An Evaluation of Crop Forecast Accuracy for Corn and Soybeans: USDA and Private Information Agencies

Thorsten M. Egelkraut, Philip Garcia, Scott H. Irwin, and Darrel L. Good

Using 1971–2000 data, we examine the accuracy of corn and soybean production forecasts provided by the USDA and two private agencies. All agencies improved their forecasts as the harvest progressed, and forecast errors were highly correlated and unbiased. The relative forecast accuracy of the agencies varied by crop and month. For corn, USDA's forecasts ranked as most accurate of the three agencies in all periods except for August during the recent period and improved most markedly as harvest progressed. For soybeans, forecast errors were very similar, with the private agencies ranking as most accurate for August and September and making largest relative improvements for August during the recent period. The USDA forecasts were dominant for October and November. Our findings identify several patterns of relative forecast accuracy that have implications for private and public decision makers.

Key Words: corn, private agencies, production forecasts, soybeans, USDA

JEL Classifications: Q11, Q13, C82, Q18

In industrial production, final output is typically known with a high degree of certainty given a set of inputs. In contrast, agricultural crop production is characterized by large variability in output corresponding to the inputs employed. This variability is often the result of changes in stochastic factors affecting agricultural production (e.g., precipitation, temperature), and makes forecasting of crop production a challenge. Uncertainty about the level of final crop production is resolved only as the growing season progresses and more information about crop conditions and crop yields becomes available. It is a well-known theoretical result that crop reports that accurately estimate the size of production before harvest can play an important role in the process of uncertainty resolution and can provide an opportunity for less risky decisions (e.g., Bradford and Kelejian).

For corn and soybeans, crop production forecasts are provided by the National Agricultural Statistics Service (NASS) of the U.S. Department of Agriculture (USDA) and numerous private agencies, of which two, Conrad Leslie and Sparks, are the most prominent. The forecasts produced by these private agencies are initially available to subscribers only but move quickly to the market, typically a

Thorsten M. Egelkraut is graduate fellow and Philip Garcia, Scott H. Irwin, and Darrel L. Good are professors, Department of Agricultural and Consumer Economics, University of Illinois, Urbana-Champaign, IL.

few days prior to the release of the NASS forecasts. These private forecasts are widely reported in the popular press, and consequently, the private as well as the USDA forecasts provide potentially valuable information to a large number of market participants. Understanding the forecast accuracy of these agencies can be important to decision makers. Identifying when crop forecasts become accurate reflections of final production—how the uncertainty about final crop production is resolved—can indicate at what point decisions that are influenced by final crop production can be made with relative certainty. Furthermore, information about the relative accuracy of these three agencies' forecasts can provide insights into which forecasts are of most value in a specific decision context. Alternatively, a high degree of correspondence in the three agencies' production forecasts could indicate little difference in their informational value. Finally, these agencies use different procedures to develop their forecasts, and to date, no research has focused on the relationship between these procedural differences and their effect on the relative accuracy of crop forecasts. Differences in the relative forecast accuracy attributable to the methods used by the agencies might point to changes in procedures that could improve overall forecast performance.

The most recent research related to the accuracy of crop production forecasts has analyzed the agencies' accuracy by indirect methods, using aggregate measures, and has not considered the relationship between the procedures used to develop the forecasts and their performance. Indirect methods assess the relative forecast accuracy of public and private agencies by measuring the reactions of corn and soybean cash and futures prices to the release of the respective reports (e.g., Fortenbery and Sumner; French, Leftwich, and Uhrig; Garcia et al.; McNew and Espinosa; Milonas; Sumner and Mueller). When private agencies' forecasts have been included in these indirect analyses, researchers have used an average of the private forecasts to reflect market expectations about final crop production to calculate how differences between these expectations and the USDA forecasts cause prices to change. The use of these procedures makes it difficult to identify the exact nature of an agency's forecast ability because findings can be influenced by noisy movements in prices that are not related to the differences in forecasts and because the use of an average of private agencies' forecasts can mask the accuracy of an individual agency.

The purpose of the paper is to evaluate the relative accuracy of the corn and soybean production forecasts provided by the USDA, Leslie, and Sparks. The assessment differs from other recent research in several dimensions. Instead of using indirect methods that rely on examining price reactions, we evaluate directly the ability of the agencies to forecast crop production. A direct approach permits the identification of the forecast accuracy of each agency and is particularly well suited for identifying underlying changes in forecast performance over time. This is important because the accuracy of the forecasts by the agencies can change, and an assessment of their relative performance could make users more aware of their relative attractiveness as a source of information. We also discuss the methods used by the USDA and private agencies to collect and interpret information. The procedures differ substantially. The USDA uses a highly systematic method to determine crop acreage and yields, whereas the other agencies use less rigorous sampling procedures. Our analysis sheds light on the relative merits of the procedures used by the agencies to develop their crop production forecasts. Finally, we use an extended and updated data set containing 30 years of observations, which permits a more detailed analysis, the use of more powerful statistical tests, and an assessment of how the agencies' forecast accuracy has changed through time.

The organization of the paper is straightforward. A brief discussion of the literature is followed by a discussion of the forecast data and the procedures used by agencies to develop their crop production forecasts. The next sections present the methods we used to assess forecast accuracy and the empirical results. This is followed by a brief summary and a discussion of the implications of our findings.

Literature

Previous studies analyzing the accuracy and value of production forecasts for corn and soybeans have most frequently evaluated either the accuracy of USDA crop reports alone or have evaluated USDA forecast performance relative to other agencies based on the effect of their release on the mean and variance of commodity cash and futures prices.[1] Clough evaluated corn forecasts by the Department of Agriculture for the period 1919–1950. He found that the uncertainty about crop production was reduced as crop size was reassessed each month. Although there were substantial decreases in uncertainty by the August and September reports, the October and November revisions were fairly close to the "final" estimates. The data used in his study did not provide any evidence of changing forecast accuracy for the time period analyzed. Gunnelson, Dobson, and Pamperin evaluated the accuracy of USDA crop production forecasts for seven row crops, including corn and soybeans, for the period 1929–1970. Over the time period analyzed, they found that the accuracy of USDA crop forecasts improved. The authors also noted that forecasts became more accurate as the growing season progressed.

Garcia et al. compared the accuracy of USDA forecasts to an average of Leslie and Sparks production forecasts for corn and soybeans for the period 1971–1992. The authors found few differences between the forecasts, with the exception that USDA corn forecasts appeared to be relatively more accurate during the early stages of the period. They concluded that, over time, the private information agencies improved their forecasting accuracy in August relative to the USDA.

Based on the notion that efficient markets incorporate new information quickly into prices, other researchers have investigated the informational value of USDA corn and soybean production forecasts by their effect on the mean and variance of subsequent cash and futures prices. Milonas analyzed the effect of the release of USDA corn, wheat, and soybean production forecasts on the cash prices of corn, wheat, soybeans, soybean oil, and soybean meal for the period 1966–1984. He found significant market reactions to the reports and further noted that first-crop forecasts produced larger price reactions than subsequent forecasts. Sumner and Mueller investigated the effect of USDA reports for corn and soybeans on futures prices for the period 1961–1982. They found evidence that USDA crop forecasts affected the respective futures prices. French, Leftwich, and Uhrig examined corn and soybean crop reports over 1969–1981 and found that the unanticipated component of USDA forecasts explained a significant amount of the variation in corn and soybean futures prices immediately following release of the reports. Garcia et al. measured corn and soybean futures price reactions and also concluded that the USDA production forecasts have substantial informational value.

In contrast, Fortenbery and Sumner, who examined the reactions of futures markets to releases of USDA production forecasts for corn and soybeans in three periods (1969–1989; 1969–1982; and 1985–1989), reported finding no evidence of larger than average price movements for the period 1985–1989. The authors concluded that USDA crop production forecasts might no longer contain new information. They provided and tested three different hypotheses to explain their findings. First, prices for corn and soybeans were at or near government support levels during the period 1985–1989 and therefore did not react to the release of crop forecasts. Second, a diminishing U.S. share of the world export market could have resulted in a smaller price effect of the release of USDA crop forecasts. Third, they proposed that the introduction of options on corn and soybean futures might have provided market participants with different means of adjusting their positions in the futures market, and price reactions might therefore no longer be observable in the futures market. Given the brief period examined, they sug-

---

[1] We focus on crop production forecasts. A recent study by Bailey and Brorsen examined USDA forecast accuracy of hog and cattle production. Their analysis, which did not make comparisons between private and USDA forecasts, suggested that the USDA forecast accuracy has improved over time.

gested further research to determine the validity of their proposed explanations and the sensitivity of their results to the data period studied. McNew and Espinosa obtained results consistent with Fortenbery and Sumner. Focusing on the period 1985–1991, they found no evidence that USDA corn and soybean production forecasts influence the level of futures prices. The authors argued that despite these findings, the USDA corn and soybean production forecasts have economic value because the reports significantly reduce the uncertainty prevailing in the market and thus validate the expectation of the market regarding the crop size.

Overall, the evidence suggests that the USDA and aggregate market forecasts have provided relatively accurate forecasts of crop production that improve through the crop year. The findings also suggest that the relative accuracy of the USDA and aggregate market forecasts is changing. In order to provide a more accurate assessment of the relative accuracy of production forecasts, we focus on the direct analysis of the forecasts using the data described in the next section.

Data

Crop production forecasts for corn and soybeans provided by the USDA, Conrad Leslie, and Sparks Companies, Inc., are analyzed for the 30-year period 1971–2000. Production forecasts are provided for August, September, October, and November of each year. The private agencies make their predictions about crop size available to their customers 5 to 7 days prior to the publication of the USDA report. The USDA releases the "final" crop production estimate in January of the year after harvest is completed. Following Clough and Garcia et al., the January estimates are used for final crop production. Most market participants consider the January estimates as being accurate and assign little value to revisions that are made a year later by USDA.

The USDA provided corn and soybean production forecasts for all months during the period evaluated. Leslie did not issue reports for November 1989, August 1990, and November 1992. Sparks did not provide forecasts for November 1972. With the exception of these cases, the resulting data series included 30 observations for each crop, agency, and month a report was released. Years in which no forecasts were issued were deleted only when direct statistical comparisons among the forecasts required paired observations.

Crop Production Forecasting Methods

USDA

The USDA uses a highly elaborate and well-documented procedure to generate its crop production forecasts. At different stages of the production process, forecasts of total acres and yield per acre by crop are developed. For corn and soybeans, the USDA generates production forecasts based on estimates of planted acreage and two types of yield indications, a farmer-reported survey and objective measurements (NASS/SMB). The planted acreage figures are obtained using a survey of farmers during the first 2 weeks of June. These acreage estimates are used in subsequent production forecasts unless acreage figures, which are monitored through the growing season, indicate a change.

The farmer-reported yield survey is conducted primarily by computer-assisted telephone interviewing (CATI), but some data are collected by mail and by face-to-face interviews. The farmers are randomly selected from a list frame and asked monthly for a subjective prediction of their final corn and soybean yields. The list frame is a noncomplete set of all corn and soybean farmers. The list changes through time, reflecting farming arrangements. The objective yield survey is based on an area-frame random sampling design, where the survey samples are selected from respondents to the USDA's June Agricultural Survey in the major producing states. The sample fields are then selected with a probability proportional to their size. The objective yields are obtained from two independently located plots in each randomly selected field. Physical counts and measurements of the number of plants and production per plant are

conducted. Yield per acre is generated for the field after standardizing for moisture content and harvest loss. Objective yield indications are derived from models based on observations over the last 5 years for the corresponding months compared with end-of-season yields. Separate monthly models are constructed by maturity stage so forecast adjustments are automatically made for early or late growing seasons.

It is important to note that accuracy of the objective yield indications can change through the growing season. Early in the season, the yield indications are influenced by assumed relationships between plant counts and fruit numbers, and an assumed fruit weight is adjusted for moisture content and harvest loss. As the season progresses, fruit counts become known. At the end of the season, plots are harvested, and yields are calculated based on actual grain weights and harvest losses.

The yield forecasts are developed monthly from August to November. The data on yields are collected during the last week of the previous month and the first few days of the survey month. Yield forecasts then reflect crop conditions at the beginning of the survey month. The crop production forecasts are based on the assumption of normal growing conditions for the remainder of the season as reflected by historical records.

The subjective and objective yield indications are combined in a multistage process employing statistical and judgmental techniques. This procedure is conducted independently in each state. The state results are then aggregated and adjusted by USDA statisticians to generate national production forecasts.

Conrad Leslie

Conrad Leslie[2] generates crop production forecasts primarily on the basis of a mail survey. The objective of the survey is not to predict USDA forecasts, but the actual size of the crops. The resulting corn and soybean production forecasts are obtained from a statistical model that incorporates at least two components: the yield information from the mail survey and the USDA acreage estimates. The precise statistical model is confidential.

The yield forecasts are "based on 1,250 'card' reports received from elevator managers, processors, grain dealers and milling correspondents—a base which differs from that often used in other estimates" (Leslie-ADM Investor Services). Most of the respondents, however, are grain elevator managers because "these observers are very sensitive to changing conditions in their operating areas ..." (private correspondence with Conrad Leslie). The questions asked in the survey include: during the growing season, compared to normal, how would you rate the condition of the crop? (The response is in terms of percentage of normal, with normal being defined as no damage from weather, insects, etc.); and near harvest, what do you think the yield is in your reporting district? In predicting corn and soybean production, Leslie further "utilizes the latest available government acreage estimates for harvest" (private correspondence with Conrad Leslie).

Sparks Companies, Inc.

Sparks Companies, Inc.,[3] also employs a model that calculates crop production as the product of acreage and yield. Sparks uses the USDA acreage estimates, which they adjust "under specific circumstances." The yield forecasts are obtained from three types of information: a yield survey, which is the dominant source; observations in the field; and any other relevant information available.

The yield survey is conducted by mail. The population surveyed consists of individuals with knowledge about agriculture: county extension agents, bankers, farmers, grain elevator managers, and input suppliers. The information about the number of questionnaires distributed is confidential, but the response rate is above 50%. Using these subjective as-

---

[2] Conrad Leslie now works in a relationship with Archer Daniels Midland, Inc.

[3] The material in this section is based on private communication with Donald Frahm of Sparks Companies, Inc.

sessments of yields, the U.S. yield forecasts are generated from the area-specific yield responses weighted by crop reporting district acreage.

In the largest crop growing states, data from field observations are also collected by "crop scouts" according to a predetermined sampling plan. The plan involves driving through an area, stopping at fixed distance intervals, assessing the field, and conducting physical counts. No particular farms are surveyed. An objective yield is calculated based on these field observations.

Sparks also uses any other relevant information that is available and believed to be reliable. This information could involve subjective opinions from professional and non-professional contacts such as USDA weekly crop condition reports. The survey yield and the objective yield, as well as the other information obtained, are then combined in a very "flexible model" that generates the final crop production forecasts.

Several points emerged from the discussion of the methods used by the three agencies that could influence the crop production forecasts. All three agencies initially base their acreage forecasts on the USDA survey of farmers conducted during the first 2 weeks of June. The USDA acreage estimates are the result of an extensive survey, and their use by Leslie and Sparks might reflect a prohibitively high cost of conducting independent surveys. Subsequently, the USDA updates their acreage estimates if subsequent monitoring indicates a change from the June survey results. Sparks acreage estimates are adjusted independently if conditions warrant, whereas Leslie uses the most current USDA acreage estimates. Because the crop production forecasts are simply the product of acreage and yields, the forecasts by the three agencies should be highly correlated, and differences should primarily be due to the differences in agencies' methods of forecasting yields. The procedures for forecasting yields differ greatly among the agencies. The USDA uses a highly systematic method based on the development of objective and subjective measures of yields from producer plots and producer interviews, which are eventually aggregated from state estimates to a national estimate. Sparks bases its yield forecasts on a survey of knowledgeable individuals, selected field observations, and related information. In contrast, Leslie uses a survey of grain market participants that focuses heavily on elevator managers. The relative effectiveness of these procedures is difficult to establish prior to the analysis but could be determined by the accuracy of the USDA objective estimates, which are generated through the growing season, and the relative accuracy of the three agencies' different procedures for generating subjective estimates of yields.

Forecast Evaluation Methods

Four measures of accuracy are employed to evaluate the forecasts provided by the three reporting agencies. Two measures are defined using percent forecast errors, $p_{t,i} = [(A_t - F_{t,i})/A_t] \times 100$, and two using absolute forecast errors, $e_{t,i} = A_t - F_{t,i}$, where A_t denotes the actual crop production defined as the January USDA estimate and $F_{t,i}$ denotes the forecast in year t and month i. The following accuracy measures are computed for each commodity and forecasting agency:

mean absolute percent error,

$$MAPE = \frac{1}{T} \sum_{t=1}^{T} |p_{t,i}|,$$

root-mean-squared percent error,

$$RMSPE = \sqrt{\frac{1}{T} \sum_{t=1}^{T} p_{t,i}^2},$$

mean absolute error,

$$MAE = \frac{1}{T} \sum_{t=1}^{T} |e_{t,i}|,$$

and root-mean-squared error

$$RMSE = \sqrt{\frac{1}{T} \sum_{t=1}^{T} e_{t,i}^2}.$$

The differences in the accuracy measures

between the forecasts from the agencies are evaluated using the Modified Diebold Mariano (*MDM*) test proposed by Harvey, Leybourne, and Newbold (1997). The procedure involves specifying a cost-of-error function, $g(e)$, of the forecast errors e with pairwise testing of the null hypothesis of equality of expected forecast performance. The test statistic, which Harvey, Leybourne, and Newbold (1997) indicate should be compared with the critical values from the Student's t distribution with $T - 1$ degrees of freedom, is computed for one-step-ahead forecasts as

$$MDM = \sqrt{\frac{T - 1}{\frac{1}{T} \sum_{t=1}^{T} (d_t - \bar{d})^2}} \bar{d},$$

where $d_t = g(e_{t,i,1}) - g(e_{t,i,2})$, $\bar{d}$ is the average difference for the period, and the null hypothesis is $E(d_t) = 0$. For example, when testing for significant differences of the *MAPE*s of two forecast agencies, $g(e_{t,i,1}) = p_{t,i,1}$ is the absolute percent forecast error of agency 1, $g(e_{t,i,2}) = p_{t,i,2}$ is the absolute percent forecast error of agency 2, and $d_t = p_{t,i,1} - p_{t,i,2}$ is the difference between the respective absolute percent forecast errors at time t.

Harvey, Leybourne, and Newbold (1997) demonstrate that the size of the *MDM* test is quite similar to the nominal significance levels for the number of forecasts used in our analysis, a finding that is insensitive to the degree of contemporaneous correlation between forecast errors and departures from normality. They argue that these characteristics are important because researchers attempting to differentiate between forecasts are often faced with a limited number of correlated forecasts that possess occasional large errors. Furthermore, for the degree of contemporaneous correlation in forecast errors in our sample, the power of the test is rather substantial and declines only marginally with departures from normality (Harvey, Leybourne, and Newbold 1998). Harvey, Leybourne, and Newbold (1997) identify other advantages of the *MDM* test, including its applicability to multiple-step-ahead forecast horizons, its nonreliance on an assumption of forecast unbiasedness,

and its applicability to cost-of-error functions other than the conventional quadratic loss. They conclude by asserting that the *MDM* test constitutes the "best available" method for determining the significance of observed differences in competing forecasts.

To examine in more detail how forecast behavior has changed through time, the individual forecast accuracy measures are also aggregated on a cumulative and a rolling basis and graphed. The cumulative measures are obtained by successively adding each year to the previous set of observations. The cumulative measures smooth graphs and provide an overall indication of directional changes in forecast accuracy. The rolling measures are computed as 5-year moving averages of the respective accuracy measures. The rolling measures display larger variability than the cumulative measures and are sensitive to poor predictions in a single year. The cumulative measures provide a description of overall patterns in forecast accuracy, whereas the rolling measures provide information on more short-run changes in forecast behavior.

Results and Discussion

The first section that follows compares the errors of the crop production forecasts released in August, September, October, and November for each crop separately, between both crops, and between the earlier and later periods, independent of which agency has generated the forecasts. This evaluation permits the identification of possible common patterns underlying the crop production forecast errors. The next section focuses on evaluating the relative performance of the individual forecasting agencies.

General Patterns in Crop Forecasts

The measures for evaluating the forecast accuracy of corn and soybean production estimates by USDA, Leslie, and Sparks are presented in Tables 1 and 2. The measures are calculated for the entire period 1971–2000 and for the subperiods 1971–1984 and 1985–2000. The length of the two subperiods was chosen

Table 1. Accuracy Measures and Ranks for Corn Production Forecasts

Month	Period	Agency	MAPE (%)	(r)	MAE (mil. bu.)	(r)	RMSPE (%)	(r)	RMSE (mil. bu.)	(r)
August	1971–2000	USDA	5.19	(1)	356.4	(1)	7.47	(1)	475.2	(2)
		Leslie	5.51	(2)	367.2	(2)	7.92	(2)	472.3	(1)
		Sparks	5.69	(3)	371.7	(3)	8.46	(3)	501.3	(3)
	1971–1984	USDA	5.33	(1)	306.1	(1)	8.22	(1)	426.7	(1)
		Leslie	6.14	(2)	357.6	(2)	9.12	(2)	481.6	(2)
		Sparks	6.94	(3)	402.9	(3)	9.89	(3)	528.3	(3)
	1985–2000	USDA	5.06	(3)	400.3	(3)	6.75	(2)	514.0	(3)
		Leslie	4.93	(2)	376.1	(2)	6.59	(1)	463.5	(1)
		Sparks	4.61	(1)	344.4	(1)	6.98	(3)	476.4	(2)
September	1971–2000	USDA	4.04	(1)	287.2	(2)	5.10	(1)	366.0	(1)
		Leslie	4.10	(2)	284.4	(1)	5.61	(2)	380.1	(2)
		Sparks	4.60	(3)	321.7	(3)	6.00	(3)	408.6	(3)
	1971–1984	USDA	3.68	(1)	224.4	(2)	4.23	(1)	256.0	(1)
		Leslie	3.71	(2)	220.9	(1)	4.70	(2)	269.1	(2)
		Sparks	4.30	(3)	265.4	(3)	5.17	(3)	315.7	(3)
	1985–2000	USDA	4.35	(1)	342.1	(2)	5.76	(1)	440.2	(1)
		Leslie	4.43	(2)	340.0	(1)	6.30	(2)	455.5	(2)
		Sparks	4.87	(3)	370.9	(3)	6.64	(3)	475.2	(3)
October	1971–2000	USDA	2.57	(1)	185.4	(1)	3.43	(1)	245.4	(1)
		Leslie	3.02	(2)	213.0	(2)	3.87	(2)	271.8	(2)
		Sparks	3.36	(3)	235.6	(3)	4.19	(3)	286.2	(3)
	1971–1984	USDA	2.42	(1)	156.6	(1)	2.84	(1)	188.1	(1)
		Leslie	2.80	(2)	174.6	(2)	3.35	(2)	211.7	(2)
		Sparks	3.47	(3)	218.4	(3)	3.85	(3)	241.4	(3)
	1985–2000	USDA	2.70	(1)	209.0	(1)	3.87	(1)	286.3	(1)
		Leslie	3.20	(2)	246.6	(2)	4.27	(2)	315.1	(2)
		Sparks	3.27	(3)	250.8	(3)	4.46	(3)	320.2	(3)
November	1971–2000	USDA	1.24	(1)	86.3	(1)	1.65	(1)	108.6	(1)
		Leslie	1.59	(2)	105.5	(2)	2.22	(2)	137.5	(2)
		Sparks	1.90	(3)	129.8	(3)	2.42	(3)	156.7	(3)
	1971–1984	USDA	1.40	(1)	91.8	(1)	1.66	(1)	111.3	(1)
		Leslie	1.64	(2)	105.1	(2)	1.93	(2)	125.6	(2)
		Sparks	2.19	(3)	140.2	(3)	2.37	(3)	154.2	(3)
	1985–2000	USDA	1.10	(1)	81.5	(1)	1.65	(1)	106.1	(1)
		Leslie	1.53	(2)	105.9	(2)	2.47	(2)	148.5	(2)
		Sparks	1.66	(3)	121.4	(3)	2.46	(3)	158.8	(3)

Note: *MAPE* is mean absolute percent error; *MAE* is mean absolute error; *RMSPE* is root-mean-squared percent error; *RMSE* is root-mean-squared error. The numbers in parentheses (*r*) reflect the ranking.

based on earlier research by Fortenbery and Sumner, who found that corn and soybean futures did not react to the release of USDA crop production forecasts after 1984, and by Garcia et al., who, when testing this hypothesis, also found a decline in the reaction of corn and soybean futures prices to the release of the USDA crop production forecasts after 1984.

For each crop, the changes in prediction errors in successive months (August, September, October, and November) reflect the resolution of production uncertainty. In general,

Table 2. Accuracy Measures and Ranks for Soybean Production Forecasts

Month	Period	Agency	MAPE (%)	(r)	MAE (mil. bu.)	(r)	RMSPE (%)	(r)	RMSE (mil. bu.)	(r)
August	1971–2000	USDA	4.93	(3)	96.5	(3)	5.96	(2)	119.7	(2)
		Leslie	4.64	(1)	89.5	(1)	5.69	(1)	110.7	(1)
		Sparks	4.75	(2)	95.4	(2)	6.13	(3)	127.0	(3)
	1971–1984	USDA	5.10	(1)	83.3	(1)	6.41	(1)	105.7	(1)
		Leslie	5.19	(2)	85.0	(2)	6.46	(2)	106.5	(2)
		Sparks	5.20	(3)	88.9	(3)	6.72	(3)	117.2	(3)
	1985–2000	USDA	4.78	(3)	108.1	(3)	5.53	(2)	130.7	(2)
		Leslie	4.12	(1)	93.7	(1)	4.87	(1)	114.5	(1)
		Sparks	4.36	(2)	101.1	(2)	5.56	(3)	135.1	(3)
September	1971–2000	USDA	3.92	(2)	79.4	(2)	4.64	(1)	97.6	(2)
		Leslie	3.49	(1)	67.9	(1)	4.65	(2)	92.8	(1)
		Sparks	4.12	(3)	81.6	(3)	5.15	(3)	103.9	(3)
	1971–1984	USDA	3.37	(2)	57.5	(2)	4.11	(1)	70.9	(2)
		Leslie	3.28	(1)	53.1	(1)	4.38	(2)	67.3	(1)
		Sparks	3.78	(3)	62.1	(3)	5.13	(3)	83.4	(3)
	1985–2000	USDA	4.39	(2)	98.5	(2)	5.06	(2)	116.0	(2)
		Leslie	3.67	(1)	80.8	(1)	4.88	(1)	110.4	(1)
		Sparks	4.41	(3)	98.7	(3)	5.12	(3)	119.0	(3)
October	1971–2000	USDA	2.50	(1)	47.6	(1)	2.95	(1)	56.5	(1)
		Leslie	2.60	(2)	51.8	(2)	3.14	(2)	63.3	(2)
		Sparks	2.94	(3)	57.7	(3)	3.57	(3)	70.8	(3)
	1971–1984	USDA	2.85	(2)	48.9	(2)	3.21	(2)	56.3	(2)
		Leslie	2.39	(1)	41.8	(1)	3.03	(1)	53.5	(1)
		Sparks	3.01	(3)	49.9	(3)	3.69	(3)	62.2	(3)
	1985–2000	USDA	2.19	(1)	46.6	(1)	2.70	(1)	56.7	(1)
		Leslie	2.78	(2)	60.6	(2)	3.23	(2)	70.9	(2)
		Sparks	2.88	(3)	64.6	(3)	3.46	(3)	77.6	(3)
November	1971–2000	USDA	1.33	(1)	24.9	(1)	1.59	(1)	29.3	(1)
		Leslie	1.61	(3)	30.4	(3)	2.02	(3)	37.6	(3)
		Sparks	1.47	(2)	27.5	(2)	1.91	(2)	34.2	(2)
	1971–1984	USDA	1.67	(1)	28.7	(1)	1.92	(1)	32.9	(1)
		Leslie	1.82	(2)	31.8	(3)	2.14	(2)	37.8	(2)
		Sparks	1.86	(3)	30.7	(2)	2.40	(3)	39.3	(3)
	1985–2000	USDA	1.04	(1)	21.9	(1)	1.27	(1)	26.0	(1)
		Leslie	1.42	(3)	29.1	(3)	1.90	(3)	37.5	(3)
		Sparks	1.16	(2)	24.9	(2)	1.39	(2)	29.3	(2)

Note: *MAPE* is mean absolute percent error; *MAE* is mean absolute error; *RMSPE* is root-mean-squared percent error; *RMSE* is root-mean-squared error. The numbers in parentheses (*r*) reflect the ranking.

the error measures for corn and soybeans decline as the growing season progresses. This result is not surprising because uncertainty about actual crop production is resolved as harvest approaches and estimates of yields and acreage become more accurate. This result is also consistent with earlier findings by Gunnelson, Dobson, and Pamperin and by Sumner and Mueller, who note that production forecasts become more accurate in later months as the uncertainty about the current year's crop production is resolved.

For the two crops, comparisons of relative forecast accuracy are best performed using the

Table 3. Correlations of Relative and Absolute Forecast Errors for the August Corn and Soybean Estimates by USDA, Leslie, and Sparks

Crop	Agency	Relative Forecast Errors			Absolute Forecast Error		
		USDA	Leslie	Sparks	USDA	Leslie	Sparks
Corn	USDA	1.00			1.00		
	Leslie	0.93	1.00		0.90	1.00	
	Sparks	0.96	0.98	1.00	0.94	0.96	1.00
Soybeans	USDA	1.00			1.00		
	Leslie	0.94	1.00		0.94	1.00	
	Sparks	0.94	0.95	1.00	0.94	0.95	1.00

percentage-based error measures since the sizes of the corn and soybean crops differ. With the exception of the September USDA forecast during the 1985–2000 period, the percent forecast errors for corn tend to be larger than for soybeans for August, September, and October for both time periods. The percent forecast errors for both crops became approximately equal in November.

Comparisons of relative accuracy between the two periods do not indicate a general pattern of forecast behavior. August *MAPEs* and *RMSPEs* for corn and soybeans are smaller during the second period than the first, with the largest period-to-period decline registered for corn. November corn production forecasts also registered slightly smaller *MAPEs* during the second period. In contrast, the errors for the September and October corn production forecasts are larger during the second period than during the early period for all agencies, with the exception of Sparks' October *MAPE*. The findings for the other months for soybeans also are mixed, with some agencies improving their forecasts for specific months, but declining for others. These findings might reflect the high degree of variability in crop production over the entire period.

For both crops, the errors in the production forecasts provided by the three agencies are highly correlated. As an example, the correlation coefficients for the August corn and soybean production forecasts are displayed in Table 3. In light of the similar magnitude of the forecast errors for many periods (Tables 1 and 2), this result suggests that the forecasting agencies all either overpredict or underpredict

by a similar amount.[4] One explanation for such a strong relationship is that Leslie and Sparks rely heavily on the USDA acreage numbers in generating their crop production forecasts, so that the differences between the crop production forecasts result primarily from differences in forecasted yields.

Relative Production Forecast Performance

In the following section, forecast rankings and statistical findings are complemented by cumulative and rolling accuracy measures. The cumulative measures incorporate changes more slowly, whereas the rolling measures react immediately to changes in forecast accuracy. The discussion centers on the *MAPEs* as the measure of forecast accuracy; the results using the other error measures are similar and lead to comparable conclusions.

The *MDM* test was employed to determine statistical differences of the forecasts provided by the reporting agencies within the two subperiods, 1971–1984 and 1985–2000. Because of the limited number of observations and the somewhat limited power of the statistical tests, we examined the patterns in differences of

---

[4] At a reviewer's request, we examined for the presence of asymmetric forecast errors and bias in the forecasts. Comparing mean percent errors, mean negative and positive percent errors (Ferris, pp. 143–45), we found no evidence of bias, which indicates that the agencies over- or underpredict by a similar amount. Regressing final crop production on the agencies' forecasts and testing the combined null hypothesis that the constant equaled zero and the slope coefficient equaled one also provided no evidence of bias. The results are available from the authors.

Figure 1a. Corn cumulative MAPEs, August

**Figure 1b.** Corn rolling MAPEs, August

**Figure 1c.** Soybeans cumulative MAPEs, August

Figure 1d. Soybeans rolling MAPEs, August

Figure 1.   (a) Corn Cumulative *MAPE*s, August; (b) Corn Rolling *MAPE*s, August; (c) Soybeans Cumulative *MAPE*s, August; (d) Soybeans Rolling *MAPE*s, August

forecast accuracy because these could have economic significance to specific decision makers.

Corn

The forecast performance of the agencies varies by month (Table 1). For August, USDA's *MAPE* for the entire sample period is smaller than that of either private agency. During the early period, the USDA provided the most accurate forecasts, with a mean absolute percentage error (*MAPE*) of 5.33%, followed by Leslie and Sparks with *MAPE*s of 6.14% and 6.94%, respectively. For this early period, the *MAPE*s were found to be significantly different between Sparks and Leslie (p = .036) and between Sparks and the USDA (p = .010), reflecting the lower accuracy of the early Sparks' corn production forecasts. For the later period, the USDA improved its accuracy, reducing the *MAPE* by 0.27 percentage points while Leslie's *MAPE* declined by 1.21 per-

centage points. Sparks displayed the largest increase in accuracy, reducing the *MAPE* by 2.33 percentage points and surpassing both Leslie and the USDA in predictive accuracy. During the later period, no statistically significant differences among the services were encountered, underscoring the improvement in forecast accuracy of Sparks relative to Leslie and the USDA.

As reflected in the cumulative and rolling accuracy measures, Sparks' improvement occurred during two major periods. From 1983 to 1991, Sparks' cumulative *MAPE* declined relative to Leslie's (Figure 1a). By 1991, the difference between Sparks' and Leslie's cumulative *MAPE*s had become almost zero. The relative improvement is also reflected in the rolling *MAPE*s (Figure 1b), where Sparks' declined at a faster rate than Leslie's. Hence, 1983–1991 was a period of improved forecasts by the private agencies. Beginning in 1992, the differences between the cumulative *MAPE*s of the private agencies and the USDA

Figure 2a. Corn cumulative MAPEs, September


Figure 2b. Corn rolling MAPEs, September

**Figure 2c.** Soybeans cumulative MAPEs, September

Figure 2d. Soybeans rolling MAPEs, September

Figure 2.   (a) Corn cumulative *MAPE*s, September; (b) Corn Rolling *MAPE*s, September; (c) Soybeans Cumulative *MAPE*s, September; (d) Soybeans Rolling *MAPE*s, September

decrease, while the differences in the cumulative *MAPE*s of the private agencies remained close to zero (Figure 1a). This relationship is also displayed in Figure 1b, which shows that the private forecasts have improved relative to the USDA and are now marginally more accurate. The past decade marks the improvement in the accuracy of the private agencies' August forecasts relative to the USDA.

For September, the USDA's *MAPE* for the entire sample, and for the two periods, is only marginally smaller than that of Leslie (Table 1). Sparks possessed the largest *MAPE*s, a result of a few poor forecasts in the early 1980s (Figure 2a,b). Each agency's forecast accuracy decreased from the early to the later period, but the reductions in the *MAPE*s were very small. Figure 2b also shows that after 1990 there are essentially no differences in corn production forecasts. These findings suggest that the three agencies provide about equally accurate forecasts in both subperiods. The statistical analysis, yielding no significant differ-

ences between the three agencies during the two subperiods, confirms this conclusion.

The USDA forecast accuracy dominates the private agencies for October and November, particularly after 1990. The USDA forecasts had the lowest *MAPE* followed by Leslie and Sparks over the entire period, with no change in ranking in the separate periods (Table 1). For the early and the later periods, Sparks and the USDA forecasts differ for October (p = .003 and .027, respectively) and for November (p = .001 and .035, respectively). Sparks and Leslie also differ for October (p = .013) and for November (p = .006) during the early period, but not for the later period. These findings indicate an improvement in Sparks' forecasts relative to Leslie's for October (Figure 3a,b) and an improvement of Leslie's forecasts relative to Sparks' for November (Figure 4a,b). Furthermore, for October and November, Leslie also differed from the USDA for the later period (p = .009 and .083, respectively), but not for the early period, indicating

Figure 3a. Corn cumulative MAPEs, October

Figure 3b. Corn rolling MAPEs, October

Figure 3c. Soybeans cumulative MAPEs, October

Figure 3d. Soybeans rolling MAPEs, October

Figure 3. (a) Corn Cumulative *MAPE*s, October; (b) Corn Rolling *MAPE*s, October; (c) Soybeans Cumulative *MAPE*s, October; (d) Soybeans Rolling *MAPE*s, October

an improvement of the USDA's forecasts relative to Leslie's.

Overall, the USDA's forecasts rank as most accurate for the September, October, and November, and for August forecasts during the early period. However the August forecasts during the later period showed the private forecasting agencies, particularly Sparks, to have improved their forecast accuracy, providing marginally superior forecasts of final corn production. Although the USDA ranks as the most accurate forecaster for September, there are virtually no differences in the magnitude of the forecast errors, particularly in recent years. The USDA forecasts dominate for October and November. It appears that with successive monthly forecasts, the USDA becomes relatively more accurate than the private agencies, particularly during the past decade.

Soybeans

The results for soybeans differ from those for corn, but there are also important similarities.

Because of the similarity in forecasts, none of the forecast error measures for the agencies were significantly different for any period. Similarly, because of the small differences in the forecasts, the rankings were susceptible to even modest forecast errors and changed more frequently than for corn. These small differences can be illustrated by considering that for the entire period, the minimum and maximum differences in *MAPE*s between the most accurate and least accurate forecast for each month were 0.29 percentage points between USDA and Leslie for August and 0.63 percentage points between Leslie and Sparks for September. Despite the small magnitude of the differences in the *MAPE*s, several patterns emerge.

For the early period, the USDA's August forecasts provided the most accurate forecasts, displaying marginally smaller *MAPE*s than those of either Leslie or Sparks (Table 2). For the later period, the forecast errors of all agencies decreased, with the private agencies achieving the largest improvement in accura-
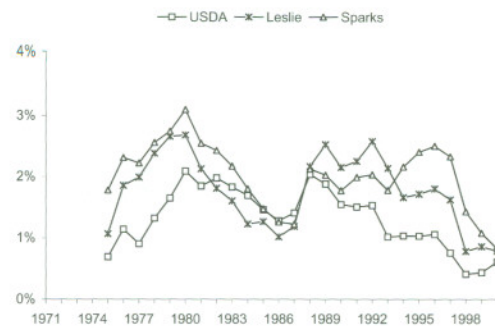
Figure 4a. Corn cumulative MAPEs, November


Figure 4b. Corn rolling MAPEs, November

**Figure 4c.** Soybeans cumulative MAPEs, November

Figure 4d. Soybeans rolling MAPEs, November

Figure 4.   (a) Corn Cumulative *MAPE*s, November; (b) Corn Rolling *MAPE*s, November; (c) Soybeans Cumulative *MAPE*s, November; (d) Soybeans Rolling *MAPE*s, November

cy. Similar to corn, both private agencies provided moderately more accurate forecasts than the USDA for the 1985–2000 period. The cumulative and rolling accuracy measures also demonstrate this development (Figure 1c,d).

For September, Leslie's forecasts resulted in the smallest *MAPE*s, followed by the USDA and Sparks. This result is in contrast to that for corn, where the USDA provided the best September forecasts. For October and November, the USDA begins to dominate the forecasts, a pattern also reflected in corn forecasts. For October, the USDA improved forecast accuracy most dramatically by reducing the *MAPE* 0.66 percentage points from the early to the later subperiod. For November, all three agencies improved their forecast accuracy, with the USDA providing the most accurate forecasts independent of the period analyzed.[5]

Overall, the forecast accuracy for soybean crop production tended to improve through the growing season for all agencies, with Leslie and the USDA generally providing the most accurate forecasts. Similar to corn, the USDA's forecasts dominated for October and November, but their accuracy for August declined relative to the other agencies during the later period.

Summary and Discussion

The accuracy of corn and soybean production forecasts was evaluated for USDA, Leslie, and Sparks. In general, percent forecast errors are larger for corn than for soybeans, and forecast accuracy improves in successive months of the crop year. The larger errors in corn reflect the

---

[5] In November, we deleted the Leslie and USDA soybean predictions because Sparks did not issue a

---

forecast. This year was associated with an extremely large percent error in soybeans experienced by the other agencies and changed the relative rankings. Our other findings were not sensitive to the small number of nonexistent forecasts by the agencies.

relative sensitivity of corn yields to weather events and its longer growing season. The improvement in forecasts in successive months of the crop year reflects the resolution of uncertainty of the size of the harvest. By October the mean absolute percent errors are generally below 3% and continue to decline in November. Our findings also indicate that forecast errors are highly correlated across the three agencies, show no sign of bias, and generally tend to over- and underpredict production by about the same amount.

The relative forecast accuracy of the agencies varied by crop and month. For corn, the USDA forecasts ranked as the most accurate of the three agencies for all periods, except for August during recent times. The USDA forecasts also improved more markedly than the private agencies as the harvest progressed. For soybeans, the forecast errors were very similar. The private agencies ranked as most accurate for August and September and, similar to corn, they also made the largest relative improvements in August during the most recent period. Also, similar to corn, the USDA soybean production forecasts are more accurate for October and November.

In general terms, the high degree of correspondence in the agencies' forecasts, particularly for soybeans, suggests that their economic value might be similar for decision makers. The small differences in these forecasts also indicate that their economic value would be more meaningful during periods of reduced inventories and higher price volatility. Furthermore, the similarity in the forecasts suggests that their value might be partially determined by the timing of the forecast release. Subscribers to the private agencies could have a very small window of opportunity before the private agencies' forecasts move into the public domain. This short period might permit them to take positions in cash and futures markets to strengthen their financial situation. Although the differences in forecasts might not be large in magnitude, several consistent patterns appear to exist, such as the dominance of the USDA forecasts for October and November for both crops and the relative improvement in accuracy of the private agencies

for both crops for August during the later period, which could have economic meaning to specific decision makers. In particular, the improvement of the private agencies' August forecast relative to the USDA forecasts signal their increasing value and could offer market opportunities at a time when the level of crop production and prices are set and marketing and pricing decisions are made.

From a crop forecast-generating perspective, the confidentiality associated with the methods of generating the private agency forecasts makes a complete assessment of the effects of the procedures on crop production forecasts difficult. Yet, several points can be made. The similarity in the forecasts across years and their improvement within the year indicate that each agency has identified salient and highly related information affecting final production. As identified, some of the similarity over time reflects the use of the USDA acreage estimates by both private agencies in their production calculations. However, forecast improvements within the crop year also indicate that, regardless of the procedure used, as the final yields becomes more precise, the agencies are able to reflect the changes appropriately. The results indicate that the three agencies are doing a reasonable job of forecasting production prospects in advance of crop maturity.

As mentioned, yield forecasts, which are the most important source of forecast differences, can be composed of subjective and objective components. When objective and subjective yield estimates are highly consistent and accurate, yield forecasts will be similar regardless of the source of or weight placed on the individual components. However, as the accuracy of the subjective and objective yield measurements differ, the overall accuracy of the yield forecast will be affected by the magnitude of the difference and the weights placed on the components. Here, the agencies differ in the use of procedures for developing subjective and objective yield measurements. USDA forecasts are calculated based on a highly systematic subjective survey of producers and a repeated rigorous objective assessment of crop development. Sparks uses a

somewhat less rigorous objective assessment procedure, complemented by a subjective survey of a more broadly defined sample of the crop-producing sector. In contrast, Leslie uses only a subjective survey of market participants. In this context, our finding that the private agencies have improved their accuracy for both August corn and soybeans production forecasts relative to the USDA suggests that during the period of most uncertainty, when the distributions of objective yields are not well established, gains in forecast accuracy can be achieved by subjective sampling of a wider variety of market participants. As harvest progresses and objective yields become relatively more precise, the value of repeated location-specific sampling procedures employed by the USDA that are designed to evaluate this component of overall yields increases. The dominance of the USDA forecasts for both crops for October and November also is consistent with this explanation.

Another possible explanation for the improved forecast accuracy of the private agencies for August is the increased use of USDA Crop Condition reports since the mid-1980s. Numerous private agencies have developed yield forecasting models based on these reports, and the resulting forecasts are now routinely reported by market news organizations. The development of such forecasting models might assist the private agencies in increasing their accuracy. Furthermore, the relative cost of information technology and computing power has decreased dramatically over the period analyzed. Hence, the resource gap between the USDA and the private agencies has narrowed, allowing the private agencies to include more data and perform more complex analyses. Future technological developments, combined with declining costs for information such as satellite images, could lead to further improvements in forecast accuracy for all agencies.

The improved performance by the private agencies for August for both crops during the most recent years and the ability of the private agencies to generate relatively accurate forecasts in soybeans suggest that it might be useful for the USDA to investigate expanding the scope of their subjective yield analysis to incorporate a wider range of market and industry participants. Such a strategy, if proved effective, might lead to improved crop production forecasts.

References

Bailey, D., and W. Brorsen. "Trends in the Accuracy of USDA Production Forecasts for Beef and Pork." *Journal of Agricultural and Resource Economics* 23(December 1998):515–25.

Bradford, D.F., and H.H. Kelejian. "The Value of Information for Crop Forecasting with Bayesian Speculators: Theory and Empirical Results." *Bell Journal of Economics* 9(Spring 1978):123–44.

Clough, M. "Changes in Corn Acreage and Production after the Early Indications." *Agricultural Economics Research* 3(October 1951):140–46.

Ferris, J.N. *Agricultural Prices and Commodity Market Analysis.* New York: McGraw-Hill, 1997.

Fortenbery, R., and D. Sumner. "The Effects of USDA Reports in Futures and Options Markets." *Journal of Futures Markets* 13(April 1993):157–73.

French, K.R., R. Leftwich, and N. Uhrig. "The Effects of Scheduled Announcements on Futures Markets." Working paper 273, Center for Research in Security Prices, Graduate School of Business, University of Chicago, Chicago, 1989.

Garcia, P., S.H. Irwin, R.M. Leuthold, and L. Yang. "The Value of Public Information in Commodity Futures Markets." *Journal of Economic Behavior and Organization* 32(April 1997):559–70.

Gunnelson, G., W. Dobson, and S. Pamperin. "Analysis of the Accuracy of USDA Crop Forecasts." *American Journal of Agricultural Economics* 54(Part I November 1972):639–45.

Harvey, D., S. Leybourne, and P. Newbold. "Testing the Equality of Prediction Mean Squared Errors." *International Journal of Forecasting* 13(June 1997):281–91.

———. "Tests for Forecast Encompassing." *Journal of Business and Economic Statistics* 16(April 1998):254–59.

Leslie-ADM Investor Services. *Crop Survey,* re-

leased Tuesday, Sept. 7, 1999, at 7:30 am CDT. Private correspondence with C. Leslie, 1999.

McNew, K., and J. Espinosa. "The Informational Content of USDA Crop Reports: Impacts on Uncertainty and Expectations in Grain Futures Markets." *Journal of Futures Markets* 14(June 1994):475–92.

Milonas, N. "The Effects of USDA Crop Announcements on Commodity Prices." *Journal of Futures Markets* 7(October 1987):571–89.

National Agricultural Statistical Service/Statistical Methods Branch (NASS/SMB). *The Yield Forecasting and Estimating Program of NASS.* Washington, DC: U.S. Department of Agriculture/Statistical Methods Branch, Estimates Division NASS Staff report SMB 98-01, June 1998.

Sumner, D., and R. Mueller. "Are Harvest Forecasts News? USDA Announcements and Futures Market Reactions." *American Journal of Agricultural Economics* 71(February 1989): 1–8.