Statistical Science
2019, Vol. 34, No. 2, 301–316
https://doi.org/10.1214/18-STS687

Producing Official County-Level Agricultural Estimates in the United States: Needs and Challenges

Nathan B. Cruze, Andreea L. Erciulescu, Balgobin Nandram, Wendy J. Barboza and Linda J. Young

Abstract. In the United States, county-level estimates of crop yield, production, and acreage published by the United States Department of Agriculture's National Agricultural Statistics Service (USDA NASS) play an important role in determining the value of payments allotted to farmers and ranchers enrolled in several federal programs. Given the importance of these official county-level crop estimates, NASS continually strives to improve its crops county estimates program in terms of accuracy, reliability and coverage. In 2015, NASS engaged a panel of experts convened under the auspices of the National Academies of Sciences, Engineering, and Medicine Committee on National Statistics (CNSTAT) for guidance on implementing models that may synthesize multiple sources of information into a single estimate, provide defensible measures of uncertainty, and potentially increase the number of publishable county estimates. The final report titled *Improving Crop Estimates by Integrating Multiple Data Sources* was released in 2017. This paper discusses several needs and requirements for NASS county-level crop estimates that were illuminated during the activities of the CNSTAT panel. A motivating example of planted acreage estimation in Illinois illustrates several challenges that NASS faces as it considers adopting any explicit model for official crops county estimates.

Key words and phrases: Agricultural surveys, auxiliary data, benchmarking, official statistics, small area estimation.

1. INTRODUCTION

The United States Department of Agriculture (USDA) has been involved in producing county-level agricultural estimates since 1917 (Iwig, 1996). Presently, county-level agricultural estimates of crop acreage, production and yield published by the USDA's National Agricultural Statistics Service (NASS) inform many crop insurance and agricultural support programs administered by other USDA agencies. Two such USDA agencies include the Farm Ser-

Nathan B. Cruze is Mathematical Statistician in Research and Development Division, USDA National Agricultural Statistics Service (NASS), 1400 Independence Avenue, SW, Washington, DC 20250, USA (e-mail: nathan.cruze@nass.usda.gov). Andreea L. Erciulescu is Senior Statistician, Westat, 1600 Research Boulevard, Rockville, Maryland 20850, USA (e-mail: AndreeaErciulescu@westat.com). Balgobin Nandram is Senior Professor of Statistics, Worcester Polytechnic Institute, Department of Mathematical Sciences, Stratton Hall, 100 Institute Road, Worcester, Massachusetts 01609, USA (e-mail: balnan@wpi.edu). Wendy J. Barboza (retired) is former Deputy Director of Research and Development Division, USDA National Agricultural Statistics Service, 1400 Independence Avenue, SW, Washington, DC 20250, USA (e-mail: wbarboza@cox.net). Linda J. Young is Chief Mathematical Statistician and Director of Research and Development Division, USDA National Agricultural Statistics Service, 1400 Independence Avenue, SW, Washington, DC 20250, USA (e-mail: linda.young@nass.usda.gov).

vice Agency (FSA), and the Risk Management Agency (RMA), which manages the Federal Crop Insurance Corporation. (A list of frequently used abbreviations is provided in the Appendix.) While both FSA and RMA have relied upon NASS crops county estimates as important inputs into decision-making processes as they oversee their respective agricultural programs, recent policy changes legislated in the Agricultural Act of 2014 (the 2014 Farm Bill) introduced a revenue loss protection program at the county level (called the Agricultural Risk Coverage County Option, or ARC-CO) that couples price estimates with yield estimates at the county level to determine a threshold revenue that triggers payment for covered commodities within the county. In the context of this recent program, NASS's published crops county estimates may now be a *determining factor* of the magnitude of disbursements individual farmers receive. If NASS suppresses county-level crops estimates for any reason, FSA must derive the corresponding quantity information by other means in order to administer its mandated programs.

Motivated in part by the need to incorporate a variety of existing and potential data sources, NASS has renewed its interest in small area estimation and model-based approaches to combining survey data with other auxiliary information. In particular, NASS convened an expert panel under the auspices of the National Academies of Sciences, Engineering and Medicine Committee on National Statistics (CNSTAT) to review its crops county estimates program and to advise on the possible role model-based estimates could serve in the production of official statistics of acreage, production and yield at the sub-state level. The report, titled *Improving Crop Estimates by Integrating Multiple Data Sources*, became publicly available in 2017 (National Academies of Sciences Engineering, and Medicine, 2017). This paper discusses some of the particular needs and challenges of the NASS crops county estimates program that were elaborated during meetings of the CNSTAT panel. The NASS survey cycle, input data sources and NASS's current method for producing sub-state crop estimates are discussed in Section 2. In Section 3, a case study for planted acreage in Illinois highlights some of the challenges NASS faces as it considers the increased use of model-based estimation in the production of its official statistics. Discussion and conclusions are offered in Section 4.

2. OFFICIAL SMALL AREA STATISTICS AND NEEDS OF THE NASS CROPS COUNTY ESTIMATES PROGRAM

Like NASS, other official statistical agencies within the United States and worldwide face demands for statistics at finer levels of disaggregation with increasing frequency. To meet these demands, models incorporating survey data and a variety of auxiliary data are often proposed to enhance the reliability of small area estimates, for example, county-level estimates, where even well-designed surveys alone may not have the sample density or timeliness needed to provide reliable estimates. Producing small area estimates that can rise to the level of *official statistics* is a task that requires methodical planning, understanding of end uses, dedicated resources and the buy-in of numerous stakeholders.

In particular, the United States Census Bureau's history with the Small Area Income and Poverty Estimates (SAIPE) program illustrates many of the challenges shared by official statistical agencies when implementing these changes. An overview of the SAIPE program can be found in Bell, Basel and Maples (2016); SAIPE was conceived based on the need for postcensal estimates of income distribution and poverty estimates at state and sub-state areas including counties, cities and school districts (United States Census Bureau, 2017). Today, SAIPE plays an important role in the allocation of Title I aid to school districts and local geographies in over 3100 counties nationwide. Its creation and evolution are documented in no fewer than six National Research Council reports, each a product of a panel like the one convened by NASS; see National Research Council (1997, 1998, 1999, 2000a, 2000b, 2007).

NASS first convened its own CNSTAT panel in November 2015, with four subsequent meetings held in 2016 and 2017. In order to more fully assess NASS's needs, the panel inquired extensively about the NASS survey cycle, currently available data, and the role of the Agricultural Statistics Board (ASB) and expert assessment used to produce NASS official statistics. The panel's activities culminated in the release of its publicly available report in late 2017.

2.1 NASS Survey Cycle

The demands placed on the NASS crops estimates program include timely estimates of planted area, harvested area, production and yield for the diverse array of crops grown throughout the entire United States.

To address this need, NASS implements a multivariate probability proportional to size (MPPS) survey design in many of its crop surveys. Bailey and Kott (1997) note that this design offers additional flexibility to target key crops grown within each state.

A partial NASS survey cycle and publication timeline is depicted in Figure 1; the width of each interval represents the approximate data collection window for each survey. NASS conducts quarterly Acreage, Production and Stocks (APS) surveys in an ongoing effort to capture activities throughout the life cycle of the crop, including planting intentions (March), early estimates of planted acreage (June), and estimates of harvest and output activities for small grains crops (September) and major row crops (December). An area frame component of the June survey provides an undercoverage adjustment for the list-based samples obtained during the September and December APS surveys. Coverage-adjusted *national and state* survey estimates are available to inform the official ASB consensus estimates for state and national activity that are released in the Small Grains Summary in late September or in the Annual Summary (for row crops) in January of the following calendar year; see USDA NASS (2014, 2015) for examples of these annual publications produced during the 2014 crop year.

In 2011, NASS fully implemented the County Agricultural Production Survey (CAPS) to augment county-level sample sizes and to standardize the data collection procedures nationwide. At the conclusion of CAPS data collection, the MPPS samples of both the APS survey and CAPS are pooled and reweighted. The resulting direct estimates for sub-state domains are essentially derived from a single *list-based* sample. Consequently, some degree of undercoverage in the survey direct estimates is still to be expected, although the extent may vary by state and by crop type. Acquisition of additional observations during CAPS data collection can potentially more than double the total number of survey responses obtained within each state relative to the APS survey sample size alone. However, this does not guarantee that the number of reports *for each commodity* will double as the sampled respondents may not grow all types of crops for which NASS produces estimates.

NASS conducts the row crops CAPS surveys in 43 states excluding the 5 New England states (shown in white) in Figure 2. The small grains CAPS is also conducted in the group of 37 states shown in light gray. The list of commodity crops targeted may differ from state to state and from year to year, subject to providing required coverage for federally mandated program crops, and satisfying the needs of other stakeholders, for example, specific state program commodities. The official county estimates for small grains (e.g., barley, oats and wheat) are published in December. The first row crops county estimates for corn, soybeans, sunflower and sorghum are published in February of the following calendar year. Row crops county estimates for additional commodities are subsequently released at intervals, concluding with the release of county estimates of potatoes in the month of October. As shown in Figure 1, the data collection window for CAPS extends beyond the release of the national and state-level official statistics. Therefore, benchmarking to previously published state acreages, production and yield is a necessary step in order to ensure consistency of estimates at all sub-state levels. While the survey is an integral part of the process, the ASB can and does incorporate other auxiliary information in the production of official crop statistics when it is available.

2.2 Survey Data and Other Auxiliary Data Types

Although the county is generally the smallest area and the level at which many USDA policies are to be administered, NASS also produces estimates at an intermediate domain between the county and state called the agricultural statistics district (ASD). The ASD is a predefined group of neighboring counties within a state. The median number of ASDs per state is 9, but small states may have only a single ASD whereas Texas has 15. (Maps of these administrative boundaries have been included in the Supplementary Material (Cruze et al., 2019).) NASS utilizes its own survey data, administrative data provided by FSA and RMA, and remote sensing data available at both the county level and the ASD level when setting official sub-state crop estimates.

NASS obtains auxiliary county-level acreage data from administrative sources. Farmers who choose to participate in FSA programs certify the acreages and crop types that they grow with FSA. FSA programs are popular but not compulsory, and the FSA planted acreage data may have some degree of undercoverage. The extent of undercoverage in the FSA data can differ by commodity, by state and even by county within state. For example, known Amish communities in Pennsylvania and other midwestern states may represent significant portions of local agricultural activity but tend not to participate in federal or commercial crop insurance programs. In NASS's current operating

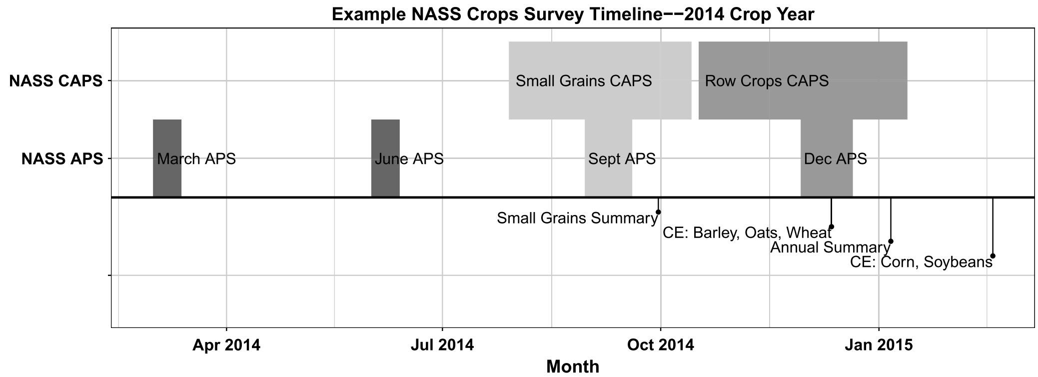

FIG. 1. *Typical NASS crop survey and publication timeline.*

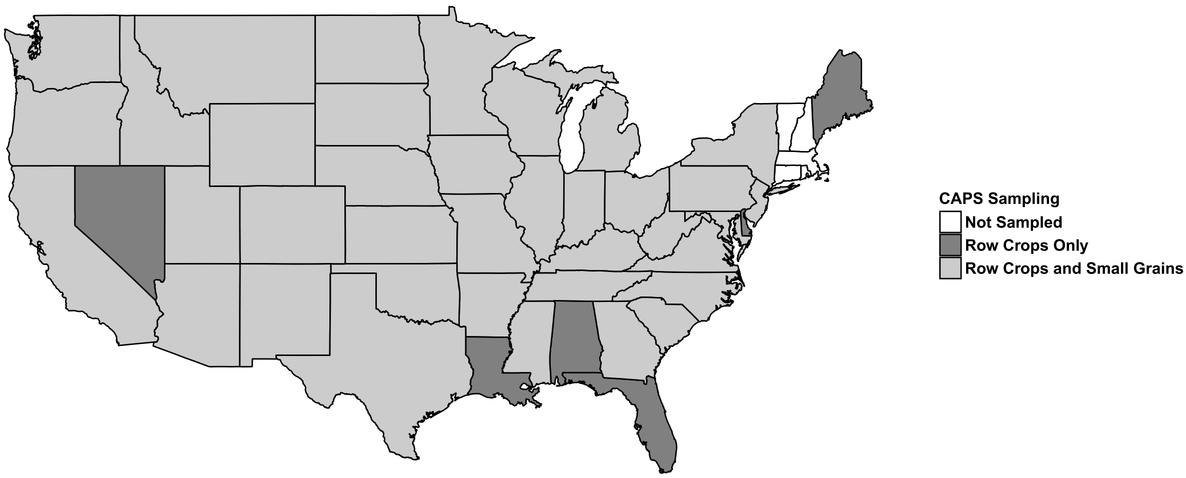

FIG. 2. *Row crops and small grains CAPS states.*

environment, the FSA planted acreage data are generally thought to represent a minimum amount of planting activity that is known to have taken place within the county lines.

As an underwriter of crop insurance policies, the RMA receives administrative data on crop area including *failed acreage* (acreage that was planted but not harvested for any reason) from various independent crop insurance agents. In ideal conditions, a high proportion of planted area would be harvested, however, a variety of conditions including local weather and more widespread natural disasters including drought, hail, high winds and flooding can contribute to the loss of crops. The RMA administrative data on failed acreage likely represent a minimum amount of crop abandonment as not all farmers will participate in a crop insurance program, and even those who do may not file a claim in the given crop year.

Described in Boryan et al. (2011), NASS's Cropland Data Layer[1] (CDL) is a remotely sensed land-cover classification product that identifies a wide variety of crops. CDL products were generated as early as 1997, but only for select states due to computational capacities at the time. The CDL first offered coverage for the conterminous United States in 2009 (Boryan, 2010, page 2). Note that the CDL achieves high classification accuracies for major crops (85% to 95%), but as Gallego (2004) and Czaplewski (1992) indicated, biases may be realized in areal estimates derived from pixel counting as a result of misclassification. As such, the CDL is released to the public as a complementary data product. Whenever the CDL is used internally for acreage estimation, NASS adjusts for the (typically downward) biases through regression modelling incorporating survey data, including the unit-level regres-

[1] NASS's CDL can be freely accessed at https://nassgeodata.gmu.edu/CropScape/.

sion models for county-level planted area discussed in Section 2.3.1.

The CDL product itself pertains directly to crop area as opposed to total production, but the crop-specific masks inform NASS's in-house remote sensing approaches to crop yield forecasting at regional and state levels, and to estimating county-level corn and soybean yields in corn belt states (Johnson, 2014). Remote sensing and other new and emerging technologies are increasingly of interest to NASS and the agricultural sector at large. Johnson (2016) studied the relationship between MODIS products and crop yields for 10 distinct commodity crops. Precision agricultural data could potentially be obtained from instruments found onboard tractors, combines and even drones. In addition to the logistical challenges in obtaining, storing, and processing such data, testimony provided to the House Committee on Agriculture highlights issues of data ownership and privacy for precision agricultural data (United States Government Publishing Office, 2015, 2016). These policy issues are still being debated and developed at the Congressional and Departmental levels.

2.3 Production of NASS Official Statistics

Incorporating multiple sources of information into official crop estimates has been a hallmark of the NASS crops estimates programs. Traditionally, the mechanism for combining possibly disparate sources of information has been through expert assessment by the ASB. As described in an earlier report by the National Research Council, the ASB "process is intended to maximize the use of what is believed to be the best information and to ensure consistency with other estimates published by USDA." (National Research Council, 2008, page 1230). The ASB and the NASS regional field office staff appointed to serve in the esimation process set smaller-than-state estimates upon review of several sources of data.

2.3.1 *Sub-state crop estimates*. At both the ASD and county level, NASS produces estimates of four parameters of interest: total planted area, total harvested area, total production and yield. Since the state estimates are determined prior to setting county estimates, NASS employs a "top-down" strategy, first setting ASD-level estimates for acreage and production subject to equation (1) and then setting corresponding county estimates subject to equation (2)

$$\text{(1)} \qquad Total_{\text{State}} = \sum_{\text{ASD}\in\text{State}} Total_{\text{ASD}},$$

$$\text{(2)} \qquad Total_{\text{ASD}} = \sum_{\text{counties}\in\text{ASD}} Total_{\text{county}}.$$

Within each administrative boundary, estimates of yield are obtained by dividing total production estimates by corresponding total harvested acreage estimates.

For each parameter to be estimated, NASS commodity statisticians have the opportunity to create a composite of the various input sources available to aid their deliberation. A prescribed set of weights is given as a starting point, and may depend on the availability of certain inputs, for example, remotely sensed data (National Academies of Sciences Engineering, and Medicine, 2017, page 21). In some sense, these composite estimates are implicit models or rules to maximize the use of the best data at hand, and the choice of weights may embody a *prior belief or knowledge* about the relative strengths of each component. The weight assigned to each component in the composite may vary by state and commodity, but final estimates are always subject to ASB's review and approval.

TABLE 1

*Summary of inputs typically reviewed for ASD and county crop estimates. Abbreviations in parenthesis identify these variables in Figure 3 below. **Denotes inputs which may not be available for all commodities or within every state*

	Planted Area	*Harvested Area*	*Production/Yield*
Inputs	Survey Planted Total (CAPS.PL) FSA Certified Acreage (FSA)	Survey Harvested Total (CAPS.HV) **PL**×Survey Harvested/Planted Ratio (CAPS.HP)	Survey Production Total (CAPS.PD) **HV**×Survey Yield Ratio (CAPS.YD)
	Battese–Fuller (B.F)**	**PL**–RMA Failed Acreage (RMA) Previous Year **HV** (PY.HV)	**HV**×Remote Sensing Yield** (RS.YD) Previous Year **PD** (PY.PD)
ASD Estimates	$\mathbf{PL}_{\mathbf{ASD}}$	$\mathbf{HV}_{\mathbf{ASD}}$	$\mathbf{PD}_{\mathbf{ASD}}$ and $\mathbf{YD}_{\mathbf{ASD}}$
County Estimates	$\mathbf{PL}_{\mathbf{county}}$	$\mathbf{HV}_{\mathbf{county}}$	$\mathbf{PD}_{\mathbf{county}}$ and $\mathbf{YD}_{\mathbf{county}}$

Letting **PL**, **HV**, **PD** and **YD** denote official estimates of planted area, harvested area, production and yield, respectively, Table 1 summarizes the process by which a commodity expert would first set ASD-level estimates for all four parameters of interest followed by county-level estimates of the same parameters. Beginning with planted area, a commodity statistician reviews the ASD-level survey direct estimate, administrative acreage data obtained from FSA and, where available, a Battese–Fuller model-based estimate of planted area. Note that the latter is an early variant of the celebrated Battese–Harter–Fuller *unit-level nested error regression model* that estimates small areas by adjusting NASS's remotely-sensed CDL pixel count data using NASS's June Area Survey data as a source of ground truth (Fuller and Battese, 1973, Walker and Sigman, 1984, Battese, Harter and Fuller, 1988). A composite estimate aids the deliberation; for planted area, FSA administrative data often carry a high weight. Finally, the official ASD-level estimate of **PL** satisfying equation (1) is established in accordance with NASS rounding rules and enforcing relationships with FSA as a lower bound.

Once ASD-level estimates for **PL** are established, the commodity expert may proceed to setting ASD-level estimates for harvested area. Two current year CAPS direct estimates may be used to help inform harvested area estimates: a survey-based harvested area total, as well as a survey-based ratio of harvested to planted area multiplied by **PL** as determined in the previous step. Additionally, an input derived from RMA administrative acreage data may be assessed along with the past year's harvested area estimate. Current crop year estimates of harvested acreage (**HV**) must be derived satisfying equation (1) and subject to the additional constraint that **HV** $\leq$ (**PL**-*RMA failed acreage*) within every area of interest; this condition ensures that the harvested area is no greater than planted area (a physical necessity) and that administrative data on abandoned acres are honored.

In a similar manner, all ASD-level estimates of production (**PD**) are set subject to equation (1), based on an assessment of survey estimates, past official estimates and remote sensing yield estimates. Note that the remote sensing estimates are only available for corn and soybeans in a limited number of states, and they do not directly incorporate any survey information at this time; see Johnson (2014) for details. The assessment of distinct inputs takes place on the production scale; all corresponding official yield estimates are derived as the ratio of **PD** and **HV** estimates. Once all ASD-level estimates for a particular commodity have been set, county estimates are set in an analogous manner, satisfying equation (2) for each total.

Figure 3 illustrates the result of this completed process for a subset of counties within a major corn producing state. County identifiers have been replaced with an index number shown on the vertical axes of each panel below. Membership in an ASD has been preserved, and counties bounded between solid black horizontal lines fall within the same ASD. Each cell in the row corresponds to a different input or resulting estimate for the indexed county. Color indicates the magnitude of these data, and the figure illustrates some of the within and between ASD variation that may be present, even in states that are major producers of a commodity. All cells to the left of the "Composite" column represent the inputs as previously enumerated in Table 1. NASS's statisticians set composite values based on available separate inputs. The "Board" statistic is approximately equal to this composite, up to manual benchmarking against state and ASD totals and enforcement of NASS rounding rules. The synthesis of separate inputs gives rise to estimates of the three totals: **PL**, **HV** and **PD**. Official **YD** estimates are the ratio of the official **PD** and **HV** totals. While NASS will attempt to set estimates in all counties where the presence of a commodity is indicated, some official county estimates may ultimately be suppressed, see, for example, counties 14, 16 and 19, in keeping with NASS county estimates publication standards. Additional figures provided in the Supplementary Material contrast the availability of these types of inputs in other states.

2.3.2 *The current publication standards.* NASS incorporates information from outside the survey design to set its county estimates, but the basis of its current publication standard is tied to features of the survey sample alone (National Academies of Sciences Engineering, and Medicine, 2017, page 28). NASS uses a compound rule for its publication standard, verifying either:

- a minimum number of reports, or
- a minimum harvested area coverage threshold.

Item-level nonresponse is permitted for production and yield, meaning that a respondent taking the survey could provide acreage information but decline to provide total production or yield. The number of reports that determine production and yield may be smaller than the number of reports used to estimate acreages.

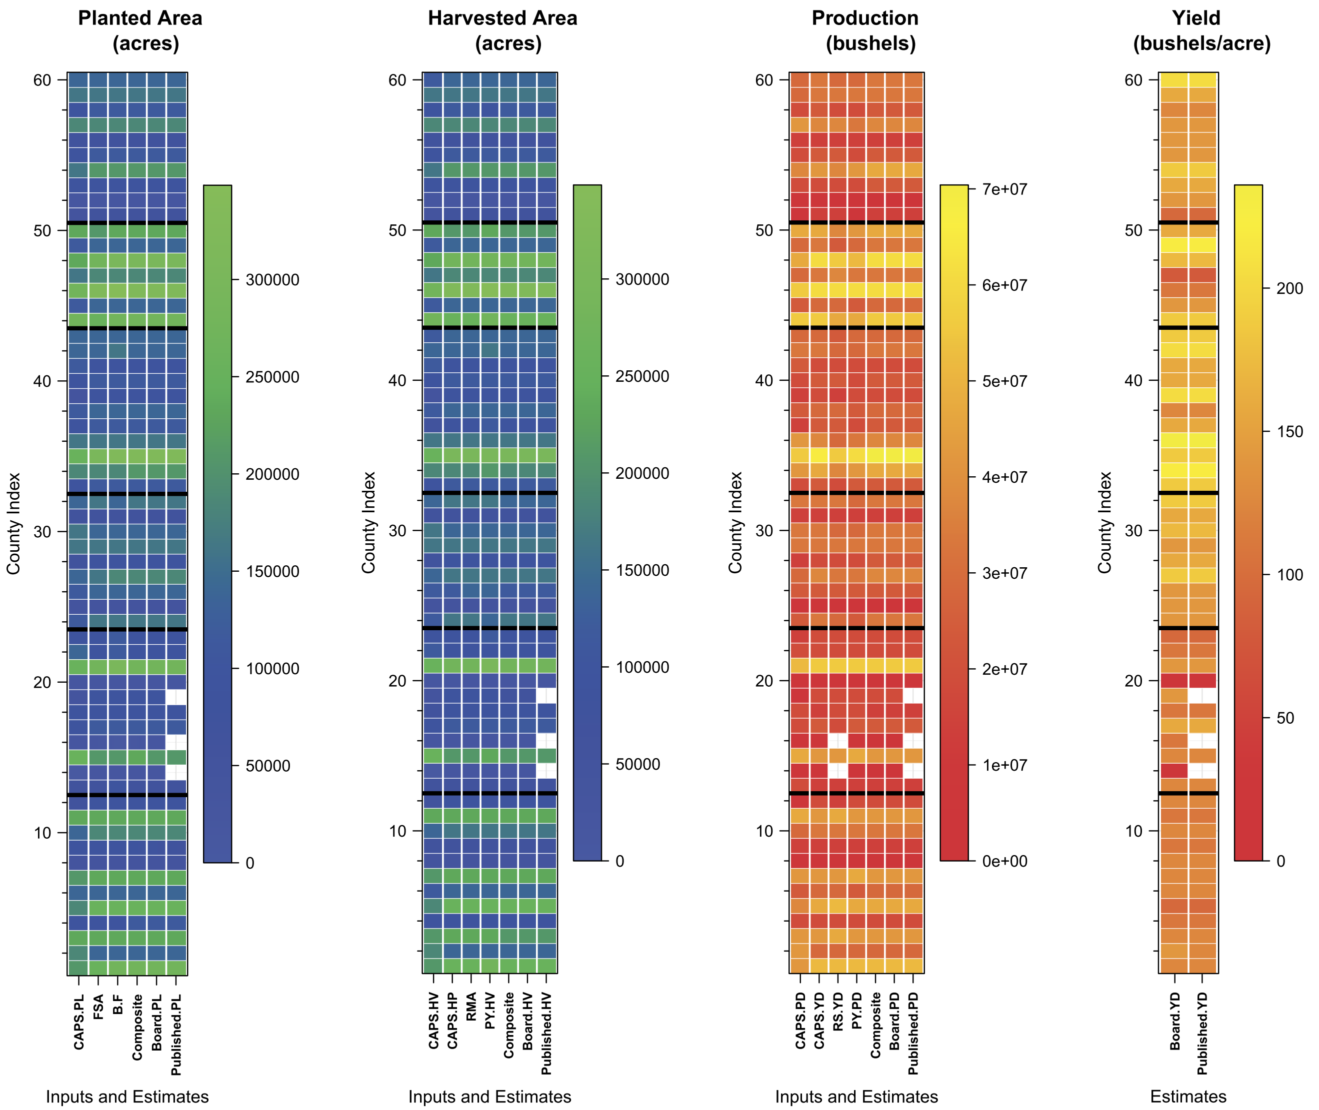

FIG. 3. *Available inputs and resulting estimates for select counties in a major corn producing state.*

NASS will first check that at least 30 valid reports of positive production or yield (respondents may report either quantity on the questionnaire) have been realized within the county, in which case the county estimate may be published. If 30 positive production or yield reports cannot be realized in the county for the commodity of interest, a second check for coverage is made to determine whether county estimates can still be published. A sum of unweighted harvested area reports *based on the (possibly smaller) number of realized yield reports* is compared to the harvested area *estimate* (the official statistic **HV** determined by the ASB); county estimates can still be published provided a minimum of 25% coverage is obtained based on at least 3 positive yield reports. NASS will either publish estimates of all parameters with respect to each commodity in the county (i.e., planted and harvested area, production, and yield) or it will suppress all estimates of that commodity. Counties that must be suppressed will be grouped into larger aggregates of counties within the state, and those aggregates must also pass the publication standard. These rules are applied for each commodity independently; it may be possible to publish estimates for one commodity in a given county but necessary to suppress estimates for a different crop within the same county. Even when NASS suppresses county estimates, the FSA and RMA must continue to administer programs in those counties. Although NASS produces aggregates of several counties or ASD-level estimates, those estimates may lack the specificity required to administer good local farm policy.

Bell and Barboza (2012) previously considered the relationship between the coefficient of variation associated with CAPS survey estimates and the coverage threshold, noting that the coverage threshold may include some counties that lack adequate precision, and it may exclude some counties with more reasonable precision from publication. One caveat in that analysis and comparison is that it only considered the coefficient of variation associated with the probability survey, whereas NASS official statistics often incorporate auxiliary information from outside the survey design. The benefit of doing so may not be reflected in measures of uncertainty from the survey alone.

2.4 CNSTAT Panel Findings and the Potential for Model-Based Estimates

The process of manual assessment of separate inputs is time consuming, and it must be repeated for each state and commodity separately. As Table 1 showed, the official estimates of harvested area are functions of the planted area estimates set in the previous step. In turn, the production estimates are a function of the harvested area estimates. Their magnitudes are determined conditionally, and any errors in this process may be propagated through this sequence of estimates and down to the county scale. Quantification of the uncertainty associated with the official estimates produced in this manner is difficult, *and presently, NASS does not publish any measures of uncertainty for its official sub-state crop estimates.* Presumably, weather, drought and soil information may be informative for estimates of crop production and yield. Aside from the remote sensing yield estimate, there is no other auxiliary data that translates any changes due to these factors directly onto the production scale for use by the commodity expert.

A primary motivation for small area estimation and model-based approaches is that it may help stabilize estimates for domains that lack sufficient precision due to small sample sizes (Rao and Molina, 2015). The CNSTAT panel considered the additional role model-based estimates could play in the production of NASS official statistics at the county and ASD levels in order to increase the number of published county estimates and better support other USDA agencies in their respective missions. In its major findings, the panel recommended the increased use of model-based estimation to improve crop estimates, which could help meet the following needs of the NASS crops county estimates program:

- enhanced reproducibility and transparency,
- quantification of associated uncertainties,
- incorporation of a wider variety of data types,
- increased number of publishable county estimates, given an appropriate publication standard.

3. CHALLENGES IN MODELING: A CASE STUDY OF PLANTED CORN IN ILLINOIS

The combination of multiple sources of data is a recurring theme throughout the NASS estimation program and at virtually all geographic levels. The traditional role of the ASB has been to assess several sources of information and come to a one-number, consensus estimate given that information. As the National Research Council noted, the ASB's estimation process ensures consistency with other estimates published by USDA. This is often achieved on the basis of point-to-point comparisons of several area estimates. In the case of planted area, this point-to-point comparison to FSA results in an unusual notion of model accuracy.

3.1 Past Research

An overview of many model-based techniques currently used in NASS production processes can be found in Young (2019). Wang et al. (2012) and Nandram, Berg and Barboza (2014) developed model-based approaches for combining multiple sources of survey data with other sources of information with the aim of informing the NASS *forecasting* process for corn and soybeans. (The forecasts precede the publication of end-of-season state estimates in the Small Grains and Annual Summaries.) These methods were refined for in-house use by the ASB and later extended to include winter wheat, and upland cotton at state and regional levels; see Adrian (2012), Cruze (2015, 2016), Cruze and Benecha (2017). Kim et al. (2018) explored a candidate small-area approach to combining NASS June Area Survey data, CDL data and FSA administrative acreage data in order to produce more precise *state* estimates of planted area for multiple commodity crops.

For sub-state crop estimates, NASS has actively pursued research in small area estimation, although much of the research predates the full implementation of CAPS in 2011. Both unit-level models, for example, Walker and Sigman (1984), Battese, Harter and Fuller (1988), Stasny, Goel and Rumsey (1991), Bellow and Lahiri (2012) and area-level approaches, for example, Bellow and Lahiri (2010, 2011, 2012), Kott (1989), Williams (2013), have been investigated. Recent developments include subarea-level approaches to estimating and benchmarking crops county estimates

(Erciulescu, Cruze and Nandram, 2016, 2017, 2018, 2019). Collectively, these efforts have touched on all four parameters of interest for major crops (including corn, soybeans, winter wheat and sorghum).

Although the potential benefits of modeling crops county estimates are well understood, several challenges have prevented NASS from adopting any one explicit model to date. A Bayesian formulation of an area-level model for planted corn acreage is developed in Section 3.2. Comparison and interpretation relative to other auxiliary data illustrate many of these challenges.

3.2 An Area-Level Model for Planted Area

Following Fay and Herriot (1979), we define a model for planted area totals within a state

$$\text{(3)} \qquad \text{Level 1}: \quad \hat{\theta}_i = \theta_i + \epsilon_i,$$

$$\text{(4)} \qquad \text{Level 2}: \quad \theta_i = \mathbf{x}_i'\boldsymbol{\beta} + u_i,$$

where $i \in \{1, \dots, m\}$ is an index over m counties and $\hat{\theta}_i$ is the survey outcome for the ith county computed from a realized sample of n_i reports of positive planted area in the county. These n_i could vary with crop type since a respondent may not grow all sampled commodities. Equation (3) describes the CAPS direct estimate as an estimate of the true planted area θ_i up to sampling error ϵ_i. It is assumed that $\epsilon_i \overset{\text{ind.}}{\sim} \mathrm{N}(0, \hat{\sigma}_{\epsilon i}^2)$, and $\hat{\sigma}_{\epsilon i}^2$ are the estimated sampling variances of the CAPS direct estimates. The Level 2 model in equation (4) describes the county-level total θ_i in terms of a linear function of observable covariates $\mathbf{x}_i$ and a county-specific random effect $u_i \overset{\text{iid}}{\sim} \mathrm{N}(0, \sigma_u^2)$ assumed to be independent from the sampling errors ϵ_i.

Planted corn area in Illinois with reference year 2014 was selected as a test case because the state is a major producer of corn, and the commodity is grown across the state. For this case study, FSA administrative acreage for corn planted area was used as a covariate at the county level. In southern parts of the US, corn may be planted as early as the beginning of March; in Illinois, typical corn planting begins by mid-April, and harvesting usually begins by mid-September (USDA NASS, 2010, page 9). These decisions can vary by year, influenced by weather conditions and a number of other factors.[2] Planting is a decision finalized earlier in the crop year, whereas the final realized harvested area, production and yield are influenced by an accumulation of events that transpire after the crop is planted and up to the date of harvest. Presence or absence of rain prior to planting could potentially accelerate or delay planting decisions in the given year. A source of precipitation data described in Vose et al. (2014) were obtained directly from the National Oceanic and Atmospheric Administration (NOAA) for inclusion in the model. The NOAA data are available for monthly intervals at the so-called climate division level and, for Illinois, there is a one-to-one correspondence between the NOAA climate division and the NASS ASD. Based on correlation analysis with NASS planted area survey data, district-level estimates of March precipitation (in inches) were included in the model; the district-level data were duplicated for counties within the same ASD. Thus, the vector of regressors for the ith county consists of $\mathbf{x}_i = (1, \text{FSA acres, NOAA (ASD-level) precipitation})'$ in this case study.

Due to observed skewness in the distribution of the county-level direct estimates $\hat{\theta}_i$, the decision was made to scale each estimate by the number of positive reports n_i and model the transformed variable $\hat{\theta}_i/n_i$. (This transformation was applied to the FSA covariate as well.) Resulting point estimates of the area totals obtained under equation (3) and equation (4) may not automatically sum to the already published state totals. Benchmarking to the state total is necessary to ensure consistency of acreage and production estimates at state, ASD and county levels.

3.2.1 *External benchmarking*. In practice, NASS requires that official county and ASD estimates satisfy external benchmarks since state totals are published prior to the publication of any sub-state crop estimates. Therefore, model-based approaches that also honor equation (1) and equation (2) are desirable. Letting a denote the target state-level estimate of planted area, $n = \sum_{i=1}^m n_i$ denote the number of all positive reports obtained from m counties in the state, and $\tilde{\theta}_i^{\mathrm{ME}} \equiv n_i\tilde{\theta}_i$ denote a modeled estimate transformed back to the scale of the ith county total, a ratio benchmarking approach was considered. This non-parametric approach applies the same corrective factor shown in equation (5) to each of the m modeled estimates $\tilde{\theta}_i^{\mathrm{ME}}$; after benchmarking, the m county-level totals agree with the established state total as in equation (6):

$$\text{(5)} \qquad \tilde{\theta}_i^{\mathrm{MERB}} = \tilde{\theta}_i^{\mathrm{ME}} * a\left(\sum_{k=1}^m \tilde{\theta}_k^{\mathrm{ME}}\right)^{-1},$$

$$\text{(6)} \qquad a = \sum_{i=1}^m \tilde{\theta}_i^{\mathrm{MERB}}.$$

[2]An excerpt from NASS's Crop Progress Report included the Supplementary Material provides a synopsis of planting and growth stages of the Illinois corn crop during the 2014 crop year.

3.2.2 *Preliminary findings*. Model-based estimates of area planted in corn for the 2014 crop year were computed for the 102 counties and 9 ASDs in the state of Illinois. The models were formulated as Bayesian hierarchical models and fit by Markov chain Monte Carlo simulation using R and JAGS software. Vague, proper prior distributions were specified for all model parameters. The prior distributions for the model parameters $\boldsymbol{\beta} \equiv (\beta_0, \beta_{\text{FSA}}, \beta_{\text{NOAA}})'$ are assumed $\text{Normal}_3(\mathbf{0}, \mathbf{diag}(10^6))$. Three choices of prior distribution for the model variance component (or transformations of the variance component) σ_u^2 were explored:

- $\sigma_u^2 \sim \text{Uniform}(0, 10^8)$,
- $\sigma_u^2 \sim \text{IG}(0.001, 0.001)$,
- $\log(\sigma_u^2) \sim \text{Uniform}(-\log(10^8), \log(10^8))$.

For each combination of prior specifications, simulations consisted of three chains, each with 10,000 Monte Carlo iterates. Discarding the first 1000 iterates as burn-in, the remaining 9000 iterates were thinned every 9 iterates to construct posterior summaries. Potential scale reduction factor convergence diagnostic statistics were inspected for each of the 102 θ_i associated with counties, the three regression parameters, and the variance components. The maximum potential scale reduction factor under uniform prior specification for σ_u^2 was 1.039, whereas the inverse gamma on the variance and the uniform on the log variance specifications resulted in several potential scale reduction factors greater than 1.1 (maximum of 11.57 and 2.224, respectively) indicating a lack of convergence. Therefore, we report only the results under the $\sigma_u^2 \sim \text{Uniform}(0, 10^8)$ prior specification below.

The posterior means (posterior standard deviations) of the regression coefficients were $\hat{\beta}_0 = 97.0$ (205.2); $\hat{\beta}_{\text{FSA}} = 0.865$ (0.037); and $\hat{\beta}_{\text{NOAA}} = -48.6$ (118.0), respectively. The coefficient on FSA acreage and its relatively small standard deviation indicates the strength of the positive linear association between FSA and NASS corn acreage data in Illinois.

Model estimates of the county planted area totals and their associated standard errors were also obtained as posterior means and standard deviations, after transforming back to the total scale through multiplication by sample size. ASD totals are sums of county-level totals; posterior means and standard deviations of ASD-level planted area were obtained by summing corresponding iterates of the chains of member counties, and then averaging and computing the standard deviations of those summed iterates.

First, models were fit assuming no external benchmark (denoted ME). Externally benchmarked estimates (MERB) were obtained by subsequently applying the ratio benchmarking step; in this process, the factor $a(\sum_{k=1}^{m} \tilde{\theta}_k^{\text{ME}})^{-1}$ was applied to all iterates of the obtained Monte Carlo samples for county totals. The 102 estimated MERB posterior means sum to the $a = 11.9$ million acre planted area total NASS published for the state of Illinois in crop year 2014 (USDA NASS, 2015, page 8). Injecting the ratio into the stored Monte Carlo samples captures the effects of benchmarking on estimated variances.

While this approach differs from the NASS ASB's strictly top-down approach, *it automatically generates point estimates for ASDs consistent with estimates of member counties while producing measures of uncertainty for counties and ASDs*. A comparison of coefficients of variation in Table 2 shows that this area-level model (whether benchmarked or not) reduces the CV of the sub-state estimate relative to the CAPS direct estimates (DE) of planted area. In this case study, reductions in CV of nearly 50% and more were realized compared to the direct estimates alone. Modest changes in CV are realized after benchmarking the county estimates to the state total.

While offering a considerable reduction in CV, the model-based approach that does not incorporate the state total as an external benchmark (ME) produces estimates that are inconsistent with the published state total and with FSA data at the county and district level. Figure 4 shows ME and MERB estimates are plotted against corresponding county- and ASD-level FSA administrative acreage data. (In the interest of not disclosing unpublished data, both axes have been redacted but 1-to-1 ratio is enforced in both panels.) Nearly all of the ME estimates fall below the plotted 45 degree line

TABLE 2
Coefficients of Variation (%) for Illinois corn planted area estimates in crop year 2014 *for* 102 *counties within* 9 *Agricultural Statistics Districts*

Level	Statistic	DE	ME	MERB
County	min	9.1	4.2	3.6
	median	19.2	7.5	7.3
	max	92.3	31.0	31.1
ASD	min	4.4	2.6	1.7
	median	6.8	2.9	2.1
	max	8.7	4.4	4.4

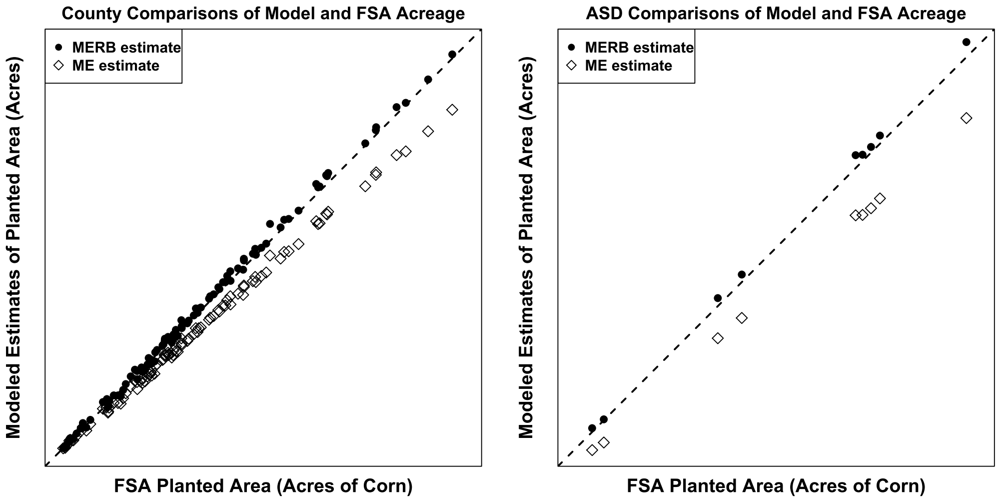

FIG. 4. *Comparison of ME and MERB estimates with FSA acreage.*

indicating that they are smaller than the corresponding FSA administrative corn acreage, whereas benchmarking improves consistency with FSA administrative acreage totals. All of the ASD-level MERB estimates cover the administrative data, but a few county-level MERB estimates may fall below the 45 degree line.

The ASB's estimation procedures entail the review of potentially disparate estimates to produce one-number official statistics. Shown in terms of number of posterior standard deviations from MERB estimates, Figure 5 depicts distance between MERB estimates and other available estimates, that is, $(\tilde{\theta}^{\text{MERB}} -$ auxiliary estimate$)/\,\text{SD}(\tilde{\theta}^{\text{MERB}})$. The sources compared at county and district level include the CAPS direct estimates (DE), the CDL acreage derived from pixel counting, the Battese–Fuller planted area estimate based on a unit-level regression of CDL and NASS June Area Survey data, the FSA planted area, and the NASS official statistics. (It must be noted that apart from the NASS official statistics, none of the other sources are benchmarked to the state total.) Positive values indicate that the MERB estimate is greater than the other source whereas negative values indicate that the MERB estimate is the smaller of the two. At both district and county levels, the spread between

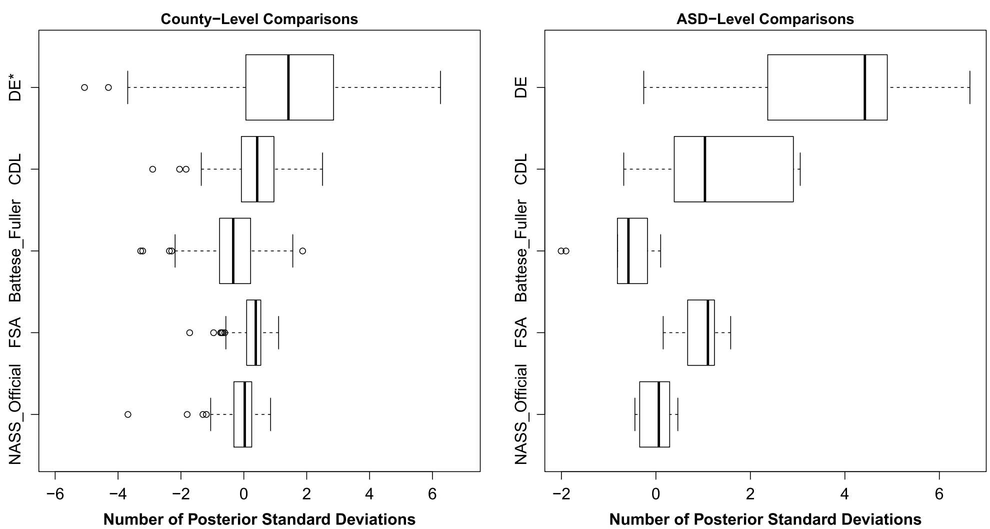

FIG. 5. *Distances from MERB posterior means in terms of number of posterior standard deviations. The comparisons are for all* 102 *counties and* 9 *ASDs in the state. *For clarity, a single outlier beyond* −6 *was omitted in the plotted county-level comparisons to direct estimates.*

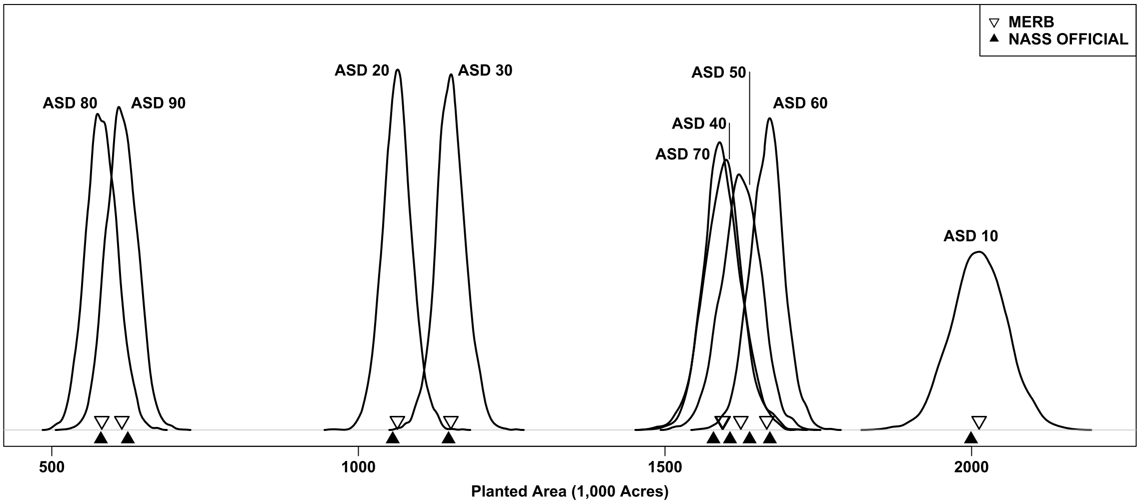

FIG. 6. *Posterior distributions for planted area in the nine agricultural statistics districts of Illinois.*

the MERB estimates and FSA acreage, Battese–Fuller estimates, or the CDL is smaller than the spread between the MERB estimates and the CAPS direct estimates. At the county level, the compactness of the FSA boxplot indicates shrinkage of the MERB estimate toward the FSA administrative data. Comparison to FSA shows negative values between the minimum and the first quartile; in terms of a literal point-to-point comparison, these counties do not cover the FSA totals, although they are within a reasonable number of posterior standard deviations. The similarities of the FSA and Official boxplots likely indicate the emphasis that the ASB places on the FSA data whenever it is available (National Academies of Sciences Engineering, and Medicine, 2017, page 21). At the ASD level, the median differences between MERB and direct estimates, FSA and CDL are all positive, perhaps indicating the accumulated effects of undercoverage, or the downward tendency due to pixel counting in the case of the CDL, when summed across counties within the same district. At the ASD-level, all MERB estimates are greater than the FSA administrative data. All 9 official ASD estimates fall within plus or minus half of one posterior standard deviation of the corresponding MERB estimates.

Posterior distributions for planted area totals in the 9 ASDs of Illinois[3] are shown in Figure 6. Posterior means under MERB and the NASS official statistics are shown for comparison. Districts 10 and 20 are the northernmost ASDs in the state; the posterior distribution for ASD 10 is remarkable in both its location (posterior mean) and its spread. The comparatively smaller planted acreage in ASD 20 may be attributed to land occupied by the Chicago metropolitan area. Similarly, smaller planted areas are found in the southernmost districts, ASD 80 and ASD 90, likely due to the presence of heavily forested lands. The remaining ASDs form a band through the central part of Illinois; ASDs 40, 50, 60 and 70 are contiguous districts, and they show similarities both in terms of estimated planted areas and in terms of their respective posterior distributions. When compared to the official statistics, the posterior means obtained under MERB show similar magnitudes to the NASS official statistics. In this case, the MERB estimates have even preserved the same rank among ASD estimates as the published NASS district-level estimates. Since the MERB approach incorporates benchmarking, modeled ASD estimates and NASS official statistics sum to the same 11.9 million acre target.

Direct estimates are obtained with respect to a sampling design. Uncertainty is inherent in NASS official statistics, although no characterization of the uncertainty has been published with the official statistics to date. In the case of the presented model, the desired quantities are characterized by entire distributions. The uncertainty associated with direct estimates may be reduced by incorporating additional information through a small area model. The ASB has traditionally incorporated auxiliary data through its informal composites; given the similarity of the NASS official statistics and the MERB point estimates in Figure 6, the increased

[3]An ASD map for Illinois is included with Supplementary Material, or it can be obtained at https://www.nass.usda.gov/Charts_and_Maps/Crops_County/boundary_maps/il.pdf.

use of modeling may offer a means of quantifying the benefit of incorporating additional information.

4. DISCUSSION AND CONCLUSIONS

As a producer of official statistics, NASS must make the best use of all available data inputs and produce high quality sub-state estimates for more than 20 commodity crops nationwide. As immediate consumers of NASS official crop county estimates, both FSA and RMA must consider the NASS official statistics in light of their own administrative records. Where incentives to participate in FSA or RMA programs may vary by state and commodity, NASS official statistics fulfill statutory obligations and provide additional insight for these agencies in the administration of their respective farm programs. NASS official county-level statistics provide important data for the US agricultural sector at large.

The NASS crops survey cycle is designed to capture crop activities from initial planting decisions to the final harvest and production. Starting in 2011, a larger probability-based sample for sub-state estimates was realized with the full implementation of CAPS. Even after strengthening the direct estimates for counties and ASDs, issues of nonresponse, variation in year-to-year planting decisions, and natural sparsity of some types of crops within state and county lines have led to the suppression of many NASS county estimates. Small area models have the potential to add value to the NASS crops county estimates program in terms of improved reproducibility and quantification of uncertainties all while harnessing a wider variety of auxiliary data types. As evidenced in the findings of the CNSTAT report, transitioning NASS's complex crops county estimates program to a model-based framework is a challenging undertaking.

The motivating case study presented in this paper has focused on a model for just one of four key parameters, planted area and, in particular, it focused on interpretation of survey and modeled estimates with respect to important current-year administrative planted acreage data source provided by FSA. A number of ongoing challenges and opportunities are noted below:

- *The best use of all available data*—At the authors' discretion, the model presented in this preliminary case study made use of just one key auxiliary data source for acreage, namely the FSA planted area data. Both the FSA acreage and the CDL acreages are highly correlated. An important question worth investigating is whether the quality of the model has been affected by foregoing the use of the CDL data, and under what circumstances it might be a preferred source of auxiliary information. Novel uses of administrative or remotely sensed data could in turn change the way expensive and time-consuming surveys are conducted in the future.
- *Modeling totals or modeling ratios*—Looking ahead to harvested area, production, and yield, NASS may use its survey data twice in its traditional composite process: once to obtain an estimate of a total, see, for example, production, and once to construct an estimate of a ratio. Where ratios are concerned in the traditional process, they are translated back to the total scale *conditional* on some fixed denominator total. The merits of incorporating one or the other into a model should be carefully weighed. Benchmarking a ratio, for example, yield, introduces additional challenges.
- *Valid zeros versus 'missing" data*—In some cases, the survey may not indicate any planted area with respect to a particular commodity, but some positive acreage may be represented by auxiliary data for the county. Due to item-level nonresponse associated with production, it is possible, in some cases, to have positive acreage survey estimates within a county but have no reports to support a production survey estimate in that same county. This particular challenge is often encountered in counties where a certain commodity crop is already sparse. The valid zero does not represent a problem for benchmarking to state totals; if a county did not contribute toward a state total, it can simply be omitted from the constraint. On the other hand, "missing" estimates could potentially affect the quality of other county estimates through their omission when benchmarking. Therefore, an appropriate synthetic estimate may be desirable in such cases.
- *Implied relationships among estimates*—Related to the valid zero, a zero estimate for planted area implies a zero estimate for harvested area, and a zero estimate for harvested area implies a zero estimate for production. In principle, any defensible (point) estimate of harvested area within a region must be no greater than the planted area estimate for that region; in practice, separately modeled estimates for planted and harvested area could result in inconsistent point estimates. The ASB has traditionally sought to cover FSA acreages in its estimation process; part of external model validation will be to determine whether auxiliary sources *complement or contravene* a modeled estimate. Since yield is a ratio of production to

harvested area, once any two of the three point estimates are known, the third is determined. Whereas model-based estimates obtained from several independent models may not automatically satisfy these relationships, NASS's traditional estimation and review process enforces these physical relationships in the official statistics. Incorporating constraints into the models or jointly modeling county-level parameters may be worthwhile.

- *A meaningful pool of covariates for production*—The inclusion of NOAA precipiation data in the case study reflected knowledge surrounding typical planting dates in Illinois. Its incorporation (at the ASD level) into the model may not be needed given the strength of the FSA administrative covariate. In the absence of administrative data on production, other characteristics including measures of soil productivity, weather, climate data, and remote sensing indices and their importance to crop phenology may be taken into consideration. Building the best pool of auxiliary data for the 43 CAPS states and the range of commodities to be supported by the crops county estimates program is a formidable task. Moreover, the best use of these data will require crop-specific knowledge about critical growing stages of each commodity and incorporating data at the appropriate temporal and spatial resolution.
- *An updated publication standard*—As NASS transitions toward increased use of model-based estimates, a publication standard that reflects the fitness for use of the estimate is essential. The panel's findings suggest that the standard should be tied to published measures of uncertainty, however, which measures of uncertainty, on which estimates, and at what stated thresholds are to be determined by the Agency. Additional matters of nondisclosure and confidentiality in the context of modeled estimates and NASS's establishment surveys must be carefully vetted.

With the release of *Improving Crop Estimates by Integrating Multiple Data Sources* in 2017, a vision for transitioning NASS crops county estimates to a program incorporating model-based crops estimates has been established. NASS looks forward to meeting these challenges in fulfillment of its mission to produce timely, accurate, and useful statistics in service to U.S. agriculture.

APPENDIX: LIST OF FREQUENTLY USED ABBREVIATIONS

Bodies within the United States Department of Agriculture (USDA)

- NASS—National Agricultural Statistics Service. Federal statistical agency housed with USDA
- FSA—Farm Service Agency. Responsible for the administration of several USDA agricultural support programs. A primary source of administrative data on *planted area*.
- RMA—Risk Management Agency. Underwriter of approved independent crop insurance policies. A source of administrative data particularly on failed acreage.

Bodies outside USDA

- CNSTAT—Committee on National Statistics. Advisory body convened under National Academies of Science, Engineering, and Medicine
- NOAA—National Oceanographic and Atmospheric Administration. Provider of weather data.

NASS Data Concepts and Notation

- ASB—Agricultural Statistics Board. The source of NASS official statistics. The ASB interprets several sources of information and publishes a consensus estimate based on its expert review of survey and auxiliary inputs.
- ASD—Agricultural Statistics District. A smaller-than-state area comprised of neighboring counties
- MPPS—Multivariate probability proportional to size. Survey sampling design underlying several key NASS crop surveys.
- APS—The quarterly Acreage, Production and Stocks survey. This MPPS design survey informs state benchmarking targets and it is a partial source of data used to compute county-level direct estimates.
- CAPS—County Agricultural Production Survey. Provides a complementary MPPS sample used in conjunction with the APS sample. The pooled APS/CAPS samples are used to generate county-level direct estimates.
- CDL—Cropland Data Layer. A NASS agricultural land cover classification product based on remotely sensed information. Another potential measurement of planted area.
- **PL, HV, PD, YD**—these denote official planted area, harvested area, production, and yield estimates, respectively.
- DE—Direct estimate. An output obtained from NASS surveys under the given survey design.
- ME—A modeled (small area) estimate.
- MERB—A modeled estimate which has been benchmarked to a given state total through ratio benchmarking.

DISCLAIMER AND ACKNOWLEDGEMENTS

The findings and conclusions in this preliminary publication have not been formally disseminated by the U.S. Department of Agriculture and should not be construed to represent any agency determination or policy. This research was supported in part by the intramural research program of the U.S. Department of Agriculture, National Agriculture Statistics Service. The work of Erciulescu was completed as a Research Associate at the National Institute of Statistical Sciences (NISS) working on NASS projects.

The work of Dr. Nandram was supported by Simons Foundation Grant 350953.

SUPPLEMENTARY MATERIAL

Supplement to "Producing Official County-Level Agricultural Estimates in the United States: Needs and Challenges" (DOI: 10.1214/18-STS687SUPP; .pdf). Supplementary information

REFERENCES

ADRIAN, D. W. (2012). A model-based approach to forecasting corn and soybean yields. In *Proceedings of the Fourth International Conference on Establishment Surveys* Amer. Statist. Assoc., Montreal, QC. https://ww2.amstat.org/meetings/ices/2012/papers/302190.pdf [Accessed: 2019-01-31].

BAILEY, J. T. and KOTT, P. S. (1997). An application of multiple list frame sampling for multi-purpose surveys. In *JSM Proceedings, Survey Research Methods Section* 496–500. Amer. Statist. Assoc., Alexandria, VA.

BATTESE, G. E., HARTER, R. M. and FULLER, W. A. (1988). An error-components model for prediction of county crop areas using survey and satellite data. *J. Amer. Statist. Assoc.* **83** 28–36.

BELL, J. and BARBOZA, W. (2012). Evaluation of using CVs as a publication standard. In *Proceedings of the Fourth International Conference on Establishment Surveys* Amer. Statist. Assoc., Montreal, QC. https://ww2.amstat.org/meetings/ices/2012/papers/302002.pdf [Accessed: 2019-01-31].

BELL, W., BASEL, W. W. and MAPLES, J. J. (2016). An overview of the U.S. census bureau's small area income and poverty estimates program. In *Analysis of Poverty Data by Small Area Estimation* (M. Pratesi, ed.) 349–378 19. Wiley, Hoboken, NJ.

BELLOW, M. E. and LAHIRI, P. (2010). Empirical Bayes methodology for the NASS county estimation program. In *JSM Proceedings, Survey Research Methods Section* 343–355. Amer. Statist. Assoc., Alexandria, VA.

BELLOW, M. E. and LAHIRI, P. (2011). An empirical best linear unbiased prediction approach to small area estimation of crop parameters. In *JSM Proceedings, Survey Research Methods Section* 3976–3986. Amer. Statist. Assoc., Alexandria, VA.

BELLOW, M. E. and LAHIRI, P. (2012). Evaluation of methods for county level estimation of crop harvested area that employ mixed models. In *Proceedings of the DC-AAPOR/WSS Summer Conference, Bethesda, MD* Amer. Statist. Assoc., Alexandria, VA.

BORYAN, C. (2010). The USDA NASS Cropland Data Layer Program: Transition from Research to Operations (2006–2009. USDA NASS Education and Outreach: Research Reports https://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/reports/Boryan_CDL_Program_Research_to_Operations_2006-2009_final.pdf [Accessed: 2019-01-31].

BORYAN, C., YANG, Z., MUELLER, R. and CRAIG, M. (2011). Monitoring US agriculture: The US Department of Agriculture, National Agricultural Statistics Service, cropland data layer program. *Geocarto Int.* **26** 341–358.

CRUZE, N. B. (2015). Integrating survey data with auxiliary sources of information to estimate crop yields. In *JSM Proceedings, Survey Research Methods Section* 565–578. Amer. Statist. Assoc., Alexandria, VA.

CRUZE, N. B. (2016). A Bayesian hierarchical model for combining several crop yield indications. In *JSM Proceedings, Survey Research Methods Section* 2045–2053. Amer. Statist. Assoc., Alexandria, VA.

CRUZE, N. B. and BENECHA, H. K. (2017). A model-based approach to crop yield forecasting. In *JSM Proceedings, Survey Research Methods Section* 2724–2733. Amer. Statist. Assoc., Alexandria, VA.

CRUZE, N. B., ERCIULESCU, A. L., NANDRAM, B., BARBOZA, W. J. and YOUNG, L. J. (2019). Supplement to "Producing Official County-Level Agricultural Estimates in the United States: Needs and Challenges." DOI:10.1214/18-STS687SUPP.

CZAPLEWSKI, R. L. (1992). Misclassification bias in areal estimates. *Photogramm. Eng. Remote Sens.* **58** 189–192.

ERCIULESCU, A. L., CRUZE, N. B. and NANDRAM, B. (2016). Small area estimates incorporating auxiliary soucres of information. In *JSM Proceedings, Survey Research Methods Section* 3591–3605. Amer. Statist. Assoc., Alexandria, VA.

ERCIULESCU, A. L., CRUZE, N. B. and NANDRAM, B. (2017). Small area estimates of end-of-season agricultural quantities. In *JSM Proceedings, Survey Research Methods Section* 541–560. Amer. Statist. Assoc., Alexandria, VA.

ERCIULESCU, A. L., CRUZE, N. B. and NANDRAM, B. (2018). Benchmarking a triplet of official estimates. *Environ. Ecol. Stat.* **25** 523–547. MR3878890

ERCIULESCU, A. L., CRUZE, N. B. and NANDRAM, B. (2019). Model-based county level crop estimates incorporating auxiliary sources of information. *J. Roy. Statist. Soc. Ser. A* **182** 283–303. MR3902644

FAY, R. E. III and HERRIOT, R. A. (1979). Estimates of income for small places: An application of James-Stein procedures to census data. *J. Amer. Statist. Assoc.* **74** 269–277. MR0548019

FULLER, W. A. and BATTESE, G. E. (1973). Transformations for estimation of linear models with nested-error structure. *J. Amer. Statist. Assoc.* **68** 626–632. MR0359188

GALLEGO, F. J. (2004). Remote sensing and land cover area estimation. *Int. J. Remote Sens.* **25** 3019–3047.

IWIG, W. (1996). The national agricultural statistics service county estimates program. In *Indirect Esitmators in U.S. Federal Programs* (W. Schaible, ed.) 129–144 7. Springer, New York.

JOHNSON, D. M. (2014). An assessment of pre- and within-season remotely sensed variables for forecasting corn and soybean yields in the United States. *Remote Sens. Environ.* **141** 116–128.

JOHNSON, D. M. (2016). A comprehensive assessment of the correlations between field crop yields and commonly used MODIS products. *Int. J. Appl. Earth Obs. Geoinf.* **52** 65–81.

KIM, J. K., WANG, Z., ZHU, Z. and CRUZE, N. B. (2018). Combining survey and non-survey data for improved sub-area prediction using a multi-level model. *J. Agric. Biol. Environ. Stat.* **23** 175–189. MR3805267

KOTT, P. S. (1989). Robust small domain estimation using random effects modeling. *Surv. Methodol.* **15** 3–12.

NANDRAM, B., BERG, E. and BARBOZA, W. (2014). A hierarchical Bayesian model for forecasting state-level corn yield. *Environ. Ecol. Stat.* **21** 507–530. MR3248537

NATIONAL ACADEMIES OF SCIENCES ENGINEERING, AND MEDICINE (2017). *Improving Crop Estimates by Integrating Multiple Data Sources*. The National Academies Press, Washington, DC.

NATIONAL RESEARCH COUNCIL (1997). *Small-Area Estimates of School-Age Children in Poverty: Interim Report 1, Evaluation of 1993 County Estimates for Title I Allocations*. The National Academies Press, Washington, DC.

NATIONAL RESEARCH COUNCIL (1998). *Small-Area Estimates of School-Age Children in Poverty: Interim Report 2, Evaluation of Revised 1993 County Estimates for Title I Allocations*. The National Academies Press, Washington, DC.

NATIONAL RESEARCH COUNCIL (1999). *Small-Area Estimates of School-Age Children in Poverty: Interim Report 3*. The National Academies Press, Washington, DC.

NATIONAL RESEARCH COUNCIL (2000a). *Small-Area Estimates of School-Age Children in Poverty: Evaluation of Current Methodology*. The National Academies Press, Washington, DC.

NATIONAL RESEARCH COUNCIL (2000b). *Small-Area Income and Poverty Estimates: Priorities for 2000 and Beyond*. The National Academies Press, Washington, DC.

NATIONAL RESEARCH COUNCIL (2007). *Using the American Community Survey: Benefits and Challenges*. The National Academies Press, Washington, DC.

NATIONAL RESEARCH COUNCIL (2008). *Understanding American Agriculture: Challenges for the Agricultural Resource Management Survey*. The National Academies Press, Washington, DC.

RAO, J. N. K. and MOLINA, I. (2015). *Small Area Estimation*, 2nd ed. Wiley, Hoboken, NJ. MR3380626

STASNY, E. A., GOEL, P. K. and RUMSEY, D. J. (1991). County estimates of wheat production. *Surv. Methodol.* **17** 211–225.

UNITED STATES CENSUS BUREAU (2017). Small Area Income and Poverty Estimates (SAIPE) Program. https://www.census.gov/programs-surveys/saipe/about.html. [Accessed: 2019-01-31].

UNITED STATES GOVERNMENT PUBLISHING OFFICE (2015). Big Data and Agriculture: Innovation and Implications–Hearing before the Committee on Agriculture, House of Representatives, 114th Congress. https://www.govinfo.gov/content/pkg/CHRG-114hhrg97412/pdf/CHRG-114hhrg97412.pdf. [Accessed: 2019-01-31].

UNITED STATES GOVERNMENT PUBLISHING OFFICE (2016). Big Data and Agriculture: Innovation in the Air–Hearing before the Subcommittee on General Farm Commodities and Risk Management of the Committee on Agriculture, House of Representatives, 114th Congress. https://www.govinfo.gov/content/pkg/CHRG-114hhrg20574/pdf/CHRG-114hhrg20574.pdf. [Accessed: 2019-01-31].

USDA NASS (2010). *Field Crops: Usual Planting and Harvesting Dates*. Agricultural Handbook No. 628. https://downloads.usda.library.cornell.edu/usda-esmis/files/vm40xr56k/dv13zw65p/w9505297d/planting-10-29-2010.pdf [Accessed: 2019-01-31].

USDA NASS (2014). *Small Grains 2014 Summary*. https://downloads.usda.library.cornell.edu/usda-esmis/files/5t34sj573/2r36v149p/qj72p9909/SmalGraiSu-09-30-2014.pdf [Accessed: 2019-01-31].

USDA NASS (2015). *Crop Production 2014 Summary*. https://downloads.usda.library.cornell.edu/usda-esmis/files/k3569432s/5q47rr167/1n79h650b/CropProdSu-01-12-2015_revision.pdf [Accessed: 2019-01-31].

VOSE, R. S., APPLEQUIST, S., SQUIRES, M., DURRE, I., MENNE, M. J., WILLIAMS, C. N., FENIMORE, C., GLEASON, K. and ARNDT, D. (2014). Improved historical temperature and precipitation time series for U.S. climate divisions. *J. Appl. Meteorol. Climatol.* **53** 1232–1251.

WALKER, G. and SIGMAN, R. (1984). The use of LANDSAT for county estimates of crop areas: Evaluation of the huddleston-ray and the battese-fuller estimators for the case of stratified sampling. *Commun. Statist. Theory Methods* **13** 2975–2996.

WANG, J. C., HOLAN, S. H., NANDRAM, B., BARBOZA, W., TOTO, C. and ANDERSON, E. (2012). A Bayesian approach to estimating agricultural yield based on multiple repeated surveys. *J. Agric. Biol. Environ. Stat.* **17** 84–106. MR2912556

WILLIAMS, M. (2013). Small area modeling of county estimates for corn and soybean yields in the US. Presentation at the Federal Committee on Statistical Methodology Research Conference. http://www.copafs.org/UserFiles/file/fcsm/C2_Williams_2013FCSM.pdf [Accessed: 2019-01-31].

YOUNG, L. J. (2019). Agricultural crop forecasting for large geographical areas. *Ann. Rev. Stat. Appl.* **6**. 173–196.