Differences in the premium rates and underlying distributions are revealed in plots of the densities implied by OLS and the mixture of normals cases. Figures 2(A) and 2(B) illustrate nonparametric kernel estimates of the densities associated with the OLS residuals.[14] Strong positive skewness is revealed in the estimates. In several cases, slight bimodality is revealed, suggesting that large, positive errors are sometimes observed—i.e., that the distributions may be a mixture of an infrequent high-variance regime and a frequent lower-variance regime. The distributions do not resemble normal densities, and thus the assumption of normality would again seem questionable. Figures 2(C) and 2(D) show the normal distributions implied by the maximum-likelihood estimates obtained under normality. Figures 2(E) and 2(F) illustrate the distributions under the mixture of normals case, which have a noticeably lower variance. This lower variance underlies the low premium rates suggested by the mixture of normals models. Figures 2(G)–2(J) graphically portray the distributions of the mixtures involving lognormal densities. In all cases, the models were dominated by a single lognormal distribution, suggesting that the distributions are very similar to those obtained under a single lognormal distribution. The distributions are nearly identical in the case of the normal/lognormal mixture.

In summary, this part of the analysis suggests that current premium rates which are based on normality are likely to be lower than the underlying price risk estimate implied by lognormality. Rates calculated in this manner, however, are based solely on historical information, and consequently may not fully reflect the uncertainty underlying market participants' actions at the time contracts are offered.

Concluding Remarks

This analysis evaluates distributional implications of modeling price uncertainty. The issue of price uncertainty has taken on increased importance with the introduction of three revenue insurance programs. In addition, changes in the farm policy environment that occurred with the 1996 Farm Bill have led to increased concerns regarding the stability of farm prices.

[14] Note that the nonparametric densities do not assume normality. OLS is a nonparametric estimation technique providing unbiased parameter estimates regardless of the underlying distribution. It has been noted, however, that least-squares estimation may make sample residuals more symmetric than the actual errors (see Huang and Bolch).

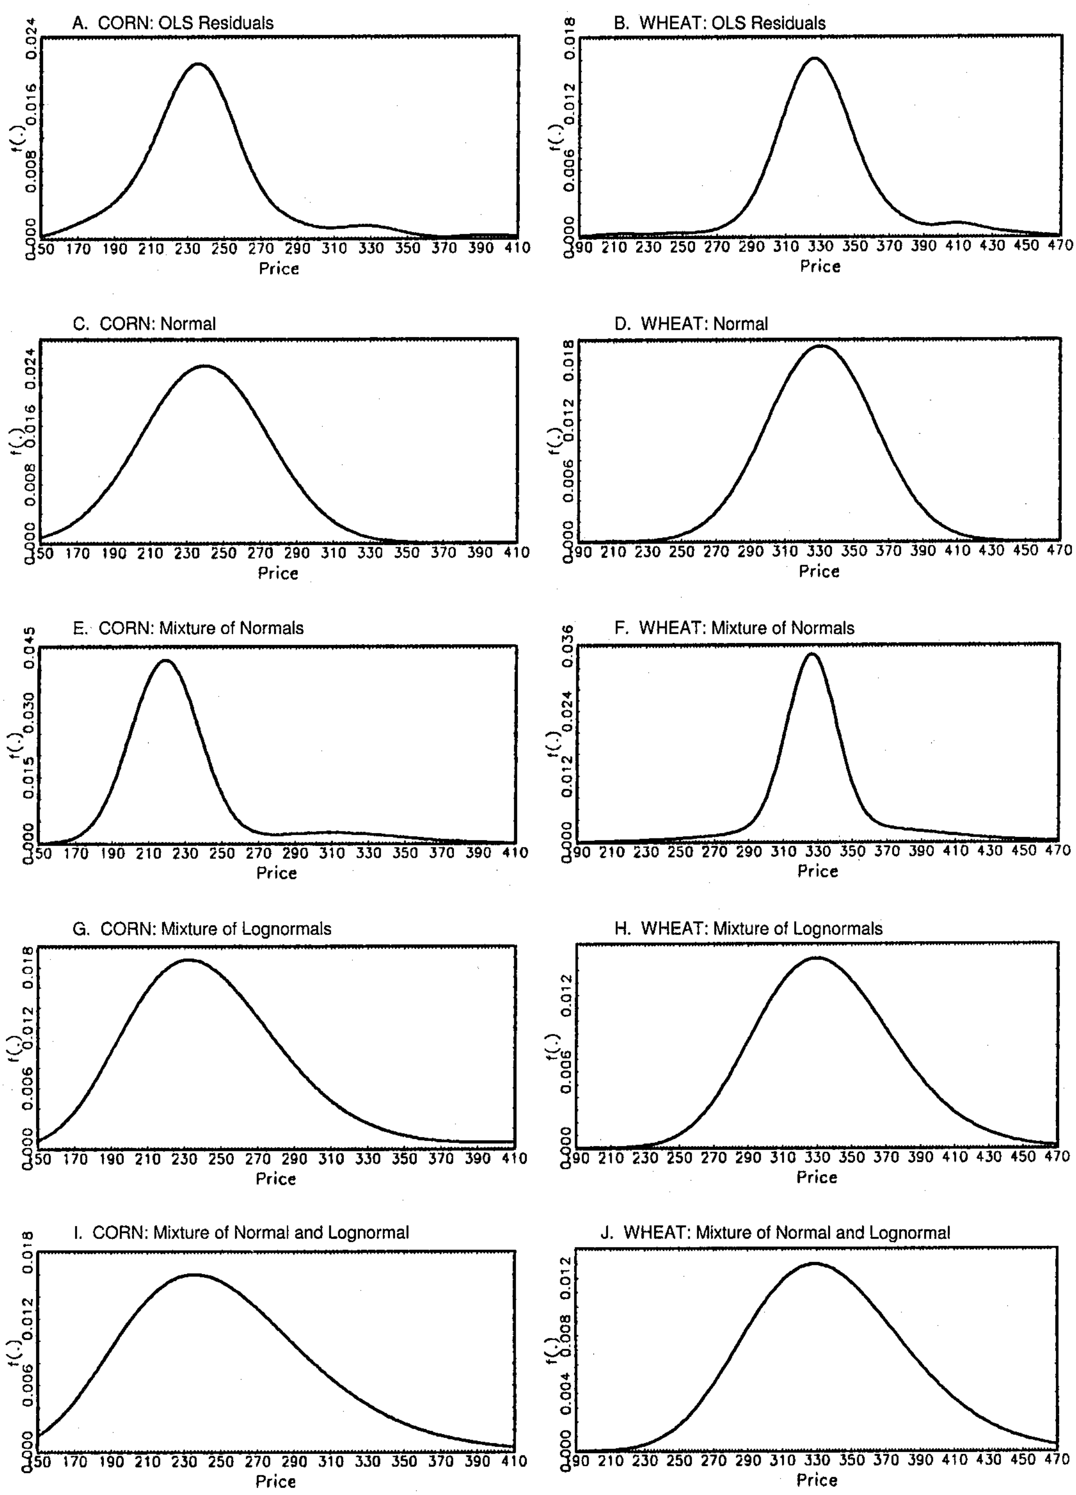

Figure 2. Estimations of price densities obtained under five alternative models for corn and wheat

An analysis of the conditional variance of corn and wheat prices revealed that variance decreases as time to maturity rises, and is highest during important growing periods. The findings of this analysis also imply that a nonconstant variance may contribute to significant departures from normality when data are aggregated over time. The results also indicate that conventional approaches to measuring price variability and rating revenue insurance may be misspecified. Our empirical results strongly reject normality. Although conditional moment tests also reject lognormality, testing results obtained from flexible mixtures of normals and lognormals provide reasonably strong support for a lognormal distribution. Insurance rates based on lognormality are considerably higher than those implied by normality.

Although our research findings have important implications for rating revenue insurance contracts, many important research issues remain. Most fundamentally, we have followed current CRC revenue insurance rating procedures and ignored yield-price correlation. Our methods could be extended to consider bivariate density estimation using mixture distributions that explicitly model such correlation. However, such an extension of our methods faces the same hurdle as nearly all insurance programs—a general lack of available yield data. In particular, nearly all crop insurance programs have been hampered by the fact that individual producer yield data are almost always scarce. Extension of the methods described here to yield models would also raise a number of other issues, including representation of regional differences in yield patterns and the appropriate geographic area for which to consider common yield distribution models. Avenues for making use of the limited data that are available within the context of mixture distribution estimation methods remain an important topic for further research.

Future research will consider additional explanatory factors (such as options premia, stocks, demand shocks, and growing conditions) which may be used to condition variance forecasts. Additional attention will also be given to modeling the complex correlation structure underlying our analysis of overlapping contracts.

[Received January 1999; final revision received September 1999.]

References

Andrews, D. W. K. "Tests for Parameter Instability and Structural Change with Unknown Change Point." *Econometrica* 61(1993):821–56.

Bera, A., and C. Jarque. "Efficient Tests for Normality, Heteroscedasticity, and Serial Independence of Regression Residuals: Monte Carlo Evidence." *Econ. Letters* 6(1980):255–59.

Black, F., and M. Scholes. "The Pricing of Options for Alternative Stochastic Processes." *J. Polit. Econ.* 81(1973):637–59.

Bridge Information Systems, Inc. "Historical Financial and Commodity Market Database." Chicago IL, January 1998.

Chicago Board of Trade. *Annual Report of the Trade and Commerce of Chicago.* Selected issues, 1899–1960.

Cornew, R. W., D. E. Town, and L. D. Crowson. "Stable Distributions, Futures Prices, and the Measurement of Trading Performance." *J. Futures Mkts.* 4(Fall 1984):531–57.

Feng, Z. D., and C. E. McCulloch. "On the Likelihood Ratio Test Statistic for the Number of Components in a Normal Mixture with Unequal Variances." *Biometrics* 50(1994):1158–62.

———. "Using Bootstrap Likelihood Ratios in Finite Mixture Models." *J. Royal Statis. Soc.* B,58(1996): 609–17.

Hall, J. A., B. W. Brorsen, and S. H. Irwin. "The Distribution of Futures Prices: A Test of the Stable Paretian and Mixture of Normals Hypotheses." *J. Finan. and Quant. Anal.* 24(1989):105–16.

Hansen, B. E. "Inference in TAR Models." *Studies in Nonlinear Dynamics and Econometrics* 2(1997). Online. Available at http://mitpress.mit.edu/e-journals/SNDE/002/articles/v2n1001.pdf. [Note: Access to subscribers only.]

Harvey, A. "Estimating Regression Models with Multiplicative Heteroscedasticity." *Econometrica* 44(1976):461–65.

Hennessy, D. A., and T. I. Wahl. "The Effects of Decision Making on Futures Price Volatility." *Amer. J. Agr. Econ.* 78(August 1996):591–603.

Hsieh, D. A. "Modeling Heteroscedasticity in Daily Foreign-Exchange Rates." *J. Bus. and Econ. Statis.* 7(July 1989):307–17.

Huang, C. J., and B. W. Bolch. "On the Testing of Regression Disturbances for Normality." *J. Amer. Statis. Assoc.* 69(June 1974):330–35.

Hudson, M. A., R. M. Leuthold, and G. F. Sarassoro. "Commodity Futures Price Changes: Recent Evidence for Wheat, Soybeans, and Live Cattle." *J. Futures Mkts.* 7(June 1987):287–301.

McLachlan, G. J. "On Bootstrapping the Likelihood Ratio Test Statistic for the Number of Components in a Normal Mixture." *Appl. Statis.* 36(1987):318–24.

Ramirez, O. A. "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." *Amer. J. Agr. Econ.* 79(February 1997):191–205.

Risk Management Agency (RMA). Unpublished crop revenue insurance product participation statistics. RMA, Washington DC, 1998.

Samuelson, P. A. "Is Real-World Price a Tale Told by the Idiot of Chance?" *Rev. Econ. and Statis.* 58(February 1976):120–23.

Stokes, J. R., W. I. Nayda, and B. C. English. "The Pricing of Revenue Insurance." *Amer. J. Agr. Econ.* 79(May 1997):439–51.

U.S. General Accounting Office. "Crop Revenue Insurance: Problems with New Plans Need to Be Addressed." Pub. No. GAO/RCED-98-111, Washington DC, April 1998.