



## Opening Up the Black Box: More on the USDA Corn Yield Forecasting Methodology

Scott Irwin and Darrel Good

Department of Agricultural and Consumer Economics  
University of Illinois

August 26, 2016

*farmdoc daily* (6):162

---

Recommended citation format: Irwin, S., and D. Good. "Opening Up the Black Box: More on the USDA Corn Yield Forecasting Methodology." *farmdoc daily* (6):162, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, August 26, 2016.

Permalink: <http://farmdocdaily.illinois.edu/2016/08/opening-up-the-black-box-more-usda-corn-yield.html>

---

We have written extensively about U.S. corn yield forecasts issued by the National Agricultural Statistics Service (NASS) of the USDA (Irwin and Good, 2006; Irwin and Good, 2011; Irwin, Sanders and Good, 2014; *farmdoc daily*, August 28, 2013; August 29, 2014; August 18, 2016). As indicated in these previous publications, NASS uses two surveys to provide indicators of corn yield prior to harvest. These include the Agricultural Yield Survey (or farmer-reported survey) and the Objective Yield Survey (or field-measurement survey) in major corn producing states. A formal Agricultural Statistics Board (ASB) of NASS statisticians is convened in each month of the forecasting cycle to review yield indicators and determine an official yield forecast. In our previous publications we focused on the survey procedures used by NASS and had relatively little to say about the methodology used by the ASB to combine the various yield indicators and determine an official yield forecast. This omission reflects the fact that NASS has historically provided very little information about this part of the forecasting methodology. Fortunately, several publications have appeared in the scientific literature during the last several years that provide important insights into the nature of the different survey yield indications and the other non-survey information used by the ASB to determine official published yield forecasts. The purpose of this article is to review the information in these scientific articles and provide a more informed understanding of the role of the ASB in generating these important market-moving forecasts.

### NASS Corn Yield Survey Indicators

We begin with a brief review of the two surveys that form the backbone of the NASS corn yield forecasting methodology—the Agricultural Yield Survey (or the farmer-reported survey) and the Objective Yield Survey (or the field measurement survey). For the August 2016 forecast, the Agricultural Yield Survey (AYS) included 22,144 operations for all crops and was conducted in 32 states for corn. The sample of farm operations surveyed was drawn from those who responded to the survey of planted acreage in June. The same operations will be interviewed each month from September through November. In the August survey, respondents were asked to identify the number of acres of corn to be harvested and to provide a forecast of the final yield.

---

We request all readers, electronic media and others follow our citation guidelines when re-posting articles from *farmdoc daily*. Guidelines are available [here](#). The *farmdoc daily* website falls under University of Illinois copyright and intellectual property rights. For a detailed statement, please see the University of Illinois Copyright Information and Policies [here](#).

The Objective Yield Survey (OYS) is designed to generate yield forecasts based on actual plant counts and measurements. The sample of fields for the OYS survey is selected from farms that reported corn planted or to be planted in the June survey of acreage. Samples are selected in the 10 principal corn producing states. A random sample of fields is drawn with the probability of selection of any particular field being proportional to the size of the tract. For the August 2016 OYS, a total of 4,544 plots were sampled for corn, soybeans, cotton, and winter wheat.

Two counting areas, or plots, are randomly selected in each field. Objective measurements (such as counts of plants and ears) are made for each plot each month during the survey cycle. When mature, the plots are harvested and yield is calculated based on actual production minus an allowance for harvest loss. During the August survey, the operator is asked to verify, field-by-field, the acreage reported in June. For corn, each of the two independently located sample plots in a field consists of two parallel 15 foot sections of row. Each plot is selected by using a random number of rows along the edge of the field and a random number of paces into the field. Enumerators count all fruit and fruiting positions in corn and, if ears have formed, a sample of ears is measured for length and circumference. Just before the field is harvested, both plots are hand harvested and weighed by the enumerator.

At each visit, the enumerator establishes a corn maturity category for the plot, ranging from 1 (no ear shoots) to 7 (mature). Prior to the blister stage the number of ears is forecast based on the number of stalks, ear shoots, or ears and both the weight per ear and harvest loss are forecast based on the 5-year average. From the blister through the dough stage, the weight per ear is forecast based on kernel row length and harvest loss is forecast based on the past 5-year average. Ear weight is measured in the dent and/or mature stage. Harvest loss is measured following harvest.

Prior to maturity and harvest, the OYS corn yield is forecast based on the forecast of the number of ears, the forecast of the weight per ear, and the forecast of harvest loss. Forecasts are based on conditions as of the survey date and projected assuming normal weather conditions for the remainder of the growing season. The OYS forecast of gross corn yield then is based on the following formula:

$$\text{Gross Yield} = [\text{number of ears} \times \text{weight per ear at 15.5\% moisture}] \div 56$$

Number of ears and ear weight are either forecast or actual and 56 is a conversion from weight to bushels.

### **The ASB**

The survey and forecasting procedures described above produce a number of indicators of the net yield of corn from the AYS and OYS. It is the job of the Agricultural Statistics Board (ASB) to combine the information from these indicators and determine the official published estimates. How does the ASB actually do this? The [latest version of the official NASS publication](#) on yield forecasting programs describes the ASB's role this way:

*“The sample surveys are designed to produce State level estimates of acreage, expected yields, final yields, and total production. The surveys are conducted by each State, and the first level of analysis is done by each State. Each Field Office does its independent appraisal of the relationships between the survey estimates and the final official statistics and forwards this information to Headquarters. While each Field Office is analyzing its survey data, statisticians in Headquarters are doing a parallel analysis of all survey data at the State, U.S., and regional levels. For the major field crops discussed in this paper, a formal Agricultural Statistics Board is convened to review regional indications and determine the official forecast or estimate. This Board is made up of 7 to 10 statisticians representing different divisions of NASS. Each Board member evaluates the regional survey indications and supporting data and determines their forecast or estimate. Each member brings their individual perspective to the review which can result in different conclusions being drawn. Through review and discussion, the Board must collectively reach a consensus and establish the National number. The Board process ensures all perspectives are examined and the national or regional forecast or estimate is the result of a thorough analysis. The summation of the individual State estimates as prepared by each State is compared to the Board number. The Headquarters statisticians will re-examine all national and*

*State data relationships and either adjust State estimates so they sum to the U.S. or change the previously determined U.S. number.”*

*“The formal meeting of the ASB to establish the final numbers and prepare the report is conducted under “lock up” conditions. Lock up begins with a complete isolation of all facilities required by the Board. All doors are locked, windows and elevators are covered and sealed, phones are disconnected, and the computer network inside “lock up” is isolated from the full network. Transmitters are not permitted and the area is monitored for electronic signals. Highly speculative data are decrypted only after the area is secure. Only after all security is in place does the Board begin final deliberations. The area remains locked up until a prescribed release time (8:30 a. m. for Crop Production) at which time the report is disseminated in electronic and paper forms.” (pp. 96-97)*

While this provides a helpful overview of the ASB in general times, it provides almost no detail on the nature of yield indications, the non-survey information, if any, used to supplement the survey indications, or how the ASB reaches a collective consensus.

In terms of statistical properties, NASS does acknowledge in its official yield forecasting publication that the AYS yield indications are biased:

*“Note that AYS data are the respondents’ expected yield. The crop may not be mature and ready for harvest for another several weeks after the respondent has been contacted. Experience has shown these responses tend to be conservative (biased down). Under drought conditions, this bias gets much larger as respondents perceptions of a crop are influenced by current weather conditions. Therefore, the interpretation phase of the review must recognize this tendency and factor it into the final deliberations.” (p. 19)*

Likewise, NASS acknowledges that OYS yield indications may also be biased:

*“Forecasts of state or regional yields are inherently subject to differ from the final, administratively determined, yield. The difference is due to weather and crop conditions yet to be encountered at the time the forecast is made, the difference between weather and crop conditions in the current year and the historic weather and crop conditions utilized to predict current yields, and systematic non-sampling errors which contribute to forecast error. NASS makes every attempt to minimize the impact of these three sources of error by means of administratively determine final values. These final, sometimes referred to as official, values are used as dependent variables to estimate an ordinary least squares equation with the averages calculated from objective yield samples as the independent variable. For example, each state will have an administratively determined, official corn yield for all previous years. To forecast the official corn yield for the current year, NASS regresses the objectively determined State level average corn yields from previous years to the corresponding official corn yields. This process provides a historical context for weather, biases, and non-sampling errors.” (p. 38)*

The fact that both yield indications are biased is interesting in and of itself, but the really important question is by how much and in what direction. We know the direction of bias for the AYS (downward) but neither the direction nor magnitude of bias for the OYS. If the biases are relatively small then this is largely a technical issue that has little practical import, but if the biases are large it becomes an important part of how the ASB generates published yield estimates.

A recent paper by Wang et al. (2011) is the first we know of to present data on actual AYS and OYS corn yield indications for an extended sample period. Figure 1 is taken directly from the Wang et al. article, and it presents: i) the OYS indications for the corn objective yield region and each report release month over 1993-2009; ii) the AYS yield indications for the same corn objective yield region and each report release month over 2001-2009; and iii) the December Agricultural Survey (DAS) corn yield indication for each year over 1996-2009. The objective region yield is simply the weighted-average yield for the objective yield states (10 since 2005) and the DAS is very close to the final yield estimates released in January. After viewing the chart, it is probably more understandable why NASS has been so cautious in disclosing this data previously. The magnitude of the bias in both yield indications is striking. The

coarseness of the scales prevents precise measurements, but the upward bias in the OYS, regardless of the release month, varies from around 10 to 15 bushels per acre (using the DAS as the benchmark). The downward bias in the AYS in August ranges from about 10 to 20 bushels, and then generally declines in September and October and is usually, but not always, fairly small in November. The pattern of bias in the AYS yield indications make sense, in that farmer-reported yields should become more accurate as the growing season progresses and harvest eventually is completed.

What this means from a practical standpoint is that the ASB is confronted with substantial bias in its two most important yield indications. For example, based on Figure 1, the final DAS objective region yield for corn in 2009 was about 175 bushels and the August OYS yield indication was around 190 bushels and the AYS yield indication was around 155 bushels, a spread of 35 bushels. A more recent paper by Nandram, Berg, and Barboza (2014) provides precise measurements. They report that the average biases for the OYS objective region corn yield indications relative to the DAS over 1993-2010 were 13.96, 11.62, 13.13, 14.99 and 14.97 bushels for August, September, October, November and December, respectively (note that we assume the bias estimates were stated in bushels but the article does not explicitly state the units). The average differences between the AYS indications and the DAS indications were respectively -12.26, -13.42, -9.81, and -4.10 bushels for August, September, October, and November. These imply AYS/OYS yield indicator spreads that average 26.22, 25.04, 22.94, and 19.09 bushels. It is important to keep in mind that these spreads are not symmetric around zero, as the AYS bias diminishes substantially by November but the OYS bias does not. Assuming these bias estimates are similar to the estimates available to the ASB when making published estimates in recent years, we can use the averages to provide a rough estimate of the OYS and AYS corn yield indications for the objective region in August 2016. The published objective region corn yield this year was 182 bushels, which combined with the bias estimates suggest an OYS yield of about 196 bushels and an AYS yield of about 170 bushels.

From this discussion, it should be obvious that the corn yield indications have to be adjusted for bias by the ASB in order to make useful published estimates. The key is the consistency of the bias across time. Figure 1 suggests there is some variation in the bias for a given report month across years and this likely contributes to the errors in the published corn yield estimates. Nonetheless, there appears to be a strong degree of consistency in the bias from year-to-year and this can be used to efficiently “de-bias” the AYS and OYS yield indications applying standard statistical techniques. The larger question is why the biases exist in the first place. Nandram, Berg, and Barboza (2014) conjecture that:

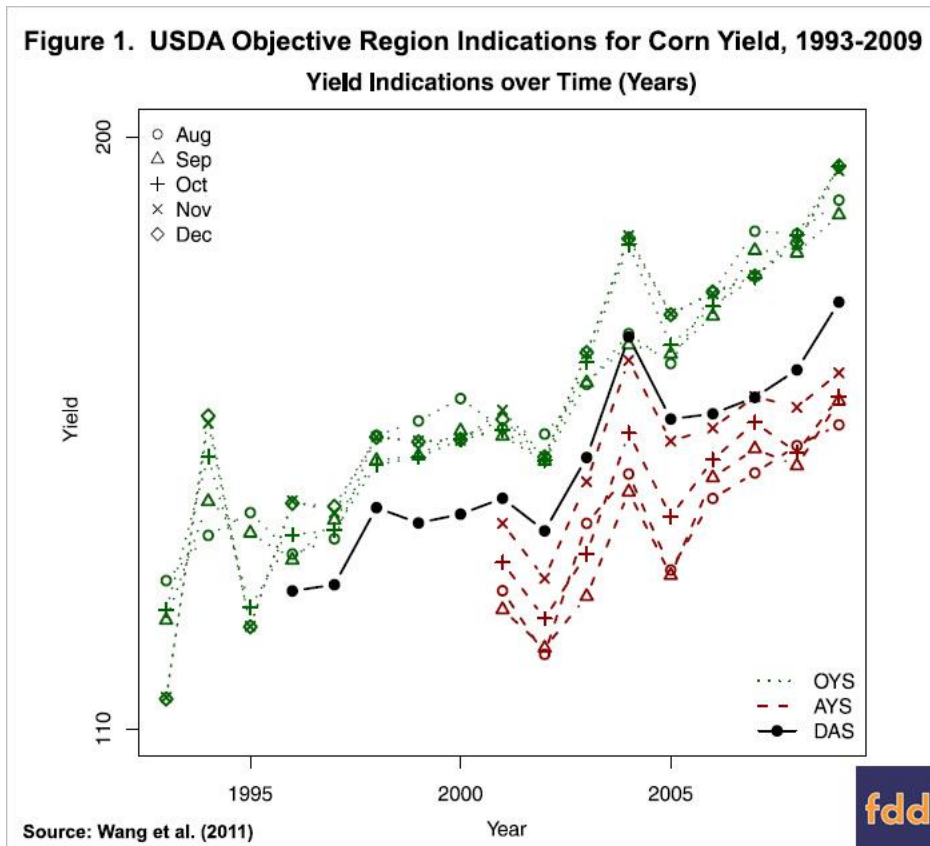
*One possible reason for the bias of the OYS indications is that the measurement process leads to a systematic overestimation of the plant density in a field. Two potential reasons for the biases in the AYS indications are pessimism on behalf of the farmers and exclusion of farms in the largest size stratum from the sampling frame. Large farms are conjectured to have higher average yields than moderately sized farms because of greater investment in advanced technology. (p. 510)*

There is also a large literature (e.g., Poate, 1988) that formally evaluates the accuracy of so-called “crop cutting” yield estimates, and the evidence shows that such techniques have a tendency to over-estimate actual yields for a variety of reasons, including variability in population density. So, it is not surprising that the OYS yield indicator is biased upwards. NASS has also been aware of the bias in OYS indications for decades and has conducted extensive research to better understand the underlying reasons for the bias, apparently without great success (Warren, 1985). The underlying reasons for the downward bias in the AYS yield indicator are hard to pinpoint beyond a general conservatism among farmers about yield prospects before harvest. Given the dramatic shrinkage in the AYS bias through November and the fact that the same farms are surveyed each month for a given year, it is unlikely that under-sampling large farm operations is one of the main reasons for the AYS bias, as suggested by Nandram, Berg, and Barboza (2014).

In light of the level and variability of the bias in AYS and OYS yield indications, it is sensible for the ASB to seek out additional non-survey information when generating published yield estimates. The Wang et al. (2011) article provides helpful information in this regard:

*Each month (August through December), estimates are produced by the ASB with the primary objective of providing the most accurate projection or end-of-season estimate given the survey*

indications and standard errors, in addition to other factors. When projecting corn yield, information on crop maturity, planting date, weather conditions, and other auxiliary factors is also taken into consideration by the ASB. (p. 85)



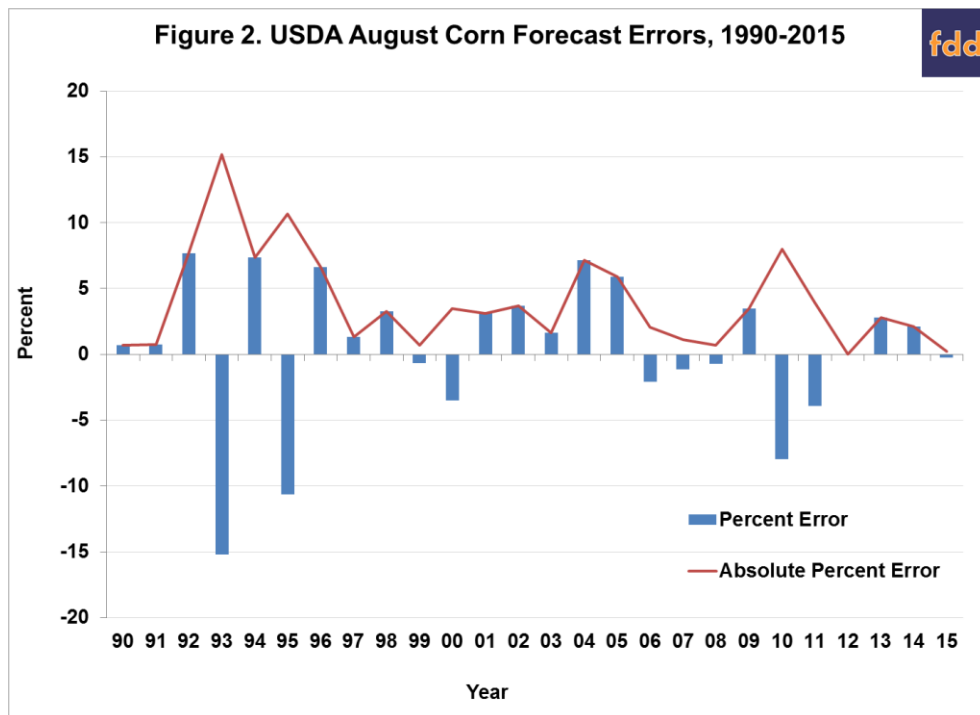
Since two of the co-authors of this paper at the time of its publication were NASS employees, we take this as authoritative evidence that the ASB uses “auxiliary” information when generating published estimates. While the details of exactly what type of auxiliary information the ASB uses when making corn estimates is not available, several recent articles by provide important clues. The stated purpose of the Wang et al. (2011), Adrian (2012), and Nandram, Berg, and Barboza (2014) articles is to basically formalize the decision-making process of the ASB in making corn yield forecasts. In order to do so, the models in the articles incorporate a variety of auxiliary information, including trend yields, planting progress, July temperature and precipitation, and crop condition ratings, which presumably reflect variables actually used by the ASB in the estimation process. Another type of auxiliary information used by the ASB in recent years is remote-sensing satellite data (Adrian, 2012; Johnson, 2014), such as the Normalized Difference Vegetative Index (NDVI), a measure of biomass density.

In sum, NASS has an active research program to try to formalize and improve the ASB methodology (e.g., Adrian, 2012; Wang et al., 2011; Nandram, Berg, and Barboza, 2014; Cruze, 2015). This provides important clues about the types of non-survey information used by the ASB, but we still don’t have any information about the weights placed on this information in ASB deliberations. There is obviously some concern within NASS about the subjective weights placed on the survey and non-survey information given discussions found in the articles cited above. For example, Adrian (2012) notes that the emphasis placed on different information, “will inevitably vary from person to person and depend upon the composition of the Board.” Our view is that an element of subjective expert judgement will likely always be a necessary component of a process as complex and difficult as forecasting crop yields. We concur with the late Bruce Gardner, who made this trenchant observation after sitting through ASB lock-up deliberations for an early 1990s *Crop Production* report:

“A NASS board in Washington then assesses all the indicators of yield, including the estimates of a month earlier. This is not done using a pre-specified formula—in which case a computer could replace the NASS board—but through a consensus of the Board members based on their experience and the full information before them.” (Gardner, 1992, p. 1068)

## ASB Track Record

When viewing the details of the ASB forecasting methodology, especially the magnitude of the bias in the AYS and OYS yield indications, it is easy to lose sight of the big picture. That is, the ASB is tasked with combining the AYS and OYS yield indications with other relevant information in order to make the best possible published estimates of corn yield, and it is the accuracy of the final published estimates that ultimately matters. So, we updated our earlier evaluation (*farmdoc daily*, August 29, 2014) of the historic accuracy of USDA published August forecasts of the U.S. average corn yield relative to the “final” yield estimate released in January (we say “final” because January estimates are sometimes revised based on the *Agricultural Census* conducted every five years). The differences between the forecasts and the final estimates in percentages over 1990-2015 are presented in Figure 2. When interpreting the errors, note that a positive error implies an under-estimate on the part of USDA and a negative error implies an over-estimate. The errors associated with the USDA corn yield forecasts are occasionally very large, such as 1993 and 1995. These examples of large errors are not surprising due to the unusual weather events that occurred in those years. It is interesting to note that USDA corn yield forecast errors in 2012 were extremely small, with the August forecast exactly equal to the final estimate, and thus, having a zero forecast error. This is surprising given the magnitude of drought conditions that prevailed in 2012 and the difficulty of forecasting corn yields under these circumstances. Accuracy was also exceptional in 2015, when the forecast error was only -0.2 percent. The plot suggests a clear downward trend in USDA corn forecast errors over time. Forecast performance after 2011 has been particularly impressive.



It is important to emphasize that the USDA track record of forecasts evaluated here is based on “real-time” forecasts. In other words, this is the record of actual forecasts produced and released to the public over time. As Tetlock and Gardner (2015) observe in their best-selling book *Superforecasting: The Art and Science of Prediction*, measurement and evaluation of a large sample of real-time quantitative forecasts is the only scientifically-valid way of determining true forecasting skill. Recently, there has been a considerable media coverage of several new “big data” firms that claim to produce superior corn yield forecasts to that of the USDA (e.g., Brokaw, 2016; Woyke, 2016). What is not emphasized is that the



evidence for these claims is based on historical simulations of model performance, or “backcasting,” a much lower and less rigorous hurdle. These firms may actually be able to beat the USDA, but real-time track records will be required to convincingly prove the point. Until then, the available evidence indicates it is hard to beat the accuracy of USDA corn yield forecasts.

## Implications

The Agricultural Statistical Board (ASB) of the National Agricultural Statistics Service (NASS), is responsible for determining the official published corn yield estimates of the USDA. While there is considerable information available about the survey methodology used by NASS to generate yield indications, until recently, very little information has been publically available about the nature of the corn yield indications, the non-survey information, if any, used by the ASB to supplement the survey indications, or how the ASB reaches a collective consensus. Several publications have appeared in the scientific literature during the last several years that help open up the ASB “black box.” These publications present for the first time historical data on the separate Agricultural Yield Survey (AYS) and the Objective Yield Survey (OYS) corn yield indications and the results are striking. The upward bias in the OYS corn yield indications, regardless of the release month, varies from around 10 to 15 bushels per acre, while the downward bias in the AYS in August ranges from about 10 to 20 bushels, and then generally declines in September and October and is usually, but not always, fairly small in November. Given the level and variability of the bias in AYS and OYS yield indications, it is sensible for the ASB to seek out additional non-survey information when generating published yield estimates. The recent publications indicate that the ASB incorporates a variety of auxiliary information, including trend yields, planting progress, July temperature and precipitation, crop condition ratings, and satellite imagery data. Unfortunately, we still don’t have any information about the weights placed on this information in ASB deliberations. While it is important to better understand all aspects of the NASS corn yield forecasting methodology, and ASB procedures in particular, in the end what matters is the accuracy of the final published estimates. On that score, the evidence suggests USDA corn yield forecasts are still hard to beat.

## References

- Adrian, D. “A Model-Based Approach to Forecasting Corn and Soybean Yields.” Fourth International Conference on Establishment Surveys, 2012.  
<http://www.amstat.org/meetings/ices/2012/papers/302190.pdf>
- Brokaw, A. “This Startup Uses Machine Learning and Satellite Imagery to Predict Crop Yields.” *The Verge*, August 4, 2016. <http://www.theverge.com/2016/8/4/12369494/descartes-artificial-intelligence-crop-predictions-usda>
- Cruze, N.B. “Integrating Survey Data with Auxiliary Sources of Information to Estimate Crop Yields.” JSM Proceedings, Survey Research Methods Section. Alexandria, VA: American Statistical Association, 2015.  
<https://www.amstat.org/sections/srms/proceedings/y2015/files/233920.pdf>
- Gardner, B.L. “How the Data We Make Can Unmake Us: Annals of Factology.” *American Journal of Agricultural Economics* 74(1992):1066-1075.
- Good, D., and S. Irwin. “USDA Corn and Soybean Acreage Estimates and Yield Forecasts: Dispelling Myths and Misunderstandings.” Marketing and Outlook Brief 2011-02, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, March 2011.
- Good, D.L., and S.H. Irwin. “Understanding USDA Corn and Soybean Production Forecasts: Methods, Performance and Market Impact over 1970-2005.” AgMAS Project Research Report 2006-01, Department of Agricultural and Consumer Economics, University of Illinois at Urbana Champaign, February 2006.
- Irwin, S., and D. Good. “Understanding “Implied Ear Weight” in USDA’s August Corn Yield Forecast.” *farmdoc daily* (6):156, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, August 18, 2016.

Irwin, S., D. Good, and D. Sanders. "Are USDA Corn Yield Forecasts Getting Better or Worse over Time?" *farmdoc daily* (4):166, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, August 29, 2014.

Irwin, S.H., D.R. Sanders, and D.L. Good. "Evaluation of Selected USDA WAOB and USDA Forecasts and Estimates in Corn and Soybeans." Marketing and Outlook Research Report 2014-01, Department of Agricultural and Consumer Economics, University of Illinois at Urbana-Champaign, January 2014.

Nandram, B, E. Berg, and W. Barboza. "A Hierarchical Bayesian Model for Forecasting State-Level Corn Yield." *Environmental and Ecological Statistics* 21(2014): 507-530.  
<http://link.springer.com/article/10.1007/s10651-013-0266-z>

NASS. "The Yield Forecasting Program of NASS." SMB Staff Report No. SMB12-01, Statistical Methods Branch, National Agricultural Statistics Service, U.S. Department of Agriculture, May 2012.  
[https://www.nass.usda.gov/Publications/Methodology\\_and\\_Data\\_Quality/Advanced\\_Topics/Yield%20Forecasting%20Program%20of%20NASS.pdf](https://www.nass.usda.gov/Publications/Methodology_and_Data_Quality/Advanced_Topics/Yield%20Forecasting%20Program%20of%20NASS.pdf)

Tetlock, P.E., and D. Gardner. *Superforecasting: The Art and Science of Prediction*. New York, NY: Crown Publishing, 2015.

Wang, J.C., S.H. Holan, B. Nandram, W. Barboza, C. Toto, and E. Anderson. "A Bayesian Approach to Estimating Agricultural Yield Based on Multiple Repeated Surveys." *Journal of Agricultural, Biological and Environmental Statistics* 17(2012): 84-106. <http://link.springer.com/article/10.1007/s13253-011-0067-5>

Warren, F.B. "Evaluation of the 1983 Corn Objective Yield Validation Survey." SRS Staff Report No. AGES850109. Washington DC: USDA, Statistical Reporting Service, 1985.  
[https://www.nass.usda.gov/Education\\_and\\_Outreach/Reports,\\_Presentations\\_and\\_Conferences/Yield\\_Reports/Evaluation%20of%20the%201983%20Corn%20Objective%20Yield%20Validation%20Study.pdf](https://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/Yield_Reports/Evaluation%20of%20the%201983%20Corn%20Objective%20Yield%20Validation%20Study.pdf)

Woyke, E. "Crystal Ball for Corn Crop Yields Will Revolutionize Commodity Trading." *MIT Technology Review*, August 9, 2016. <https://www.technologyreview.com/s/602096/crystal-ball-for-corn-crop-yields-will-revolutionize-commodity-trading/>