

Issues and Strategies for Aggregate Supply Response Estimation for Policy Analyses

Octavio A. Ramirez, Samarendu Mohanty, Carlos E. Carpio, and Megan Denning

We demonstrate the use of the small-sample econometrics principles and strategies to come up with reliable yield and acreage models for policy analyses. We focus on demonstrating the importance of proper representation of systematic and random components of the model for improving forecasting precision along with more reliable confidence intervals for the forecasts. A probability distribution function modeling approach, which has been shown to provide more reliable confidence intervals for the dependent variable forecasts than the standard models that assume error term normality, is used to estimate cotton supply response in the Southeastern United States.

Key Words: nonnormality, probability distribution function, supply response

JEL Classifications: Q11, Q18, C32

Aggregate supply response estimation has been a key applied research issue for at least 30 years, because supply equations are critical components of increasingly more sophisticated U.S. and even world-level partial equilibrium trade models that are built, maintained, and continuously improved for policy analyses. The most common strategy for estimating the supply response is to build separate econometric models for yield and acreage. In most applications, these need to be estimated at the state, regional, or national level, which means that only relatively small samples are available for estimation. Although 80 to 100 years of

data might be available at the highest levels of aggregation, significant structural changes often constrain the researcher from using all available data to build the supply models. In regional-level analyses like the one exemplified in the present article, it is not uncommon to have to build yield and acreage response models on the basis of fewer than 30 observations.

We discuss key issues and exemplify the use of the small-sample econometrics principles and strategies that are needed to come up with reliable yield and acreage models under such circumstances. These are systematically presented throughout the following exemplary applications involving cotton supply response in the Southeastern United States.

Methods and Procedures

Yield Response Model Specification and Estimation

Conceptually, aggregate yields depend on the own crop price, the prices of competing

Octavio A. Ramirez is professor, New Mexico State University, Las Cruces, NM. Samarendu Mohanty is associate professor, Texas Tech University, Lubbock, TX. Carlos E. Carpio is graduate research assistant, North Carolina State University, Raleigh, NC. Megan Denning was a former graduate student at Texas Tech University.

This work was supported by a USDA-CSREES grant. The authors thank two anonymous reviewers of the *Journal* for their useful comments and insights.

crops, input costs, weather conditions during the growing season (particularly rainfall), and the number of acres planted. Our econometrics strategy begins by adopting the probability distribution function (pdf) modeling approach suggested by Ramirez, Misra, and Field. Specifically, yields at year t are specified as a linear function of all these variables plus a flexibly distributed error term (U_t):

$$(1) \quad Y_t = X_t B + U_t = B_1 + B_2 CP_t + B_3 SP_t \\ + B_4 VCP_t + B_5 RF_t + B_6 ACI_t \\ + B_7 AP_t + B_8 t + B_9 t^2 \\ + B_{10} t^3 + U_t,$$

where CP_t represents the effective cotton price received by producers in this region, SP_t represents the effective competing crop prices received by producers (regional soybean price in the case of the southeast), VCP_t stands for variable costs of cotton production, RF_t represents growing season rainfall, ACI_t stands for the number of southeast cotton acres enrolled in crop insurance, AP_t represents the number of cotton acres planted in that region, and t is a time trend.

The effective cotton and soybean price variables were constructed to form a close approximation of the regional price realizations, considering the prices received by the producers participating in government programs as well as those obtained by nonparticipating producers, and the total variable costs of cotton production were used as a proxy for input costs. For details on the effective price calculations and data sources, see Denning (pp. 71–75).

Econometric theory states that these variable approximations (i.e., measurement errors) might result in biased and inconsistent model parameter estimates and recommends taking corrective action (e.g., using some form of instrumental variable procedure) to at least achieve consistency. This, however, makes no sense, considering the degrees of freedom available for estimation in this and most other supply response applications. Instead, we recommend accepting the fact that

unbiasedness and consistency can not be ensured and attempting to estimate a flexible model that provides for the best possible representation of the data at hand. This, of course, means that the parameter estimates need to be interpreted with caution, but the supply and supply distribution forecasts obtained from such a model are likely to be adequate for policy analyses.

The flexible error term (U_t) in Equation (1) is specified as in Ramirez, Misra, and Field:

$$(2) \quad U_t = (Y_t - X_t B) \\ = \sigma_t [\sinh(\theta Z) - F(\theta, \mu)] / [\theta G(\theta, \mu)], \\ \text{where } Z \sim N(\mu, 1),$$

$$(3) \quad \sigma_t = B_{11} + B_{12} t,$$

$$(4) \quad F(\theta, \mu) = E[\sinh(\theta Z)] \\ = \exp(\theta^2/2) \sinh(\theta \mu), \quad \text{and}$$

$$(5) \quad G(\theta, \mu) = \{[\exp(\theta^2) - 1] \\ \times [\exp(\theta^2) \cosh(-2\theta \mu) \\ + 1] / 2\theta^2\}^{1/2},$$

where $\sigma_t^2 > 0$, $-\infty < \theta < \infty$, and $-\infty < \mu < \infty$ are distributional parameters; \sinh , \cosh , and \exp denote the hyperbolic sine and cosine and the exponential function, respectively; and Z is an independent normally distributed random variable. The pdf model above is estimated by finding the parameter values that maximize the following concentrated log-likelihood function:

$$(6) \quad LL = \sum_{t=1}^T \ln(G_t) - 0.5 \times \sum_{t=1}^T H_t^2, \\ \text{where: } G_t = [\sigma_t^2 / G(\theta, \mu) (1 + R_t^2)]^{-1/2}, \\ H_t = [\sinh^{-1}(R_t) / \theta] - \mu, \\ R_t = \{\theta U_t / [\sigma_t^2 / G(\theta, \mu)]^{1/2}\} + F(\theta, \mu),$$

where $\sinh^{-1}(x) = \ln\{x + (1 + x^2)^{1/2}\}$ is the inverse hyperbolic sine function, and σ_t^2 , $F(\theta, \mu)$, and $G(\theta, \mu)$ are as defined in Equations (3), (4), and (5). Ramirez, Misra, and Field also show that in this probability distribution function model:

$$(7) \quad \begin{aligned} E[Y_t] &= \mathbf{X}_t \mathbf{B}, \\ \text{Var}[Y_t] &= \sigma_t^2 = (B_{11} + B_{12}t)^2, \\ \text{Skew}[Y_t] &= S(\theta, \mu), \\ \text{Kurt}[Y_t] &= K(\theta, \mu), \end{aligned}$$

where $S(\theta, \mu)$ and $K(\theta, \mu)$ involve combinations of exponential and hyperbolic sine and cosine functions of θ and μ . In other words, the mean of the yield distribution shifts through time in response to changes in the explanatory variables in $\mathbf{X}_t \mathbf{B}$ [i.e., in Equation (1) above] its standard deviation (σ) is a linear function of time, and its skewness and kurtosis are determined by the parameters θ and μ . If $\theta \neq 0$ and μ approaches zero, the Y_t distribution becomes symmetric, but it remains kurtotic. Higher absolute values of θ cause increased kurtosis. If $\theta \neq 0$ and $\mu > 0$, Y_t has a kurtotic and right-skewed distribution, and $\mu < 0$ results in a kurtotic and left-skewed distribution. Higher absolute values of μ produce increased skewness.

In practice, under normality, θ would approach zero, and the pdf model above would collapse into a normal distribution with mean $\mathbf{X}_t \mathbf{B}$ and variance σ_t^2 . Therefore, the null hypothesis of normality versus the alternative of asymmetric nonnormality is $H_0: \theta = \mu = 0$ versus $H_a: \theta \neq 0, \mu \neq 0$. The null hypothesis of symmetric nonnormality versus the alternative of asymmetric nonnormality is $H_0: \theta \neq 0, \mu = 0$ versus $H_a: \theta \neq 0, \mu \neq 0$. In addition, the null hypothesis of homoskedasticity can be tested by $H_0: B_{12}$ versus $H_a: B_{12} \neq 0$. More sophisticated heteroskedasticity alternatives (such as $\sigma_t = B_{11} + B_{12}t + B_{13}t^2$) could, of course, be hypothesized and evaluated. However, we advise against it because of the small sample size.

Our proposed yield model estimation strategy then is to begin with the flexible but parsimonious pdf model described above and conduct four critical misspecification tests to make sure that the parsimony in the mean and variance specifications is not preventing us from identifying a more suitable representation for the data-generating process. The first is the well-known "RESET3"

test for functional form and omitted variables (Maddala). This test will evaluate if the simple linear specification for the conditional mean of the yield distribution hypothesized in Equation (1) could likely be improved on. The second is the also-common "White" test for heteroskedasticity. The third is the autocorrelation function analysis outlined by Bartlett complemented with the Q -test statistic for the null hypothesis of error term independence proposed by Box and Pierce. The fourth is the D'Agostino-Pearson omnibus nonnormality test. The last three tests are applied to the normalized and standardized residuals:

$$(8) \quad \hat{Z}_t = \langle \sinh^{-1}\{[\theta G(\theta, \mu)U_t/\sigma_t] + F(\theta, \mu)\}/\theta \rangle - \mu,$$

where $U_t = Y_t - \mathbf{X}_t \mathbf{B}$, $G(\theta, \mu)$ and $F(\theta, \mu)$ are as defined above, and all parameters are replaced by their estimated values. Note that Equation (8) is the inverse of Equation (2); therefore, if the initially assumed error term specification is correct, \hat{Z}_t should be an independently, identically, and normally distributed random variable and should not reject the null hypotheses of homoskedasticity, independence, and normality. Rejection of the null hypothesis in any of these tests would point to the possibility of improving the model in that particular aspect.

If none of the previously discussed hypotheses is rejected, one may proceed to the final stage of model specification, which involves identifying and excluding (i.e., setting to zero) all statistically insignificant parameters. Because standard maximum-likelihood algorithms only provide potentially imprecise numerical approximations of the elements of the Hessian matrix used to compute standard error estimates, we recommend that parameter significance be evaluated through the more reliable likelihood ratio tests. Also, because the small sample size would tend to increase the size of the type-2 error (i.e., the probability of rejecting H_0 when it is, in fact, false) at any given level of significance (α), we recommend to use a liberal α of 0.20.

As a final check, because of the possibility

of high correlations between some of the parameter estimators, we suggest conducting a joint likelihood ratio test to verify the statistical insignificance of all excluded parameters as a set and to again perform the four previously discussed misspecification tests on the final model residuals.

Acreage Response Model Specification and Estimation

The only difference between the initial specification of the yield and the acreage response models is in the conditional mean and variance functions [i.e., Equations (1) and (3)], specifically:

$$(9) \quad A_t = B_1 + B_2RC_t + B_3RS_t \\ + B_4RVCR_t + B_5EDP_t \\ + B_6DUM_{74} + B_7DUM_{86} \\ + B_8DUM_{97} + U_t,$$

$$(10) \quad \sigma_t = B_9DUM_{65} + B_{10}DUM_{74} \\ + B_{11}DUM_{86} + B_{12}DUM_{97},$$

where RC_t and RS_t represent measures of the expected net returns to upland cotton and soybean, respectively; $RVCR_t$ is a measure of the perceived variability of upland cotton net returns relative to soybean net returns; EDP_t is the effective diversion payment; and DUM_{65} , DUM_{74} , DUM_{86} , and DUM_{97} represent dummy variables for 1965–1973, 1974–1985, 1986–1996, and 1997–2001, respectively. For details about the calculation of the measures of expected net returns and of the perceived relative variability of net returns and a discussion of data sources, see Denning (pp. 76–78).

As Equation (1), Equation (9) is an ad hoc approximation of the unknown process relating explanatory variables to the conditional mean of the probability distribution of the number of cotton acres planted in the southeast during a particular year. In addition, although the set of variables included in this equation can be justified on the basis of economic theory (Denning, pp. 57–66), three of them are unobservable perceptions or expectations and have to be approximated.

Because of the small sample size, we recommend against evaluating too many alternative approximation methods, because one of them would be likely to provide seemingly outstanding results by mere chance. Instead, we suggest thinking about the expectation generating mechanisms thoroughly, from both a theoretical and a practical perspective, and coming up with one or two sensible methods to approximate these variables. In the present study, we opted for the procedure described in Denning (pp. 76–77). Also because of these unavoidable approximations, we again needed to set aside biasedness and inconsistency concerns, as well as the goal of estimating a “true” structural relation, and aim instead to estimate a flexible model that provides for a good representation of the basic characteristics of the unknown statistical process generating the data.

In regards to Equation (10), note that the typical zero-one dummy variable arrangement would result in DUM_{65} and DUM_{97} encompassing only six and five observations, respectively. This is obviously an insufficient amount of data to reliably estimate variability and, as illustrated later, can lead to econometrically unacceptable results. In addition, one could argue that the transition to an alternative “regime” triggered by a policy shift may take several years.

To address those two issues, in the present article, except for DUM_{97} , we use a “terminal” dummy variable arrangement of [1 0.75 0.5 0.25 0] instead of [1 1 0 0 0], and, except for DUM_{65} , we use an “initial” arrangement of [0 0.25 0.5 0.75 1] instead of [0 0 1 1 1]. This provides for smoother shifts across policy periods and, as illustrated later, the linkages established between the parameters corresponding to the different periods alleviates the econometric problem encountered when using the standard arrangement.

The last adjustment needed in the case of the acreage response model is for the fact that it might exhibit autocorrelation. Ramirez and Somarriba explained how to adjust the pdf model described above to account for autocorrelation. As in the case of normally distributed errors, if $U_t = (Y_t - \mathbf{X}_t\mathbf{B})$ is autocorre-

lated, it needs to be transformed into an independently distributed variable (U_t^*), as follows:

$$(11) \quad U_t^* = U_t - \rho_1 U_{t-1} - \rho_2 U_{t-2} - \dots \\ - \rho_L U_{t-L} (t = L + 1, \dots, T),$$

where L is the order of autocorrelation identified through the testing process described above. Note that Equation (11) is simply another equation entering the concentrated log-likelihood function [Equation (6)] where U_t has to be replaced by U_t^* , so that the autocorrelation coefficients ($\rho_1, \rho_2, \dots, \rho_L$) are jointly estimated with all other model parameters. In small-sample applications, however, one does not want to disregard the information contained on the first L observations.

As in the case of normally distributed errors, this can be avoided by also transforming the first L residuals (U_1, U_2, \dots, U_L) into independently distributed variables ($U_1^*, U_2^*, \dots, U_L^*$) using the well-known factors (which are also functions of the autocorrelation coefficients) and including these first L transformed residuals in the likelihood function as well (Judge et al., p. 294). All remaining aspects of the error-term specification for the acreage response model are as in the yield response model, because estimation is also accomplished by finding the parameter values that maximize the so-modified, concentrated log-likelihood function [Equation (6)]. We also follow the same steps for model specification and testing.

Supply Response Predictions

We start this section with the premise that, in the context of policy analysis, a good supply response prediction must include a reliable confidence interval for that prediction. Analysts are much better prepared to inform policymakers if they have a well-founded feeling of how precise their predictions are; and policymakers are beginning to demand that, in addition to point forecasts, "advisors" provide them with a range of possible outcomes and their associated probabilities (i.e., with the

probability distributions associated with their forecasts).

In addition to their flexibility, Ramirez, Misra, and Field's pdf models have been shown to provide more reliable confidence intervals for the dependent-variable forecasts than the standard models that assume error-term normality. These confidence intervals are, of course, obtained from the underlying probability distribution of the forecast. In the Results section, we will provide examples of how to derive these more reliable probability distributions and confidence intervals.

Before going into the exemplary applications, however, it is critical to discuss the issue of an unconditional versus a conditional forecast. For the purposes of the present study, we assumed that the analyst is interested in a 1-year-ahead dependent-variable forecast. This means that the values to be taken by some of the explanatory variables in the model might be unknown. Typically, the analyst will use forecasts for these unknown explanatory variable values to obtain a 1-year-ahead dependent-variable prediction. The question then is how to come up with a reliable (i.e., statistically sound) confidence interval for that prediction.

In short, in the general case to be illustrated here, where the model's error term and some of the explanatory variables are nonnormally distributed, there is no simple way to accomplish the former. Once its parameters (σ, θ , and μ) have been estimated, the probability distribution of the error term [U_t in Equation (2)] can be simulated using the procedures outlined by Ramirez, Misra, and Field and added to the yield forecast. The resulting distribution of the forecast, however, would then be conditional on the explanatory variables taking their exact predicted values, which is far from realistic.

The only theoretically sound alternative to obtain the "unconditional" distribution of this dependent-variable forecast is to also model the joint probability distribution of the explanatory variables which values themselves had to be forecasted, simulate joint realizations from this distribution 1 year into the future, and add the product of the simulated values

times the corresponding estimated model parameters plus the simulated error-term values to the dependent-variable forecast. Mathematically, in the case of the yield response model

$$(12) \quad Y_s = \hat{B}_1 + \hat{B}_2 CP_s + \hat{B}_3 SP_s + \hat{B}_4 VCP_s + \hat{B}_5 RF_s + \hat{B}_6 ACI_k + \hat{B}_7 AP_k + \hat{B}_8 t + \hat{B}_9 t^2 + \hat{B}_{10} t^3 + U_s,$$

where \hat{B}_i stands for the i th parameter estimate, the subscript s indicates a simulated value (jointly simulated in the case of the explanatory variables), and t has to be set at $T + 1$ for a 1-year-ahead forecast. The subscript k indicates that the number of cotton acres planted and the number enrolled in crop insurance are assumed known at the time of the forecast and, therefore, do not have to be modeled, forecasted, and simulated.

The question now is how to jointly model and simulate the distribution of the other explanatory variables, recognizing that they might not be normally distributed, heteroskedastic, and/or autocorrelated. Ramirez, Misra, and Field also helped address this question by providing a multivariate equivalent of the pdf model discussed above, which is estimated by finding the parameter values that maximize the following concentrated joint-likelihood function:

$$(13) \quad MNLL = \left\{ -\frac{n}{2} \times \ln|\Sigma| - 0.5 \times \sum_{j=1}^m [\ln(|\psi_j|)] + \sum_{t=1}^T \sum_{j=1}^m [\ln(G_{jt})] - 0.5 \times \sum_{t=1}^T \sum_{j=1}^m \left\{ \left[H_t * \left(\Sigma^{-1} \right) * H_t \right] \right\} \right\},$$

where m is the number of models to be jointly estimated; Σ is an $m \times m$ matrix with unit diagonal elements and nondiagonal elements ρ_{jk} (the cross-equation correlation coefficients); G_{jt} ($j = 1, \dots, m$; $t = 1, \dots, T$) is as defined in Equation (6) if Y_j is not normally

distributed or $G_{jt} = \sigma_{jt}^{-1}$ if Y_j is normally distributed; and H_t is a $1 \times m$ row vector with elements H_{jt} ($j = 1, \dots, m$) also defined in Equation (6) if Y_j is not normally distributed and $H_{jt} = (Y_{jt} - X_{jt}\beta_j)/\sigma_{jt}$ if Y_j is normally distributed. The operator $*$ stands for matrix multiplication; and $\cdot*$ indicates element-by-element matrix multiplication.

The multivariate log-likelihood function Equation (13) links m univariate log-likelihood functions through a cross-error term correlation matrix Σ . Estimates for the elements of Σ are needed for a proper simulation of the joint pdf characterizing the m explanatory variables being forecasted with this "auxiliary" model. Heteroskedasticity in any particular variable (Y_j) can be handled through σ_{jt} , as in Equation (3). Autocorrelation can also be addressed as explained above for the univariate pdf model case. Note that some variables could be heteroskedastic but not autocorrelated and vice versa, and some could be i.i.d.

In the exemplary case of yield response, a joint pdf model for cotton and soybean prices, variable costs of cotton production and rainfall are needed. We recommend pure time-series models, meaning that the means be initially specified as third-degree polynomial functions of time, the standard deviations (σ_{jt}) to be linear functions of a time trend, and the error term assumed to follow a first-order autoregressive process. The misspecification testing procedures described for the yield and acreage response models will also be applied in the case of these auxiliary models.

Results

Yield Response

Three alternative initial and final yield response models and related test statistics under the assumption of normality are presented in Table 1. Initial Model I is as specified in Equations (1) and (3) (with θ and $\mu = 0$). The problem with this model is that the estimate for B_4 , the parameter associated with the variable costs of cotton production (VCP), is positive and highly significant. This, of course, makes no economic sense and, therefore, is empiri-

Table 1. Alternative Initial and Final Yield Response Models and Related Test Statistics under the Assumption of Normality

	Initial Model I			Initial Model II			Final Model II			Initial Model III			Final Model III		
	Par. Est.	S.E. Est.	P.V.	Par. Est.	S.E. Est.	P.V.	Par. Est.	S.E. Est.	P.V.	Par. Est.	S.E. Est.	P.V.	Par. Est.	S.E. Est.	P.V.
B_1	-251.03	144.376	0.093	-91.995	0.641	0.641	-21.912	46.586	0.641	-128.75	166.553	0.446	-127.56	109.127	0.251
B_2	0.266	2.674	0.922	2.380	0.377	0.377				3.313	2.310	0.162	2.206	1.295	0.098
B_3	-23.318	19.264	0.237	-20.855	0.338	0.338				-26.919	19.679	0.182			
B_4	2.814	1.180	0.024												
B_5	6.448	4.824	0.192	5.340	0.367	0.367				9.420	4.584	0.049	10.925	4.689	0.026
B_6	-1.360	0.631	0.040	-0.927	0.296	0.296	-1.035	0.571	0.079	-1.033	0.564	0.078	-1.536	0.420	0.001
B_7	2.031	0.421	0.000	1.751	0.002	0.002	1.974	0.212	0.000	1.401	0.549	0.016	1.194	0.404	0.006
B_8	-9.419	14.159	0.512	0.977	0.966	0.966				9.299	16.927	0.587	11.253	3.361	0.002
B_9	95.131	108.923	0.390	193.491	0.281	0.281	213.927	34.679	0.000	96.686	102.525	0.354			
B_{10}	-20.930	26.238	0.432	-51.203	0.189	0.189	-55.687	12.799	0.000	-28.278	21.394	0.197			
B_{11}	25.086	15.758	0.123	21.672	0.285	0.285	7.780	9.022	0.395	61.477	7.148	0.000	64.730	7.526	0.000
B_{12}	2.615	1.115	0.027	3.208	0.040	0.040	4.470	0.955	0.000	0.829	0.096	0.000	0.873	0.102	0.000
	Cal. V.	Crt. V.		Cal. V.	Crt. V.		Cal. V.	Crt. V.		Cal. V.	Crt. V.		Cal. V.	Crt. V.	
<i>WHI</i>	12.676	14.68		6.784	13.36		5.695	7.78		7.795	13.36		5.539	9.24	
<i>RES</i>	0.082	2.53		0.761	2.52		0.572	2.49		0.104	2.52		0.045	2.50	
<i>SKE</i>	0.345	1.65		0.453	1.65		0.480	1.65		2.377	1.65		2.616	1.65	
<i>KUR</i>	0.717	1.65		0.531	1.65		0.427	1.65		1.852	1.65		2.159	1.65	
<i>DAG</i>	0.632	4.61		0.437	4.61		0.413	4.61		9.077	4.61		11.507	4.61	
<i>LEV</i>		-175.15			-177.98			-179.12			-179.09			-180.99	

Notes: Par, Par. Est., S.E. Est. and P.V. stand for parameter, parameter estimate, standard error estimate, and p value, respectively; the B_i 's ($i = 1, \dots, 12$) are as defined in Equations (1) and (3); *RES* and *WHI* refer to the RESET 3 and the White test, also discussed above; *SKE*, *KUR*, and *DAG* stand for the skewness, kurtosis, and the omnibus normality tests; Cal. and Crt. value refer to their calculated and critical table values ($\alpha = 0.1$), respectively; and *LFV* refers to the maximum value reached by the concentrated log-likelihood function. For the critical values, note that under H_0 : homoskedasticity, the White tests statistic is distributed as a χ^2_s , where s is the number of suspect variables in the test; under H_0 : adequate functional form, the RESET 3 statistic is distributed as an $F_{r,r,k}$; under H_0 : no skewness and H_0 : no kurtosis, *SKE* and *KUR* are standard normal; and under H_0 : normality, *DAG* is χ^2_6 .

cally unacceptable. This problem remains after excluding all statistically insignificant variables from Initial Model I. Initial Model II excludes the explanatory variable VCP, whereas Final Model II sets all nonsignificant parameters ($\alpha = 0.20$; i.e., $p \geq 0.20$) in Initial Model II equal to zero. Note that the likelihood-ratio test statistic [$2(179.12 - 177.98) = 2.28 < \chi^2_{(4,0.1)} = 7.78$] confirms the joint nonsignificance of all excluded parameters.

The misspecification tests applied to the residuals from Final Model II do not indicate an inadequate functional form, heteroskedasticity, or nonnormality (Table 1), which suggests that this model might be empirically acceptable. Note, however, that the estimated standard deviation equation is $\sigma_t = 7.78 + 4.47t$. This would suggest that yield variability, as measured by the standard deviation of the yield distribution, has increased 14-fold, from 12.25 pounds per acre in 1965 ($t = 1$) to 173.17 pounds per acre in 2002 ($t = 37$). These empirically unacceptable estimates illustrate a common "precision" problem when trying to estimate volatility trends with small time-series samples.

To address this problem, we recommend using a Gauss Constrained Maximum-Likelihood (CML) procedure to maximize Equation (6) subject to a constraint in the rate of volatility increase, to an empirically sensible percentage; such a constraint significantly lessens the statistical model's ability to represent the observed data. Initial Model III constraints the rate of increase in σ during 1965–2002 to 50%. Then, the estimated standard deviation equation becomes $\sigma_t = 61.48 + 0.83t$ (i.e., $\sigma_0 = 61.48$ and $\sigma_{37} = 92.20$), which is an empirically sensible amount of volatility increase during that time period. Further, the likelihood-ratio test statistic [$2(179.09 - 177.98) = 2.22 < \chi^2_{(1,0.1)} = 2.71$] suggests that this linear restriction on σ_t is not rejected by the data. Therefore, Initial Model III was adopted as the starting point for the analysis.

Final Model III sets all nonsignificant parameters ($\alpha = 0.20$; i.e., $p \geq 0.20$) in Initial Model III equal to zero, and the likelihood-ratio test statistic [$2(180.99 - 179.09) = 3.80 < \chi^2_{(3,0.1)} = 6.25$] confirms the joint nonsignif-

icance of all excluded parameters. Note, however, that once this volatility-increase restriction is imposed, the magnitude of the three nonnormality test statistics goes up substantially, which indicates that the model residuals are positively skewed, kurtotic, and jointly nonnormal at far beyond the 90% certainty level (Table 1).

The nonnormal yield pdf models estimated using Ramirez, Misra, and Field procedures are presented in Table 2. Note that, in the full model (Initial Model I), the parameter estimate corresponding to VCP (B_4) was again positive but this time was statistically insignificant at the 20% level. B_{10} and B_{12} , the parameter estimates corresponding to t^3 in the yield response function [Equation (1)] and to t in the standard deviation equation were also insignificant. Final Model I was obtained by setting these three parameters equal to zero.

Note that Final Model I is homoskedastic—that is, once error-term nonnormality is taken into account, the null hypothesis of a time-trending standard deviation cannot be rejected at any reasonable level of statistical significance (Table 2). Table 2 also presents the misspecification tests applied to both the "regular" residuals [$U_t = (Y_t - \mathbf{X}_t\mathbf{B})$] and after those same residuals have been standardized and normalized to obtain \hat{Z}_t according to Equation (8). Although neither shows signs of functional form or heteroskedasticity problems, there is strong statistical evidence that U_t is positively skewed, kurtotic, and jointly nonnormal ($\alpha = 0.01$).

On the other hand, the fact that the \hat{Z}_t series does not reject the skewness, kurtosis, and the joint omnibus nonnormality tests of D'Agostino et al. ($\alpha = 0.20$) suggests that the nonnormal pdf model properly accounts for the nonnormality in U_t . Error-term nonnormality is also evidenced by the single-parameter statistical significance of θ and μ , whereas a positive μ estimate confirms right-skewness (Table 2). A likelihood-ratio test of Initial (normal) Model II versus Initial (nonnormal) Model I [$2(177.98 - 171.93) = 12.10 > \chi^2_{(2,0.01)} = 9.21$] provided additional support for the alternative hypothesis of error-term nonnormality and suggested that the nonnormal

Table 2. Alternative Initial and Final Nonnormal Yield pdf Models and Related Test Statistics

Par.	Initial Model I			Final Model I		
	Par. Est.	S.E. Est.	P.V.	Par. Est.	S.E. Est.	P.V.
B_1	-436.68	111.115	0.001	-348.61	84.437	0.000
B_2	3.999	1.410	0.009	5.246	0.876	0.000
B_3	-37.921	9.717	0.001	-38.028	9.626	0.000
B_4	0.827	0.697	0.246			
B_5	12.200	2.971	0.000	11.572	3.060	0.001
B_6	-1.997	0.218	0.000	-1.882	0.193	0.000
B_7	2.282	0.318	0.000	2.084	0.309	0.000
B_8	30.943	10.033	0.005	29.172	6.777	0.000
B_9	-77.579	53.173	0.156	-46.637	17.444	0.012
B_{10}	6.139	9.860	0.539			
B_{11}	86.134	41.954	0.050	93.382	24.460	0.001
B_{12}	0.317	1.748	0.858			
θ	0.752	0.193	0.001	0.744	0.184	0.000
μ	2.708	0.881	0.005	2.747	0.857	0.003
		Correct.			Correct.	
	Resid.	Resid.	Crt. V.	Resid.	Resid.	Crt. V.
<i>WHI</i>	8.557	6.441	14.68	4.167	3.109	12.02
<i>RES</i>	0.051	0.243	2.53	0.039	0.276	2.52
<i>SKE</i>	3.360	-0.214	1.65	3.539	-0.189	1.65
<i>KUR</i>	2.347	-1.590	1.65	2.565	-1.344	1.65
<i>DAG</i>	16.798	2.575	4.61	19.106	1.843	4.61
<i>LFV</i>	-171.93			-172.70		

Notes: θ and μ are the nonnormality parameters in Equations (2) through (7); Resid. refers to the usual model residual [$U_i = (Y_i - X_iB)$], and Correct. Resid. refers to a residual that has been standardized and normalized [i.e., theoretically corrected for heteroskedasticity and nonnormality through Equation (8)]. All other abbreviations are as defined in Table 1.

pdf model is a far superior statistical representation of the data-generating process.

Having established error-term nonnormality in the empirical yield response models, it is interesting to compare the differences between the final specifications under the nonnormal pdf model versus the standard normal-error

model assumption (i.e., Table 1 vs. Table 2). First, in retrospect, it becomes clear that an improper assumption of error-term normality can result in an erroneous diagnosis of heteroskedasticity, because the former can become a statistical artifact for lessening residual nonnormality. Second, there can be substantial differences in the set of explanatory variables that is identified as statistically significant, particularly when the final normal-error model included an incorrectly specified variance function.

Also notice the marked differences between the predicted yield trends (holding all other explanatory variables at their sample mean levels) from Final Normal Model II, Final Normal Model III (i.e., with a restricted standard deviation function), and the Final Nonnormal Model (Figure 1). Most cotton yield experts would agree that those obtained




Figure 1. Yield Time Trends under Three Alternative Model Specifications




Figure 2. Probability Distributions of the 2002 Southeastern U.S. Yield Forecasts under the Normal and Nonnormal Models




Figure 3. Observed Values vs. the 80% Confidence Intervals for the Predictions from the Normal and Nonnormal Southeastern U.S. Cotton Yield Models

under the nonnormal pdf model are more likely compatible with the actual yield trends during the past four decades.

A key product of a yield response model is yield forecasts and confidence intervals for these predictions. It is important to point out that such confidence intervals need to take into account the two major elements of uncertainty affecting the predictions: the model's error term, as reflected on its estimated distribution, and the uncertainty about the true values of the model's parameters, which is captured by the joint distribution of their estimators. In all but ordinary least-squares-estimated models, these confidence intervals have to be obtained by numerical simulation of both the error-term and the model parameter estimators' distribution. Basically, the intervals are derived from the simulated distributions of the yield forecasts. For details on how to perform these simulations, please refer to Ramirez, Misra, and Nelson.

Figure 2 shows the distributions for the year 2002 cotton yield forecasts for the Southeastern United States under the final normal and the nonnormal pdf models, with the assumption that all explanatory variables take their known year 2002 values. Note that the distribution of the nonnormal model's forecast (standard deviation of 97.6) is much narrower than the one for the normal model's forecast (standard deviation of 108.5), which confirms that the first provides for a more precise statistical representation of the yield data.

At this point, it is useful to also warn the reader against interpreting the standard devi-

ations as comparable measures of dispersion when comparing normal with nonnormal distributions. For example, in this case, the 11.2% difference suggested by the ratios of their standard deviations is not indicative of the much wider degree of dispersion of the distribution of the forecast from the normal model. We have actually encountered situations where nonnormal distributions that are narrower across the board exhibit higher standard deviations than an evidently more disperse normal distribution.

The simulated data underlying Figure 2 can be used to compute confidence intervals. An 80% interval for the nonnormal forecast, for example, is obtained by finding the yield amounts that leave $\alpha/2 = 10\%$ of the simulated values below them and $\alpha/2 = 10\%$ above them. Within-sample confidence bands can be constructed by joining the lower and the upper bounds of the $T = 37$ $(1 - \alpha)$ confidence intervals for the T within-sample predictions. The constructed 80% confidence bands for the predictions from the normal and nonnormal error-term models are presented in Figure 3, in comparison to the observed yield values.

Previous research using much larger sample sizes (Ramirez and Fadiga; Ramirez, Misra, and Nelson) has shown that the confidence bands from Ramirez, Misra, and Field's nonnormal pdf model outperform those from a similar model that assumes error-term normality, in the sense that they are generally narrower and more closely adhere to theoretical expectations. In this case, the average width

of the 37 intervals comprising the normal model's band is 224.8 pounds per acre versus 194.6 pounds per acre for the nonnormal band. Theoretically, one would expect that 80% confidence bands would leave an average of $0.10 \times 37 = 3.7$ observations above them and 3.7 below them. In this application, the normal bands leave four observations below them and three above them, whereas the nonnormal bands leave three below and four above them.

In short, when the error-term is nonnormally distributed, using confidence intervals from a nonnormal pdf model provides two empirical advantages. First, in general, the intervals would tend to be significantly narrow, which suggests that the corresponding prediction is more precise. Second, under repeated sampling, they would be more likely to adhere to theoretical expectations—that is, to leave the “required” number of yield realizations within, below, and above them. These improvements would certainly be valuable to analysts and policymakers.

Confidence intervals developed as above will work well on reflecting the behavior of the yield data within-sample. They will, however, be inappropriate to represent the stochastic behavior of the dependent variable out-of-sample (i.e., into the future), because the values to be undertaken by at least some of the explanatory variables become unknown. As discussed above, theoretically suitable (i.e., empirically sound) confidence intervals for out-of-sample predictions have to incorporate a third element of uncertainty affecting these predictions: the randomness associated with the explanatory variable forecasts used to obtain the dependent variable predictions. Because VCP was excluded from the final nonnormal yield pdf model, a joint pdf model for cotton and soybean prices and rainfall is needed for this purpose.

For this purpose, the pure time-series model discussed above was estimated, and the final model was obtained by setting all nonsignificant parameters ($\alpha = 0.20$) equal to zero and confirming joint nonsignificance through a likelihood ratio test [$2(151.28 - 150.50) =$

$1.56 < \chi_{(3,0.10)}^2 = 9.24$].¹ Likelihood-ratio tests indicate that all three sets of nonnormality parameters (μ_i and θ_i , $i = 1, 2, 3$) are statistically significant at the 5% level, which suggests that the Southeastern U.S. cotton and soybean price and growing season rainfall distributions are all nonnormal. The positive μ_i estimates indicate right skewness in all three cases. Note that, strictly speaking, it is the error-term distributions that are nonnormal. Econometrically, the dependent variable and the error-term distribution only differ from each other by a known intertemporally shifting constant (X, B). Thus, as long as X, B can be considered fixed, they would have the same variance, skewness, and kurtosis. Properly speaking, then, we conclude that the conditional distributions are nonnormal.

In the case of the cotton and soybean prices, in addition to being nonnormal, the error-term distributions are increasing in variance through time (from 8.81 to 13.04 cents per pound for cotton and from \$1.16 to \$1.71 per bushel for soybeans) and show mild positive autocorrelation. The rainfall distribution is homoskedastic, as expected, but it shows a statistically significant negative autocorrelation. Because we are not sure that negative autocorrelation is contradicted by meteorological theory, we decided to retain this unusual model characteristic.

The 2002 joint price and the rainfall distributions are simulated as outlined above (for more details see, Ramirez, Misra, and Field). Their means and standard deviations are 58.87 and 13.98 cent per pound (cotton price), \$6.10 and \$1.81 per bushel (soybean price), and 21.87 and 2.69 inches (rainfall). As anticipated, they are all right-skewed and kurtotic, with the one corresponding to soybean price exhibiting the highest degree of nonnormality. The unconditional probability distribution for the 2002 yield forecast was then simulated using Equation (12).

As expected, the means of these two simulated distributions were nearly identical (540.46 vs. 540.13 pounds per acre). The stan-

¹ Detailed results of the model can be obtained from the author(s) on request.

standard deviation of the unconditional distribution (137.72 pound per acre), however, was over 40% larger than the standard deviation of the conditional distribution (97.64 pounds per acre). This, of course, was caused by the added uncertainty about the 2002 cotton and soybean prices and rainfall. Although they both retained the right-skewness and kurtosis of the nonnormal yield pdf model error term (which is a component in their simulation), their skewness and kurtosis coefficients were substantially different. However, the large difference in their standard deviations was still reflected in misleadingly narrow confidence intervals from the conditional distribution. The width of the 80% confidence interval, for example, was 198.96 under the conditional distribution of the yield forecast, versus 298.67 under the unconditional distribution.

The empirical implications of such a difference are evident. In short, all of these sources of uncertainty must be considered when constructing confidence intervals for a model's forecast, and error-term nonnormality needs to be accounted for as well. As illustrated above, neglecting any of these factors can result in confidence intervals that would not accurately reflect the future probabilistic behavior of the dependent variable and, therefore, undermine the credibility of the analyst's inferences in the long run.

Acreage Response

Many of the empirical issues surrounding the specification, estimation, and simulation of acreage response models are similar to those encountered when building yield response models. Therefore, in this section, we will focus on the issues that turned out to be somewhat different in our applications. One of these issues is that, even with the "overlapping" dummy (policy period) variable coding system explained in the Methods and Procedures, the maximum-likelihood estimate for the dummy variable coefficient estimating the standard deviation of the conditional acreage distribution for the 1997–2001 period approaches zero in the case of the normal-error model. This is clearly an empirically unac-

ceptable aberration due to the small number of observations available to estimate that volatility parameter, in combination with the also small sample size available for estimating all other model parameters.

The proposed solution for this problem is a restriction to the standard deviation parameters [B_9 , B_{10} , B_{11} , and B_{12} in Equation (10)] so that no individual coefficient estimate deviates from the average of the four estimates by more than 50% of this average. This is also implemented through Gauss-CML. The initial and final normal-error acreage response models estimated under this restriction and the analogous nonnormal pdf models (which do not require the restriction) are presented in Table 3.

Note, first, that the residuals from the initial normal model show strong signs of nonnormality. Also notice that the estimated autocorrelation coefficient (0.537) is highly significant and similar in magnitude to the ones obtained under the nonnormal models. The estimates for the standard deviation parameters (B_9 , B_{10} , B_{11} , and B_{12}) and their corresponding standard errors suggest that B_9 and B_{10} and B_{11} and B_{12} might not be statistically different from each other. A model that incorporates this restriction (i.e., $B_9 = B_{10}$ and $B_{11} = B_{12}$) and sets the three parameters that are not statistically significant in the initial normal model (B_2 , B_4 , and B_6) equal to zero (not presented) exhibits a maximum-likelihood function value of -205.35 . A likelihood ratio test [$2(205.35 - 203.63) = 3.44 < \chi^2_{(5,0.10)} = 9.24$] easily justifies these restrictions.

The correlation coefficient estimate in that restricted model, however, drops to 0.341 and becomes statistically insignificant ($p = 0.175$), which leads to the final normal model presented in Table 4. The estimate for $B_9 = B_{10}$ is substantially lower than either of the separate estimates obtained in the initial model, whereas the estimate for $B_{11} = B_{12}$ is much higher than either of the individual B_{11} or B_{12} estimates in that model, and the 50% mean-spread restriction on these coefficients becomes unbinding and is no longer needed. In short, the final normal model in Table 4 ap-

Table 3. Initial and Final Acreage Models and Related Test Statistics under the Assumptions of Error-Term Normality and Nonnormality

Par.	Initial Normal Model				Initial Nonnormal Model				Final Nonnormal Model			
	Par. Est.	S.E. Est.	P.V.	Crt. V.	Par. Est.	S.E. Est.	P.V.	Crt. V.	Par. Est.	S.E. Est.	P.V.	Crt. V.
B_1	1530.87	169.586	0.000	55.618	1424.19	165.531	0.000	0.000	1481.31	95.128	0.000	0.000
B_2	2.269	2.309	0.335		1.932	1.035	0.073	0.073	2.348	1.048	0.033	0.033
B_3	-15.816	8.359	0.070	1.123	-11.711	5.469	0.042	0.042	-13.373	3.118	0.000	0.000
B_4	10.624	18.939	0.580		-2.752	6.161	0.659	0.659	0.000			
B_5	-1.778	0.876	0.053	0.777	-1.407	0.204	0.000	0.000	-1.453	0.260	0.000	0.000
B_6	52.148	467.456	0.912		-55.579	337.644	0.871	0.871	0.000			
B_7	819.540	453.013	0.082	196.417	535.422	356.973	0.146	0.146	609.600	140.596	0.000	0.000
B_8	2568.61	519.632	0.000	361.829	2486.31	260.475	0.000	0.000	2513.35	91.396	0.000	0.000
B_9	162.724	118.459	0.181	22.395	555.751	370.847	0.146	0.146	458.710	118.221	0.001	0.001
B_{10}	193.055	81.711	0.026	22.395	301.883	176.470	0.099	0.099	458.710	118.221	0.001	0.001
B_{11}	377.690	84.094	0.000	110.003	643.225	237.516	0.012	0.012	458.710	118.221	0.001	0.001
B_{12}	273.707	212.611	0.209	110.003	214.410	79.172	0.012	0.012	458.710	118.221	0.001	0.001
ρ	0.537	0.204	0.014		0.521	0.050	0.000	0.000	0.516	0.049	0.000	0.000
θ	0.000				1.499				1.499			
μ	0.000				-0.143	0.267	0.597	0.597	0.000			
	Cal. V.				Cal. V.	Corr.			Cal. V.	Corr.		
	(Resid.)				(Resid.)	Resid.			(Resid.)	Resid.		
<i>WHI</i>	1.562	14.68		9.24	1.643	2.735	14.68	14.68	0.309	1.331	12.02	12.02
<i>RES</i>	0.074	2.53		2.50	0.276	0.187	2.53	2.53	0.124	0.292	2.52	2.52
<i>SKE</i>	2.353	1.65		1.65	0.311	3.317	1.65	1.65	0.394	4.729	1.65	1.65
<i>KUR</i>	2.538	1.65		1.65	-0.472	3.483	1.65	1.65	-0.276	4.374	1.65	1.65
<i>DAG</i>	11.976	4.61		4.61	0.319	23.132	4.61	4.61	0.231	41.494	4.61	4.61
<i>LFV</i>		-203.63		-206.26		-197.38				-200.60		

Notes: B_i ($i = 1, \dots, 12$) are the parameters in Equation (9), and all other abbreviations are as in Tables 1 and 2.

Table 4. Initial and Final Acreage Response Models and Related Test Statistics Under the Assumption of Error-Term Normality and Homoskedasticity

Par.	Initial Model			Final Model		
	Par. Est.	S.E. Est.	P.V.	Par. Est.	S.E. Est.	P.V.
B_1	1388.57	318.255	0.000	1497.87	326.548	0.000
B_2	2.186	2.024	0.290	0.000		
B_3	-9.277	9.059	0.315	-8.647	5.867	0.151
B_4	37.720	22.191	0.101	0.000		
B_5	-1.897	1.025	0.076	-1.650	1.039	0.123
B_6	-326.292	593.008	0.587	0.000		
B_7	579.152	530.408	0.285	968.105	389.781	0.019
B_8	1961.93	728.016	0.012	2138.01	746.575	0.008
B_9	287.767	34.967	0.000	306.629	37.288	0.000
B_{10}	287.767	34.967	0.000	306.629	37.288	0.000
B_{11}	287.767	34.967	0.000	306.629	37.288	0.000
B_{12}	287.767	34.967	0.000	306.629	37.288	0.000
ρ	0.709	0.170	0.000	0.747	0.172	0.000
	Cal. V. (Resid.)	Crt. V.		Cal. V. (Resid.)	Crt. V.	
<i>WHI</i>	3.1886	14.68		5.9202	9.24	
<i>RES</i>	0.1846	2.53		0.7119	2.50	
<i>SKE</i>	3.3485	1.65		3.7124	1.65	
<i>KUR</i>	3.5290	1.65		3.9302	1.65	
<i>DAG</i>	23.6664	4.61		29.2284	4.61	
<i>LEV</i>	-209.86			-212.08		

Note: Abbreviations are as in Table 1.

pears to be the best specification that can be obtained while assuming normality.

The initial nonnormal model, in contrast, can be estimated without the 50% mean-spread restriction on the standard deviation coefficients (Table 3). Furthermore, once the assumption of normality is relaxed, these four standard deviation coefficients are no longer statistically different from each other. The final nonnormal model is homoskedastic and excludes the variability of net revenues from cotton relative to sorghum ($RVCR_t$) and DUM_{65} , which were insignificant in the initial model. A likelihood ratio test [$2(200.60 - 197.38) = 6.44 < \chi^2_{(3,0.10)} = 9.24$] provides the statistical basis for these restrictions. In addition, a likelihood ratio test [$2(203.63 - 197.38) = 12.50 > \chi^2_{(3,0.01)} = 11.34$] indicates that the initial nonnormal model is quite a superior statistical representation of the acreage response data than the initial normal model (the third restriction in $\chi^2_{(3,0.01)}$ is the 50% mean-spread restric-

tion required in the initial normal model). Also note that, although the final models are not nested and therefore cannot be compared through a likelihood ratio test, the maximum value reached by the likelihood function of the nonnormal model (-200.60) is much higher than the one achieved by the normal model (-206.26).

Having established that the nonnormal model specification is statistically superior, one must conclude that the normality assumption in this case leads to misspecification of other error-term components, in particular the variance and autocorrelation functions. The more than fourfold increase in acreage variability after 1986 and the lack of autocorrelation in acreage response implied by the final normal model would be hard to accept for an empirical policy analysis.

Final evidence of the empirical advantage of using the nonnormal pdf model is provided in Figure 4. The average width of the 33 with-




Figure 4. Observed Values vs. the 90% Confidence Intervals for the Predictions from the Normal and Nonnormal Southeastern U.S. Cotton Acreage Models




Figure 5. Observed Values vs. the 90% Confidence Intervals for the Predictions from the Normal Homoskedastic and Nonnormal Southeastern U.S. Cotton Acreage Models

in-sample 90% confidence intervals under the nonnormal model is 826.6 lbs per acre, versus 1212.20 in for the normal model intervals. Although the predicted acreages are not that different (3,560 versus 3,881 pounds per acre), the terminal (year 2002) interval is more than twice as wide under the normal model.

In the case of this acreage response sample, for instance, the 90% nonnormal bands leave one observation below and one above them, versus the theoretically required $0.10 \times 33/2 = 1.65$ —that is, the nonnormal bands could actually be wider and still accommodate the observed data. The nonnormal bands, in contrast, leave two observations below and two above them (i.e., they might actually need to be even wider). Note that the normal bands are actually narrower than the nonnormal bands during the first half of the sample and then widen considerably because of the extraordinary increase in error-term variance during the second half. What is happening is that, in the final normal model, most of the variability in the independent variables is being exploited to explain the acreage observations during the first half of the sample.

Considering the above, we decided to re-estimate the normal model under the assumption of homoskedasticity (Table 4). The final homoskedastic model was obtained by setting B_2 , B_4 , and B_6 equal to zero (note that, although B_4 is almost significant at the 10% level, the parameter estimate exhibits the wrong sign). These three restrictions were justified by a likelihood-ratio test [$2(212.08 - 209.86) =$

$4.44 < \chi^2_{(3,0.20)} = 4.64$] at the 20% level. Note that, although the autocorrelation coefficient estimate was high (0.747) and significant at the 1% level, the two remaining explanatory variables (RS and EDP) were only moderately significant ($p = 0.151$ and 0.123 , respectively)—that is, the explanatory power of the final normal homoskedastic model is mostly driven by two policy-shift dummies and its autocorrelation function.

Although the null hypothesis of homoskedasticity was rejected [$2(209.86 - 203.63) = 12.46 > \chi^2_{(3,0.01)} = 11.34$] at the 1% level, we believe that this model is preferable to the final normal model in Table 3, and we would feel more comfortable using it for empirical analyses. The most suitable model, however, is the final nonnormal one. Figure 5 compares the 90% confidence bands for the within-sample predictions from the final normal homoskedastic versus the final nonnormal pdf model bands. The nonnormal model clearly retains its superiority in forecasting precision. These forecasts and confidence intervals are, of course, conditional on known values of the explanatory variables. The additional procedures needed to obtain unconditional forecasts and confidence intervals for acreage response are similar to those of yields.

Supply Response Forecasts and Confidence Intervals

Note that, in constructing the yield forecasts and confidence intervals above, it was as-

sumed that the number of acres planted was known. This would be correct if the supply response was being predicted right after planting. Then, the supply response forecast and its probability distribution would simply be obtained by multiplying the yield forecast and confidence intervals by the known acreage. Otherwise, more steps would be needed to finalize the task of obtaining theoretically sound and efficient cotton supply response forecasts for the Southeastern United States.

First note that, because the final nonnormal pdf models are recursive (i.e., acres planted at time t affects yields but yields do not contemporaneously affect acreage), both the yield forecast and the simulations to construct the probability distribution of that forecast would depend on the acreage response predictions and simulations. In other words, acreage response for year $T + 1$ would have to be predicted and the unconditional distribution of that forecast simulated first. Each simulated value (AP_s) from that acreage forecast distribution would then have to be used for simulating the probability distribution of the year $T + 1$ yield forecast [i.e., in place of AP_k in Equation (12)], and the resulting simulated yield value (Y_s) would be multiplied by the same AP_s to obtain a simulated realization (S_s) from the distribution of the supply response forecast.

Last, note that, in this case study, we have assumed that the error terms of the yield and acreage response equations are contemporaneously uncorrelated. If this does not hold, the previously described procedure for generating the probability distribution of the supply response forecast would have to be modified. Specifically, the two equations would have to be estimated jointly, using Equation (13). The estimated contemporaneous correlation coefficient would then be used to simulate the joint error term distribution corresponding to the yield and acreage response equations. All other steps would remain as explained above.

Concluding Remarks

In practice, one can think of supply response modeling at two levels. The first is focusing

on the “proper” modeling of the average yield and acreage responses only (i.e., on the systematic components of the models). We defined “proper” as a model that is the best possible statistical representation of the data at hand given the limitations of econometric estimation procedures in small sample applications and is also compatible with economic theory and other scientific knowledge about yield and acreage behavior. The econometric estimation strategies illustrated in the present article can be useful for this. The resulting models should produce adequate but not very precise predictions and, because their random components (i.e., error terms) could be misspecified, any confidence intervals for these predictions need to be interpreted with caution.

The next level is striving to also have a proper representation of the random components of the models. Although this is econometrically involved, it might substantially improve forecasting precision and allow for more reliable confidence intervals for the forecasts. In practice, this generally means that the analysts’ conclusions are more likely to be on target. Thus, the extra effort should be justified given the magnitude of the economic implications from the decisions being made on the basis of these forecasts.

References

- Bartlett, M.S. “On the Theoretical Specification of Sampling Properties of Autocorrelated Time Series.” *Journal of the Royal Statistical Society Series B* 27(1946):27–41.
- Box, G.E.P., and D.A. Pierce. “Distribution of Residual Autocorrelations in Autoregressive-Integrated Moving Average Time Series Models.” *Journal of the American Statistical Association* 65(December 1970):1509–26.
- Denning, M. “Producer Supply Response for Cotton in the United States.” M.S. thesis. Texas Tech University, 2002.
- Judge, G.G., W.E. Griffiths, R. Carter Hill, H. Lutkepohl, and T.-C. Lee. *The Theory and Practice of Econometrics*. New York: John Wiley & Sons, 1985.
- Maddala, S.G. *Introduction to Econometrics*, 2nd ed. New York: Macmillan, 1992.

- Ramirez, O.A., and E. Somarriba. "Risk and Returns of Diversified Cropping Systems under Nonnormal, Cross and Autocorrelated Commodity Price Structures." *Journal of Agricultural and Resource Economics* 25(2000):653-68.
- Ramirez, A.O., and M. Fadiga. "Forecasting Agricultural Commodity Prices with Asymmetric-Error GARCH Models." *Journal of Agricultural and Resource Economics* 28(2003):71-85.
- Ramirez, O.A., S. Misra, and J. Field. "Crop Yield Distribution Revisited." *American Journal of Agricultural Economics* 85(2003): 108-20.
- Ramirez, O.A., S. Misra, and J. Nelson. "Efficient Estimation of Agricultural Time Series Models with Nonnormal Dependent Variables." *American Journal of Agricultural Economics* 85(2003):1029-40.

