



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

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

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

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

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

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Quantile Regression Estimates of Confidence Intervals for WASDE Price Forecasts

Olga Isengildina-Massa, Scott H. Irwin, and Darrel L. Good

This study uses quantile regressions to estimate historical forecast error distributions for WASDE forecasts of corn, soybean, and wheat prices, and then compute confidence limits for the forecasts based on the empirical distributions. Quantile regressions with fit errors expressed as a function of forecast lead time are consistent with theoretical forecast variance expressions while avoiding assumptions of normality and optimality. Based on out-of-sample accuracy tests over 1995/96–2006/07, quantile regression methods produced intervals consistent with the target confidence level. Overall, this study demonstrates that empirical approaches may be used to construct accurate confidence intervals for WASDE corn, soybean, and wheat price forecasts.

Key Words: commodity, evaluating forecasts, government forecasting, judgmental forecasting, prediction intervals, price forecasting

Introduction

Price volatility causes many agricultural firms to rely on forecasts when making decisions. As a result, the U.S. Department of Agriculture (USDA) devotes substantial resources to agricultural situation and outlook programs. A prominent example of USDA forecasting efforts is the World Agricultural Supply and Demand Estimates (WASDE) program, which provides monthly forecasts for major crops, both nationally and globally. Unlike all other WASDE estimates, WASDE price forecasts are published in the form of an interval. Interval forecasts, in contrast to point estimates, represent a range of values in which the realized value of the series is expected to fall with some prespecified probability (Diebold, 1998). WASDE price forecasts are generated using a balance sheet approach, with published intervals reflecting uncertainty associated with prices in the future (Vogel and Bange, 1999). For example, the October 2007 WASDE forecast of the 2007/08 marketing year average farm price was \$2.90–\$3.50 per bushel for corn, \$7.85–\$8.85 per bushel for soybeans, and \$5.80–\$6.40 per bushel for wheat. However, the confidence level associated with the published interval is not revealed. One of the challenges in calculating the forecast intervals and specifying an associated confidence level is the fact that these are consensus forecasts and therefore cannot be described by a formal statistical model.¹

Olga Isengildina-Massa is assistant professor, Department of Applied Economics and Statistics, Clemson University. Scott H. Irwin is the Laurence J. Norton Chair of Agricultural Marketing and Darrel L. Good is professor emeritus, both in the Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign. Funding support from the U.S. Department of Agriculture under Cooperative Agreement No. 43-3AEK-5-80076 is gratefully acknowledged. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the view of the U.S. Department of Agriculture. The authors thank the editor, Joseph A. Atwood, and two anonymous reviewers for their comments which helped improve the paper substantially.

Review coordinated by Joseph A. Atwood.

¹ According to Vogel and Bange (1999), “The process of forecasting price and balance sheet items is a complex one involving the interaction of expert judgment, commodity models, and in-depth research by Department analysts on key domestic and international issues” (p. 10).

Bayesian learning models emphasize the importance of information uncertainty in market participants' interpretation and reaction to information releases (Kandel and Pearson, 1995; Hautsch and Hess, 2007). This is consistent with repeated calls in the agricultural economics literature for increased use of interval and probability forecasting (Timm, 1966; Teigen and Bell, 1978; Bessler and Kling, 1989). However, application and analysis of interval and probability forecasts have received relatively little attention. Sanders and Manfredo (2003) examined one-quarter-ahead WASDE interval forecasts of livestock prices from 1982–2002 and found hit rates (the proportion of time the interval contains the subsequent actual price) ranging from 48% for broilers to only 35% for hogs. Isengildina, Irwin, and Good (2004) showed that monthly WASDE interval forecasts of corn and soybean prices during the 1980/81–2001/02 marketing years had hit rates ranging from 36% to 82% for corn and from 59% to 89% for soybeans, depending on the forecast month. According to the authors, these hit rates were low compared to target confidence levels—about 80% pre-harvest and 90% post-harvest—as revealed by a survey of USDA forecast analysts conducted in their study (p. 994).

While numerous procedures have been proposed to calculate confidence limits generated by statistical forecasting models (e.g., Chatfield, 1993; Prescott and Stengos, 1987; Bessler and Kling, 1989), these procedures provide little guidance for forecasts based on a combination or consensus process rather than formal models, as is the case with WASDE forecasts. In reviewing the prediction interval literature, Chatfield (1993) observes that the use of empirically based methods should be considered as a general-purpose alternative when theoretical formulae are not available or there are doubts about model assumptions. Chatfield also notes the empirical method "... is attractive in principle, however, it seems to have been little used in practice, presumably because of the heavy computational demands" (p. 127). He suggests that since computational demands have become much less of a burden, this method should be reexamined.

Empirical methods are based on the notion that confidence limits for future forecasts may be estimated by evaluating historical forecast errors. Williams and Goodman (1971) were the first to apply an empirical method to constructing confidence limits for economic forecasts. Their approach consisted of splitting the data into two parts—a training set and a validation set. The economic model was then estimated using the training set, and confidence limits were computed for the validation set. The model was reestimated each year by adding an additional observation to the training set in order to compute the confidence limits for the next year's forecast in the validation set. The key assumption of this method is that future forecast errors belong to approximately the same distribution as past forecast errors.² As argued by Williams and Goodman, this assumption is less restrictive than the standard assumption that a forecasting model describes the series adequately in the future. Therefore, by accumulating forecast errors through time, one can obtain an empirical distribution of forecast errors to estimate confidence limits for future forecasts. The benefit of this method is that it can be applied in a straightforward manner to any type of error distribution, including fat-tailed and/or asymmetric distributions.

Empirical methods of constructing forecast confidence intervals have been used successfully in a variety of fields (Murphy and Winkler, 1977; Stoto, 1983; Keilman, 1990; Zarnowitz, 1992; Shlyakhter et al., 1994; Jorgensen and Sjoberg, 2003). One of the main limitations of empirical methods is the necessity for a reasonably large sample of forecasts to reliably estimate confidence intervals. Consequently, empirical methods have been most widely used

² It is worth noting that most theoretical variance expressions are based on the same assumption.

in areas where data limitations are less common, such as weather, population, and software development forecasting.

Taylor and Bunn (1999a, b) suggested a new approach to empirical interval estimation that addresses the small sample problem by pooling data across time and forecasting horizons and estimating forecast error distributions via quantile regression. The authors developed forecast error quantile models which are functions of lead time, as suggested by theoretically derived variance expressions. The use of quantile regression avoids the normality and optimality assumptions underlying theoretical forecast variance expressions. Further, this approach relaxes the assumption that error distributions for each forecasting month are independent, since forecast errors tend to decline from the beginning to the end of the forecasting cycle as more information becomes available.

The purpose of this paper is to investigate the use of quantile regression for estimating empirical confidence limits for WASDE interval forecasts of corn, soybean, and wheat prices using data from 1980/81–2006/07. Our research results will provide valuable information which can be used to accurately estimate confidence limits for WASDE price interval forecasts as well as other interval forecasts.

Data and Descriptive Analysis


The USDA releases monthly WASDE interval price forecast reports for corn, soybeans, and wheat, usually between the 9th and 12th of the month. As shown in figure 1, the first price forecast for a marketing year is released in May preceding the U.S. marketing year (September through August for corn and soybeans and June through May for wheat). Estimates are typically finalized by September (for wheat) or November (for corn and soybeans) of the following marketing year. Thus, 18 forecast updates of marketing year average corn and soybean prices and 16 forecast updates of marketing year average wheat prices are generated in the WASDE forecasting cycle. According to one of the USDA analysts surveyed by Isengildina, Irwin, and Good (2004), who was involved in compiling WASDE corn and soybean price interval forecasts, the intervals are symmetric as “each month a midpoint is forecast using the U.S. and global supply and use and then a range is put on each side of the midpoint” (p. 997). While the forecasts are published in the form of an interval, the probability with which the realized price is expected to fall within the forecast interval is not specified.

Tables 1–3 present descriptive statistics for WASDE interval price forecasts for corn, soybeans, and wheat, respectively, over the 1980/81–2006/07 marketing years.³ During the study period, the first (May prior to harvest) forecast intervals averaged \$0.39/bushel for corn, \$1.27/bushel for soybeans, and \$0.46/bushel for wheat. In relative terms, May forecast intervals for wheat were the narrowest, representing about 14% of the average forecast price, compared to 18% for corn and 22% for soybeans.⁴ By November after harvest, these average intervals narrowed to \$0.36/bushel for corn, \$0.90/bushel for soybeans, and \$0.25/bushel for wheat. The relative magnitude of post-harvest wheat forecast intervals was about half the size

³ Tables 1 and 2 present descriptive statistics for 17 monthly forecasts of corn and soybean prices, and table 3 presents descriptive statistics for 14 monthly forecasts of wheat prices because the last “forecast” provides the final estimate for each commodity.

⁴ Isengildina, Irwin, and Good (2004) provide survey evidence that WASDE price intervals are symmetric—i.e., a midpoint is forecast and then an equal interval is added to each side of the midpoint. Therefore, the average forecast price is computed in this study by taking an average of the midpoint of forecast prices for each month.

U.S. Marketing Year for Corn and Soybeans



U.S. Marketing Year for Wheat




Figure 1. The 2006/2007 WASDE forecasting cycle for corn, soybeans, and wheat relative to the U.S. marketing year

Table 1. Descriptive and Accuracy Statistics for WASDE Corn Price Interval Forecasts, 1980/81–2006/07 Marketing Years

| Month of Forecasting Cycle (k) | Average Price (\$/bu.) | Average Interval (\$/bu.) | Average Interval (%) | MAE (\$/bu.) | MAPE (%) | Hit Rate (%) |
|--|------------------------|---------------------------|----------------------|--------------|----------|--------------|
| Pre-harvest: | | | | | | |
| 1 May _{t} | 2.29 | 0.39 | 17.55 | 0.32 | 13.34 | 37.04 |
| 2 June _{t} | 2.31 | 0.39 | 17.25 | 0.31 | 13.03 | 29.63 |
| 3 July _{t} | 2.35 | 0.39 | 16.90 | 0.27 | 11.20 | 44.44 |
| 4 August _{t} | 2.39 | 0.38 | 16.38 | 0.23 | 9.33 | 55.56 |
| 5 September _{t} | 2.39 | 0.37 | 15.94 | 0.21 | 8.71 | 55.56 |
| 6 October _{t} | 2.38 | 0.37 | 15.97 | 0.19 | 7.53 | 55.56 |
| Average | 2.35 | 0.38 | 16.67 | 0.25 | 10.52 | 46.30 |
| Post-harvest: | | | | | | |
| 7 November _{t} | 2.38 | 0.36 | 16.01 | 0.14 | 5.77 | 74.07 |
| 8 December _{t} | 2.37 | 0.34 | 15.01 | 0.12 | 4.73 | 81.48 |
| 9 January _{$t+1$} | 2.38 | 0.30 | 12.96 | 0.09 | 3.85 | 85.19 |
| 10 February _{$t+1$} | 2.38 | 0.25 | 11.00 | 0.08 | 3.28 | 81.48 |
| 11 March _{$t+1$} | 2.38 | 0.20 | 8.74 | 0.07 | 2.74 | 74.07 |
| 12 April _{$t+1$} | 2.39 | 0.14 | 6.36 | 0.05 | 2.18 | 74.07 |
| 13 May _{$t+1$} | 2.39 | 0.10 | 4.56 | 0.04 | 1.85 | 70.37 |
| 14 June _{$t+1$} | 2.38 | 0.07 | 3.47 | 0.04 | 1.59 | 70.37 |
| 15 July _{$t+1$} | 2.38 | 0.04 | 1.98 | 0.02 | 0.93 | 74.07 |
| 16 August _{$t+1$} | 2.38 | 0.01 | 0.42 | 0.02 | 0.64 | 44.44 |
| 17 September _{$t+1$} | 2.38 | 0.00 | 0.19 | 0.01 | 0.47 | 48.15 |
| Average | 2.38 | 0.17 | 7.34 | 0.06 | 2.55 | 70.71 |

Notes: Average price is calculated by averaging the midpoints of forecast intervals. Errors are calculated as the difference between the final (November _{$t+1$}) estimate and the midpoint of the forecast interval. Hit rate is the proportion of times the interval contained the final estimate. $N = 27$ marketing years.

of corn and soybean price intervals, with respective November averages for wheat, corn, and soybeans of 7%, 16%, and 15%. These forecast intervals usually collapsed to point estimates in May after harvest for wheat and soybeans, and in August after harvest for corn. Mean absolute errors (MAEs) and mean absolute percent errors (MAPEs) in tables 1–3 represent the average distance of the final estimate from the forecast midpoint for each month. These statistics demonstrate that in May prior to harvest, the magnitude of the average percent error was the lowest for wheat (10.08%), indicating narrower intervals may be required to provide accurate wheat price forecasts relative to soybeans (MAPE = 11.37%) and corn (MAPE = 13.34%). No trends in the magnitude of interval forecasts over time were detected. Thus, intervals in the same months did not become smaller (or larger) from the beginning to the end of the sample period.

Interval forecast accuracy is typically described in terms of the hit rate (i.e., the proportion of time the forecast interval included the final value). As observed in tables 1–3, hit rates for individual months ranged from 30% to 85% for corn, 26% to 81% for soybeans, and 37% to 78% for wheat. Prior to harvest, hit rates for corn and wheat price forecast intervals were lower, both averaging 46%, compared to 67% for soybeans. These findings imply that, on

Table 2. Descriptive and Accuracy Statistics for WASDE Soybean Price Interval Forecasts, 1980/81–2006/07 Marketing Years

| Month of Forecasting Cycle (k) | Average Price (\$/bu.) | Average Interval (\$/bu.) | Average Interval (%) | MAE (\$/bu.) | MAPE (%) | Hit Rate (%) |
|--|------------------------|---------------------------|----------------------|--------------|----------|--------------|
| Pre-harvest: | | | | | | |
| 1 May _{t} | 5.72 | 1.27 | 21.57 | 0.71 | 11.37 | 51.85 |
| 2 June _{t} | 5.73 | 1.22 | 20.86 | 0.68 | 10.76 | 55.56 |
| 3 July _{t} | 5.77 | 1.19 | 20.23 | 0.59 | 9.47 | 66.67 |
| 4 August _{t} | 5.89 | 1.19 | 19.58 | 0.48 | 7.67 | 81.48 |
| 5 September _{t} | 5.96 | 1.07 | 17.46 | 0.44 | 6.85 | 77.78 |
| 6 October _{t} | 5.93 | 0.97 | 15.86 | 0.43 | 6.83 | 70.37 |
| Average | 5.83 | 1.15 | 19.26 | 0.55 | 8.82 | 67.28 |
| Post-harvest: | | | | | | |
| 7 November _{t} | 5.93 | 0.90 | 14.70 | 0.37 | 5.89 | 70.37 |
| 8 December _{t} | 5.94 | 0.79 | 12.91 | 0.28 | 4.59 | 81.48 |
| 9 January _{$t+1$} | 5.92 | 0.68 | 11.34 | 0.23 | 3.74 | 77.78 |
| 10 February _{$t+1$} | 5.91 | 0.59 | 9.79 | 0.19 | 3.13 | 81.48 |
| 11 March _{$t+1$} | 5.89 | 0.44 | 7.36 | 0.12 | 1.97 | 81.48 |
| 12 April _{$t+1$} | 5.91 | 0.26 | 4.40 | 0.10 | 1.62 | 77.78 |
| 13 May _{$t+1$} | 5.93 | 0.00 | 0.00 | 0.07 | 1.21 | 25.93 |
| 14 June _{$t+1$} | 5.94 | 0.00 | 0.00 | 0.06 | 0.94 | 40.74 |
| 15 July _{$t+1$} | 5.95 | 0.00 | 0.00 | 0.03 | 0.55 | 48.15 |
| 16 August _{$t+1$} | 5.95 | 0.00 | 0.00 | 0.02 | 0.27 | 62.96 |
| 17 September _{$t+1$} | 5.95 | 0.00 | 0.00 | 0.01 | 0.14 | 66.67 |
| Average | 5.90 | 0.33 | 5.50 | 0.13 | 2.19 | 64.98 |

Notes: Average price is calculated by averaging the midpoints of forecast intervals. Errors are calculated as the difference between the final (November _{$t+1$}) estimate and the midpoint of the forecast interval. Hit rate is the proportion of times the interval contained the final estimate. $N = 27$ marketing years.

average, corn and wheat price interval forecasts prior to harvest contained the final price estimate only 46% of the time. After harvest, the hit rates for all commodities increased, averaging 71% for corn, 65% for soybeans, and 65% for wheat price interval forecasts. All three commodities demonstrated some very low hit rates late in the forecasting cycle. For example, hit rates for corn price interval forecasts averaged 44% and 48% in August and September after harvest; soybean hit rates averaged 26%, 41%, and 48% from May–July after harvest; and wheat hit rates averaged 41% and 37% in May and June after harvest. This loss in accuracy late in the forecasting cycle is associated with prematurely collapsing forecast intervals to point estimates.

An important basic assumption of the empirical approach to estimating confidence limits is that the distribution of forecast errors remains stable over time. The validity of this assumption for corn, soybean, and wheat price forecast errors is tested by dividing the sample into two parts—from 1980/81–1994/95 and from 1995/96–2006/07—and examining whether the first two moments of forecast error distributions differed between the two subperiods. We conducted analysis for both unit errors and percentage errors. Unit errors were calculated as the difference between the final estimate (November for corn and soybeans, and September

Table 3. Descriptive and Accuracy Statistics for WASDE Wheat Price Interval Forecasts, 1980/81–2006/07 Marketing Years

| Month of Forecasting Cycle (k) | Average Price (\$/bu.) | Average Interval (\$/bu.) | Average Interval (%) | MAE (\$/bu.) | MAPE (%) | Hit Rate (%) |
|---|------------------------|---------------------------|----------------------|--------------|----------|--------------|
| Pre-harvest: | | | | | | |
| 1 May _{t} | 3.31 | 0.46 | 13.98 | 0.35 | 10.08 | 40.74 |
| 2 June _{t} | 3.32 | 0.46 | 13.73 | 0.32 | 9.45 | 37.04 |
| 3 July _{t} | 3.28 | 0.44 | 13.40 | 0.23 | 6.80 | 59.26 |
| Average | 3.30 | 0.45 | 13.70 | 0.30 | 8.78 | 45.68 |
| Post-harvest: | | | | | | |
| 4 August _{t} | 3.30 | 0.43 | 12.97 | 0.18 | 5.27 | 66.67 |
| 5 September _{t} | 3.30 | 0.36 | 10.82 | 0.14 | 4.10 | 74.07 |
| 6 October _{t} | 3.33 | 0.31 | 9.18 | 0.12 | 3.43 | 77.78 |
| 7 November _{t} | 3.34 | 0.25 | 7.49 | 0.10 | 2.84 | 66.67 |
| 8 December _{t} | 3.34 | 0.21 | 6.16 | 0.08 | 2.25 | 70.37 |
| 9 January _{$t+1$} | 3.34 | 0.17 | 4.99 | 0.06 | 1.79 | 70.37 |
| 10 February _{$t+1$} | 3.34 | 0.12 | 3.55 | 0.05 | 1.44 | 70.37 |
| 11 March _{$t+1$} | 3.33 | 0.10 | 3.00 | 0.04 | 1.08 | 77.78 |
| 12 April _{$t+1$} | 3.33 | 0.07 | 2.22 | 0.03 | 0.93 | 66.67 |
| 13 May _{$t+1$} | 3.34 | 0.00 | 0.06 | 0.02 | 0.50 | 40.74 |
| 14 June _{$t+1$} | 3.34 | 0.00 | 0.00 | 0.02 | 0.51 | 37.04 |
| Average | 3.33 | 0.18 | 5.50 | 0.07 | 2.19 | 65.32 |

Notes: Average price is calculated by averaging the midpoints of forecast intervals. Errors are calculated as the difference between the final (September _{$t+1$}) estimate and the midpoint of the forecast interval. Hit rate is the proportion of times the interval contained the final estimate. $N = 27$ marketing years.

for wheat) and the midpoint of the forecast interval. Percentage errors were calculated as the difference between the final estimate (November for corn and soybeans, and September for wheat) and the midpoint of the forecast interval divided by the final estimate.

Based on independent sample t -tests (not shown), there is no consistent evidence of statistically significant differences in the means of errors and percentage errors of price forecasts between the two subperiods (except July through September after harvest in corn). Test results based on Levene's F -statistic (not shown) revealed no statistically significant difference in the variances of errors and percentage errors for each forecasting month between the two subperiods (except September after harvest in corn, and May and June prior to harvest in wheat). This evidence suggests forecast error distributions of monthly WASDE corn, soybean, and wheat price forecasts were generally stable over time. Even though in most cases results were consistent across both types of errors, percentage errors typically demonstrated smaller differences between the two subperiods. The use of percentage errors may be preferred to unit errors when price levels change (as they did for all three commodities after 2006). In this case, intervals based on unit errors will be understated relative to intervals based on percentage errors. We therefore use percentage errors to calculate empirical forecast intervals throughout the remainder of the paper.

Quantile Regression Models

Koenker and Bassett (1978) developed quantile regression as an extension of the linear model for estimating rates of change in not just the mean but all parts of the distribution of a response variable. Consider the simple case of the constant-only model $y_t = \beta_0 + e_t$, where β_0 is a constant parameter and e_t is an i.i.d. random error term. Koenker and Bassett note that the τ th quantile of y_t can be derived from a sample of observations as the solution $\beta_0(\tau)$ to the following minimization problem:

$$(1) \quad \min \beta_0 \left[\sum_{t|y_t \geq \beta_0} \tau |y_t - \beta_0| + \sum_{t|y_t < \beta_0} (1 - \tau) |y_t - \beta_0| \right].$$

As a means for finding the τ th sample quantile, this minimization problem readily extends to the more general case where y_t is a linear function of explanatory variables (\mathbf{X}). The estimates are semi-parametric in the sense that no parametric distributional form is assumed for the random part of the model, although a parametric form is assumed for the deterministic part of the model. The conditional quantiles denoted by $Q_y(\tau | \mathbf{X})$ are the inverse of the conditional cumulative distribution function of the response variable, $F_y^{-1}(\tau | \mathbf{X})$, where $\tau \in [0, 1]$ denotes the quantiles (Koenker and Machado, 1999). As an example, for $\tau = 0.90$, $Q_y(0.90 | \mathbf{X})$ is the 90th percentile of the distribution of y conditional on the values of \mathbf{X} . An approximation of the full probability distribution can be produced from the quantile estimates corresponding to a range of values of τ ($0 < \tau < 1$). For symmetric distributions, the 0.50 quantile (or median) is equal to the mean μ .

Taylor and Bunn (1999a, b) suggested using quantile regressions for generating prediction intervals of forecasts based on exponential smoothing. The authors show that quantile regressions with fit errors expressed as a function of forecast lead time are consistent with theoretical forecast variance formulas. In the present context, quantile regression models based on this approach can be specified for a given commodity as follows:

$$(2) \quad Q_{tk}(\tau) = \beta_0 + \beta_1 k_t + \beta_2 k_t^2 + \varepsilon_{tk},$$

where $Q_{tk}(\tau)$ is the τ th conditional quantile of the distribution of WASDE forecast errors in marketing year t and forecast month k , and k is indexed from 1–16 for corn and soybeans and 1–14 for wheat (see tables 1–3).⁵ These quantile regression estimates can be used to compute confidence limits for commodity price forecasts for marketing year t and forecast month k as follows. First, we select desired quantiles based on the target confidence level; i.e., for an 80% confidence level, $\tau = 0.1$ and $\tau = 0.9$ should be used for the lower and upper bounds, respectively. Second, we estimate parameters for the selected quantiles for equation (2) with historical forecast errors through $t - 1$ (training set) as the dependent variable and forecast month and forecast month squared as independent variables. Third, we plug in specific k in the estimated equation to compute the desired quantiles, $Q_{tk}(\tau)$, for each forecast month k . The computed quantiles indicate the distance of the lower and upper bounds of the confidence interval from forecast midpoint and should be added to the midpoint to construct such an interval.⁶

⁵ The last two months (17 and 18 for corn and soybeans, and 15 and 16 for wheat) were not included in the analysis because the errors were usually zero, so the distributions were impossible to estimate.

⁶ A specific example is discussed later, in the “Accuracy Evaluation” subsection.

Both heteroskedasticity and autocorrelation are likely to be estimation issues in the quantile regressions specified in equation (2). Heteroskedasticity is likely to be present since the variance of forecast errors is decreasing over the forecasting cycle. Autocorrelation is likely to be present due to the overlapping nature of forecast horizons for a given marketing year. Five covariance estimators were considered to test the sensitivity of estimation results to alternative assumptions about the error distributions: Huber sandwich (valid for independent but non-identical errors), residual bootstrap, XY -pair bootstrap (valid when errors and explanatory variables are not independent), Markov chain marginal bootstrap, and modified Markov chain marginal bootstrap (robust to heteroskedasticity). While there was little difference in the quantile estimation results across the covariance estimators, we present results using the XY -pair bootstrap with 100 replications, as it was the most conservative method (with the largest standard errors for hypothesis testing).⁷

A benefit of the quantile regression approach is that other factors impacting forecast error distributions may be included in the analysis. Economic theory indicates the size of the forecast error in each marketing year may be related to the “tightness” of underlying supply and demand conditions. These supply and demand conditions are often summarized by the stocks/use ratio (e.g., Westcott and Hoffman, 1999). For instance, historical stocks/use ratio estimates during the period of study for corn ranged from 5% in 1995 to 66% in 1985. It is reasonable to hypothesize that forecast errors are larger during periods of low stocks/use ratios and vice versa.⁸ The expanded quantile regression model is specified as:

$$(3) \quad Q_{tk}(\tau) = \beta_0 + \beta_1 k_t + \beta_2 k_t^2 + \beta_3 SU_{tk} + \varepsilon_{tk},$$

where SU_{tk} is the WASDE stocks/use estimate for marketing year t and forecast month k .

Detailed estimation results are presented only for particular quantiles in order to conserve space.⁹ As noted earlier, confidence levels associated with WASDE interval price forecasts are not published. Isengildina, Irwin, and Good (2004) conducted a survey of USDA analysts involved in compiling WASDE corn and soybean price interval forecasts to determine confidence levels associated with the forecasts. Analyst responses were variable across respondents (by as much as 30% in the beginning of the season) and over the forecasting cycle (from 65% in May prior to harvest to 95% in April after harvest). The average confidence level reported by USDA analysts prior to harvest was 81% for corn and 78% for soybeans; the average confidence level after harvest was 91% for corn and 87% for soybeans. Based on this information, and assuming USDA wheat analysts provide interval forecasts for similar confidence levels, $\tau = 0.10$ and $\tau = 0.90$ quantiles are estimated prior to harvest and $\tau = 0.05$ and $\tau = 0.95$ quantiles after harvest. These quantile estimates are then used to generate upper and lower bounds of 80% and 90% confidence intervals pre- and post-harvest, respectively, for each commodity. All quantile regressions were estimated using *Eviews 6.0* econometric software.

⁷ Note that only quantile regression parameter estimates are used for computing empirical confidence intervals; thus, unbiasedness of estimated coefficients is the most important assumption for this approach. Koenker (2005, p. 74) argues that parameter estimates of the quantile regression are unbiased even in cases when the i.i.d. assumption is not satisfied. Therefore, the potential impact of heteroskedasticity and autocorrelation in the errors will likely be limited to efficiency issues for hypothesis testing.

⁸ Other model specifications were also explored. The price level and a measure of volatility, calculated as an absolute value of percentage difference of the forecasted price from the average of the previous five years' prices, were considered. However, these alternative specifications failed to improve model performance. The results for the additional alternative specifications are not presented here but are available from the authors upon request.

⁹ The complete set of quantile estimation results is available from the authors upon request.

Table 4. Quantile Regression Estimation Results for Corn, Soybeans, and Wheat WASDE Price Forecast Errors, Forecast Month Only, 1980/81–2006/07 Marketing Years

| Commodity/Estimate | Quantile | | | |
|--------------------|------------|------------|------------|------------|
| | 0.05 | 0.10 | 0.90 | 0.95 |
| Corn: | | | | |
| Constant | −0.318*** | −0.267*** | 0.236*** | 0.249*** |
| k | 0.035*** | 0.030*** | −0.025*** | −0.019*** |
| k^2 | −0.001*** | −0.001*** | 0.001*** | 0.000 |
| Pseudo- R^2 | 0.357 | 0.260 | 0.285 | 0.300 |
| Quasi-LR Statistic | 329.735*** | 238.440*** | 257.263*** | 183.603*** |
| Soybeans: | | | | |
| Constant | −0.307*** | −0.204*** | 0.225*** | 0.290*** |
| k | 0.036*** | 0.023*** | −0.022*** | −0.025*** |
| k^2 | −0.001*** | −0.001*** | 0.001*** | 0.000 |
| Pseudo- R^2 | 0.302 | 0.235 | 0.275 | 0.316 |
| Quasi-LR Statistic | 253.164*** | 236.240*** | 266.222*** | 225.015*** |
| Wheat: | | | | |
| Constant | −0.218*** | −0.174*** | 0.226*** | 0.248*** |
| k | 0.032*** | 0.025*** | −0.033*** | −0.036*** |
| k^2 | −0.001*** | −0.001*** | 0.001*** | 0.001*** |
| Pseudo- R^2 | 0.420 | 0.316 | 0.318 | 0.438 |
| Quasi-LR Statistic | 327.255*** | 248.856*** | 229.637*** | 331.868*** |

Notes: Single, double, and triple asterisks (*, **, ***) denote statistical significance at the 10%, 5%, and 1% levels, respectively. The month of the forecasting cycle is denoted by k . The number of observations for corn and soybeans is 459, and for wheat is 378.

In-Sample Results

Table 4 presents quantile regression estimation results for models based only on forecast month for corn, soybeans, and wheat over the 1980/81–2006/07 marketing years. The quantile regression approach offers the benefit of pooling data across months and years, thereby substantially increasing the statistical power of the empirical approach to forecast interval estimation. Specifically, quantile regressions estimated over the 1980/81–2006/07 marketing years use 459 observations for corn and soybeans, and 378 observations for wheat, while standard empirical methods would use only 27 observations (one per marketing year) to estimate distributions of forecast errors. All but two of the estimated coefficients (k^2 for the 0.95 quantile in corn and soybeans) are significant at conventional levels. The constants for different quantiles describe the shape of the forecast error distribution by indicating the distance from the forecast midpoint to the respective quantile; thus, the sign is negative for $\tau < 0.5$ and positive for $\tau > 0.5$.

The coefficient on the forecast month k in table 4 describes an inverse relationship between forecast month and forecast error; i.e., forecast errors become smaller as more information becomes available over the forecasting cycle (k becomes larger). The sign is therefore positive for $\tau < 0.5$ and negative for $\tau > 0.5$. The coefficient on k^2 describes the nonlinearity in the relationship between forecast error and forecast month k . The negative sign for $\tau < 0.5$

and positive sign for $\tau > 0.5$ of the k^2 coefficient indicate that forecast errors are decreasing at an increasing rate as forecast month k increases. This pattern is illustrated in figure 2, which shows forecast errors for all forecast months over the 1980/81–2006/07 marketing years as well as quantile regression estimates for the 0.10 and 0.90 quantiles based on the same data. Figure 2 demonstrates how forecast errors become smaller over the forecasting cycle in a nonlinear fashion and the ability of quantile regression to capture this pattern in its estimated coefficients.

Table 4 indicates that using only forecast month as an explanatory variable explains from 25% to 34% of the variation in forecast errors at the identified quantiles in corn and soybeans and from 33% to 49% of the variation in wheat. As shown by the quasi-likelihood ratio statistics, the explanatory power of all estimated models is statistically significant.

Figures 3–5 present estimated coefficients and 95% confidence bounds for the regression model [equation (2)] across the full range of quantiles ($\tau \in [0, 1]$). These figures demonstrate the different impact of the quantile regression variables on various quantiles of the forecast error distribution. Thus, we observe the positive relationship with the constant and k^2 , and the negative relationship with k . The impact of k and k^2 is different for all quantiles except for the upper quantiles in corn and wheat, where the 0.80 and 0.90 quantiles are affected in about the same way. The fact that the magnitude of the coefficients on k and k^2 is greater for the tails and zero for the median suggests changes in the variability but not the expected value of the error as the forecasts move through the forecasting cycle. This may be tested formally by comparing estimated slopes at different points of the error distributions. Wald tests (Koenker and Bassett, 1982) at the 0.05, 0.10, 0.90, and 0.95 quantile levels reject the null hypothesis of slope equality, so conditional quantiles are not identical.

In addition to the main set of results discussed above, this study investigated the impact of stocks/use ratios on forecast error distribution as described in equation (3). It was hypothesized that during the periods of low stocks/use ratios, which reflect underlying supply and demand conditions, forecast errors may be larger than during periods of high stocks/use ratios. However, these alternative estimation results (not shown) reveal that the stocks/use variable has very little impact on the forecast error distributions. For this reason, we conduct out-of-sample evaluation in the next section for quantile regression models which are functions of forecast lead time only.

Out-of-Sample Results

In-sample results suggest quantile regression is a useful tool for generating empirical confidence intervals for WASDE corn, soybean, and wheat price forecasts. In order to rigorously assess the potential of the quantile regression approach to improve upon published WASDE price forecasts, we evaluate out-of-sample performance. The first 15 observations (1980/81–1994/95) from the training set were used to generate confidence limits for the 16th year (1995/96), or the first observation of the validation set; the first 16 observations were used to generate confidence limits for the 17th year (1996/97), and so on. The training set consists of 15 (1980/81–1994/95) through 26 (1980/81–2005/06) observations, and the validation set consists of 12 observations (1995/96–2006/07). The target confidence level prior to harvest is 80% and after harvest is 90%.

Empirical confidence intervals are constructed using the quantile regression estimates based on all previous forecast errors. For example, the out-of-sample confidence intervals for 2007/08 forecasts can be constructed using the estimates based on the 1980/81–2006/07 marketing years

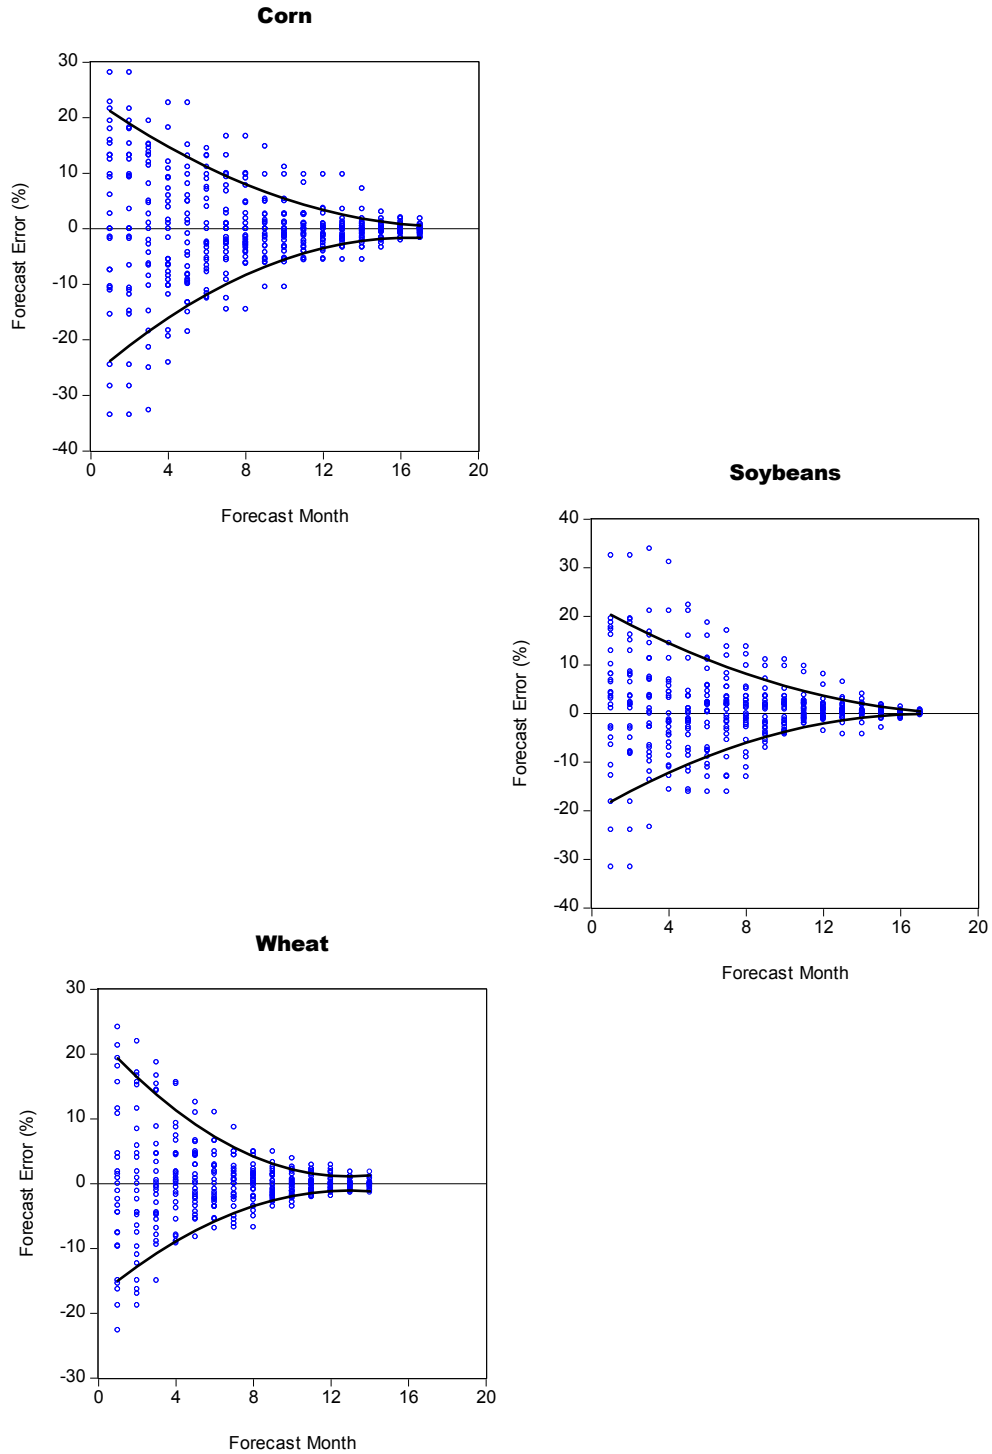


Figure 2. Errors by forecast month and estimated 0.10 and 0.90 quantiles for WASDE corn, soybean, and wheat price forecasts, 1980/81–2006/07 marketing years




Figure 3. Quantile regression coefficient estimates and 95% confidence bounds for WASDE corn price forecast errors, 1980/81–2006/07 marketing years

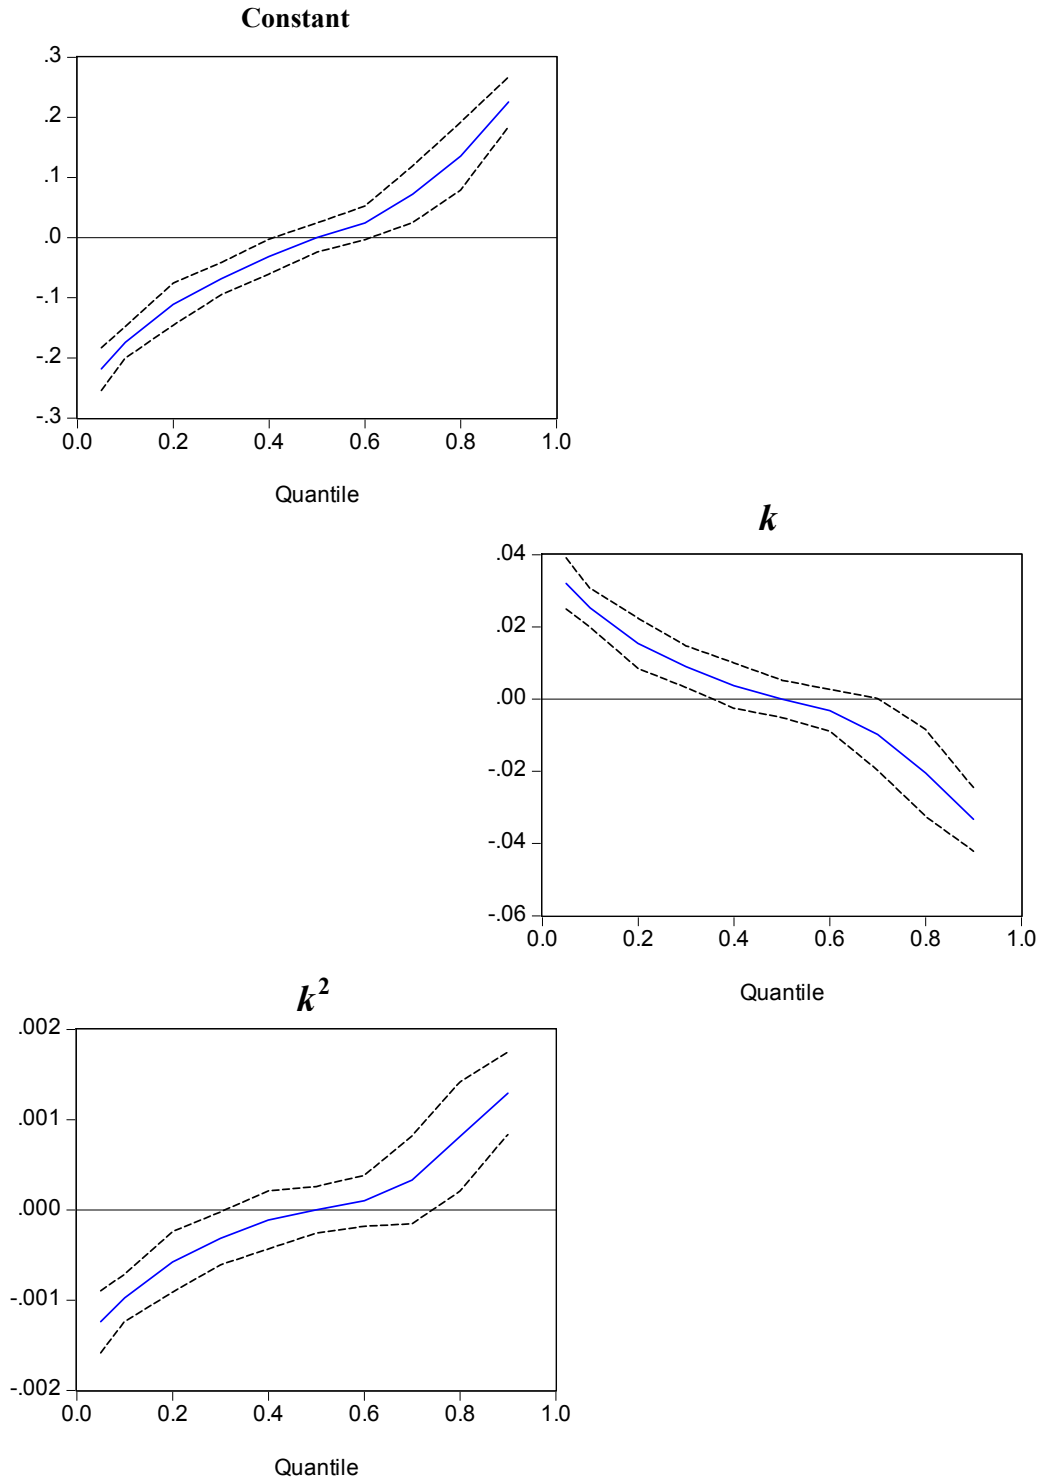


Figure 4. Quantile regression coefficient estimates and 95% confidence bounds for WASDE soybean price forecast errors, 1980/81–2006/07 marketing years




Figure 5. Quantile regression coefficient estimates and 95% confidence bounds for WASDE wheat price forecast errors, 1980/81–2006/07 marketing years

presented in table 4. The 80% confidence limits for May prior to harvest ($k = 1$) corn price forecasts are $\{-0.267 + 0.030 \times 1 - 0.001 \times 1^2 = -0.24\}$ for $\tau = 0.10$, and $\{0.236 - 0.025 \times 1 + 0.001 \times 1^2 = 0.21\}$ for $\tau = 0.90$. Based on these results, 24% of the forecast midpoint should be subtracted and 21% of the midpoint should be added to the midpoint to construct an 80% confidence interval. For a \$3.40/bushel average price, the estimated interval would be \$2.59 to \$4.12/bushel. These confidence intervals were constructed for each forecast month of the validation set, resulting in 576 out-of-sample forecast intervals. The performance of these forecasts was compared to published WASDE forecast intervals in terms of accuracy and informativeness. Several measures of forecast accuracy are used to describe whether the interval contained the final estimate at the desired level of confidence. Forecast informativeness is evaluated based on the tradeoff between accuracy and interval width.

Accuracy Evaluation

Hit rates describe the proportion of times forecast intervals contain the final or “true” value (y_t) and may be defined as an indicator variable (I_t^k) such that:

$$(4) \quad I_t^k = \begin{cases} 1 & \text{if } y_t \in [l_{t/k}(\alpha), u_{t/k}(\alpha)], \\ 0 & \text{if } y_t \notin [l_{t/k}(\alpha), u_{t/k}(\alpha)], \end{cases}$$

where $[l_{t/k}(\alpha), u_{t/k}(\alpha)]$ are the lower and upper limits of the interval forecast for y_t made at time k with confidence level α . The closer the hit rate to the stated confidence level, the more accurate the forecast. Forecast coverage is based on the expectation of the indicator variable (I_t^k), and examines whether the proportion of times the forecast interval includes the true value is equal to the target (stated) confidence level. Thus, forecast coverage may be examined by testing the hypothesis $H_0: E(I_t^k) = \alpha$ against $H_A: E(I_t^k) \neq \alpha$. If H_0 is not rejected and the interval hit rate is equal to the stated confidence level, forecasts are said to be calibrated. Since the indicator variable I_t^k has a binomial distribution (Christoffersen, 1998), the likelihood function under the null hypothesis H_0 is:

$$(5) \quad L(\alpha) = (1 - \alpha)^{n_0} \alpha^{n_1},$$

where L is a likelihood function. Under the alternative hypothesis H_A , the likelihood function is denoted by:

$$(6) \quad L(p) = (1 - p)^{n_0} p^{n_1},$$

where n_1 and n_0 are the number of times an interval was “hit” (1) or “missed” (0) in the indicator sequence I_t^k . Then, forecast coverage may be tested via the likelihood-ratio test:

$$(7) \quad LR_c = -2 \ln \left(\frac{L(\alpha)}{L(\hat{p})} \right) \xrightarrow{asy} \chi^2(1),$$

where $\hat{p} = n_1 / (n_0 + n_1)$ is the maximum-likelihood estimator of p . This test is described by Christoffersen (1998) as an unconditional coverage test.¹⁰

¹⁰ Christoffersen (1998) also proposed additional tests that examine interval forecast independence and forecast coverage conditional on independence. However, due to a small number of observations, these tests cannot be applied reliably to the prediction part of our sample (1995/96–2004/05).

Table 5. Out-of-Sample Accuracy Statistics for Empirical versus Published Confidence Intervals for WASDE Corn Price Forecasts, 1995/96–2006/07 Marketing Years

| Month of Forecasting Cycle (<i>k</i>) | Published Intervals | | | Quantile Regression Intervals | | | | |
|---|---------------------|---------------------------|-----------------------------|-------------------------------|---------------------------|------------------------|------------------------|-----------------------------|
| | Hit Rate (%) | Average Interval (\$/bu.) | Unconditional Coverage Test | Hit Rate (%) | Average Interval (\$/bu.) | Average Below (\$/bu.) | Average Above (\$/bu.) | Unconditional Coverage Test |
| Pre-harvest: | | | | | | | | |
| 1 May _{<i>t</i>} | 42 | 0.40 | 8.46*** | 75 | 1.03 | -0.48 | 0.55 | 0.18 |
| 2 June _{<i>t</i>} | 33 | 0.40 | 12.26*** | 75 | 0.92 | -0.44 | 0.49 | 0.18 |
| 3 July _{<i>t</i>} | 50 | 0.40 | 5.36** | 83 | 0.82 | -0.39 | 0.43 | 0.09 |
| 4 August _{<i>t</i>} | 58 | 0.40 | 2.92* | 75 | 0.72 | -0.35 | 0.37 | 0.18 |
| 5 September _{<i>t</i>} | 67 | 0.40 | 1.17 | 75 | 0.63 | -0.31 | 0.32 | 0.18 |
| 6 October _{<i>t</i>} | 58 | 0.40 | 2.92* | 83 | 0.54 | -0.27 | 0.27 | 0.09 |
| Average | 51 | 0.40 | | 78 | 0.78 | -0.37 | 0.41 | |
| Post-harvest: | | | | | | | | |
| 7 November _{<i>t</i>} | 92 | 0.40 | 0.04 | 92 | 0.63 | -0.31 | 0.32 | 0.04 |
| 8 December _{<i>t</i>} | 92 | 0.38 | 0.04 | 100 | 0.54 | -0.27 | 0.27 | N/A |
| 9 January _{<i>t+1</i>} | 100 | 0.34 | N/A | 100 | 0.46 | -0.23 | 0.23 | N/A |
| 10 February _{<i>t+1</i>} | 92 | 0.28 | 0.04 | 100 | 0.38 | -0.19 | 0.19 | N/A |
| 11 March _{<i>t+1</i>} | 92 | 0.22 | 0.04 | 100 | 0.32 | -0.16 | 0.16 | N/A |
| 12 April _{<i>t+1</i>} | 67 | 0.13 | 4.83** | 83 | 0.26 | -0.13 | 0.13 | 0.50 |
| 13 May _{<i>t+1</i>} | 75 | 0.11 | 2.22 | 92 | 0.21 | -0.11 | 0.10 | 0.04 |
| 14 June _{<i>t+1</i>} | 75 | 0.09 | 2.22 | 83 | 0.16 | -0.08 | 0.08 | 0.50 |
| 15 July _{<i>t+1</i>} | 83 | 0.05 | 0.50 | 83 | 0.12 | -0.07 | 0.06 | 0.50 |
| 16 August _{<i>t+1</i>} | 42 | 0.00 | 16.99*** | 92 | 0.09 | -0.05 | 0.04 | 0.04 |
| 17 September _{<i>t+1</i>} | 58 | 0.00 | 8.20*** | 100 | 0.07 | -0.04 | 0.03 | N/A |
| Average | 79 | 0.18 | | 93 | 0.29 | -0.15 | 0.15 | |

Notes: Single, double, and triple asterisks (*, **, ***) denote statistical significance at the 10%, 5%, and 1% levels, respectively. Published WASDE intervals are symmetric, so the average below and the average above are equal to one-half of the average interval. Empirical price forecast intervals were calculated using percentage errors from the 1980/81 marketing year forward. Since pre-harvest months are characterized by a much greater degree of uncertainty associated with production of the commodity relative to the post-harvest months, target confidence level is 80% prior to harvest and 90% after harvest.

Results of the accuracy tests for out-of-sample forecast intervals computed using quantile regression are reported in tables 5–7. As was observed in tables 1–3 for the entire sample, published WASDE forecasts had relatively low hit rates in the prediction subsample, 1995/96–2006/07, although significant improvement in forecast accuracy was observed in corn price forecast intervals after harvest. The hit rates for published intervals averaged 51% for corn, 65% for soybeans, and 44% for wheat prior to harvest. Empirical confidence intervals had much higher hit rates, averaging 78% for corn, 75% for soybeans, and 58% for wheat prior to harvest. These hit rates were statistically different from the target confidence level of 80% in 2 out of 15 cases, or about 13% of the time. For comparison, published intervals' hit rates were statistically different from the assumed target level in 9 out of 15 cases, or 60% of the time.

After harvest, the hit rates for published intervals averaged 79% for corn, 56% for soybeans, and 71% for wheat. After-harvest hit rates for empirical confidence intervals averaged

Table 6. Out-of-Sample Accuracy Statistics for Empirical versus Published Confidence Intervals for WASDE Soybean Price Forecasts, 1995/96–2006/07 Marketing Years

| Month of Forecasting Cycle (k) | Published Intervals | | | Quantile Regression Intervals | | | | |
|--|---------------------|---------------------------|-----------------------------|-------------------------------|---------------------------|------------------------|------------------------|-----------------------------|
| | Hit Rate (%) | Average Interval (\$/bu.) | Unconditional Coverage Test | Hit Rate (%) | Average Interval (\$/bu.) | Average Below (\$/bu.) | Average Above (\$/bu.) | Unconditional Coverage Test |
| Pre-harvest: | | | | | | | | |
| 1 May _{t} | 58 | 1.09 | 2.92* | 75 | 2.55 | -1.32 | 1.23 | 0.18 |
| 2 June _{t} | 67 | 1.08 | 1.17 | 75 | 2.25 | -1.16 | 1.09 | 0.18 |
| 3 July _{t} | 67 | 1.05 | 1.17 | 75 | 1.97 | -1.01 | 0.96 | 0.18 |
| 4 August _{t} | 75 | 1.03 | 0.18 | 75 | 1.71 | -0.87 | 0.84 | 0.18 |
| 5 September _{t} | 67 | 0.92 | 1.17 | 75 | 1.47 | -0.75 | 0.73 | 0.18 |
| 6 October _{t} | 58 | 0.83 | 2.92* | 75 | 1.25 | -0.63 | 0.62 | 0.18 |
| Average | 65 | 1.00 | | 75 | 1.87 | -0.96 | 0.91 | |
| Post-harvest: | | | | | | | | |
| 7 November _{t} | 67 | 0.81 | 4.83** | 83 | 1.56 | -0.76 | 0.81 | 0.50 |
| 8 December _{t} | 75 | 0.73 | 2.22 | 92 | 1.32 | -0.62 | 0.70 | 0.04 |
| 9 January _{$t+1$} | 75 | 0.68 | 2.22 | 92 | 1.10 | -0.50 | 0.60 | 0.04 |
| 10 February _{$t+1$} | 83 | 0.59 | 0.50 | 92 | 0.90 | -0.39 | 0.51 | 0.04 |
| 11 March _{$t+1$} | 83 | 0.40 | 0.50 | 92 | 0.72 | -0.30 | 0.43 | 0.04 |
| 12 April _{$t+1$} | 83 | 0.31 | 0.50 | 83 | 0.57 | -0.22 | 0.35 | 0.50 |
| 13 May _{$t+1$} | 17 | 0.00 | 35.66*** | 92 | 0.43 | -0.15 | 0.28 | 0.04 |
| 14 June _{$t+1$} | 25 | 0.00 | 28.58*** | 92 | 0.31 | -0.09 | 0.21 | 0.04 |
| 15 July _{$t+1$} | 25 | 0.00 | 28.58*** | 75 | 0.21 | -0.05 | 0.16 | 2.22 |
| 16 August _{$t+1$} | 42 | 0.00 | 16.99*** | 67 | 0.13 | -0.02 | 0.11 | 4.83** |
| 17 September _{$t+1$} | 42 | 0.00 | 16.99*** | 75 | 0.07 | 0.00 | 0.06 | 2.22 |
| Average | 56 | 0.32 | | 85 | 0.67 | -0.28 | 0.38 | |

Notes: Single, double, and triple asterisks (*, **, ***) denote statistical significance at the 10%, 5%, and 1% levels, respectively. Published WASDE intervals are symmetric, so the average below and the average above are equal to one-half of the average interval. Empirical price forecast intervals were calculated using percentage errors from the 1980/81 marketing year forward. Since pre-harvest months are characterized by a much greater degree of uncertainty associated with production of the commodity relative to the post-harvest months, target confidence level is 80% prior to harvest and 90% after harvest.

93% for corn, 85% for soybeans, and 94% for wheat. These hit rates were statistically different from the target level of 90% in 1 out of 33 cases, or about 3% of the time. Published confidence intervals' hit rates were statistically different from the assumed target level in 12 out of 33 cases, or 40% of the time. Overall, these results demonstrate a dramatic improvement in accuracy for empirical confidence intervals relative to published intervals.

Informativeness Evaluation

The results presented in the previous section concentrated only on the issue of interval accuracy as described by the hit rates. While obviously important, interval accuracy may not be the only argument in a forecast user's utility function. For example, tables 5–7 demonstrate that while quantile regression intervals are more accurate, they are also much wider. If the forecaster's objective is to provide the most accurate interval that also has the narrowest

Table 7. Out-of-Sample Accuracy Statistics for Empirical versus Published Confidence Intervals for WASDE Wheat Price Forecasts, 1995/96–2006/07 Marketing Years

| Month of Forecasting Cycle (<i>k</i>) | Published Intervals | | | Quantile Regression Intervals | | | | |
|---|---------------------|---------------------------|-----------------------------|-------------------------------|---------------------------|------------------------|------------------------|-----------------------------|
| | Hit Rate (%) | Average Interval (\$/bu.) | Unconditional Coverage Test | Hit Rate (%) | Average Interval (\$/bu.) | Average Below (\$/bu.) | Average Above (\$/bu.) | Unconditional Coverage Test |
| Pre-harvest: | | | | | | | | |
| 1 May _{<i>t</i>} | 33 | 0.54 | 12.26*** | 50 | 1.09 | -0.47 | 0.62 | 5.36** |
| 2 June _{<i>t</i>} | 33 | 0.54 | 12.26*** | 50 | 0.95 | -0.41 | 0.54 | 5.36** |
| 3 July _{<i>t</i>} | 67 | 0.54 | 1.17 | 75 | 0.82 | -0.35 | 0.46 | 0.18 |
| Average | 44 | 0.54 | | 58 | 0.95 | -0.41 | 0.54 | |
| Post-harvest: | | | | | | | | |
| 4 August _{<i>t</i>} | 75 | 0.54 | 2.22 | 92 | 0.86 | -0.35 | 0.51 | 0.04 |
| 5 September _{<i>t</i>} | 83 | 0.44 | 0.50 | 92 | 0.72 | -0.30 | 0.42 | 0.04 |
| 6 October _{<i>t</i>} | 92 | 0.38 | 0.04 | 100 | 0.59 | -0.24 | 0.34 | N/A |
| 7 November _{<i>t</i>} | 75 | 0.29 | 2.22 | 92 | 0.47 | -0.20 | 0.27 | 0.04 |
| 8 December _{<i>t</i>} | 75 | 0.23 | 2.22 | 83 | 0.37 | -0.16 | 0.21 | 0.50 |
| 9 January _{<i>t+1</i>} | 75 | 0.19 | 2.22 | 100 | 0.29 | -0.12 | 0.16 | N/A |
| 10 February _{<i>t+1</i>} | 75 | 0.11 | 2.22 | 92 | 0.22 | -0.10 | 0.12 | 0.04 |
| 11 March _{<i>t+1</i>} | 83 | 0.10 | 0.50 | 92 | 0.16 | -0.07 | 0.09 | 0.04 |
| 12 April _{<i>t+1</i>} | 67 | 0.07 | 4.83** | 92 | 0.12 | -0.06 | 0.06 | 0.04 |
| 13 May _{<i>t+1</i>} | 42 | 0.00 | 16.99*** | 100 | 0.09 | -0.04 | 0.05 | N/A |
| 14 June _{<i>t+1</i>} | 42 | 0.00 | 16.99*** | 100 | 0.08 | -0.04 | 0.04 | N/A |
| Average | 71 | 0.21 | | 94 | 0.36 | -0.15 | 0.21 | |

Notes: Single, double, and triple asterisks (*, **, ***) denote statistical significance at the 10%, 5%, and 1% levels, respectively. Published WASDE intervals are symmetric, so the average below and the average above are equal to one-half of the average interval. Empirical price forecast intervals were calculated using percentage errors from the 1980/81 marketing year forward. Since pre-harvest months are characterized by a much greater degree of uncertainty associated with production of the commodity relative to the post-harvest months, target confidence level is 80% prior to harvest and 90% after harvest.

width, both accuracy (measured by hit rates) and informativeness (interval width) should be taken into account. Thus, informativeness may be evaluated by comparing the width of intervals that have the same accuracy or by comparing the accuracy of forecasts that have the same width.

Ideally, informativeness should be evaluated for competing interval forecasts across a range of target confidence levels. This is not possible in the present context because WASDE forecast intervals are not specified for different target confidence intervals. To provide some evidence in this regard, it is assumed the observed hit rates for WASDE forecasts in the out-of-sample period (reported in tables 5–7) are equal to the target confidence levels of USDA analysts: 51% pre-harvest and 79% post-harvest for corn, 65% pre-harvest and 56% post-harvest for soybeans, and 44% pre-harvest and 71% post-harvest for wheat.¹¹ Quantile-based empirical intervals for these target confidence levels are then computed using the methods outlined earlier. Since both the WASDE and empirical intervals are assumed to have the same confidence level (accuracy), informativeness can be evaluated by comparing the width of intervals.

¹¹ Note that this assumption is made for the purpose of this illustration only, as the survey evidence indicates the target confidence levels were around 80% pre-harvest and 90% post-harvest.




Figure 6. Accuracy-informativeness tradeoff of published vs. empirical confidence intervals computed using quantile regression for WASDE price forecasts, 1995/96–2006/07 marketing years

Figure 6 shows the average accuracy (hit rates) and informativeness (interval width) of the computed empirical confidence intervals and WASDE forecast intervals for all three commodities from 1995/96–2006/07. Presenting the comparisons in this format allows a type of dominance test. Specifically, a forecast user would prefer a forecast having a higher average hit rate and a narrower average interval. The only case of dominance is pre-harvest corn, where WASDE forecast intervals have a higher average hit rate and slightly smaller intervals. In the remaining cases, higher average hit rates are accompanied by wider intervals. At the same time, there is an interesting pattern in the results. WASDE forecast intervals during pre-harvest tended to be more informative than empirical forecast intervals based on quantile regression since WASDE produced equally accurate forecasts with smaller intervals in soybeans and more accurate forecasts with equally wide intervals in corn and wheat. After harvest, the findings are just the opposite, with quantile regression tending to outperform USDA forecast intervals for all commodities—i.e., quantile regressions produced results with about the same accuracy with smaller intervals in soybeans and more accurate results with about the same intervals in corn and wheat.

The lack of dominance of one method over the other suggests the answer as to which forecast is “best” will depend on the tradeoff between accuracy and informativeness in a forecast user’s utility function. If some forecast users put a heavy emphasis on accuracy, they would prefer more accuracy and less informativeness. In contrast, some forecast users may prefer more informativeness and less accuracy. Some earlier attempts to evaluate this tradeoff were made by Yaniv and Foster (1995) and applied by Isengildina, Irwin, and Good (2004). However, their results are based on university students’ preferences, which may not adequately reflect WASDE forecast users’ preferences. Finally, even if the empirical approach is not dominant at the confidence levels implied by historical WASDE hit rates, the approach still may be valuable to forecast users who want to compute intervals for different confidence levels or who believe WASDE intervals are too narrow compared to target confidence levels, as some evidence suggests.

Summary and Conclusions

A prominent example of USDA forecasting efforts is the World Agricultural Supply and Demand Estimates (WASDE) program, which provides monthly forecasts for major crops, both for the United States and the world. WASDE price forecasts (unlike all other WASDE estimates) are published in the form of an interval, reflecting uncertainty associated with future prices. However, the confidence level associated with the published interval is not revealed. One of the challenges in calculating WASDE price forecast intervals and specifying an associated confidence level is due to the fact that these are consensus forecasts, which cannot be described by a formal statistical model. Such forecasts cannot use the confidence interval formulas derived for statistical models, but may instead rely on empirically based methods.

The basic empirical method was first introduced by Williams and Goodman (1971) and is based on the notion that by accumulating forecast errors through time, one can obtain an empirical distribution of forecast errors. One of the main limitations of the empirical method is the heavy data requirement for forecast error distribution estimation. Taylor and Bunn (1999a, b) suggest a new approach to empirical interval estimation which addresses the small sample problem by pooling data across time and forecasting horizons and estimating forecast error distributions. The authors then develop forecast error quantile models that are functions

of lead time, as suggested by theoretically derived variance expressions. This paper explored the use of quantile regression for estimating empirical confidence limits for WASDE forecasts of corn, soybean, and wheat prices.

Following Taylor and Bunn, quantile regressions for corn, soybean, and wheat forecast errors from 1980/81–2006/07 were specified as a function of forecast lead time measured as the forecast month from the beginning to the end of the forecasting cycle. The estimated coefficients indicate the distance from the forecast midpoint to a particular point of error distribution. A benefit of the quantile regression approach is that other factors impacting forecast error distribution may be included in the analysis. This study hypothesized that during the periods of low stocks/use ratios, which reflect the underlying supply and demand conditions, forecast errors may be larger than during the periods of high stocks/use ratios. However, little impact of the stocks/use variable on the forecast error distributions was found.

The quantile regression approach to calculating forecast intervals was evaluated based on out-of-sample performance, where the first 15 observations (1980/81–1994/95) were used to generate confidence limits for the 16th year (1995/96), the first 16 observations were used to generate confidence limits for the 17th year (1996/97), and so on. Hit rates for empirical confidence intervals averaged 78% for corn, 75% for soybeans, and 58% for wheat prior to harvest. These hit rates were statistically different from the target confidence level of 80% in 2 out of 15 cases, or about 13% of the time. After harvest, hit rates for empirical confidence intervals averaged 93% for corn, 85% for soybeans, and 94% for wheat. These hit rates were statistically different from the target level of 90% in 1 out of 33 cases, or about 3% of the time. Our findings suggest empirical confidence intervals calculated using quantile regressions may significantly improve the accuracy of WASDE corn, soybean, and wheat price forecasts.

This study also investigated the relative characteristics of published WASDE and quantile forecast intervals in the context of the tradeoff between accuracy (hit rates) and informativeness (interval width). The only case of dominance is pre-harvest corn, where WASDE forecast intervals have a higher average hit rate and slightly smaller intervals. In the remaining cases, higher average hit rates are accompanied by wider intervals. Because of the lack of dominance of one method over the other, the answer as to which forecast is “best” will depend on the tradeoff between accuracy and informativeness in a forecast user’s utility function. Despite the lack of dominance, the empirical approach still may be valuable to forecast users who want to compute intervals for different confidence levels or who believe WASDE intervals are too narrow compared to target confidence levels, as some evidence suggests.

The results of this study can be extended to calculation of confidence intervals for price forecasts associated with other WASDE commodities. Furthermore, the quantile regression approach to calculating empirical confidence intervals discussed here could be used to generate confidence intervals for non-price WASDE categories, such as export forecasts, which are not currently published in interval form.

[Received September 2009; final revision received October 2010.]

References

- Bessler, D. A., and J. L. Kling. “The Forecast and Policy Analysis.” *Amer. J. Agr. Econ.* 71(1989):503–506.
- Chatfield, C. “Calculating Interval Forecasts.” *J. Bus. and Econ. Statis.* 11(1993):121–135.
- Christoffersen, P. F. “Evaluating Interval Forecasts.” *Internat. Econ. Rev.* 9(1998):841–862.
- Diebold, F. X. *Elements of Forecasting*. Cincinnati, OH: South-Western College Publishing, 1998.

- Hautsch, N., and D. Hess. "Bayesian Learning in Financial Markets: Testing for the Relevance of Information Precision in Price Discovery." *J. Finan. and Quant. Analy.* 42(2007):189–208.
- Isengildina, O., S. H. Irwin, and D. L. Good. "Evaluation of USDA Interval Forecasts of Corn and Soybean Prices." *Amer. J. Agr. Econ.* 84(2004):990–1004.
- Jorgensen, M., and D. I. K. Sjoberg. "An Effort Prediction Interval Approach Based on the Empirical Distribution of Previous Estimation Accuracy." *Information and Software Technology* 45(2003):123–136.
- Kandel, E., and N. D. Pearson. "Differential Interpretation of Public Signals and Trade in Speculative Markets." *J. Polit. Econ.* 103(1995):831–872.
- Keilman, N. W. *Uncertainty in National Population Forecasting: Issues, Backgrounds, Analyses, Recommendations*, Vol. 20. Amsterdam: Swets and Zeitlinger, Publications of the Netherlands Interdisciplinary Demographic Institute (NIDI) and the Population and Family Study Centre (CBGS), 1990.
- Koenker, R. *Quantile Regression*. New York: Cambridge University Press, 2005.
- Koenker, R., and G. Basset. "Regression Quantiles." *Econometrica* 46(1978):33–50.
- . "Robust Tests for Heteroscedasticity Based on Regression Quantiles." *Econometrica* 50(1982):43–62.
- Koenker, R., and J. A. F. Machado. "Goodness of Fit and Related Inference Processes for Quantile Regression." *J. Amer. Statist. Assoc.* 94(1999):1296–1310.
- Murphy, A. H., and R. L. Winkler. "Reliability of Subjective Probability Forecasts of Precipitation and Temperature." *J. Royal Statist. Society Series C* 26(1977):41–47.
- Prescott, D. M., and T. Stengos. "Bootstrapping Confidence Intervals: An Application to Forecasting the Supply of Pork." *Amer. J. Agr. Econ.* 9(1987):266–273.
- Sanders, D. R., and M. R. Manfredo. "USDA Livestock Price Forecasts: A Comprehensive Evaluation." *J. Agr. and Resour. Econ.* 28(2003):316–334.
- Shlyakhter, A. I., D. M. Kammen, C. L. Broido, and R. Wilson. "Quantifying the Credibility of Energy Projections from Trends in Past Data." *Energy Policy* 22(1994):119–130.
- Smith, S. K., and T. Sincich. "Stability over Time in the Distribution of Population Forecast Errors." *Demography* 25(1988):461–474.
- Stoto, M. A. "The Accuracy of Population Projections." *J. Amer. Statist. Assoc.* 78(1983):13–20.
- Taylor, J. W., and D. W. Bunn. "A Quantile Regression Approach to Generating Prediction Intervals." *Management Sci.* 45(1999a):225–237.
- . "Investigating Improvements in the Accuracy of Prediction Intervals for Combinations of Forecasts: A Simulation Study." *Internat. J. Forecasting* 15(1999b):325–339.
- Teigen, L. D., and T. M. Bell. "Confidence Intervals for Corn Price and Utilization Forecasts." *Agr. Econ. Res.* 30(1978):23–29.
- Timm, T. R. "Proposals for Improvement of the Agricultural Outlook Program in the United States." *J. Farm Econ.* 48(1966):1179–1184.
- Vogel, F. A., and G. A. Bange. "Understanding USDA Crop Forecasts." Misc. Pub. No. 1554, USDA/National Agricultural Statistics Service and Office of the Chief Economist, World Agricultural Outlook Board, Washington, DC, 1999.
- Westcott, P. C., and L. A. Hoffman. "Price Determination for Corn and Wheat: The Role of Market Factors and Government Programs." USDA Tech. Bull. No. 1878, USDA/Economic Research Service, Washington, DC, 1999.
- Williams, J. C., and B. D. Wright. *Storage and Commodity Markets*. Cambridge: Cambridge University Press, 1991.
- Williams, W. H., and M. L. Goodman. "A Simple Method for the Construction of Empirical Confidence Limits for Economic Forecasts." *J. Amer. Statist. Assoc.* 66(1971):752–754.
- Yaniv, I., and D. P. Foster. "Graininess of Judgment Under Uncertainty: An Accuracy-Informativeness Trade-off." *J. Experimental Psych.: General* 124(1995):424–432.
- Zarnowitz, V. *Business Cycles: Theory, History, Indicators, and Forecasting*. Chicago, IL: University of Chicago Press, 1992.