Improving Crop Estimates by Integrating Multiple Data Sources

DETAILS

117 pages | 6 x 9 | PAPERBACK
ISBN 978-0-309-46529-8 | DOI: 10.17226/24892

CONTRIBUTORS

Mary Ellen Bock and Nancy J. Kirkendall, Editors; Panel on Methods for Integrating Multiple Data Sources to Improve Crop Estimates; Committee on National Statistics; Division of Behavioral and Social Sciences and Education; National Academies of Sciences, Engineering, and Medicine

PREPUBLICATION COPY, UNCORRECTED PROOFS

Improving Crop Estimates by Integrating Multiple Data Sources

Panel on Methods for Integrating Multiple Data Sources to Improve Crop Estimates

Mary Ellen Bock and Nancy J. Kirkendall, *Editors*

Committee on National Statistics

Division of Behavioral and Social Sciences and Education

A Consensus Study Report of
The National Academies of
SCIENCES • ENGINEERING • MEDICINE

ADVANCE COPY

NOT FOR PUBLIC RELEASE BEFORE

Wednesday, October 11, 2017
11:00 am EDT

THE NATIONAL ACADEMIES PRESS
Washington, DC
www.nap.edu

PREPUBLICATION COPY, UNCORRECTED PROOFS

THE NATIONAL ACADEMIES PRESS 500 Fifth Street, NW Washington, DC 20001

This activity was supported by a cooperative agreement between the National Academy of Sciences Committee on National Statistics and the United States Department of Agriculture, National Agricultural Statistics Service (agreement number 58-3AEU-4-0068). Support for the work of the Committee on National Statistics is provided by a consortium of federal agencies through a grant from the National Science Foundation (award number SES-1024012). Any opinions, findings, conclusions, or recommendations expressed in this publication do not necessarily reflect the views of the organization or agency that provided support for the project.

International Standard Book Number-13:
International Standard Book Number-10:
Digital Object Identifier: https://doi.org/10.17226/24892.

Additional copies of this report are available for sale from the National Academies Press, 500 Fifth Street, NW, Keck 360, Washington, DC 20001; (800) 624-6242 or (202) 334-3313; http://www.nap.edu.

Printed in the United States of America

Suggested citation: The National Academies of Sciences, Engineering, and Medicine. 2017. *Improving Crop Estimates by Integrating Multiple Data Sources*. Washington, DC: The National Academies Press. doi: https://doi.org/10.17226/24892.

PREPUBLICATION COPY, UNCORRECTED PROOFS

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

The **National Academy of Sciences** was established in 1863 by an Act of Congress, signed by President Lincoln, as a private, nongovernmental institution to advise the nation on issues related to science and technology. Members are elected by their peers for outstanding contributions to research. Dr. Marcia McNutt is president.

The **National Academy of Engineering** was established in 1964 under the charter of the National Academy of Sciences to bring the practices of engineering to advising the nation. Members are elected by their peers for extraordinary contributions to engineering. Dr. C. D. Mote, Jr., is president.

The **National Academy of Medicine** (formerly the Institute of Medicine) was established in 1970 under the charter of the National Academy of Sciences to advise the nation on medical and health issues. Members are elected by their peers for distinguished contributions to medicine and health. Dr. Victor J. Dzau is president.

The three Academies work together as the **National Academies of Sciences, Engineering, and Medicine** to provide independent, objective analysis and advice to the nation and conduct other activities to solve complex problems and inform public policy decisions. The National Academies also encourage education and research, recognize outstanding contributions to knowledge, and increase public understanding in matters of science, engineering, and medicine.

Learn more about the National Academies of Sciences, Engineering, and Medicine at **www.nationalacademies.org**.

PREPUBLICATION COPY, UNCORRECTED PROOFS

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

Consensus Study Reports published by the National Academies of Sciences, Engineering, and Medicine document the evidence-based consensus on the study's statement of task by an authoring committee of experts. Reports typically include findings, conclusions, and recommendations based on information gathered by the committee and the committee's deliberations. Each report has been subjected to a rigorous and independent peer-review process and it represents the position of the National Academies on the statement of task.

Proceedings published by the National Academies of Sciences, Engineering, and Medicine chronicle the presentations and discussions at a workshop, symposium, or other event convened by the National Academies. The statements and opinions contained in proceedings are those of the participants and are not endorsed by other participants, the planning committee, or the National Academies.

For information about other products and activities of the National Academies, please visit www.nationalacademies.org/about/whatwedo.

PREPUBLICATION COPY, UNCORRECTED PROOFS

PANEL ON METHODS FOR INTEGRATING MULTIPLE DATA SOURCES TO IMPROVE CROP ESTIMATES

MARY ELLEN BOCK (*Chair*), Department of Statistics (emerita), Purdue University
JULIE GERSHUNSKAYA, Bureau of Labor Statistics
MALAY GHOSH, Department of Statistics, University of Florida
MICHAEL GOODCHILD, Department of Geography (emeritus), University of California, Santa Barbara
CHAD HART, Department of Economics, Iowa State University (through October 2016)
OLGA ISENGILDINA MASSA, Department of Agricultural and Applied Economics, Virginia Tech University
SUSAN OFFUTT, Department of Economics (retired), U.S. Government Accountability Office
S. LYNNE STOKES, Department of Statistical Sciences, Southern Methodist University
JON WAKEFIELD, Department of Statistics, University of Washington
ROBERT YOUNG, Department of Public Policy, American Farm Bureau Federation

NANCY J. KIRKENDALL, Study Director
BRIAN HARRIS-KOJETIN, Director, Committee on National Statistics
CONSTANCE E. CITRO, Senior Scholar
GLENN D. WHITE, JR. Senior Program Officer
MARY ANN KASPER, Senior Program Assistant

COMMITTEE ON NATIONAL STATISTICS

ACKNOWLEDGEMENTS

This report is the product of contributions from many colleagues, whom I thank for their time, generosity, and expert guidance. The report also benefited strongly from the skillful guidance of the Study Director, whom I wish to acknowledge and thank.

As chair, I first acknowledge with great appreciation the efforts of my fellow panel members. Despite their many professional commitments, every panel member donated countless hours and shared extensive expertise to make this report possible. The report reflects the collective expertise, varied perspectives, diverse backgrounds, and commitment of all panel members: Julie Gershunskaya, Bureau of Labor Statistics; Malay Ghosh, Department of Statistics, University of Florida; Michael Goodchild, Department of Geography (emeritus), University of California, Santa Barbara; Chad Hart, Department of Economics, Iowa State University; Olga Isengildina Massa, Department of Agricultural and Applied Economics, Virginia Tech University; Susan Offutt, Department of Economics (retired), US Government Accountability Office; S. Lynne Stokes, Department of Statistical Sciences, Southern Methodist University; Jon Wakefield, Department of Statistics, University of Washington; and Robert Young, Department of Public Policy, American Farm Bureau Federation

The project was sponsored by the National Agricultural Statistics Service (NASS), United States Department of Agriculture. The bulk of the panel's first four meetings were open meetings consisting of many presentations and discussions as noted in more detail below. The meetings were well attended by NASS staff, who actively engaged in discussions with the panel, helping us to appreciate the work NASS does. I would especially like to acknowledge Dr. Nathan Cruze, the panel's primary contact person with NASS. He orchestrated and organized the many presentations to the panel by NASS staff, prepared or collected background material needed by the panel, and shared his deep knowledge about NASS and especially its modeling activities.

The panel would like to express their thanks to NASS senior management, first Joe Riley (former Administrator) then Hubert Hamer (current Administrator) for their encouragement and support and Linda Young, Director of Research who helped to make this effort possible. The panel thanks the following NASS staff and contractors for their detailed presentations (often more than one), helping the panel to learn about NASS surveys, systems, and models: Edwin Anderson (cash rents); Mark Apodaca (NASS list frame building and maintenance); Jeff Bailey (cash rents survey and model results); Wendy Barboza (linking NASS and FSA data); Dan Beckler (linking NASS and FSA data); Nathan Cruze (model based research in general and specific, cash rents and crops); Lindsay Drunasky (crops surveys); Andreea Erciulescu (NASS/NISS) (model based research in general and specific, cash rents and crops); Noemi Guindin (a potential decision support system linking yield to weather changes); Lance Honig (crops, publication standards, setting estimates, data timelines); David Johnson (remote sensing yield estimates); Troy Joshua (cash rents survey); Rick Mueller (remote sensing); Balgobin Nandram (Worcester Polytechnic Institute) (crops modeling); Gerald Tillman (crops and cash rents surveys); and Linda Young (modeling and direction)

The panel would like to express special thanks to representatives of the Farm Service Agency (FSA) and Risk Management Agency for their help in understanding their use of NASS data and features of the administrative data provided by their agencies. Speakers included Michael Alston, RMA; Joy Harwood, FSA; Terry Hickenbotham, FSA; Rich Iovanna, FSA; John Jinkins, FSA; Phil Sronce, FSA; and David Zanoni, RMA.

The panel would also like to thank Vince Breneman and Roger Glassen representatives of Economic Research Service and Patrick Flanagan, Natural Resources Conservation Service for presentations that helped the panel understand some of the USDA geospatial data and its use. The panel would also like to thank representatives of other federal statistical agencies for their advice, especially on modeling issues; Wes Basel and William Bell from the U.S. Census Bureau and Alan Dorfmann, National Center for Health Statistics. Special thanks to Mary Kay Thatcher, Farm Bureau for her presentation about precision agriculture; to David Marker, Westat for his presentation concerning sample survey publication standards used by statistical agencies; to Gordon Reichert, Statistics Canada for his presentation about their new remote sensing model-based estimates for yield during the growing season. To Keith Coble, Mississippi State for his insights on yield modeling; Noel Cressie, University of Wollongong, Australia for advice on spatio-temporal data fusion for satellite remote sensing; and to Christian Frankenberg, Environmental Science and Engineering, California Institute of Technology for a discussion of solar induced chlorophyll fluorecsense (SIF) as measured by a new satellite and to Jerry Hatfield, Agricultural Research Service for his presentation about crop productivity and potential for modeling with new SIF measurements. These individuals attended open meetings and generously gave of their time to present material to inform the panel's deliberations. These presentations stimulated extensive discussion of the issues covered in this report.

The panel could not have conducted its work efficiently without the capable staff of the National Academies of Sciences, Engineering, and Medicine. The panel was fortunate to have as its Study Director Dr. Nancy Kirkendall. She brought to the panel her extensive experience in government agencies including the former directorship of the Statistics and Methods Group of the U.S. Energy Information Administration as well as an outstanding history of directing National Academy studies. Our work could not have been completed without her extraordinary dedication and many contributions. She provided technical and substantive insights, drafted and revised almost all the sections of our report, pulling together the wide range of panel expertise, and kept the project on track, also serving as the hub for communication. We were also fortunate to have the counsel of Connie Citro, Director of the Committee on National Statistics (CNSTAT). We also want to thank the CNSTAT Senior Program Officer Glenn White who gave an extra ear and eye to our deliberations; Mary Ann Kasper who provided excellent administrative and logistical support to the panel; and Kirsten Sampson-Snyder, Division of Behavioral and Social Sciences and Education, for expertly coordinating the review process.

The panel also recognizes the many federal agencies that support the Committee on National Statistics directly and through a grant from the National Science Foundation. Without their support and their commitment to improving the national statistical system, the panel's work that is the basis of this report would not have been possible.

This Consensus Study Report was reviewed in draft form by individuals chosen for their diverse perspectives and technical expertise. The purpose of this independent review is to provide candid and critical comments that will assist the National Academies of Sciences, Engineering, and Medicine in making each published report as sound as possible and to ensure that it meets the institutional standards for quality, objectivity, evidence, and responsiveness to

the study charge. The review comments and draft manuscript remain confidential to protect the integrity of the deliberative process.

We thank the following individuals for their review of this report: Zachary S. Brown, Department of Agricultural and Resource Economics, North Carolina State University; Cynthia Z.F. Clark, Independent consultant, McLean, VA; Keith Coble, Department of Agricultural Economics, Mississippi State University; Joseph W. Glauber, Markets, Trade and Institutions Division, International Food Policy Research Institute, Washington, DC; Barry Goodwin, Department of Agricultural and Resource Economics, North Carolina State University; Scott H. Holan, Department of Statistics, University of Missouri; Xiaofei Li, Department of Agricultural Economics, Mississippi State University; Prabhu L. Pingali, Tata-Cornell Institute for Agricultural and Nutrition, Cornell University; and Eric V. Slud, Statistics Program, University of Maryland College Park and Mathematical Statistics, Center for Statistical Research and Methodology, U.S. Census Bureau.

Although the reviewers listed above provided many constructive comments and suggestions, they were not asked to endorse the conclusions or recommendations of this report nor did they see the final draft before its release. The review of this report was overseen by Sarah M. Nusser, Center for Survey Statistics and Methodology, Iowa State University and Christopher A. Sims, Department of Economics, Princeton University. They were responsible for making certain that an independent examination of this report was carried out in accordance with the standards of the National Academies and that all review comments were carefully considered. Responsibility for the final content rests entirely with the authoring committee and the National Academies.

Mary Ellen Bock, *Chair*
Panel on Methods for Integrating Multiple Data Sources to Improve Crop Estimates

CONTENTS

EXECUTIVE SUMMARY

The National Agricultural Statistics Service (NASS) is the primary statistical data collection agency within the U.S. Department of Agriculture (USDA). NASS conducts hundreds of surveys each year and prepares reports covering virtually every aspect of U.S. agriculture. Among the small-area estimates produced by NASS are county-level estimates for crops (planted acres, harvested acres, production, and yield by commodity) and for cash rental rates for irrigated cropland, nonirrigated cropland, and permanent pastureland. Key users of these county-level estimates include USDA's Farm Services Agency (FSA) and Risk Management Agency (RMA), which use the estimates as part of their processes for distributing farm subsidies and providing farm insurance, respectively.

Virtually all statistical agencies are seeing declines in survey response rates, and NASS is no exception (see National Research Council, 2013). As a result, the number of counties for which NASS has been able to publish estimates for crop acreage, yield, and cash rents has been declining over time. This has caused challenges for key data users that must develop alternative estimates for those counties without official NASS estimates. Those alternative estimates are not as justifiable as the estimates from sound probability samples or consistent, high-quality, model-based approaches.

PURPOSE OF THIS STUDY

In September 2014, NASS entered into a cooperative agreement with the Committee on National Statistics of the National Academies of Sciences, Engineering, and Medicine to assess county-level crop and cash rents estimates, and to offer recommendations on methods for integrating data sources to provide more precise county-level estimates of acreage and yield for major crops and of cash rents by land use. Multiple sources of data that could be used for this purpose are potentially available, including NASS surveys, data from other agencies, and automated field-level data collected by farm equipment. The panel was asked to consider technical issues involved in using these data sources, such as methods for integrating the data, the assumptions underpinning the use of each source, the robustness of the resulting estimates, and the properties of desirable estimates of uncertainty.

The current NASS approach to integrating multiple data sources is through its Agricultural Statistics Board (ASB). While the current process follows specific steps and guidelines, it is inherently subjective and neither transparent nor reproducible. This could be improved if the ASB were provided with high-quality model-based estimates that synthesize multiple data sources.

NASS itself has observed that if it could develop and adopt model-based approaches to integration of multiple data sources, it could bring the agency into conformance with the statistical standards promulgated by the U.S. Office of Management and Budget (OMB) (2006).[1] These standards lay out the steps to be taken to ensure that published statistical information is transparent, reproducible, and reliable.

The panel's recommendations are based largely on presentations and deliberations at its four open meetings, during which it heard from many NASS staff about the agency's surveys, frames, remote sensing indications, auxiliary data sources, the process used to prepare estimates, and work on developing and implementing model-based estimates; staff of other USDA agencies about administrative data, use of NASS estimates, and the potential for use of remote sensing information; representatives of other statistical agencies that have developed and implemented model-based approaches to small-area estimation; and from academics with expertise in model development.

KEY RECOMMENDATIONS

The panel has chosen to present its major recommendations in terms of a vision for NASS in 2025. This vision has three components.

First, in the future NASS prepares for the ASB county-level estimates based on models that incorporate multiple data sources, as well as uncertainty measures for the estimates.

> **NASS should evolve the ASB role from one of integrating multiple data sources to one of reviewing model-based predictions; macro-editing; and ensuring that models are continually reviewed, assessed, and validated.** (Recommendation 2-1)
>
> **NASS should achieve transparency and reproducibility by developing, evaluating, validating, documenting, and using model-based estimates that combine survey data with complementary data in accordance with Office of Management and Budget standards.** (Recommendation 2-2)
>
> **NASS should adopt and use the following publication standard:**
>
> - **County-level estimates may be withheld to protect confidentiality.**
> - **County-level estimates may be withheld because NASS deems them unreliable for any use, based on its measure of uncertainty.**
> - **All other county-level estimates will be published, along with their measure of uncertainty.** (Recommendation 2-3)
>
> **NASS should develop and publish uncertainty measures for county-level estimates.** (Recommendation 2-4)

Second, the NASS list frame is a georeferenced farm-level database, serving as a sampling frame for surveys and facilitating the use of farm data in statistical analysis.

[1] See especially Standard 4.1.

NASS should adopt FSA's Common Land Unit (CLU) as its basic spatial unit. (Recommendation 2-8)

NASS should be prepared to maintain alternative geospatial field-level boundaries (e.g., resource land units [RLUs] and precision agriculture measurements) in its databases to facilitate completing the geospatially referenced farm-level database. (Recommendation 2-9)

NASS should update its list frame to make use of CLU information as geospatial building blocks by developing linkages between the NASS list frame and FSA and RMA administrative data. (Recommendation 2-10)

Third, NASS acquires all georeferenced administrative and remotely sensed and ground-gathered data relevant to developing estimates and uses these data to complement its survey data.

NASS should collaborate with RMA to obtain relevant individually identifiable acreage and production data and to conduct comparisons with NASS data for the same entity. (Recommendation 3-1)

NASS should collaborate with farmer cooperatives to ensure that it is one of the government agencies with which farmers can choose to share their relevant precision agriculture data. (Recommendation 3-3)

NASS should develop a precision agriculture reporting option for the County Agricultural Production Survey (CAPS)/Acreage, Production, and Stocks (APS) survey system. Farmers who reported relevant precision agriculture data would either not receive an additional survey form or receive one that was simplified and easy to use. (Recommendation 3-5)

NASS should explore collaboration with other USDA agencies that are actively involved in remote sensing applications to obtain access to data with finer spatial resolution, and possibly also to share in the costs of processing those data. (Recommendation 3-8)

NASS should keep abreast of emerging data sources; how they are used; and how they might be used to improve county estimates, especially of yield. Based on a careful evaluation, NASS might consider purchasing data. (Recommendation 3-9)

Finally,

NASS should undertake a staged, systematic effort to implement the vision presented in Chapter 2 of this report. (Recommendation 5-1)

1

Introduction

The National Agricultural Statistics Service (NASS) is the primary statistical data collection agency within the U.S. Department of Agriculture (USDA). As noted on its website,[1] the agency conducts hundreds of surveys annually and prepares reports covering virtually every aspect of U.S. agriculture. Among the small-area estimates produced by NASS are county-level estimates for crops (planted acres, harvested acres, production, and yield by commodity) and for cash rental rates for irrigated cropland, nonirrigated cropland, and permanent pastureland. Key users of these county-level estimates include USDA's Farm Services Agency (FSA) and Risk Management Agency (RMA), which use the estimates as part of their processes for distributing farm subsidies and providing farm insurance, respectively.

In September 2014, NASS entered into a cooperative agreement with the Committee on National Statistics of the National Academies of Sciences, Engineering, and Medicine to assess these county-level estimates and offer recommendations for their improvement. The statement of task for the study is shown in Box 1-1.

BOX 1-1 Statement of Task

An ad hoc panel under the auspices of the National Research Council will review, assess, and make recommendations on methods for integrating multiple data sources to improve county-level crop estimates produced by the National Agricultural Statistics Service (NASS). The goal is to provide more precise estimates with appropriate measures of uncertainty for current county-level estimates of acreage, yield, and cash rents for major crops. Multiple sources of data are potentially available for county-level crop estimates, including NASS surveys, data from other agencies, and automated field-level information collected by farm equipment dealers. The panel will explore methods for combining the information from these and other sources to produce more precise county-level estimates with valid measures of uncertainty. Issues to consider include the methods used to integrate the information, the assumptions underpinning each approach, the robustness of the estimates to a failure of one or more assumptions, and other technical issues. In addition, the panel will consider the suitability of using each data source and the properties of desirable estimates of uncertainty. The panel will produce a final report with findings and recommendations at the conclusion of the study.

[1]See https://www.nass.usda.gov/About_NASS/index.php [July 11, 2017].

The study panel was specifically directed to consider annual county-level estimates for cash rents, crop acreage, and yields. These county-level estimates are published after NASS has published national- and state-level estimates for the year. The challenge of integrating multiple data sources to improve crop estimates needs to be accomplished in a way that brings NASS into conformance with the statistical standards promulgated by the U.S. Office of Management and Budget (OMB) (2006),[2] which specify the steps to be taken to ensure that published statistical information is transparent, reproducible, and reliable.

NASS COUNTY-LEVEL ESTIMATES AND THEIR EVOLVING USES

USDA began regularly reporting estimates of final acreage, yields, and production for principal crops in 1866. Today, NASS is responsible for estimating acreage, production, and yield for most commodities grown in the United States. National- and state-level estimates are among OMB's primary economic indicators. NASS has been publishing county-level crop and livestock inventories since 1917. Originally, these county-level estimates were prepared with federal funding, an arrangement that subsequently evolved into partnerships involving state funding via cooperative agreements. Crop surveys usually were sponsored by states, with NASS statisticians in the state offices defining the samples and processes used and developing the estimates.

NASS's county-level estimates support the efficient functioning of agricultural markets by providing information about the supply of, demand for, and use of commodities. Participants in agricultural markets rely on such information to make decisions: for producers, about what to grow and how to manage inventories; for processors and traders, about how to organize production and determine sales; and for retailers and consumers, about how to anticipate costs and assess the availability of food. When market participants share a common understanding of the fundamentals of supply and demand, market transactions accurately reflect the value of commodities to those along the supply chain and help ensure that food is grown, processed, and consumed at the lowest cost to the nation.

Complementing the data on agricultural production, NASS has long estimated rental rates for farmland at the state level. Between 1950 and 1994, state-level cash rents were estimated using a list survey of real estate appraisers. Beginning in 1994, state-level cash rents were estimated primarily by using the June Area Survey to ask farmers directly about their rental agreements. The 2008 farm bill extended the requirement for state-level rental estimates by mandating that NASS provide mean rental rates for all counties with at least 20,000 acres of cropland plus pasture.

Cash rents generally reflect the value of farmland as an input to agricultural production. These estimates are used by farm operators and land owners in negotiating rental agreements; by bankers in making farm operating and ownership loans; and by real estate agents, analysts, financial advisors, and extension services, among others. As with the crop production estimates, information on cash rents contributes to the efficiency of farmland markets by providing participants in these markets with data on supply and demand.

Over the past century, then, NASS has played a key role in providing information to participants in agricultural markets. At the same time, national farm policy has evolved, as has the use of agricultural statistics in administering federal programs. Since the 1930s, the federal government has provided support to the farm sector. Initially, this support entailed direct

[2]See especially Standard 4.1.

intervention in markets whereby supply was controlled through restrictions on what acreage farmers could plant, and prices were set nationally. With the rise in the importance of global trade and attendant policy liberalization, the federal government's approach became more market oriented and less interventionist. This trend culminated in the 1996 farm bill's direct "decoupled" payment scheme, based on past, not current, production. The 2002 farm bill reintroduced price-based "countercyclical" payments so that aid is provided when current revenues fall. This aid also is increasingly conditioned on the circumstances of each farmer, with payments being more closely related to the farmer's own revenue experience. Thus, the need for accurate and disaggregated information on acreage, yield, and output to enable better alignment of payments with farmers' actual outcomes is increasingly pressing.

One of the mainstays of farm support under the 2014 farm bill is Agricultural Risk Coverage (ARC)-County Option (CO), delivered by FSA. Payments to farmers are based on the difference between a guaranteed level and an estimate of county revenue outcomes computed as the product of NASS season average price data and NASS season-end estimates of county yields. FSA also administers the Conservation Reserve Program (CRP), whereby the federal government makes a payment to a farmer in exchange for removing environmentally sensitive land from agricultural production and planting species that will improve environmental health and quality. Before 2008, FSA administered the CRP using rental rates estimated by appointed FSA boards of directors in each state. In 2009, FSA began using results from the NASS statistical survey of county average rental rates for cropland and pastureland to establish the rental rates used for the CRP. Federally subsidized crop insurance programs, such as the ARC-CO and group risk options, use NASS county yields in setting the guaranteed yield levels for policies based on outcomes in a farmer's county (as opposed to the farmer's personal experience).

NASS took steps to respond to these administrative requirements for information at the county level. First, the agency initiated the County Agricultural Production Survey[3] (CAPS), beginning in a few states in 2011 and implemented in all eligible states in 2012. Today, CAPS provides survey data with which to estimate acreage and production for selected crops at the county level for use in state and federal programs in 44 states. CAPS replaced the nonprobability samples that previously had been selected and administered by state offices. Today, the NASS county crop estimates program is defined jointly by NASS, RMA, and FSA. States are partners in the process and may add commodities to the program to address local interests. In addition, the biannual Cash Rents Survey was initiated in most states in 2008 to provide county-level estimates for use by FSA in administering the CRP.

In summary, in the last decade NASS has responded to the increasing uses of county-level data by introducing large-scale probability surveys to provide the required information. Despite these initiatives, however, farmers have raised questions with program agencies about the use of estimates derived from these surveys because they feel that a county average does not represent their own farm experience. As demonstrated in a presentation to the study panel by FSA, relatively small variation in a county yield estimate can result in relatively large changes in subsidy payments and possibly in a payment's not being triggered at all.[4] Substantive variability

---

[3]See http://www.nass.usda.gov/Surveys/Guide_to_NASS_Surveys/County_Agricultural_Production [October 2015].

[4]Ideally, the conditions for triggering or determining the size of an ARC payment would be exogenous to the individual receiving the payment. Using county estimates based only on the sample risks making the payments dependent on the actions of the beneficiaries when the survey response rate is low. As a consequence, both the

in payment rates from one county to the next also can be driven by these yield differences. The monetary significance of a county yield estimate therefore focuses attention on how it is derived and, for an individual farmer, on how well it corresponds to that farmer's own yield outcomes.

Another issue is difficulties that have emerged with shrinkage in the number of counties for which survey-based estimates are available, particularly in some areas of the country with high participation in federal commodity programs. When NASS county estimates are not available, FSA must resort to other, secondary sources of information that may not reflect a county's experience with fidelity.[5] For example, as noted by Johansson and colleagues (2017) "In the event that there is neither a NASS county estimate nor enough data to estimate an RMA county yield, the FSA State Committee will determine the county yield using best available data, including such possibilities as the NASS or RMA yield for a neighboring county, the NASS district yield estimate, or 70 percent of the transitional yield (or t-yield). NASS districts include multiple counties, which may make the yield determination too high for some counties and too low for others."

NASS strives to provide yield estimates for as many counties as possible while ensuring the quality of these estimates for use by market participants and federal agencies. NASS applies a publication standard to determine those counties for which the number of survey responses is sufficient to support estimation. This standard requires that there be at least 30 valid responses[6] for a crop in a county, or if there are fewer, they collectively must account for at least 25 percent of acreage by crop. If this standard is not met, a county's estimate is not published. Unfortunately, survey response rates have been declining over time, and the result has been a decrease in the number of counties for which estimates meet the standard. For the Cash Rents Survey, for example, the response rate fell from a U.S. average of 75.9 percent in 2014 to 69.1 percent in 2016 (U.S. Department of Agriculture-National Agricultural Statistics Service, 2014, 2016). Reflecting this decline, NASS published data for 2,879 counties in 2014 and 2,597 in 2016. For the CAPS, the response rate for row crops was 62.8 percent in 2011 and 56.4 percent in 2016, while the response rate for small grains was 69.1 percent in 2011 and 63.4 percent in 2016.[7] A decline in responses in a county increases the likelihood that an estimate cannot be published.

FSA is required by law to make payments at the county level, regardless of the availability of NASS estimates. If NASS cannot provide the necessary data, FSA must turn to other channels. This situation can be particularly problematic in counties with significant participation in ARC-CO. The ARC-CO program provides revenue loss coverage at the county level. ARC-CO payments are issued when the actual county revenue (as estimated by the product of NASS market prices and NASS end-of-season estimates of county yield) of a covered commodity is less than the preset ARC-CO guarantee for that commodity. If NASS county-level data are not available, RMA or state committee yield is used instead. An example of this situation occurred in the Northern Plains in 2015. Response rates in 2015 for surveys on production of small grains (wheat, barley, and oats) are shown in Figure 1-1a. Counties in which fewer than half of those surveyed responded—circumstances in which NASS is least likely to

---

credibility of the estimates and the perceived fairness of the size and distribution of payments are called into question.

[5]See http://nfu.org/wp-content/uploads/2016/03/ARC-County.pdf [July 11, 2017].

[6]NASS attempts to publish estimates for approximately 3000 counties, requiring much larger sample sizes that are much more costly than producing estimates only for the nation and state.

[7]Based on a May 2016 presentation to the panel by Lindsay Drunasky, updated in April 2016 by Nathan Cruze. Drunasky's presentation also included response rate maps by state showing patterns of response.

publish an estimate—are concentrated in the Dakotas, Kansas, and western Minnesota. As seen in Figures 1-1b and 1-1c, these are areas in which ARC-CO payments are potentially significant. Figure 1-1b shows the range of county wheat revenue earned from the market in the absence of any federal payments; this map shows that the Upper Plains states had numerous counties in which that revenue fell at or below historical averages. Figure 1-1c indicates the significance of ARC-CO payments when they are added to market revenue, showing that revenue in many of those Upper Plains counties subsequently topped historical averages. The Congressional Research Service has noted "wide discrepancies" in the yield estimates[8] for adjacent counties, some with and some without NASS estimates, which presumably only exacerbated producers' concerns about the credibility of the estimates used to calculate ARC-CO payments (Congressional Research Service, 2017).

FIGURE 1-1a Small grains county estimates survey response rates

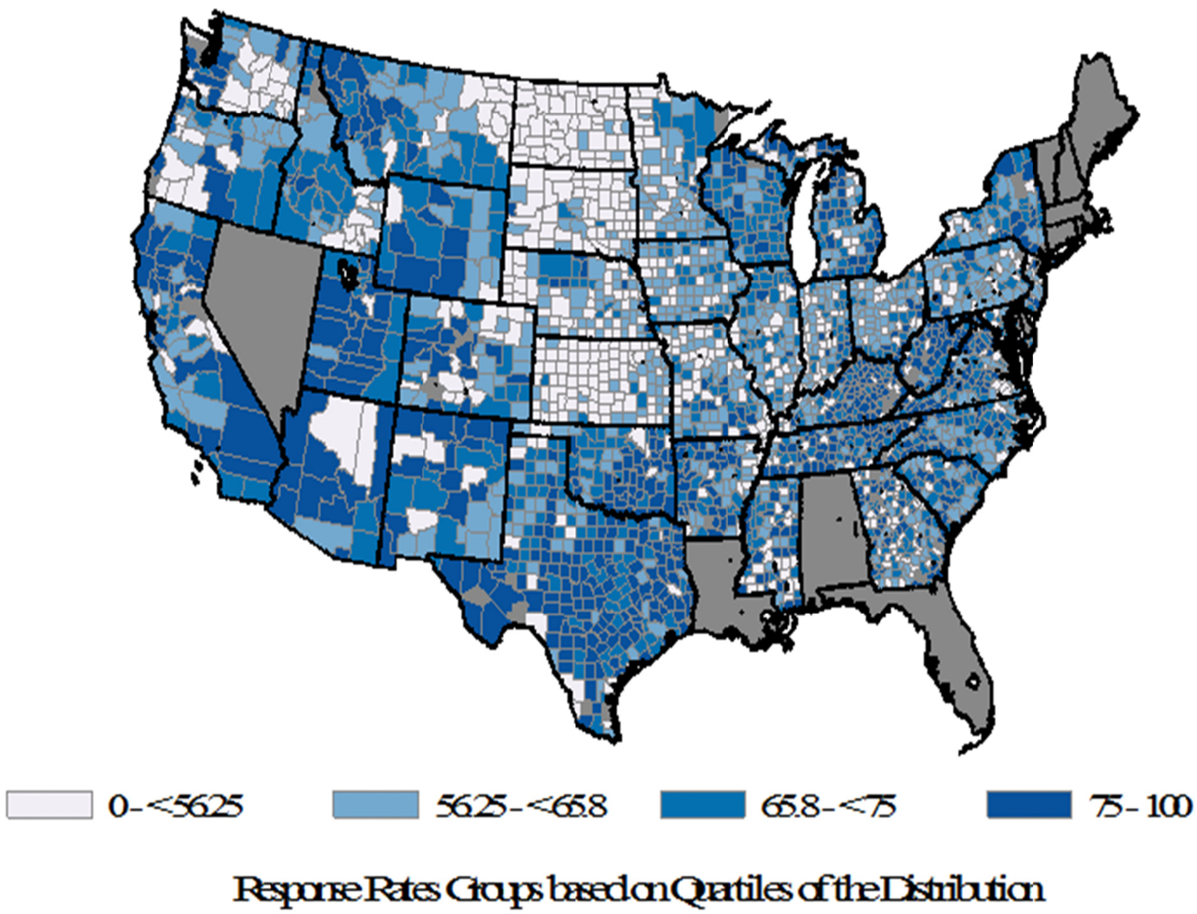

[8]Wide discrepancies in yield estimates may trigger legitimate discrepancies in ARC payments. The equations used are very sensitive.

FIGURE 1-1b 2015 crop wheat revenue without ARC-CO safety net

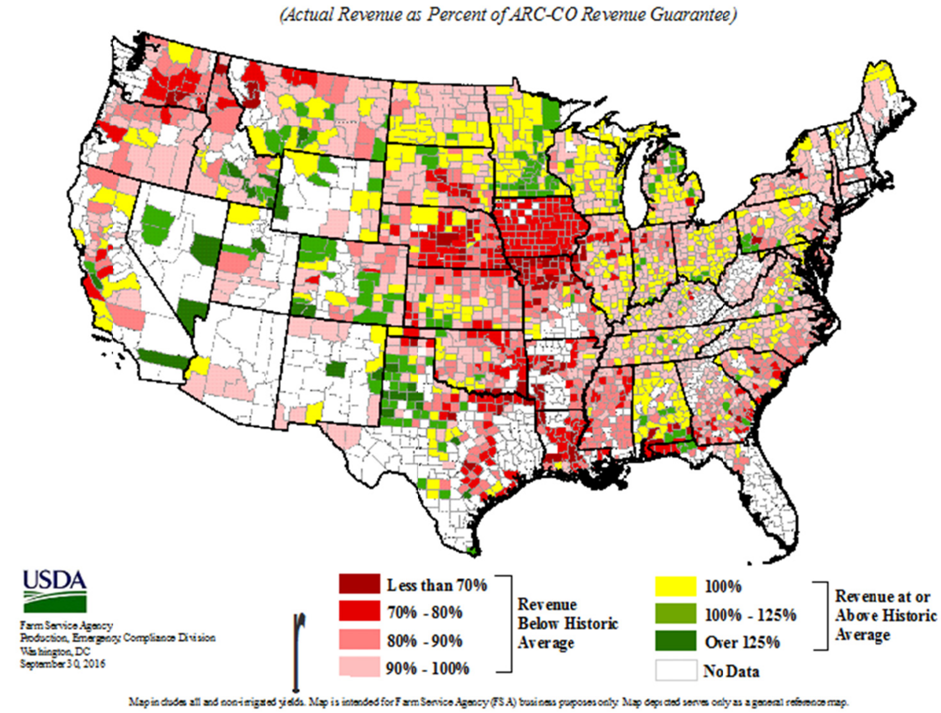

FIGURE 1-1c 2015 crop wheat revenue with ARC-CO safety net

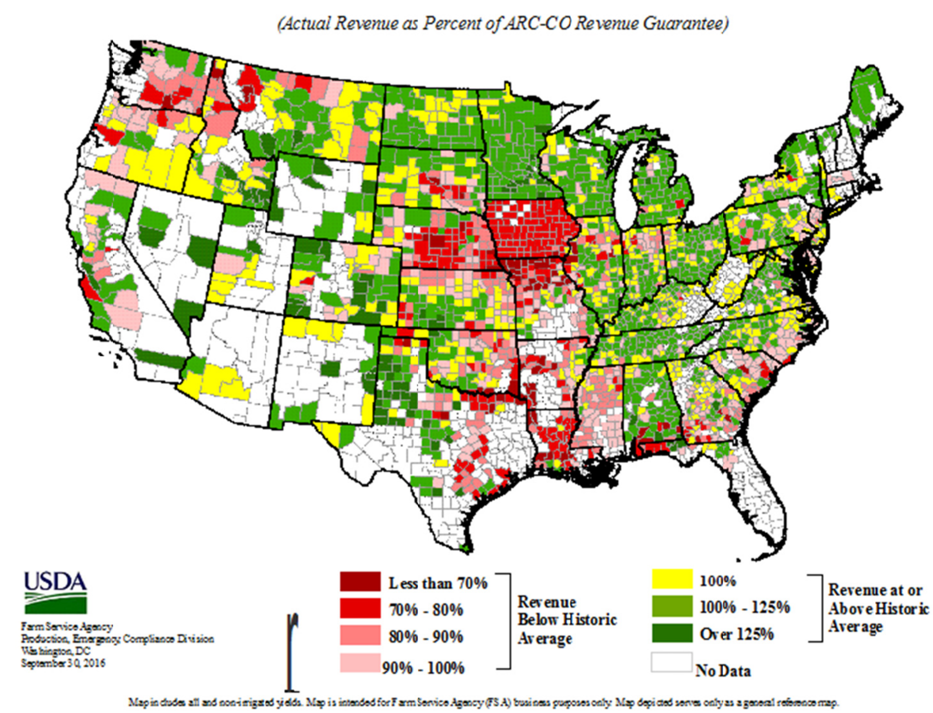

NASS has again responded by working to improve survey response rates. These efforts have included evaluating farms on the list frame to ensure that they are still operating as recorded, publishing articles and pamphlets describing the importance to farmers of participating in NASS surveys (Schleusener, 2016), and considering ways to make it easier for farmers to respond. Nonetheless, almost all federal statistical agencies are seeing a drop in survey response rates, so NASS faces a systemic problem not unique to agriculture. (See, for example, National Academy of Sciences, Engineering, and Medicine [2017] for a full discussion of threats to the survey paradigm and National Research Council [2013] for a discussion of declining rates of response to social science surveys.)

At the same time, it is important to note that survey results by themselves may not determine the official NASS county estimates for acreage, yield, or rental rates. The Agricultural Statistics Board (ASB), part of NASS, reviews all state-, district-, and county-level recommendations made by state and federal staff. ASB establishes official estimates based on a collaborative process involving NASS field offices that collect and process the data and headquarters statisticians who provide oversight and coordination. The federal statisticians consider direct survey results and their associated quality measures, as well as a range of additional information, including data from previous surveys and censuses, and, when appropriate, administrative data collected from farmers by FSA and RMA, remote sensing data on crop production, and estimates from statistical models. For the cash rents estimates, separate "indications"[9] are drawn from the Cash Rents Survey, past surveys and censuses, and a model that combines 2 years of data from the Cash Rents Survey with information on relevant crop condition and agronomic factors. Regional field office statisticians derive estimates using this information and then provide justification to the Board should their recommendations differ from the survey indications. The panel learned that in 2014, 34,452 county-level estimates were made, with 939 (8.2%) ASB changes to indications from the Cash Rents Survey. Yet despite recruiting other information to augment the county survey results, NASS still may not judge derived estimates sufficiently valid to be published, which in the case of county yields prevents their use in determining farm payments.

The decline in the number of counties for which NASS can publish statistically valid direct survey estimates has created challenges for its key customers, FSA and RMA. As noted earlier, these agencies rely on NASS estimates to the extent possible, but if NASS estimates are not available, they rely on other data. These include the NASS survey estimates for the Agricultural Statistics District (a group of counties including the one in question), production data from RMA, and subjective assessments of local FSA officials.

Given these circumstances, NASS, as it has in the past, is considering how to respond to the needs of program agencies for more comprehensive and reliable estimates of county yields and cash rents. While attempting to raise survey response rates, NASS also looks to other sources of information that may better inform official estimates developed through the ASB process. As evidenced by the nature of the indications already being considered by ASB, survey data can be augmented in a number of ways. Expanding the use and sophistication of statistical models is one obvious way to bolster the ability to publish county estimates and to do so in a systematic and transparent way. Models can potentially make use of information on crop output obtained from such sources as satellite data; field-level observations gathered by farm

[9]"Indication" is the NASS term for a preliminary calculation or judgment about the value of a variable. The output of the Board process that considers all indications is published and is termed the official NASS "estimate."

machinery; administrative data on farm operations reported to USDA program agencies; and other auxiliary information, such as weather and soil quality.

NASS also is facing a changing context for the development of crop estimates. Satellites and advanced farm machinery, respectively, are emerging as aerial and ground-based sources of information on agronomic conditions. Not only can such data be used by NASS in its estimation process, but they increasingly are being used by private-sector entities to derive their own estimates for commercial sale. Satellites and drones can generate readings on plant health, local weather, and crop conditions on a daily basis, all of which can be combined to provide real-time estimates of crop progress. Likewise, increases in the sophistication and capabilities of farm machinery have made it possible to gather field-level data on input application and crop output, which are examples of measurements used in precision agriculture, a crop management approach utilizing data from GPS systems and sensors on such equipment as tractors and combine harvesters. Both remotely sensed and ground observations can enhance analysts' ability to monitor, predict, and measure crop conditions and outputs. There are challenges to effective use of such data, however, including how to combine these data to produce accurate crop estimates. In the case of precision agriculture measurements, access to the information may be hindered by questions of ownership: the farmer, or the firm that sells and services the machinery, or the company that analyzes the precision agriculture data as a service. Ultimately, the credibility of NASS in developing crop estimates for public use may depend on its ability to acquire and interpret these data as private-sector competition grows. Accordingly, this report considers how administrative, remotely sensed, and ground-gathered data might be integrated into NASS models to yield more accurate and timely estimates.

STUDY APPROACH

In the course of this study, the panel held four in-person meetings that included public sessions to collect information and one closed meeting to finalize this report. The panel also held many teleconferences. In addition, the panel requested and received numerous briefings from NASS on a wide range of topics, including the CAPS and the Cash Rents Survey; sample design and data collection; issues related to nonresponse and its impact; the publication standard and how it is used; the availability of auxiliary information, particularly remote sensing and FSA/RMA administrative data; the Board process for deriving final county estimates; and NASS research into the development of small-area estimates to combine data sources.

The panel heard as well from representatives of FSA and RMA who described how they use NASS data in their programs, along with the Acreage and Crop Reporting Streamlining Initiative, which is improving the quality and availability of RMA and FSA data. Representatives of other agencies within USDA described their research on the use of new satellite data to predict crop yields and the development of spatial data on land characteristics and linkages to farm data from FSA. The panel also learned about a previous study of publication standards used by other statistical agencies and entities to publish direct survey data. The panel heard from the U.S. Census Bureau and the U.S. Bureau of Labor Statistics, statistical agencies that use small-area estimation methods as part of their publication programs. Another presentation described Statistics Canada's adoption of a model-based approach using remote sensing data to replace a midseason survey estimate. Finally, the panel heard from yield modelers who are experts in midseason forecasts of yield.

The panel believes that with this report, it has satisfied its statement of task (Box 1-1) to the extent possible. Terms in the statement of task include "will explore more methods" and a list of "issues to consider." While the panel could not pursue model development for each specific NASS estimate, general guidelines are offered in Appendix C. The panel has also advised NASS staff to continue their current cash rents and crop estimates modeling efforts as described in Chapter 5, Projects 2 and 3. The methods being employed in those efforts can be expanded to incorporate spatial terms. Once Common Land Units (CLUs) have been linked to NASS farms, unit-level models (or either farms or CLUs) can be pursued, as discussed in Appendix C.

ORGANIZATION OF THE REPORT

The panel presents its recommendations for improvement in terms of a vision for NASS in 2025 that is described in Chapter 2. This vision is motivated by the imperatives for a modern statistical agency to be responsive to its data users, to adopt the most robust technologies for data management, and to utilize nontraditional data sources and statistical methods.

Chapter 3 describes the multiple data sources that are available to enhance NASS's county-level estimates of planted acres, harvested acres, production, and yield by commodity; describes how these sources are currently used; and suggests improvements for the future. Included among these multiple data sources are administrative data from FSA and RMA, satellite and other remote sensing data, and data from precision agriculture measurements.

Chapter 4 reviews the data sources that might inform the estimation of cash rents. It also describes a model that is currently used as input to the ASB process and explores how this model might be carefully assessed to determine how ASB could use it in a transparent and reproducible manner.

Finally, Chapter 5 suggests how NASS could implement the vision set forth in Chapter 2 by describing specific projects that would move the agency forward. These projects are linked to the panel's recommendations, and are intended to involve different skills and different stages of implementation.

The report also includes four appendixes. Appendix A describes the survey methodology of the Cash Rents Survey and the CAPS. Appendix B details steps NASS might follow should it choose to adopt the Routine External Evaluation Protocol (REEP) as part of an external evaluation, as suggested by Dorfman (2017). Appendix C explains small area time and spatial modeling approaches for multiple sources of data. Finally, Appendix D provides biographical sketches of the panel members.

2

A Vision of NASS in 2025

This chapter presents the panel's vision for NASS in 2025. This vision has three components. First, NASS prepares its county estimates using a transparent and well-documented process, publishing measures of uncertainty along with point estimates. Second, the NASS list frame is a georeferenced farm-level database, serving as a sampling frame for surveys and facilitating the use of farm data in statistical analysis. Third, NASS acquires all relevant georeferenced administrative and remotely sensed and ground-gathered information and uses this information to complement its traditional survey data.

THE VISION

The panel's vision for NASS in 2025 is motivated by the imperatives for a modern statistical agency to be responsive to its data users, to adopt the most robust technologies for data management, and to utilize nontraditional data sources and statistical methods. As discussed in Chapter 1, the end uses of NASS data have expanded beyond informing market decisions to determining federal subsidy payments. The advent of geographic information systems (GIS) technology for data collection, management, and analysis has particular relevance for the location-specific circumstances of agriculture. And the emergence of remote sensing and precision agriculture measurements as means of acquiring timely georeferenced meteorological and agronomic data offers the opportunity to augment survey and administrative data and to enhance statistical modeling. Moving to make the vision presented in this chapter a reality will pose many challenges for NASS. Nonetheless, the adoption of new ways of doing business is necessary to maintain the agency's credibility and to promote efficiency in its operation. The following subsections describe in turn each of the three components of the panel's vision for NASS in 2025.

Vision of a Transparent, Well-Documented Process

Achieving the first component of the panel's vision will put NASS in compliance with the Standards and Guidelines for Statistical Surveys promulgated by the U.S. Office of Management and Budget (OMB) (2006). OMB Standard 4.1, Developing Estimates and Projections, lays out the steps to be taken to ensure that published statistical information, whether from surveys or from models that use survey data, is transparent, reproducible, and reliable.

> Standard 4-1: Agencies must use accepted theory and methods when deriving direct survey-based estimates, as well as model-based estimates and projections that use survey data. Error estimates must be calculated and disseminated to

> support assessment of the appropriateness of the uses of the estimates or projections. Agencies must plan and implement evaluations to assess the quality of the estimates and projections (U.S. Office of Management and Budget, 2006).

In this report, the terms "measures of uncertainty" and "uncertainty measures" are used to describe error estimates that support assessment of the appropriateness of potential uses of the estimates. (For a discussion of uncertainty and models, see National Research Council [1991a, pp. 89–96]). Box 2-1 provides the OMB guidelines related to Standard 4.1.

BOX 2-1 Office of Management and Budget Guidelines for Statistical Surveys

Section 4 provides standards and guidelines for production of estimates and projections. The following guidelines represent best practices that may be useful in fulfilling the goals of Standard 4.1, Developing Estimates and Projections.

Guideline 4.1.1: Develop direct survey estimates according to the following practices:

1. Employ weights appropriate for the sample design to calculate population estimates. However, an agency may employ an alternative method (e.g., ratio estimators) to calculate population estimates if the agency has evaluated the alternative method and determined that it leads to acceptable results.
2. Use auxiliary data to improve precision and/or reduce the error associated with direct survey estimates.
3. Calculate variance estimates by a method appropriate to a survey's sample design taking into account probabilities of selection, stratification, clustering, and the effects of nonresponse, post-stratification, and raking. The estimates must reflect any design effect resulting from a complex design.

Guideline 4.1.2: Develop model-based estimates according to accepted theory and practices (e.g., assumptions, mathematical specifications).

Guideline 4.1.3: Develop projections in accordance with accepted theory and practices (e.g., assumptions, mathematical specifications).

Guideline 4.1.4: Subject any model used for developing estimates or projections to the following:

1. Sensitivity analysis to determine if changes in key model inputs cause key model outputs to respond in a sensible fashion;
2. Model validation to analyze a model's performance by comparing the results to available independent information sources; and
3. Demonstration of reproducibility to show that, given the same inputs, the model produces similar results.

Guideline 4.1.5: Prior to producing estimates, establish criteria for determining when the error (both sampling and nonsampling) associated with a direct survey estimate, model-based estimate, or projection is too large to publicly release the estimate/projection.

Guideline 4.1.6: Document methods and models used to generate estimates and projections to help ensure objectivity, utility, transparency, and reproducibility of the estimates and projections. (For details on documentation, see Section 7.3). Also, archive data and models so the estimates/projections can be reproduced.
For more information on developing model-based estimates, see FCSM Statistical Policy Working Paper 21, Indirect Estimators in Federal Programs.

SOURCE: Excerpted from U.S. Office of Management and Budget (2006).

For its direct survey-based indications, NASS uses accepted theory and methods, and sampling error is estimated. However, NASS's published estimates are not necessarily those derived from its surveys. Instead, as summarized in Chapter 1 and described in more detail later in this chapter, NASS uses the Agricultural Statistics Board (ASB) process or expert judgment to determine the official estimate given a variety of auxiliary indications. The Board process is neither transparent nor reproducible. The panel believes that to follow the OMB standards, NASS needs to develop statistical models to replace the subjective Board process. In addition to developing and testing models, NASS would have to lay out the plan for ongoing routine evaluation procedures, develop and publish measures of uncertainty for county estimates, revise its publication standards to incorporate sound criteria and rules for publication of estimates, change the role of the ASB to rely on appropriate model-based estimates for transparency, provide documentation on its website of the survey and model methodologies that lead to county estimates, and plan and implement evaluations to assess those methodologies. The panel is confident that acceptable models for crop estimates and for cash rents can be developed. Especially for crop estimates, the auxiliary information available to NASS provides strong predictors.

Vision of a Georeferenced Farm-Level Database

Development of a georeferenced farm-level list frame, the second component of the vision, would be done most expeditiously by adopting the geospatial convention already in use by the Farm Services Agency (FSA) and Risk Management Agency (RMA) (both important sources of administrative farm data). FSA defines the Common Land Unit (CLU) as an individual, contiguous farming parcel, which is the smallest unit of land that has a permanent, contiguous boundary; common land cover and land management; and a common owner and/or common producer association.[1] FSA digitizes CLUs into GIS shapefiles or geodatabases and populates the associated farm data. NASS would identify the CLUs that make up each NASS farm,[2] requiring changes to the list frame to accommodate the georeferenced CLUs. NASS would have to determine how to collect or identify CLU (or equivalent) information for farms that are not on FSA or RMA lists.

[1]See the 2017 Common Land Unit information sheet at https://www.fsa.usda.gov/Assets/USDA-FSA-Public/usdafiles/APFO/support-documents/pdfs/clu_infosheet_2017_Final.pdf [June 2017].
[2]Note that the CLUs that make up a NASS farm can change from year to year with leasing arrangements.

Vision of Using Data from Other Sources to Complement Survey Data

Achieving the third component of the vision will require progress on the second component and a nimble approach to surveys. Having georeferenced CLUs as the basic unit for the list frame will allow linkage to FSA and RMA administrative records, as well as to GIS data generated by remote sensing and precision agriculture measurements. To the extent that NASS can acquire information already provided by farmers to the U.S. Department of Agriculture (USDA), the respondent burden of surveys will potentially be reduced. Moreover, crop acreage data might be available earlier than current surveys permit, although the question of timing is complicated by agencies' differing schedule requirements. Still, greater reliance on administrative data could allow redirection of scarce survey resources to farms not in FSA or RMA records or to areas where survey response is low. Although the administrative data are subsets of the universe, they can be used in statistically valid ways as described at the end of this chapter. As for observations from remote sensing and precision agriculture measurements, they can be used as additional information in statistical models along with survey data and support unit-level modeling. Indeed, the ASB process already considers such sources, although not apparently using a formal statistical analysis.

Summary of the Vision

A commitment to achieving the vision outlined above will require changes in the way NASS manages and structures its operations. Figure 2-1 illustrates the panel's concept of the ultimate configuration of the NASS process for developing county-level estimates. As discussed, multiple data sources, along with the accompanying measures of uncertainty, will be integrated through formal statistical models that will provide the basis for ASB's decisions. ASB will exercise its judgment in deciding when model estimates can be improved by consideration of unforeseen events (e.g., a drought or a hurricane) or of systemic changes (e.g., rapid farm consolidation) that are not well captured in the models. ASB will determine whether a county estimate should not be published. As a means of providing quality control, ASB will drive a feedback loop with analysts that suggests modifications to improve model performance and interpretation. Documentation of these models and records of ASB decisions will be available to the public so that increased transparency bolsters confidence in the robust nature of the estimation process. A fuller discussion of the steps needed to achieve this vision follows. The section below describes how the current NASS environment affects the vision components, while the following section describes key factors in implementing the use of models needed for the vision. Two succeeding sections describe how to make the case for a change to more transparent use of models and the issue of selecting basic spatial units for a geospatially referenced database. The final section describes NASS surveys in the future.

CURRENT NASS ENVIRONMENT

Traditionally, government statistical agencies avoided the use of models and favored survey-based estimators having sound statistical properties with respect to the sample design theory. The most prominent virtue of the purely sample-based approach is the perceived objectivity of resulting estimates, stemming mainly from the avoidance of making further assumptions that require validation, as well as from the relative simplicity and tractability of the formulas used to develop estimates.[3] Since important government policy decisions often rely on

FIGURE 2-1 Summary of the vision of NASS in 2025

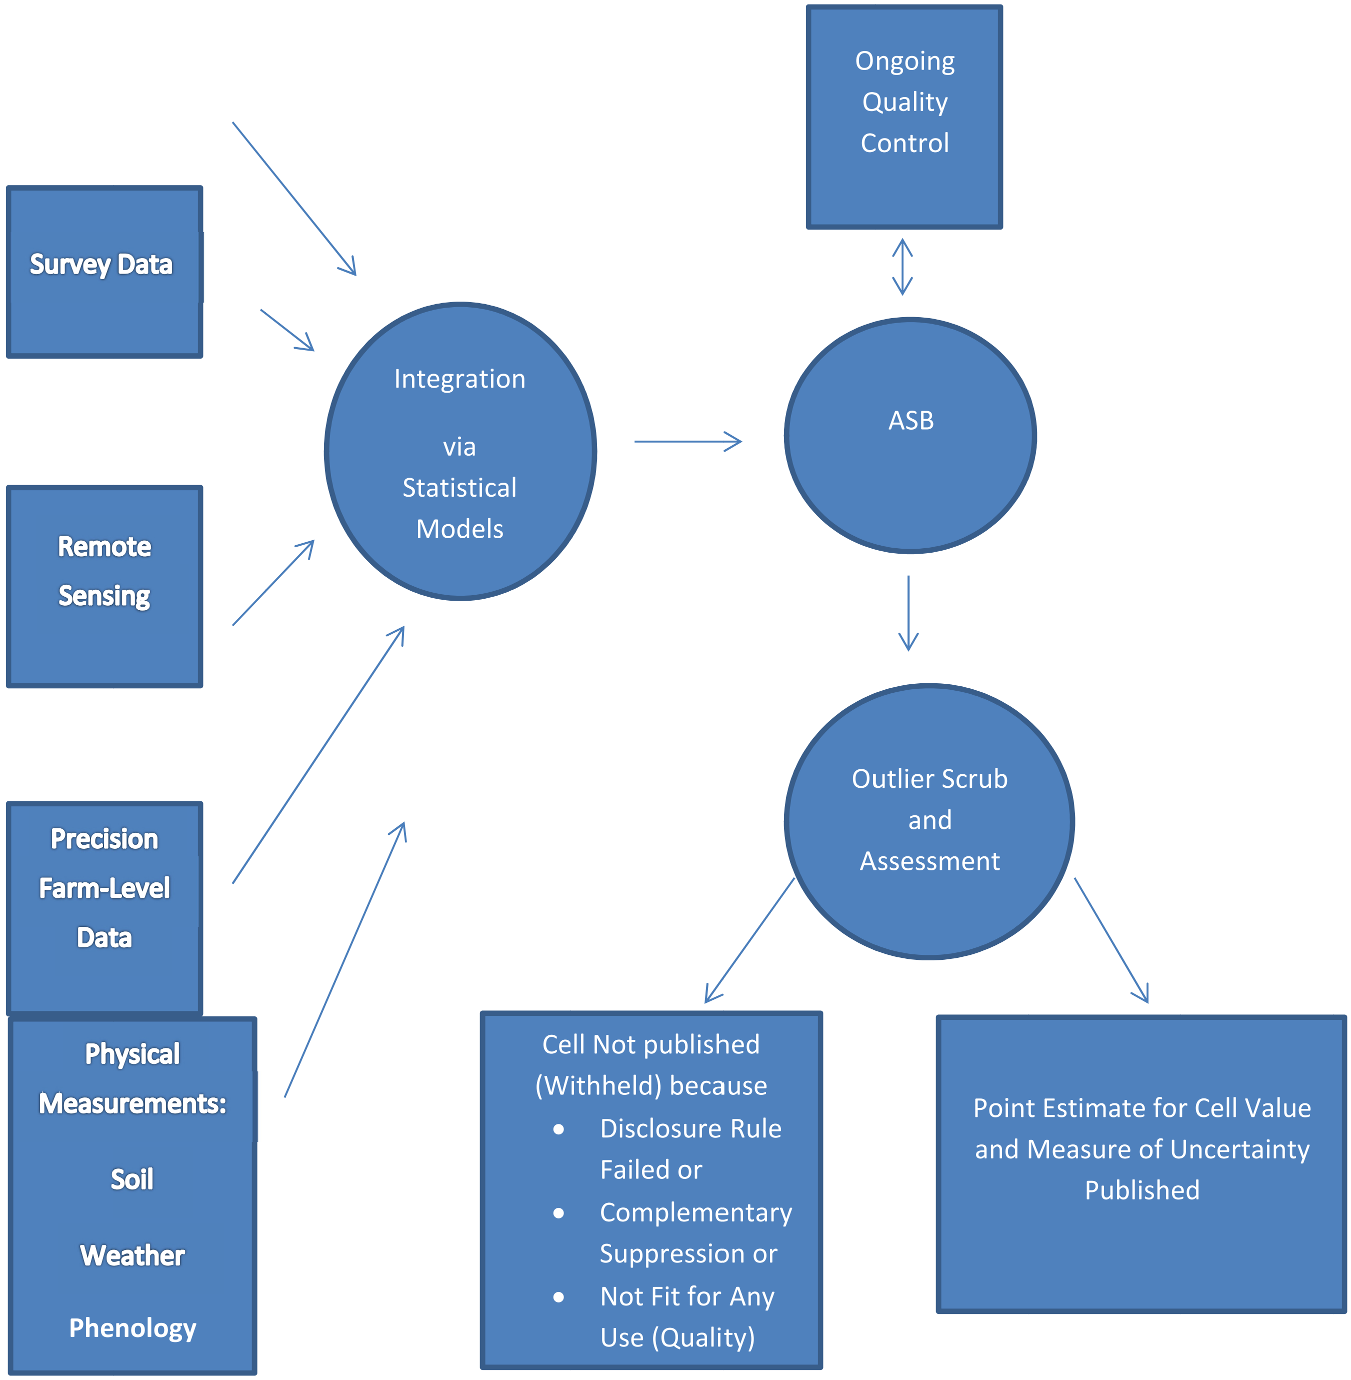

NOTE: ASB = Agricultural Statistics Board; FSA = Farm Services Agency; RMA = Risk Management Agency.

[3]See for example, Kalton (2002).

these numbers, the impartiality of the estimation process is of particular significance for statistical agencies.

In recent years, however, as a result of growing demand for estimates at finer levels of detail, together with challenges with meeting the budget constraints of complex surveys on the one hand and advances in both theoretical and practical aspects of the discipline of small-area estimation (SAE) on the other, statistical agencies have gradually become more open to the use of models as a transparent and sensible way to achieve fusion of survey-based information with other data sources in pursuit of greater precision in estimation.

It is now generally accepted that, while the pure direct sample-based approach is appropriate for large samples, incorporating additional information with the help of statistical models is essential in smaller domains where samples are scarce. For a number of years, models have been employed in such important programs as Small Area Income and Poverty Estimates (SAIPE) (discussed further below) and Small Area Health Insurance Estimates (SAHIE) of the U.S. Census Bureau and for small domains in the Current Employment Statistics (CES) survey conducted by the U.S. Bureau of Labor Statistics (BLS), as well as surveys conducted by Statistics Canada and others.

Like other statistical agencies, NASS has preferred to rely on traditional sample survey–based approaches, and many within NASS are concerned with the challenges associated with validating complex assumptions. At the same time, especially for the crop estimates program, NASS has been fortunate to have a variety of high-quality alternative data sources available in addition to survey data. The NASS approach to using these sources of information has been to charge ASB (described below) to consider all preliminary "indications" in some detail according to a set of rules and to make a judgment as to the most accurate "estimate." This process may result in estimates that are perceived to be accurate, but is neither transparent nor reproducible.

The process of integrating multiple data sources according to a set of rules (however accomplished) is already a departure from the purely survey-based approach as it is based on a set of implicit assumptions pieced together in a deliberative process such as that used by ASB (see Box 2-2). The challenge with a "deliberative model" is that it is neither transparent nor reproducible.

After giving a basic definition of a model in Box 2-2, this section describes the current ASB process for estimation, as well as four formal statistical models developed by NASS, three of which are already used to provide input to the ASB process.

BOX 2-2 What Is a Model?

A model is a means of organizing relevant information to aid in understanding a real-world phenomenon. Currently, the Agricultural Statistics Board uses an informal model when it combines information from surveys, remote sensing, and past outcomes to produce estimates of county yields. This report urges greater use of statistical models, which are mathematical equations that relate an outcome to a set of inputs or determinants in a setting of uncertainty. Models are necessarily a simplified representation of the essential causes or characteristics of the outcome of interest. Here, crop yield estimates at the county level are known to be influenced by randomness in weather and in farming techniques, among other factors. Statistical modeling provides a systematic way to predict yield by quantifying its relationship to key determinants, taking into account the uncertainty inherent in agricultural processes. The advantage of using statistical—rather than ad hoc, informal—models is that they are transparent (the equation can be written down with precision) and reproducible (another user can duplicate the results).

SOURCE: Prepared by the panel.

The Agricultural Statistics Board Process

For national- and state-level estimates, ASB meets in a carefully controlled “lock-up” to review all available indications (survey, administrative, and/or model-based) and information and prepare estimates for dissemination. A similar ASB process, but not conducted in a “lock-up,” is used to set county-level estimates. Regional field office (RFO) statisticians collect survey data, perform editing, and conduct further analysis of the data. Estimates are established using survey indications as well as alternative sources of information.

NASS requires consistency of its published estimates at all geographic levels: national, principal region, state, Agricultural Statistics District (a group of neighboring counties within the same state, denoted ASD), and finally county. National, principal region, and state estimates for acreage and production are based on surveys from the NASS list frame as well as data from the June Area Survey (an area sample intended to account for undercoverage of the list frame). Estimates for crop acreage and production generally are set first by ASB for the total produced for the region containing the principal states in which a commodity is grown. State totals in that region must add to that regional total and are set second. Then the national estimate is derived from the principal region total and the total for the remaining states producing the commodity. While no national estimate of cash rental rates is published, data obtained from the June Area Survey are used to publish state-level estimates.

For both estimation programs, NASS employs a “top-down” approach when establishing its substate estimates, meaning that national and state estimates generally are published before ASD and county estimates, even while additional data collection may be ongoing. To incorporate survey and auxiliary information into its crop estimates, NASS computes composites of available inputs for its acreage and production totals. The ASD-level composites are ratio adjusted to the published state total, and then rounding rules are enforced in producing the official ASD statistics. Subsequently, the county composite estimates are ratio benchmarked to the *rounded* ASD totals; rounding rules are then enforced again at the county level to produce official county estimates. At each level, official NASS yield estimates are obtained as the ratio of the final production estimate to the final harvested area estimate.

In a similar manner, benchmarking is performed on total rented acres and total rent expenditures in support of estimates of cash rental rates. Upon review of all information, ASD totals are adjusted to state totals, and county totals are in turn adjusted to ASD totals, subject to NASS rounding rules. Note that while this important consistency check is performed, NASS does not publish the separate acres rented or total rent estimates, only the cash rental rates.

For crops, alternative data sources include administrative data on planted acreage by crop from FSA and failed acreage by crop from RMA, and remote sensing estimates for acreage and yield for some commodities. NASS estimation software creates a weighted average indication from the benchmarked survey, administrative, and remote sensing indications using recommended weights provided by ASB. RFO staff use these composite indications as a starting point from which to set estimates for acreage, yield, and production for each commodity. For cash rents, the direct survey estimate is considered most important, and the only alternative indication available is a model-based result. Estimates established by RFO staff are submitted for review to ASB, with headquarters and RFO staff working from their respective duty stations. The RFO provides justification to ASB in cases in which recommended estimates deviate from survey results. ASB members review the estimates for accuracy and consistency across state boundaries, verify that proper procedures were followed throughout the process, and establish official estimates.

A report of the National Research Council (2008, p. 1230) states that 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 that are published by USDA." That report offers the following recommendation: "Recommendation 6.9: NASS and ERS [USDA's Economic Research Service] should provide more clarification and transparency of the estimation process" (p. 126). Adrian (2012) comments on the ASB process and its subjectivity in preparing in-season yield forecasts. He also suggests the use of statistical models as a way for NASS to achieve transparency and reproducibility.

Especially for crops, ASB is charged with assessing indications from multiple diverse data sources and producing coherent estimates for multiple geographic areas. ASB was established in 1905,[4] long before automated methods of combining multiple data sources were available. The panel believes that the current ASB process for crop estimates relies on expert judgment and is neither reproducible nor transparent, and that NASS needs to pursue the development of methods based on formal statistical models that produce measurably accurate county-level estimates. Once such formal procedures were available, vetted, and accepted, the ASB function would be to carefully assess results for reasonableness. It is a widely accepted notion that estimates need to be reviewed before they are released to the public. The final review of well-defined estimates is commonly performed by analysts in other statistical agencies.[5] The process for reviewing estimates needs to follow clear guidelines and result in a limited number of well-documented adjustments to otherwise automatically produced estimates, thus ensuring the transparency of the process.

Recommendation 2-1. NASS should evolve the ASB role from one of integrating multiple data sources to one of reviewing model-based predictions; macro-editing; and ensuring that models are continually reviewed, assessed, and validated.

Note that this recommendation maintains a role that involves judgment by ASB. Maintenance of this role acknowledges, as observed by Bunn and Wright (1991, p. 508) in their literature review, that in practice, expert ex post adjustment of model-based forecasts may be more accurate than automatic adoption of a model-based estimate. Typically, the expert judgment is better when the model contains specification error or structural change makes the current situation different from the past. Bunn and Wright (1991, p. 508) also observe, however,

[4]For a history of ASB, see U.S. Department of Agriculture-National Agricultural Statistics Service (2007).

[5]For example, in BLS's CES, analysts review sample-based estimates to see whether there are any outliers. If extreme estimates are found, analysts discuss the estimates as a team that includes analysts from the national office and from the states, and make decisions based on these discussions. Discussions are usually carried out via email within a very short time frame. As a first step in their review, analysts look at micro records to see whether an outlying estimate is due to a single or a few atypical observations or is supported by a larger sample. If adjustments are made to the estimates, the size of and reasons for those adjustments are documented.

that the expert judgment approach is "hard to justify to others and may undermine the credibility of the whole process." Hence this review process itself should be as transparent as possible, with changes made only for well-defined and documented reasons. The NASS Estimation Manual describes data sources and comparisons on which the NASS statisticians and analysts rely in reviewing the estimates. Making this manual publicly available might be a good start toward improving the transparency of the ASB review process.

Development of Models by NASS

NASS has been active in the pursuit of useful models for many years. Those NASS models that are particularly relevant to the preparation of county estimates are described below. The panel's vision of NASS in 2025, particularly the use of models to synthesize multiple data sources to achieve transparency and reproducibility, is consistent with the views of many researchers within NASS, as is evident from comments in papers by NASS staff (see, for example, Adrian, 2012; Busselberg, 2011; and Cruze, 2015a, 2015b).

The four modeling approaches described below either provide indications for the current ASB process or have the potential to provide such input. These approaches can be considered as the basis for the development of improved methods.

The first approach is based on **the composite indication**, a linear combination of basic indications prepared from (1) direct survey estimates, (2) administrative data, (3) remote sensing data, and (4) past values of indications based on surveys. The panel was told that each of the basic indications is first ratio benchmarked so the sum over counties is equal to the previously published state total. Iwig (1996) describes the precursor to the current approach. In the early 1990s, the weights were determined individually and subjectively by state offices. NASS described the current approach to the panel at its November 2015 meeting. The currently used weights[6] (Table 2-1) are set by ASB and are considered to be relatively stable over time. They are changed when needed. For example, if remote sensing estimates were improved, the weight on that indicator would be increased. ASB uses the composite indication as a starting point in developing official NASS estimates.

TABLE 2-1 Weights Used to Compute the Composite Indication

Indication/Data Source	(1) Direct Survey	(2) Planted Indication Times Direct Survey Ratio	(3) FSA	(4) Planted Indication Minus RMA Failed	(5) Remote	(6) Past Data
Planted acres (remote available)	0.15		0.65		0.2	
Planted acres (remote not available)	0.2		0.8			
Harvested acres	0.14	0.7		0.08		0.08
Production (remote sensing available)	0.12	0.78			0.05	0.05
Production (remote sensing not available)	0.15	0.8				0.05

NOTE: FSA = Farm Services Agency; RMA = Risk Management Agency.

[6] Note that weights for each composite indication sum to 1 over data sources.

In Table 2-1 the columns reflect variables that are inputs to the composite estimator. Column (1) contains the direct survey estimate for the indication shown in each row. Column (3) is the FSA estimate for planted acres; column (5) is the remote sensing estimate (planted acres and production); and column (6) contains the estimate published the previous year. Column (2) is a variable calculated as the product of a direct survey estimate for a ratio (harvested acres over planted acres or production over harvested acres) and the composite estimator for planted acres or harvested acres (as appropriate) from a row above. Column (4) is the product of the direct survey estimate for the ratio of harvested acres to planted acres and the composite estimator for planted acres from a row above.

The rows reflect the indications that will be determined during the ASB process: planted acres, harvested acres, and production. They are computed sequentially. The indication for planted acres is computed as a linear combination of the direct survey indication, the FSA indication, and the remote sensing indication (if one exists for that crop and county). The indication for harvested acres is computed as a linear combination of the direct survey estimate, the product of the just-established indication for planted acres and the direct survey indication for the ratio of harvested to planted acres, and the difference between the direct indication for planted acres and the RMA failed acres. The indication for production is computed as the linear combination of the direct survey indication for production; the product of the direct survey indication for yield and the just-established indication for harvested acres; and the remote sensing estimate, if one exists. If the remote sensing estimate does not exist, the past year's production estimate for the county is used in the composite.

The weights shown in Table 2-1 provide an indirect NASS assessment of the quality of various input data sources. FSA data have the highest weight in calculation of planted acres. Column (2) variables—computed as the product of a direct survey indication for a ratio and an acreage value from a previous step—have the highest weight in computing harvested acres and production, respectively. Remote sensing indications (when available) receive little weight for production but higher weight than the direct survey estimate for planted acres.

The composite estimate could be viewed and used as a model-based estimate if weights were derived statistically. For example, under the assumptions that the components are unbiased and mutually independent, the minimum variance of the composite estimator is achieved by using the weights that are inversely proportional to the variances of respective component indications. A similar approach is currently used for domain estimation in the CES (Gershunskaya, 2012). The above assumptions concerning lack of bias and independence, as well as the form of the target estimator, may be overly restrictive. More flexible classes of models exist (for example, the classical Fay-Herriot model [Fay and Herriot, 1979] or adaptations) that can be used to effectively "mimic" the current ASB estimation process.

The second approach is based on **the NASS Cropland Data Layer** (CDL) (a remote sensing product), which provides pixels classified by land use (including a range of crops) as the basis for independent county-level indications of planted acreage by crop and yield for corn and soybeans in the Midwest cornbelt. Schaible (1996, Chapter 6) describes the remote sensing estimates of acreage that have been produced since the early 1990s. The key changes are that today's Landsat satellites have smaller pixel sizes, and NASS uses the FSA CLU data as the training dataset for classifying pixels; in 1993 the June Area Survey data were used for that purpose. NASS makes use of the Battese and Fuller (1981) model, along with classified pixels from the CDL and segment mapping data from the June Area Survey, to produce remote sensing indications of planted acres by crop.

According to Johnson (2014) a data mining software product, Rulequest's Cubist,[7] is used to develop a predictive model for yield. Model inputs include weekly measurements of the Normalized Difference Vegetation Index (NDVI) during the growing season, temperature data from satellites over the past 6 years, and 6 years of past data on yield to establish a model for county yield. The current-year NDVI and temperature data are used as model inputs to predict current-year yield. (See Chapter 3 for additional detail on NASS's approaches for preparing indications for planted acres and yield using satellite remote sensing variables.)

As described above, these estimates are input to ASB's composite indication for some crops and some counties. The low weight given to remote sensing indications for yield in developing the composite indication indicates that ASB does not regard them as particularly accurate in general, although they may be of value for some crops in some regions.

The third approach is based on the **Berg, Cecere, and Ghosh (2014) model for cash rents**. NASS implemented this model and used the results, along with direct survey estimates (the primary indication), as an indication for cash rents in the ASB process during 2013, 2014, and 2016. The model was derived on the basis of conducting the Cash Rents Survey annually. Revised language in the 2014 farm bill, however, specified that the Cash Rents Survey would be conducted "[no] less frequently than once every other year." The survey was not conducted in 2015, and the 2016 production run of the cash rents model input 2016 and 2014 survey data. The Cash Rents Survey will be conducted in 2017.

The cash rents model of Berg, Cecere, and Ghosh (2014) uses two univariate (Fay and Herriot, 1979) area-level models—one for the average of the cash rents in two survey years and the other for the difference. The authors used least squares to estimate the model parameters. The model accounts for the correlation between cash rents in the two survey years and the assumed equality of variances in these two years. The authors constructed a single index to capture the impact of several sources of auxiliary information, including (1) the county-level total dollar value of agricultural production from the 2007 Census of Agriculture; (2) county yields published by NASS for 2004–2009 to reflect the quality of the land in the county (separately for irrigated and nonirrigated cropland, where possible, and a separate hay yield index); and (3) three National Commodity Crop Productivity Indices (NCCPIs) as developed by the National Resources Conservation Service (NCCPI-corn, NCCPI-wheat, and NCCPI-cotton). The authors also used a two-stage process to benchmark the ASD estimates to the corresponding aggregated state estimates and the county-level estimates to the corresponding benchmarked ASD estimates.

After converting the cash rents model to make use of the surveys conducted in 2014 and 2016, NASS provided the results to ASB as an indication in 2016. Potential enhancements and an evaluation of the adaptation of the model to account for data collected every two years are described in Chapter 4. One challenge is that the panel does not know how the model results are actually used by ASB.

The fourth approach is based on **NASS approaches to end-of-season model-based crop estimation**, summarized by Erciulescu, Cruze, and Nandram at the panel's January 2017 meeting. Previous efforts by the same authors include Cruze et al. (2016) and Erciulescu et al. (2016). In their January 2017 summary to the panel, they noted two types of models. The first is unit-level models (Battese and Fuller, 1981; Battese et al., 1988). These models have not yet been actively pursued by NASS. Linkage of the CLU information to the NASS list frame should

[7]Described by the company's website as a machine learning tool that automatically determines the best-fitting piecewise linear model to predict a continuous outcome variable. See https://www.rulequest.com/cubist-info.html [August, 2017].

make the development of unit-level models more feasible. The second type of model comprises area (e.g., county-level) models, such as those addressed by Fay and Herriot (1979), and subarea-level models, discussed by Fuller and Goyeneche (1998) and Torabi and Rao (2014). NASS's current efforts have made use of subarea models (for counties) so that benchmarking (to ASDs) can be incorporated into the modeling process.

Erciulescu, Cruze, and Nandram have explored estimation of planted acres, harvested acres, and yield in four states Iowa, Illinois, Indiana, and Kansas). The parameters of interest in these models are the unknown true planted acres, harvested acres, and yield, written as a linear combination of auxiliary variables, including FSA planted acres and a National Oceanic and Atmospheric Administration (NOAA) weather variable, as well as random effects. Production is defined as the product of harvested acres and yield. The direct survey estimate divided by the sample size, conditional on unknown parameters, is assumed to follow a normal distribution, with mean equal to the parameter of interest (the regression equation) and variance related to the sample design. Appropriate assumptions also are made about the conditional distributions of parameters, including variances.

Erciulescu, Cruze, and Nandram used a Bayesian model fit using Markov chain Monte Carlo (MCMC) methods and compared the model results with survey estimates, estimates prepared by ASB, and either FSA planted acres or remote sensing estimates where applicable. Ratio benchmarking was applied to posterior iterations. Point estimates and estimated variances were based on ratio-adjusted posterior iterations.

The authors state that the model-based point estimates agree quite well with the ASB estimates, and benchmarking constraints are satisfied. They state further that the estimated variances/covariances are smaller than those of the direct survey indications. The authors are continuing to work though some of the challenges associated with use of their model, but results are promising. The model as implemented appears to be a somewhat expanded version of a Fay-Herriot (1979) model that relies most heavily on survey data, with the intent that the auxiliary information will help reduce the mean square error of the estimate relative to that of the survey.

This strategy may lead to models that are useful, and its consideration by ASB is warranted. However, it does not provide a way to integrate multiple data sources. Of all the auxiliary indications available with which to develop the required estimates for crops, the only one used in this model at present is FSA planted acres. The authors account for skewness (only in the survey data) by using the sample sizes as a weighting variable. See Chapter 5 for more detail on potential improvements for the modeling of crop estimates.

While NASS has actively pursued the development of models with a clear goal of improved transparency and reproducibility, the panel finds that substantial work is still needed to fulfill the vision set forth in this chapter. As has been observed for the Census Bureau's SAIPE, which is the first of the major small-area models employed by statistical agencies,

> The development of small-area estimates of income and poverty is a major effort that includes data acquisition and review, database development, geographic mapping and geocoding of data, methodological research, model development and testing and documentation and evaluation of procedure and outputs. Since the production of small-area poverty estimates supports a range of important public policies for federal, state and local governments—including the allocation of funds—it is essential that the Census Bureau have adequate staff and other resources for all components of the estimation program, including evaluation and

> documentation. It is the responsibility of any agency that produces model-based estimates to conduct a thorough assessment of them, including internal and external evaluations of alternative model formulations. (National Research Council, 2000a, p. 167)

These observations apply equally to NASS as it pursues its small-area estimation program. Note that achieving this goal requires access to statistical/modeling talent through in-house staff, collaboration with other USDA and government agencies, and collaboration with academic researchers.

> **Recommendation 2-2. NASS should achieve transparency and reproducibility by developing, evaluating, validating, documenting, and using model-based estimates that combine survey data with complementary data in accordance with Office of Management and Budget standards.**

The initial models described as part of the third and fourth approaches detailed above are promising starts, and the committee provides advice on potential improvements to these models in Chapters 4 and 5.

IMPLEMENTING THE USE OF MODELS

NASS's adoption and full use of models will require identifying candidate models; updating the NASS publication standard; evolving the role of ASB; documenting the models; internally evaluating and externally validating the models; and perhaps most important, preparing convincing documentation and communication using measures of uncertainty to demonstrate that the selected model-based approaches provide high-quality, defensible estimates. This convincing evidence will be the key to ultimate adoption and use of the models by ASB.

The Committee on National Statistics has previously addressed how to establish and/or maintain the quality of model-based estimates in at least three National Research Council studies: on the Census Bureau's SAIPE program (National Research Council, 2000a); on microsimulation modeling of retirement income (National Research Council, 1997); and on microsimulation models (National Research Council, 1991a, 1991b). All of these reports recommend that these models be subject to internal validation, external validation, and sensitivity and uncertainty analysis, recommendations that apply to any important model.

The subsections in this section show how the NASS publication standard is related to uncertainty analysis and confidentiality and then describe these issues in general, as well as that of documentation. The discussion in this section illustrates these issues and others, such as internal and external evaluation, using the SAIPE program (National Research Council, 2000a), which produces model-based estimates for school districts, counties, and states. The main objective of this now well-established small-area modeling program is to provide updated estimates of income and poverty for use in the administration of federal programs and the allocation of federal funds to local jurisdictions. These estimates combine data from administrative records, postcensal population estimates, and the decennial census with direct estimates from the American Community Survey (ACS) to provide consistent and reliable single-year estimates. These model-based single-year estimates are more reflective of current conditions relative to multiyear survey estimates.[8]

The Census Bureau's county estimates of poor school-age children are produced using a county and a state regression model (National Research Council, 2000b, pp. 47–54). The Census Bureau uses the Fay-Herriot regression models, similar to models being considered by NASS. The National Research Council (2000b, pp. 54–61) report concludes that a comprehensive evaluation of these two components of the estimation procedure should include both internal and external evaluations, as detailed below.

Internal Evaluation

Internal evaluations investigate the validity of a model's underlying assumptions and features, and typically are based on examination of the residuals from a regression model considering both functional form and error characteristics. In the evaluation of SAIPE documented in National Research Council (2000a, pp. 60–65), categories of counties were specified, and comparisons were conducted for each category. Functional form was assessed according to

- linearity between dependent and predictor variables,
- constancy of linear relationships over time, and
- inclusion or exclusion of predictor variables.

Error characteristics were assessed according to

- normality (symmetry and moderate tail length),
- homogeneous variances (variability constant across counties and not dependent on predictor variables), and
- the absence of outliers.

Methods for internal model evaluation (e.g., the deviance information criterion [DIC], cross-validation and out-of-sample comparisons) are described in the literature and should be part of the standard toolkit of researchers who work on developing models for NASS county estimates. (See, for example, the review of recently developed methods for model selection and checking in Pfeffermann [2013] and a discussion of issues related to goodness-of-fit approaches in mixed-effect generalized linear models in Tang et al. [2014].)

The panel has little evidence regarding the extent to which NASS has pursued detailed internal model evaluations for its current models. It may be that analysts conduct some such evaluations, but do not include details of these evaluations in published papers. Such details are appropriately included in model documentation.

[8]See https://www.census.gov/did/www/saipe/about/index.html [July 11, 2017].

External Validation

External evaluation of potential models can shed light on their usefulness. NASS might conduct external model evaluation during preliminary stages of analysis, at the time of research, and for periodic monitoring of models already selected.

Comparison with a census, when available, is a standard approach to external validation. For example, BLS's CES relies on the ability to compare estimates with the Quarterly Census of Employment and Wages (QCEW) levels over a number of years of estimation to provide assurance concerning the validity of the model estimates. As another example, the National Research Council (2000b) report suggests that external evaluation of SAIPE could employ comparison with decennial estimates (with the understanding that the census is not perfect) or assessment by stakeholders with local knowledge.[9] The SAIPE report was prepared before the decennial census long form was replaced by the ACS. The availability of data from the ACS has improved the accuracy of SAIPE estimates, but external validation by comparison with the decennial census is now no longer an option.

In a report documenting its assessment of the predecessor of SAIPE, the National Research Council (1980, p. 10) states, "An estimate is considered accurate if it is close to the value of the parameter it is estimating, which is typically unknown." The report goes on to say that "ideally, an estimating procedure (or estimator) should meet four criteria: 1) low average error; 2) low average relative error; 3) few extreme relative errors; and 4) absence of bias for subgroups." In a more recent assessment of SAIPE (National Research Council, 2000a, pp. 194–195), evaluation measures were summarized and illustrated by application to the 1990 census. The four measures used here were (1) the average absolute difference (the overall absolute model–census difference in terms of numbers of poor school-age children), (2) the average proportional absolute difference (the overall model–census difference in terms of percentage errors for counties), (3) the category algebraic difference (for selected categories of counties, such as population size and metropolitan status), and (4) the category average proportional difference. In this example, the evaluation was conducted for the key models under consideration.

A census does not always provide a truly perfect standard because, like other surveys, a census has limitations related to nonresponse, undercoverage, classification error, and response error. However, a census can provide a useful benchmark for estimates that represent the same time period as the census. The **Census of Agriculture**, conducted every 5 years in years ending in 2 or 7, is appropriate for assessing the consistency of NASS county-level estimates for harvested acres and production (or yield as estimated by the ratio derived from production divided by harvested acres) for the year of the census. The Census of Agriculture provides no information with which to assess cash rents or planted acres. The Census of Agriculture has several drawbacks:

- It is conducted only once every 5 years.
- It is affected by undercoverage, nonresponse, and misclassification.
- It cannot be used in real-time estimation.

[9] National Research Council (2000b, p. 57) describes the review by local stakeholders of 1990 decennial census estimates. The Census Bureau identified groups of counties for which estimates appeared to be high or low and contacted knowledgeable local people, such as state demographers and state data center staff, to review them and provide input.

- It has good but not complete coverage of county estimates published by NASS.

The **published ASB estimates** are appropriate for use in assessing the quality and reasonableness of NASS model-based county-level estimates, especially at the research stage:

- ASB estimates are used for comparison for harvested acres in Bellow and Lahiri (2012) and for corn and soybean yields in Williams (2013).
- For certain groups of states, ASB estimates were used for exploratory analysis and shown to be within the confidence limits of model-based forecasts in most cases (Wang et al., 2012).
- Nandram and colleagues (2014, p. 509) use ASB estimates as response variables in models, thus highlighting the level of trust in those estimates.

Comparisons of model estimates with ASB estimates are likely to be necessary to convince ASB members that the model results are worthy of consideration.

Finally, Dorfman (2017) proposes that the **Routine External Evaluation Protocol (REEP)** be used as part of external evaluation. He argues that, as the demand for small-area estimation is ever-increasing, statisticians have the responsibility to develop a protocol aimed at setting criteria and enabling detection of "a tipping point" at which estimates produced using small-area models cannot be regarded as satisfactory. This determination is particularly difficult to make in groups of areas in which there is a small or no sample and that may behave differently from larger areas. For NASS, REEP would be based on comparing county model estimates with the corresponding direct sample-based estimates, with the latter being based on a large enough sample. The procedure would be implemented by selecting a random sample of counties within which a supplementary sample of farms would be selected and used to prepare direct estimates. Appendix B provides steps NASS might consider if it decided to pursue REEP.

NASS Publication Standard

The current NASS publication standard is related specifically to the publication of estimates derived from probability sample surveys. A county estimate from a sample survey can be published only if the survey had 30 nonzero respondents for the commodity represented by the cell, or fewer than 30 if the reported acreage of those respondents was at least 25 percent of the total county acreage for that cell. Bell and Barboza (2012) studied the effect of the currently used publication criteria on the coefficient of variation (CV) of the estimates and found that these criteria do not always lead to estimates with adequate precision. In addition, the panel was told by one FSA data user that in his opinion, cells with at least 30 respondents are of acceptable accuracy. However, he thought this was not the case for some of those cells with a smaller number of respondents. Although NASS might benefit from a fresh look at the part of its publication standard that deals with fewer than 30 respondents, the panel observes that the major disadvantage of the current standard is that it was designed specifically to be used for direct survey estimates and cannot apply to any other type of estimate. NASS needs a publication standard that applies to estimates derived from models as well as those derived from surveys.

The panel has chosen to recommend a generic publication standard that bases a publication decision on protection of respondent confidentiality and measures of uncertainty for point estimates. This proposed standard reflects NASS goals of protecting the confidentiality of respondents to surveys and publishing high-quality data. It also reflects the panel's judgment that the user is key to deciding whether an estimate is of sufficiently high quality for his or her own use. To make such judgments possible, the estimate needs to be accompanied by a measure of uncertainty. Box 2-3 provides more detail on the quality of NASS estimates.

BOX 2-3 The Quality of NASS Estimates

The quality of a NASS estimate depends on how well it corresponds to the real-world outcome. Statistical models can be evaluated with regard to their ability to predict the event of interest, and statisticians employ a variety of measures to characterize a model's "goodness of fit." A high-quality model will reliably return predictions that come close to the actual result. But how close is close enough? The answer is that the user determines how accurate a prediction must be for it to provide a satisfactory basis for a real-world decision. For example, greater accuracy is likely required when a subsidy payment depends on a NASS estimate of county yield than when the same estimate is used to compare multiyear trends in yield across counties. The NASS decision to publish an estimate (or not) presumes the agency understands the requirements of its users. Alternatively, NASS could provide information about the uncertainty of an estimate, thereby allowing the user to make his or her own judgment about accuracy and fitness for use in a specific context.

Recommendation 2-3. NASS should adopt and use the following publication standard:

- **County-level estimates may be withheld to protect confidentiality.**
- **County level estimates may be withheld because NASS deems them unreliable for any use, based on its measure of uncertainty.**
- **All other county-level estimates will be published, along with their measure of uncertainty.**

Measures of Uncertainty

To comply with the OMB standard, NASS must develop and publish measures of uncertainty for each of its published county-level estimates. NASS has measures of uncertainty for its direct survey indications, and indications based on statistical models also are expected to be accompanied by measures of uncertainty. Measures of uncertainty for estimates of planted acres derived from remote sensing are provided by Schaible (1996, Chapter 7). However, NASS told the panel that measures of uncertainty currently are not available for remote sensing indications even though the estimates are derived from models. Nor are measures of uncertainty for indications based on administrative data currently available. Measures of uncertainty are key to assessing the various indications and determining how they can best be used. For data users, measures of uncertainty are key to assessing estimates for fitness for use.

In addition to uncertainty measures derived as part of model estimation, uncertainty measures can be derived using such validation approaches as leave one or several out cross-validation. A popular Bayesian approach is the construction of posterior predictive p-values.

Recommendation 2-4. NASS should develop and publish uncertainty measures for county-level estimates.

Confidentiality

Like other statistical agencies, NASS collects survey data under a pledge of confidentiality and uses statistical disclosure limitation methods and suppression to protect the data. The result is a fair amount of data suppression, and this will likely continue to be the case as long as NASS relies on surveys and its current confidentiality criteria. Suppression sends a message to respondents that their data are protected, but researchers can and have provided estimates of those values to the public. (See, for example, Lokupitiya et al. [2007], in which estimates of suppressed values of area planted and yield for all counties in the United States for a period of 16 years are derived.) There has been research on confidentiality approaches that do not involve suppression, but the panel has not seen any that are ready for implementation in an operational environment. NASS would benefit from considering the implementation of new research methods for confidentiality protection as they evolve and are demonstrated to be useful.

The unanswered question is how/whether suppression analysis is required with estimates derived from auxiliary data sources (not surveys) and/or models. FSA[10,11] and RMA[12] both collect administrative data for their programs. They differ in the extent to which they view their data as being confidential, but both agencies share individually identifiable data with other USDA agencies.

Recommendation 2-5. NASS should work with its partner agencies, FSA and RMA, to determine whether to apply confidentiality protection to individually identifiable administrative data used by NASS for statistical purposes and if so, how to provide that protection. For data for which confidentiality protection is needed, NASS may propose that it apply the same confidentiality rules as those applied to NASS survey data collected under a pledge of confidentiality.

For estimates derived from models, the answer is less clear. Disclosure limitation methods are intended to ensure that the contribution to the cell total of the largest survey respondent cannot be estimated too closely from the published total. This assurance is of particular concern when the population is small. For example, in a county dominated by one

[10]FSA form 578 states, "The following statement is made in accordance with the Farm Security and Rural Investment Act of 2002, (Pub. L. 107-171). The information will be used to determine to whom program benefits will be paid. Furnishing the requested information is voluntary; however, failure to furnish the correct and complete information will result in a determination of ineligibility for program benefits. This information may be provided to other agencies IRS, Department of Justice, or other State and Federal law enforcement agencies, and in response to a court magistrate or administrative tribunal. The provisions of criminal and civil fraud statutes, including 18 USC 286, 287, 371, 641, 651, 1001; 15 USC 714m; and 31USC 3729, maybe applicable to the information provided."

[11]FSA's annual acreage spreadsheet provides acreage data by crop by FSA farm without using a threshold or dominance rule. However, if a farm has multiple owners, the values of acreage provided may be multiplied by the ownership fraction. Additionally, the definition of "county" used in the spreadsheet is the county where the farmer submitted his or her application rather than the county where the farm is located. These actions may help protect confidentiality.

[12]The panel was told that RMA applies a confidentiality analysis when preparing its county-level production estimates.

farm, any estimate for the county is an estimate for the farm. Officially, suppression would be needed only if the estimate were based on data provided to the agency under a pledge of confidentiality and the farm did not provide a waiver. If the estimate were based entirely on nonconfidential data sources, such as publicly available data, it could be published. If the farm provided the data under a pledge of confidentiality, could the agency then publish an estimate derived entirely from other sources? Officially the answer would be yes, but how would the farmer view this publication?

The experience of other agencies in applying confidentiality protection to modeled results is instructive. The SAIPE program, for example, produces model-based estimates for the counts and percentages of school-age children in poverty in school districts and counties based on data from the ACS, aggregated income tax records, and other administrative data. Although estimates are sent to the Census Bureau's Disclosure Review Board every year for a disclosure assessment prior to release, no estimates have been suppressed because of confidentiality concerns. According to the Disclosure Review Board, "given that this data product is fully synthetic, the synthesis itself is the disclosure avoidance protection technique. If this product reported information from the underlying microdata (e.g., counts, or regression estimates) without going through the synthetic model(s), then we would need to check that each cell passes a cell count threshold."

The SAIPE example relates to population statistics (about people) instead of farm-level statistics (about farms that may be of widely differing sizes). However, the concept of "suitably synthetic" may be as useful as "publicly available" in assessing the need for suppression.

Another instructive example is the CES conducted by BLS, which uses a confidentiality rule in conjunction with a threshold to determine whether model-based estimates can be published. If the confidentiality rule fails (so the direct survey estimate cannot be published) and sample coverage is less than a given threshold, the model-based estimate can be published. Cells that fail the confidentiality rule and have sample coverage above the threshold are not published unless waivers (letters of consent) are obtained from the largest respondent(s).

NASS will have to determine whether estimates that are derived indirectly from respondent-level data—such as estimates based on remote sensing—are subject to confidentiality rules. For example, even though FSA data are used in classifying pixels, and June Area Survey data are used in the Battese-Fuller regression estimate for planted acres, remote sensing estimates of planted acres by crop for a county would difficult to link either to a respondent to the June Area Survey or the FSA administrative data. The current remote sensing yield estimates rely on having a time series (6 years) of previously published county yield estimates. If a county estimate were suppressed in any of the previous 6 years, a remote sensing estimate would not be available for that county. Hence, one might conclude that these remote sensing estimates are suitably synthetic and are not a confidentiality risk.

Documentation

To its credit, NASS encourage both staff and contractors to publish papers documenting proposed models and how well they work. It also conducts occasional staff-level evaluations of competing models that are presented at conferences and published (e.g., Adrian, 2012; Bellow, 2007; Crouse, 2000; Cruze, 2015b; Williams, 2013). Panel members have been part of expert groups invited by NASS to review aspects of models during their development. This is a valuable approach to gaining expert insight. However, the fact that expert reviews have taken place does not appear to be documented by NASS. At present, there is little official NASS documentation for models, even though some model-based estimates are being used as indicators by ASB. At present, NASS appears to rely on published proceedings and journal papers for documentation.

In its evaluation of the SAIPE model, National Research Council (2000a, p. 167) makes the following statement:

> An integral part of the evaluation effort is the preparation of detailed documentation of the modeling procedures and evaluation results. No small-area estimates should be published without full documentation. Such documentation is needed for analysts both inside and outside the Census Bureau to judge the quality of the estimates and to identify areas for research and development to improve the estimates in future years.

Recommendation 2-6. NASS should prepare documentation for any indication submitted to ASB for review. Such documentation should include a description of the input data sources and the methodology used, the estimates produced, a summary of the quality of the estimates, details of assessments, and descriptions of previous and planned validations. The documentation should be easily accessible from the NASS website.

MAKING THE CASE FOR CHANGE

One of the issues NASS will have to address is the development of a decision process for the use of models. If a model performs well for most publication cells, a decision might be to rely on that model for all county estimates. An alternative approach might be a decision process with a series of steps leading to selection of specific model-based indications to serve as the estimate in different situations.

It is important to establish a procedure (consisting of steps and criteria) that could be applied to determine which county indications are suitable for use as official NASS estimates. The CV or other measure of uncertainty of alternative model-based indications and the direct sample-based indication could be used to decide which indication to use as the official estimate in specific instances.

Researchers need to be able to convince others within NASS, as well as users outside of NASS, about the value of model-based approaches. Researchers need to expect to prepare convincing evidence concerning the quality of model-based versus design-based estimates, especially if model-based results are to be used in lieu of design-based estimates or if survey estimates are not publishable.

At a minimum, the mean difference, mean absolute difference, mean relative difference, and mean absolute relative difference and assessments of outlying observations among model-based county indications, ASB county-level estimates, and other county indications need to be prepared and summarized in meaningful ways. It is important that preparation and use of such statistics be a regular part of ASB's quality assurance efforts.

Encouraging collaboration between ASB members and NASS modelers will be a key way to facilitate communication between the two groups. Modelers need to understand how ASB reviews and assesses indications from multiple data sources. This knowledge may help them identify convincing ways to present model-based estimates to ASB, and may also help them determine different ways in which the various data inputs can be used in the models.

Recommendation 2-7. A representative of NASS modelers should be a full member of ASB. Additionally, modelers should be able to observe ASB deliberations, especially when model-based results are part of the deliberations.

GEOSPATIALLY REFERENCED DATABASE: SELECTING BASIC SPATIAL UNITS

In the panel's vision for the future of NASS, the list frame will provide a high-quality georeferenced farm-level database. The first challenge in developing such a database is to select meaningful and appropriate spatial units of analysis that are also feasible for NASS to obtain and use. These spatial units will be the building blocks for the analysis of location-specific attributes.

Problems arise whenever a major project must contend with datasets that use different spatial units. For NASS, for example, remote sensing data are based on pixels that vary in size and position depending on the particular satellite and sensor, while aerial photography uses its own set of pixels, and survey data often are collected for irregularly shaped basic units such as farms. This problem of "spatial mismatch" can be addressed with GIS tools, but this approach introduces an inevitable level of uncertainty. These problems of spatial mismatch and uncertainty could be reduced for NASS if a basic spatial unit could be adopted, and all data could be expressed as attributes of one or more basic spatial units. This is the strategy adopted by the Census Bureau, where the block is the basic unit, and all other units, such as block groups and census tracts, are aggregations of blocks. The obvious candidate for NASS is FSA's CLU, aggregated into farms, counties, districts, and so on.

Challenges and Advantages

Use of this strategy entails four immediate challenges. First, CLUs are not aggregates of the pixels used in satellite remote sensing, and while the errors introduced by estimating CLU properties from pixels are small when pixels are much smaller than CLUs, the errors become substantial for such sources as the Moderate Resolution Imaging Spectroradiometer (MODIS), with its 250 m pixels. Second, there are substantial gaps between CLUs (roads, water, nonagricultural areas). The Census Bureau deals with this issue by creating pseudoblocks outside settled areas. In NASS's case, it would be simpler to ignore non-CLU areas entirely. Third, boundary files exist for counties that cross CLUs. This issue could be addressed if USDA adopted a practice of snapping CLU boundaries to county boundaries, and functions exist in GIS to do this automatically. Finally, for CLUs to function as basic spatial units, it would be desirable to regard them as permanent to the extent possible, freezing them and minimizing any changes over time.[13]

While challenges are associated with using the CLU as a basic spatial unit, the approach offers advantages as well: CLUs were intended to serve as a standardized GIS data layer [14] that

[13]Although CLUs are relatively stable, they can change, primarily as a result of FSA farm reconstitutions—changes that are due to combining or dividing tracts or farms, usually because of changes in ownership or operator. (Neill, 2011).

[14]See http://www.esri.com/news/arcuser/0402/usda.html [July 12, 2017].

would ultimately include all farm fields, rangeland, and pastureland in the United States; they have unique identification numbers;[15] they are updated annually by FSA;[16,17] and they undergo a certification process.[18] They also are widely used within USDA and are readily available to NASS. As discussed below, NASS has not used CLUs as the basis for its list frame because of differences in farm definitions among NASS, FSA, and RMA.

> **Recommendation 2-8. NASS should adopt FSA's Common Land Unit (CLU) as its basic spatial unit.**

RMA reporting by field location (CLU) began in 2010, and in 2016 RMA started requiring that all farms report field location on its acreage report (U.S. Department of Agriculture, 2016). Bulletin MGR-16-005 states,

> As a service to their insureds, many AIPs [Approved Insurance Providers] have adopted technologies that facilitate the identification of field location. To better accommodate these new technologies, RMA developed the resource land unit (RLU) data standard in consultation with AIPs. Reporting by RLU enables field location reporting without specifically reporting a FSA Farm, Tract, and Field Number. In reinsurance year 2016, RLUs were authorized as a means to identify field location where CLUs did not exist.

Beginning in 2017, "the insured can meet the acreage reporting requirement for field location identification by:

- If known and accurate, providing the FSA Farm, Tract, and Field Number, including the FSA administrative State and County; or otherwise
- Clearly identifying the field location(s) and associated boundaries using AIP map-based reporting or other mapping resources which clearly delineate the fields' location."

While it is the panel's understanding that most farms report by CLU, they can report by RLU. Reporting by RLU is being encouraged by AIPs[19] and is likely to increase in the future.

Although the panel is recommending that CLUs be used as the NASS basic spatial unit, NASS should plan to maintain RLUs and, when they become available, precision agriculture reports of field location, in its databases. (See Chapter 3 for additional discussion of precision agriculture.) These data will be particularly useful for identifying GIS field locations for all farms.

---

[15]See https://www.fsa.usda.gov/programs-and-services/aerial-photography/imagery-products/common-land-unit-clu/index [July 12, 2017].

[16]See https://www.fsa.usda.gov/Internet/FSA_Notice/cm_787.pdf [July 12, 2017].

[17]See https://www.fsa.usda.gov/Assets/USDA-FSA-Public/usdafiles/APFO/support-documents/pdfs/clu_infosheet_2016.pdf [July 12, 2017].

[18]There are two different categories of CLUs—certified and noncertified. Noncertified CLUs are the "original" CLUs generated by FSA (usually in the county FSA offices). Certified CLUs are those that have undergone a process whereby FSA confirms the acreages and boundaries with owners and operators.

[19]See http://www.proag.com/News/Real-Time-Revolutionizes-USDAs-Crop-Insurance-Verification-Program-2017-04-03/5851 [May 8, 2017].

Recommendation 2-9. NASS should be prepared to maintain alternative geospatial field-level boundaries (e.g., resource land units [RLUs] and precision agriculture measurements) in its databases to facilitate completing the geospatially referenced farm-level database.

There are many advantages to incorporating CLUs into NASS list frames. First, CLUs/RLUs will provide unique linkages among FSA, RMA, and NASS reported data, providing valuable information on current-year planted acres by crop that would improve NASS surveys (the panel's ideas are discussed at the end of this chapter). Second, FSA and RMA update the CLUs/RLUs annually, providing information about changes in farm ownership and operation that NASS can use to keep its list frame up to date. Third, having a georeferenced database means that NASS farm-level data can be linked to such parcel-specific independent variables as weather, soil quality, remote sensing information (NDVI, ground temperature), and distance to markets that are potentially very useful for farm- or unit-level models.

The following subsection describes NASS frames and the panel's assessment of the impact of the geospatial building blocks on those frames. It also describes the different definitions of "farm" used by NASS, FSA, and RMA and the impact of those definitional issues on the goal of incorporating CLUs into the NASS list frame.

IMPACT ON NASS FRAMES

NASS makes use of two sampling frames—an area frame and a list frame.[20] The area frame is defined as the entire land mass of the United States and ensures complete coverage of the U.S. farm population. The list frame is a roster of known farmers and ranchers and includes a profile of each operation indicating its size and the commodities it historically has produced. The main strengths of the area frame are its completeness and stability. Its weaknesses are its inefficiency for crops grown in small acreages or geographic regions and its cost to build and collect data. The list frame can be sampled more efficiently (at the commodity-specific level, if necessary), and data can be collected using less expensive methods (mail, telephone, web). The list frame does not, however, provide complete coverage of all farms and requires continual updating (as do most business registers) since farming arrangements are constantly changing. NASS uses its area frame to adjust for undercoverage of the list frame. For the 2012 Census of Agriculture, the adjustment for frame undercoverage for the number of farms was 12.3 percent, and the coverage adjustment for land was 3.4 percent.[21]

County-level indications (prior to benchmarking) for both crops and cash rents are based on probability sample survey data collected using samples from the list frame. Benchmarking of

[20]This description was drawn from information provided by NASS. See, for example, the description of the use of sampling frames in https://www.nass.usda.gov/Education_and_Outreach/Understanding_Statistics/Yield_Forecasting_Program.pdf [August 17, 2017].

[21]Appendix A of the NASS Census of Agriculture publication provides measures of census undercoverage, nonresponse, and classification error. See https://www.agcensus.usda.gov/Publications/2012/Full_Report/Volume_1,_Chapter_1_US/usappxa.pdf [March 2017]. For the following major crops, it shows frame undercoverage to be as follows: corn for grain = 4.0 percent of farms and 1.6 percent of acreage; winter wheat = 3.4 percent of farms and 1.8 percent of acreage; soybeans for beans = 3.8 percent of farms and 1.8 percent of acreage; cotton = 3.1 percent of farms and 1.3 percent of acreage; etc.

county-level estimates to agree with previously published state totals accounts for list frame undercoverage, and is done by using an identical multiplicative factor applied to each county.[1]

The success of NASS surveys and the Census of Agriculture depends on the quality of the sampling frames (area and list). The list frame contains basic data: name, contact information, location, and demographic information. For updating of the list frame, NASS uses annual lists of participating farms from FSA and RMA, outside source lists, and web scraping. It also has approval to use federal tax information at the record level. However, if a farm is identified from the tax information, it cannot be added to the frame until it has also been identified in another source and that other source information is used on the list frame. Maintaining the list frame is a major ongoing effort for NASS regional offices. Frame maintenance and updating is coordinated in the St. Louis office. The National Agriculture Classification Survey is used to help ensure that the list frame is as complete as possible for the Census of Agriculture.

NASS told the panel during its October 2016 meeting that, in researching issues associated with declining response rates, it had established that a fair amount of nonresponse is due to “deadwood”—farmers that are no longer in business—on its list frame and in surveys. In an initial check by field offices, NASS found that 40-50% of farm identified through an algorithm as deadwood, were actually no longer in business. NASS plans to expand its verification effort. NASS also told the panel that it has decided to link FSA CLUs to the list frame to help maintain the status of farms. The panel enthusiastically supports this decision and views this linkage as key to achieving the vision of the NASS list frame being a premier farm parcel database, serving as a sampling frame for NASS surveys and as a valuable tool for analysis.

Recommendation 2-10. NASS should update its list frame to make use of CLU information as geospatial building blocks by developing linkages between the NASS list frame and FSA and RMA administrative data.

The area frame is used to assess the completeness of the list frame. If farms not on the list frame are found, they are not added to the list frame because doing so would negate the independence of the list and area frames, thereby violating a key assumption underlying the dual-frame survey estimation and census capture–recapture adjustment for undercoverage.

The list frame is kept in a relational database developed in the mid-1990s called the Enhanced List Management Operations (ELMO) database. NASS has chosen to characterize farms by “operation” and by “people” (owners, operators, others). As a result, ELMO has three key tables: a person table (PID); an operations table (OIDS); and a person/operations table (POIDS), linking each person with all the operations with which he or she is involved. An individual person may be involved in multiple operations, and multiple people may be involved in each operation. The PID uniquely identifies people involved in farming, particularly owners and operators, and includes name, contact information, and demographics. The OIDS uniquely identifies an operation, and includes business name, physical address, location (county), and identification number. The POIDS includes identification number, current operating status, and whether the unit is a target for operator-dominant surveys such as the County Agricultural Production Survey (CAPS). For operator-dominant surveys, sampling units are defined by the “target name” for the farm or ranch, usually a person’s name. The target name reports for all of the operations in which he or she is involved.
“target name” for the farm or ranch, usually a person’s name. The target name reports for all of the operations in which he or she is involved.

[1]This process accounts for undercoverage only if certain assumptions hold.

The Census of Agriculture and some minor, commodity-specific surveys use a frame based on the operation-dominant list (also a link between the PID and OIDS, but with operation dominant). For the census and these surveys, a selected operation reports only for that operation.

Farms in the United States are complex and are defined by different USDA agencies in different ways (see Box-2-4 for the definitions of "farm" used by NASS, FSA, and RMA). NASS and USDA's Economic Research Service (ERS) have jointly established a National Academies of Sciences, Engineering, and Medicine Panel on Improving Data Collection and Reporting about Agriculture with Increasingly Complex Farm Business Structures to evaluate complex farm structures. This panel recently started its work, and its report is expected in 2018.

BOX 2-4 What is a "Farm?"

National Agricultural Statistics Service (NASS) Definition

- **Official Legal Definition:** *A* farm *is any place from which $1,000 or more of agricultural products were produced and sold, or normally would have been sold, during the census year.*
- Farms are maintained on NASS's list frame as "operations," with one or more individuals (operators) associated with each operation.
- One operator is designated as the "target operator" and is the sampling unit for many (but not all) NASS surveys. The operation is the sampling unit for the Census of Agriculture.
- Operations and operators self-identify; there are no relevant statutes, regulations, or requirements.

Farm Services Agency (FSA) Definition

- **Definition of "Farm":** *A* farm *is made up of tracts that have the same owner and the same operator. Land with different owners may be combined if all the land is operated by 1 producer with all of the following elements in common and substantially separate from that of any other tracts: (1) Labor, (2) Equipment, (3) Accounting system, and (4) Management. Note: Land on which other producers provide their own labor and equipment, but do not meet the definition of an operator, shall not be considered a separate farm.* (FSA 3-CM Handbook)
- **Definition of "Field":** *A* field *is part of a farm which is separated from the balance of the farm by: (1) permanent boundaries, such as fences, (2) permanent waterways or woodlands, (3) croplines in cases where farming practices make it probable that such cropline is not subject to change, or (4) other similar features.* (7 CFR 718.2)
- These definitions of "field" are maintained as geospatial units that are called Common Land Units (CLUs).

Risk Management Agency (RMA) Definition

RMA does not define "farm." It collects information from entities that purchase sponsored insurance from Approved Insurance Providers. These entities submit acreage data similar to FSA data, and beginning in 2017 also report by CLU (or in some cases by Resource Land Unit [RLU]).

SOURCE: Prepared by the panel.

These differences in the definition of "farm" make linkage of the data reported to NASS, FSA, and RMA more challenging. For example, NASS told the panel that a project in support of the 2012 Census of Agriculture to link the FSA payments database and the NASS list frame resulted in only a 63 percent match rate with 6 percent possible matches. At the panel's second meeting, NASS representatives described an effort undertaken in 2010 to match FSA farms to NASS farms in Nebraska (see Barboza and Iwig, 2010). This effort was quite time-consuming, and the initial match rate achieved was 80 percent, with a 45 percent one-to-one farm match rate. After review, the match rate rose to 97 percent, with a 53 percent one-to-one match rate.

Difficulties in matching the NASS list frame with the FSA list are due to the complexity of farm operations and the fact that NASS and FSA farms are not necessarily the same. NASS identifies an "operator" willing to report for as much farmland as possible, but it is not clear to the panel how NASS ensures that the selected person knows and reports for exactly the farm operations NASS expects. The available geospatial information on farm operations should facilitate such assessments. The NASS study in Nebraska showed that there were 2.4 FSA farms for every NASS farm. As NASS well knows, achieving this vision will not be easy; however, it will be extremely valuable. Once it has been achieved, NASS will have gained the ability to keep a high proportion of the frame (better than 95 percent of planted acres in the United States) up to date by regularly checking matched FSA (and ultimately RMA) farm records for changes; will have detail for matched farms concerning planted and failed acres by crop for the current year and be able to use these data to fine tune its surveys to collect only the missing information; and will be able to develop unit-level model-based approaches that can take advantage of such geospatial information as soil productivity and weather.

NASS has said that it has little experience with linking to RMA data. However, the panel was told about the Acreage and Crop Reporting Streamlining Initiative (ACRSI), a multiagency project within USDA. FSA and RMA have been the key players in this project to date, and have developed a common reporting system for acreage reports intended to allow farmers who participate in both programs to complete one acreage report for use by both programs. While the ACRSI project is still evolving, it has made great strides in achieving a common automated system whereby producers can submit their FSA/RMA acreage reports through a common site. Data reported to FSA are available electronically to AIPs, and data reported to RMA by AIPs are available electronically to FSA. FSA, RMA, the Natural Resources Conversation Service (NRCS), and NASS will all be able to use the reported information for their respective agency programs. For example, FSA would use the information for program purposes, RMA would use it for crop insurance purposes if the producer purchased crop insurance, and NASS would use it for statistical purposes.

Achieving complete linkage with these administrative sources will be time-consuming, and NASS needs to plan an approach that can build to success. The first step will be planning how to accommodate the CLU data and FSA and RMA identifiers (IDs) in ELMO (or in files easily accessed by ELMO). Second, NASS may have information on matched FSA and NASS farms from previous work. For several years, for example, the June Area Survey has made use of FSA CLUs, and remnants of the 2010 record linkage effort in Nebraska may be available. Third, NASS's traditional matching approaches using owner/operator names can provide a link for the least complex farms—those for which there is one NASS farm for one FSA farm. The Nebraska matching experiment indicated that there was an "easy" one-to-one exact match for almost 50 percent of farms. NASS needs to develop revised survey forms for matched farms that take advantage of administrative data to simplify reporting and reduce respondent burden. With the promise of a simple reporting form, perhaps NASS can work with all respondents to encourage them to provide the FSA IDs with which they are associated. Completing the match through manual reviews would be possible but time-consuming. Over time, the geographic information available for FSA and RMA farms should make it possible to track down the NASS farms to identify matches. Expensive manual efforts at matching would best be dedicated to achieving matches for the largest farms.

While conducting this linkage may be complex, maintaining the linked database should be relatively simple. FSA and RMA update their data annually. The panel heard that FSA makes some changes in CLU definitions over time, possibly because of changes in ownership or land use. At its second meeting, the panel heard from ERS analysts who have linked FSA and other USDA data and maintained that linkage over time. They estimated that no more than 10 percent of CLUs are changed from one year to the next.

NASS SURVEYS IN THE FUTURE

Well-executed NASS surveys have been the major tool for supplying useful information about crop production in the United States for decades. The new sources of information examined in this report will not replace the need for the surveys or survey infrastructure NASS has developed any time soon, and they will continue to be as important as ever for producing estimates. For the new data sources to have maximum utility, however, some changes to the survey machinery will be necessary. For example, methods will need to be developed for associating the new information with the units on the sampling frame. In this way, the auxiliary information can be used to improve either the sample design or the estimation process or both. This section outlines a strategy for how this might be accomplished. It includes subsections on using Internet survey forms, imputing for nonresponse, and improving estimation using alternative estimation methods.

Once this has been accomplished, alternative sources of information can provide a variety of potential improvements to survey methods, several of which are discussed in this section (see Chapter 3 for detailed discussion). For example, allowing alternative sources of data to substitute for operator-supplied data can reduce the burden on operators, potentially increasing response rates or even improving accuracy. These changes will require a long planning horizon, experimentation, and testing. The available sources of data are likely to continue to change as new technologies develop. Therefore, a system that builds in continuous testing and change will likely be required to take advantage of the new technologies as they appear.

The NASS surveys that feed into the county-level crop and cash rents estimates are described in Appendix A. These surveys have been providing useful county-level estimates for crops since 2012 and for cash rents since 2009. As noted previously, however, declining response rates have resulted in a decline in the number of counties for which estimates can be published. NASS has initiated efforts to evaluate and address issues related to nonresponse. The panel expects that as NASS moves toward achieving the vision set forth in this chapter, it will continue to rely on high-quality surveys but will need to become more nimble, regularly assessing the quality of auxiliary data by county and ultimately focusing surveys in those areas in which auxiliary data are not sufficient.

Internet Survey Form to Facilitate Reporting by FSA/RMA Linked Farms

As more of the FSA and RMA data are linked to the list frame, it may be useful for NASS to develop an easy-to-use Internet survey form to facilitate reporting by farms that also participate in FSA and RMA programs. The administrative data identify planted acres, crops, and some failed acres, but do not provide information on yield. Since the planted acres by crop would be known from FSA/RMA data, the Internet form would collect only data not available from administrative sources. The availability of such a form might encourage farmers to identify their FSA numbers to NASS so they could make use of the streamlined form.

Recommendation 2-11. NASS should develop a user-friendly Internet form to be completed by FSA and RMA participants matched to the NASS list frame in lieu of completing the County Agricultural Production Survey (CAPS)/Acreage, Production, and Stocks (APS) surveys. NASS could use the availability of such a form as an incentive for farmers to help in providing NASS/FSA linkage information.

Imputation

The demand for small-area estimates, resource constraints, and the availability of high-quality auxiliary data indicate the potential benefit of incorporating FSA/RMA farm-level data on planted acres, when available, into NASS's imputation methodology for CAPS and APS. The panel received a copy of a 2006 internal paper documenting a NASS experiment that involved using data on FSA planted acres by crop for imputation. The experiment demonstrated that this approach to imputation was not feasible at the time (because of difficulty in matching farms and disagreement between FSA and NASS data). However, this was before Barboza and Iwig (2010) had documented the differences between FSA and NASS farms. Once a reasonable number of FSA farms and their CLUs have been linked with the NASS list frame, using current-year information from FSA on planted acres by crop for imputation will be feasible.[23] Over time as more FSA farms are linked, NASS may realize significant improvements in accuracy due to the inclusion of this current-year information. The panel was told that cash rents, especially rented acres, are fairly stable over time, which may indicate that a simple model of imputation for rented acres based on historical data could be beneficial.

As precision agriculture data become available, they also may prove valuable in imputation for nonresponse. The current sample-based data collection and estimation method could be used, but nonresponse on items on planted and harvested acreage could be imputed with data from the actual nonresponding farm or from a "similar" farm. Either of these approaches might require permission to access the data from a unit that is not a responder to the survey, depending on how the data were obtained. If it were possible to obtain access to such data on an opt-out basis, this might be a way to address this issue. If only opt-in permission is possible, it may be unrealistic to expect that a nonresponding unit would be willing to extend permission.

[23]The panel was told that crop switching causes considerable concern during imputation. Imputation relies on past-year data for a missing farm to select crops for imputation. However, past-year data may not be related to the crops that the farmer plants in the current year.

Recommendation 2-12. As farm-level administrative data and precision agriculture data become available, NASS should consider imputing for nonresponse based on these auxiliary sources of information.

Improving Estimation

Both FSA and RMA data define subpopulations of all farms, and NASS could potentially obtain access to relevant data from precision agriculture databases on an ongoing basis for a substantial fraction of the total farm population. These auxiliary data for well-defined subpopulations could be used to provide more efficient estimators relative to the current NASS sampling and estimation methods.

Data from precision agriculture databases, FSA program participants, and RMA purchasers of insurance clearly are not probability samples and are not likely to be representative of the total farm population. Thus the data could not be treated as a direct replacement for or even an augmentation of the probability sample. A topic of current research interest in survey methods is how to make use of data from nonprobability samples (see Baker et al., 2017). The availability of data through new technologies and social media has inspired this research, and it is important for NASS to stay abreast of these developments. Still, there are methods for employing data from nonprobability samples that can be used in ways that are familiar to those accustomed to probability sampling methodology. All of these methods require that it be possible to link the data from the nonprobability samples to the list frame. These methods make use of the data either to improve the sample design or to serve as auxiliary information to improve estimation derived from a probability sample. Two such methods are described below.

First, suppose that all of the information normally collected by one of NASS's surveys is available for a subset of the population from the precision agriculture database, and it can be matched to the NASS list frame. (The same methodology holds for FSA and RMA participants linked to the list frame.) Then the frame can be divided into two strata distinguished by whether the needed data are available—stratum *N* (not available) and N' (available)—and the fraction of units for which the data are not available is denoted by p_N. Then the resources for sampling can all be directed to the stratum with no available data, and the estimator of the total is just the sum of the known total from stratum N' and the estimated total from stratum *N*. If the design effects for the new and original designs are similar and the finite population sampling correction (fpc) is negligible, then the relative precision for this estimator is

$$RP\left(\hat{t}_{y,new}, \hat{t}_{y,orig}\right) = \frac{Var(\hat{t}_{y,new})}{Var(\hat{t}_{y,orig})}$$
$$= \frac{{P_N}^2 {S_{yN}}^2}{{S_y}^2}$$

where S_y^2 and S_{yN}^2 denote the variance of the acreage in the entire population and the variance of the acreage of the part of the population in stratum *N*. This relative precision can be substantial if the coverage of the precision agriculture database is large.

Now suppose that the variable available from the auxiliary database is not the same as that needed for estimation, but is correlated with it. So, for example, suppose that only acreage aggregated over crops is available rather than acreage for the specific crop of interest. These data can still be used in estimation, and will provide increased precision under certain circumstances. Denote the desired specific acreage from farm i as y_i. An indicator r_i denotes whether (= 1) or not (= 0) auxiliary data are available for farm i, and x_i denotes the farm's available (e.g., aggregated acreage) value. Then a ratio estimator can be constructed that uses as auxiliary data the variable $r_i x_i$. That is, the estimator is defined as

$$\hat{t}_{y,ratio} = t_{rx} \frac{\hat{t}_y}{\hat{t}_{rx}},$$

where $t_{rx} = \sum_{i=1}^{N} r_i x_i$ is the total of the auxiliary data over all farms for which these data are available, and $\hat{t}_y$ and $\hat{t}_{rx}$ are the estimated totals of the specific crop acreage and the available auxiliary data for the population making up the probability sample. If the proportion of the population providing auxiliary data ($(1 - p_N)$) and the correlation between x and y are high enough, the variance of $\hat{t}_{y,ratio}$ can be substantially lower than that of $\hat{t}_y$ (Liu et al., 2017). If the data available from the auxiliary sources are considered to be equivalent to *yi* (or even of higher quality), but the matching of frame to database may contain errors, $\hat{t}_{y,ratio}$ may still perform adequately. With the auxiliary data being treated as supplemental to the sample, the variance, rather than the bias, is increased when errors occur in either matching or the accuracy of x_i as a proxy for y_i.

Recommendation 2-13. Once alternative data have been linked to the NASS list frame for a sufficiently large percentage of farms, alternative estimation methods that make use of the linked data should be evaluated.

3

Multiple Data Sources for Crops: Challenges and Opportunities

The NASS Agricultural Statistics Board (ASB) considers multiple data sources when developing its county-level crop estimates. The NASS county estimates program produces annual estimates for planted acres, harvested acres, production, and yield (production divided by harvested acres). The main source of data used to develop these estimates is the County Agricultural Production Survey (CAPS)/Acreage, Production, and Stocks (APS) survey system,[1] which provides information for all required estimates.

These survey-based estimates are considered highly reliable because they are developed using appropriate statistical procedures (Johansson et al., 2017). Over time, however, NASS has been able to publish estimates for fewer counties because of falling response rates. NASS is fortunate to have high-quality auxiliary data available to inform these estimates and is challenged in determining how best to use these data in a manner that is transparent and reproducible. As discussed in Chapter 2, these auxiliary data sources include recorded data, such as administrative data from the Farm Services Agency (FSA), the Risk Management Agency (RMA), and potentially precision agriculture measurements, as well as data observed using remote sensing tools. This chapter describes current and emerging alternative data sources potentially available to NASS. It explains how they are currently used by NASS, if they are so used, and presents the panel's recommendations for enhancements. The focus is on how NASS can achieve the goal of saving resources and reduce the response burden on farmers while improving its estimates.

The chapter begins by examining the challenges and opportunities with respect to NASS's use of administrative data from FSA and RMA. It then turns to the potential use of a new and potentially very important source of recorded data to which NASS currently lacks access—precision agriculture measurements—and highlights the benefits and challenges of incorporating this information into NASS's county-level estimation process. This is followed by a discussion of data observed by satellite and other remote sensing tools. Satellite remote sensing data are available through government agencies or private vendors; other sources of remote sensing data include data collected by drones and aircraft (available through private vendors) or flux towers in fields (available through the Agricultural Research Service [ARS]). These data also are not currently used by NASS.

Until NASS has high-quality auxiliary data on production or yield, it will need to rely on its surveys, using auxiliary data to improve certain estimates for certain crops/regions. The panel sees two possibilities for achieving the goal of high-quality auxiliary data on production or yield. First, if RMA moves the date for collecting production data forward, as it is currently considering, NASS may be able to use these data directly to improve its estimates in some areas.

[1]See Appendix A for a description of NASS surveys.

Second, NASS has developed its satellite remote sensing program on a shoestring budget. If NASS can improve its satellite remote sensing estimates of yield and expand the number of commodities covered, it will have many options available for achieving this goal.

ADMINISTRATIVE DATA FROM THE FARM SERVICES AGENCY (FSA) AND RISK MANAGEMENT AGENCY (RMA)

FSA Administrative Data

FSA manages a number of programs to benefit farmers, using NASS data in program administration and collecting important administrative data. NASS county-level yield estimates are used in the Agricultural Risk Coverage program, NASS county-level production estimates are used in the Marketing Loan Benefits program, and NASS county-level cash rents estimates are used in the Conservation Reserve Program, as described in Chapter 1.

To apply for participation in one of these FSA programs, producers must file an annual application (form FSA-578) that certifies (i.e., “signs up”) their acreage for participation. The FSA-578 is due to FSA by July 15 for many crops. Updates to failed acres are provided within 15 days of the failure. The FSA data are reported by Common Land Unit (CLU), roughly a field, as described in Chapter 2. Generally, the information in which NASS is interested on the FSA-578 form includes the specific crop(s) (by type) grown, the number of acres planted in that crop, the FSA farm number (and sometimes tract and field numbers), the intended usage of the crop, and whether irrigation was available. FSA provides county-level aggregates of these data by both physical acres (where the acres are located) and administrative acres (the county that administers FSA programs). FSA also developed and maintains CLU[2] information for each FSA-578 farm. In 2017, FSA data covered about 95 percent of planted acres in the United States for major commodities.

RMA Administrative Data

RMA and the Federal Crop Insurance Corporation (FCIC) (a wholly owned government corporation that administers the federal crop insurance program) design insurance products, develop standards, and establish premium rates for federal crop insurance. All crop insurance policies are sold to farmers by private crop insurance agents, or Approved Insurance Providers (AIPs),[3] who are paid by FCIC to market and service policies.

Michael Alston of RMA told the panel that in 2015, the two largest types of FCIC insurance were revenue policies (75%) and actual production history policies (19%). Revenue policies protect producers against loss of revenue due to price fluctuations and yield loss due to natural causes. Actual production history and yield protection policies insure producers against yield losses due to natural causes. The remaining 6 percent of policies were based on county results rather than individual farms and used adjusted gross revenue, with a fixed dollar amount of insurance. RMA uses county-level production history and county-level estimates of yield from

[2]See https://www.fsa.usda.gov/Assets/USDA-FSA-Public/usdafiles/APFO/support-documents/pdfs/clu_infosheet_2017_Final.pdf [June 2017].

[3]A list of AIPs in 2017 can be found at https://www3.rma.usda.gov/tools/agents/companies/indexCI.cfm [July 12, 2017].

NASS. In 2017, RMA data covered about 85 percent of planted acres of major commodities in the United States.

Farm product insurance agents sell the same federal insurance products at rates that depend on a producer's productivity and location. For a given producer, premium rates are the same regardless of the insurance provider. Insurance providers compete by providing customer service, part of which includes preparing the data for reporting to RMA. (The data items are similar to what is reported to FSA, as well as to NASS in the CAPS.) In January 2017, David Zanoni from RMA told the panel that AIPs have various systems for collecting these data, including precision agriculture measurements. Examples include subcontracting with a software vendor to process machine files straight off the tractor; building their own software to do the same; using custom farm management information systems that allow producers or agents to "draw" their acreage through geographic information systems (GIS) and attach the necessary crop data on mobile software (i.e., an agent uses an iPad in the field to draw the acreage report); and agents buying commercial GPS systems and mapping the fields themselves as a service to the producers. RMA specifies the data that must be reported with insurance applications, but does not specify how the data should be collected from (or for) farmers. RMA has developed a web service application that AIPs can use to report boundary information. It is based on arcGIS software using a commonly accepted agriculture data format, GeoJSON, for reporting by resource land unit (RLU).[4]

Acreage data are reported to insurance agents/companies for many crops by July 15. Production data are reported to insurance agents/companies by about the end of May following the planting year. There is a lag in the reporting of the data by insurance agents/companies to RMA.

NASS Use of FSA and RMA Data

Definitions of planted acres, failed acres, and production are similar enough among NASS, FSA, and RMA that data from one agency can be used to compute the data for another. CLUs provide geospatial identifiers for most planted acres in the United States, with attribute data available from the FSA-578 form. The Acreage and Crop Reporting Streamlining Initiative (ACRSI) has resulted in greater consistency between RMA and FSA in lists of commodities and reporting dates. It also has resulted in the use of a common reporting system for the collection of FSA and RMA data and RMA's adoption of CLU/RLU as its approach for identifying the geospatial location of farm fields.

The FSA data account for all farms that participate in FSA programs but not the universe of farms. FSA collects data needed for classifying crop use (grain and seed, forage, or by cropping practice) early in the year and does not require farmers to report the final usage of the crops they planted, while NASS's data are intended to reflect actual crop use. Also, if irrigation is available for a field, FSA records the acres as irrigated regardless of whether any water was applied to the crop during the growing season. The most challenging difference between the FSA and NASS data is that FSA and NASS farms are not identical (see also Chapter 2).

ASB currently uses FSA data on planted acres as part of its process for estimating planted acres in two ways. First, FSA data on county-level planted acres serve as a minimum level for

[4]See https://www.rma.usda.gov/ftp/Publications/M13_Handbook/2017/approved/RES_LAND_UNIT_SUMM.PDF [June 2017].

NASS planted acres because they represent data for part of the population (a very large part in some areas for some crops). Second, FSA data on planted acres by county benchmarked to previously published state totals are an input to the composite indicator described in Chapter 2.

This discussion illustrates that administrative data from FSA and RMA represent high-quality acreage and production data collected by these agencies during farm program administration. The main challenge with these data is that they are incomplete, as they cover only certain parts of the required NASS estimates and not for all crops, and they are not timely. For example, FSA collects data on planted and failed acres by crop from farmers that participate in its programs (about 95% of total planted acres, with coverage varying by crop and region). FSA data on planted acres by crop are invaluable in preparing NASS estimates for planted acres, but the data on failed acres (potentially useful for estimating harvested acres) are viewed as being incomplete, especially at the time the NASS county-level estimates are developed. RMA collects data on planted acres by crop, failed acres, and production from farms that purchase RMA insurance products (about 85% of total planted acres, with coverage varying by crop and region). NASS uses RMA data on failed acres in estimating harvested acres. NASS views RMA data on failed acres as more complete than those reported by FSA because farmers must report failed acres to receive insurance payments. While precision agriculture measurements, discussed later in this chapter, are another excellent source of recorded data, they also are unlikely to provide complete coverage of all planted acres. The high coverage rates of FSA and RMA data make them valuable to NASS. It should be observed, however, that these data are based on farm policy, and specifically subsidies. A reduction in benefits could also cause participation to drop.

Timing is another important issue for the use of these data. Table 3-1 illustrates that while acreage data from both FSA and RMA are available in the fall as crops progress and have proven highly beneficial for producing acreage estimates (Good and Irwin, 2016), producers are not required to report their production to their insurance agents/companies until they sign up for program participation for the following year (typically the end of May). This timing makes it impossible for NASS to use these administrative data for county production and yield estimates, which typically are published in mid-February. However, NASS production and yield estimates are used for farm payments only after the marketing year is over (May for wheat, July for cotton, August for corn and soybeans). Thus it appears that changes in the timeline for data collection by RMA and/or for release of NASS's county-level estimates could greatly simplify the task of providing county estimates without sacrificing their availability for intended uses later in the year.

TABLE 3-1 Timeline for Reporting and Use of County-Level Corn and Soybean Data, September–August Marketing Year

Month	NASS Input	FSA Input	RMA Input	NASS Pub.	FSA Pub.	FSA Use
Jun	JAS[a], APS[b]					
July		FSA-578[c]				
Aug			Acreage[d]		Acreage[e]	
Sept	APR				Acreage	
Oct					Acreage	
Nov					Acreage	
Dec	CAPS[f]/APS				Acreage	
Jan					Acreage	

Feb			County estimates[g]	
Mar	APS			
Apr				
May				
Jun		Production[h]		
July				
Aug				
Sep				ARC payments[i]

[a]The June Area Survey collects data on acres planted for corn and soybeans in the current year.
[b]The quarterly Acreage, Production, and Stocks surveys capture the flow of planting through harvest for many crops.
[c]From FSA-578, annual acreage report, due June 30 for small grains, July 15 for all other crops. If planted after that date, within 15 days of planting. Change reported within 15 days of a loss.
[d]Acreage reporting due to insurance agent/company at same time as FSA. There is a lag in reporting to RMA. Generally RMA has about 80% of the acreage covered by mid-August. Sales closing is generally March 15; production reported to agent/company within 45 days after sales closing.
[e]Acreage spreadsheet published monthly with increasing completeness over the season at https://www.fsa.usda.gov/news-room/efoia/electronic-reading-room/frequently-requested-information/crop-acreage-data/index.
[f]The County Agricultural Production Survey collects data on planted acres, harvested acres, and yield or production at the end of the growing season.
[g]County estimates published by NASS.
[h]Production, the reporting deadline for producers to send production data to agents/companies, is near the end of May (generally 45 days after sales closing, usually March 15). There may be a delay in sending the data to RMA.
[i]Payments under the Agriculture Risk Coverage program.
NOTE: APS = Acreage, Production, and Stocks; ARC = Agricultural Risk Coverage; CAPS = County Agricultural Production Survey; FSA = Farm Services Agency; JAS = June Area Survey; NASS = National Agricultural Statistics Service; RMA = Risk Management Agency.

The panel understands that RMA is rethinking the reporting of its production data, and the report deadline may be moved forward. RMA officials told the panel that they doubted any new reporting data would be timely enough to enable NASS to use the data in preparing its county-level estimates absent a change in the current publication schedule. However, as discussed in Chapter 2, should RMA production data become available in an appropriate time frame, NASS could consider using these data in conjunction with its own survey data to improve production and yield estimates. RMA data should be particularly helpful to NASS in counties with high RMA coverage. Linkage of RMA data to the list frame also would enhance the utility of the RMA production data, making it feasible to use these data for imputation and estimation as described in Chapter 2. Even if RMA production data were not available in time for the release of NASS county-level estimates on the current schedule, it might be possible for NASS to provide additional estimates for county-level yield/production that previously were not publishable by NASS standards later in the year, after the RMA production data became available and before these estimates were needed by FSA. As described below in the section on near-term enhancements to NASS remote sensing indications, RMA farm-level production data, with location available from the farms' CLUs, may be useful in training or ground-truthing remote sensing estimates of yield.

> **Recommendation 3-1. NASS should collaborate with RMA to obtain relevant individually identifiable acreage and production data and to conduct comparisons with NASS data for the same entity.**

Future Enhancements to NASS Use of FSA and RMA Data

> **Recommendation 3-2. NASS should collaborate with FSA and RMA in the development of an approach for using RMA data in conjunction with NASS survey data to prepare crop estimates for counties for which NASS estimates are withheld. These estimates should be available in time to be used by FSA.**

RMA has extensive experience in yield modeling and county yield forecasting that may be useful for NASS to explore. A report by Coble and colleagues (2010) provides an extensive review and summary of the literature on crop yield modeling. The challenges highlighted by these studies include evidence of yield heteroscedasticity over time (e.g., Atwood et al., 2002, 2003; Gallagher, 1987; Harri et al., 2009; Hazell, 1984; Traxler et al., 1995; Yang et al., 1992) and a trend toward increases in mean yields due to technological advances (e.g., Goodwin and Ker, 1998; Zhu et al., 2008). Furthermore, McCarl and colleagues (2008) argue that variability in yield is nonstationary and influenced by climate. Of particular relevance to NASS's county-level yield estimation goals are the studies examining the spatial heterogeneity of yield distributions. Glauber (2004) and Babcock (2008) argue that this heterogeneity can be explained partly by varying weather/climate conditions and resource endowments (e.g., soil quality, topography). The authors demonstrate how publicly available data, such as monthly total rainfall, mean temperature, and mean Palmer Drought Index, can be used to model heterogeneity in yield. They further argue "that a significant amount of the county-to-county variation in rates may be explained by differences in soil types, elevation, slope, production systems employed and other data that should be readily collectible and either stable across time or with changes that can be documented" (Coble et al., 2010, p. 75).

To address estimation challenges entailed in modeling yield, the Statplan database, which contains historical insurance records from 1948 forward, was constructed (Coble et al., 2010, p. 15). It appears that access to this database could be very useful for NASS's yield modeling efforts. According to Coble and colleagues (2010, p. 29), "RMA uses credibility weighting to smooth rates among adjoining counties. Credibility in the RMA process is a function of net acres insured so a historical record of net acres insured is retained in the Statplan process….RMA refers to the surrounding counties as the county group or credibility complement."

As described in Coble et al. (2010, p. 38), the "RMA Type/Practice Rating Methodology Interim Underwriting Guidelines" describe the procedure RMA currently uses to aggregate counties with similar type/practice factors:

> …in deriving the TpFactors for irrigated and non-irrigated practices in the western States, grouping a smaller number of geographically clustered counties within the state is more typical since the average rainfall changes significantly over shorter distances. In contrast, the Eastern States have rainfall patterns that are

> more stable across greater distances and grouping more counties, whole states or multiple states may be appropriate. Determination of the county groupings for use in developing type/practice factors is largely left to the subjective judgment of the RMA regional offices.

A credibility procedure similar to that described above is then applied.

This approach of grouping counties with similar type/practice factors is sharply different from the NASS approach of developing estimates based on administrative boundaries, such as county, Agricultural Statistics District (ASD), and state. Figure 3-1 illustrates that differing values of key variables (in this case cash rents) do not necessarily follow ASD boundaries (the dark lines in the figure) and that key variables may be highly heterogeneous across counties. Adopting the approach of grouping counties with similar characteristics for purposes of model fitting or estimation regardless of the administrative units (ASD, state) to which they belong might result in improved estimates. Keith Coble, in his discussion with the panel in January 2016, noted that soil type is a key variable for grouping of farm-level data to estimate yield. He noted that including spatial relationships smooths outliers. Another way to account for such variability is to explicitly include spatial variation in models (see Appendix C).

NASS would benefit from exploring the impact of soil productivity on yield by defining estimation groups based on the National Commodity Crop Productivity Index (NCCPI).[5]

FIGURE 3-1 2016 Illinois cash rent paid per-acre non-irrigated cropland

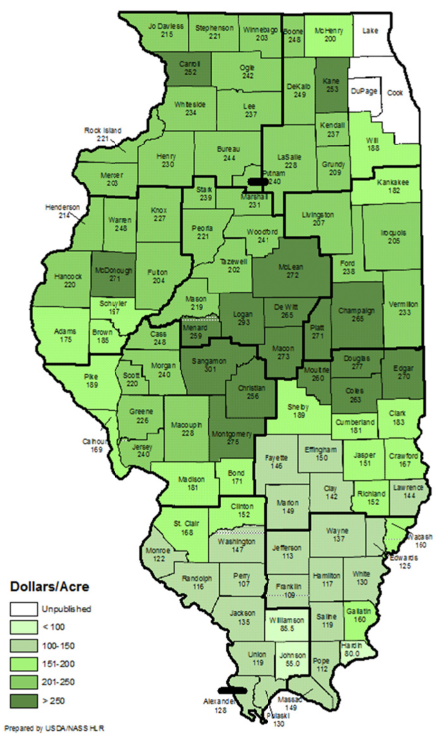

---

[5]The NCCPI is an index produced by the Natural Resource Conservation Service (NRCS), which rates land according to its potential for production based on soil and other physical characteristics. There are some commodity-specific indices, but only the general one is considered here. In other words, the same covariate applies for corn and for soybeans for any county that produces both crops. For more information about the NCCPI, see the user guide (U.S. Department of Agriculture-Natural Resource Conservation Service, 2012).

A NEW SOURCE OF RECORDED DATA: PRECISION AGRICULTURE

Precision agriculture is a site-specific crop management approach that depends on measuring field variability in crops. Precision agriculture is enabled by GPS systems, computers, and sensors built into modern combine harvesters and other equipment. Penetration of tractor guidance using GPS had grown to about 50 percent of planted acres for such crops as corn, rice, and peanuts by 2013 (Schimmelpfennig, 2016). Penetration of GPS soil and yield mapping technologies lagged behind, with adoption rates of about 25 percent. At a minimum, and relevant to NASS crop estimates, these GPS technologies can provide the GIS boundary of the field planted, area planted, and yield, all by commodity. Precision agriculture measurements are becoming highly valuable.

According to Blake Hurst (2015) of the Missouri Farm Bureau Federation in testimony to Congress in 2015:

> Most combines traveling across fields in the Midwest this fall had a GPS receiver located in the front of the cab. Although agriculture has been experimenting with this technology for a decade or so, only now is the industry starting to consider all the uses of this transformative technology….If 1,000 machines randomly spread across the Corn Belt were recording yield data on the second day of harvest, that information would be extremely valuable to traders dealing in agricultural futures. Traders have traditionally relied on private surveys and U.S. Department of Agriculture yield data. These yield estimates are neither timely nor necessarily accurate. But now, real-time yield data is available to whoever controls those databases. The company involved says it will never share the data. Farmers may want access to that data, however, and they may not be averse to selling the information to the XYZ hedge fund either, if the price is right—but that's only possible if farmers retain ownership and control of the data.

Hurst went on to describe the efforts of a number of farm and commodity organizations, agriculture service providers, and agriculture technology providers (ATPs) to collaborate around several issues. First was the development of a set of principles for data privacy and security for farm data. Second was the design of a transparency evaluator, an automated tool to assist farmers in assessing the data protection/sharing features in the "fine print" in contracts. Third was the organization of an agriculture data repository that would protect data but also make them available for specific uses approved by the farmer.

At its January 2016 meeting, the panel heard from Mary Kay Thatcher, American Farm Bureau, who provided an update on these initiatives. She reported that additional organizations have signed on to these collaborative efforts, agreeing with the principles developed. The transparency evaluator has been completed and allows farmers to compare and contrast specific issues within contracts. It provides a seal ATPs can use that demonstrates their adherence to the data principles and commitment to data transparency. Finally, the Ag Data Coalition was established to provide a farmer-controlled, secure, and flexible data repository, established as a cooperative. In 2017, the collaborating entities joined forces with the Grower Information Services Cooperative to announce the AgXchange, an independent data repository available commercially through the Growers Ag Data Cooperative.[6]

[6]See http://agdatacoalition.org/newsroom [August 2017].

In the past, only NASS had access to most U.S. ground (as opposed to aerial) farm-level data for a crop. But some of these ground farm-level data now are available through precision agriculture measurements to other private organizations that provide decision-making tools and analysis for growers submitting the precision agriculture data. Monsanto Corporation and Dow AgroSciences (2016) are two such companies that may offer farm management tools, site-specific weather information, and seed planting prescriptions to individual farmers who submit their precision agriculture data. Current data sharing and privacy agreements allow the parent company to use the aggregated data obtained from its farm customers. While the company may not have access to data from all or even most growers of a crop, it may have enough information to develop credible estimates, and it is not clear how these estimates might be used. NASS estimates have always been freely available to everyone on the official release date. If a private organization had better or equivalent earlier estimates, it could use them to speculate in commodity markets.

The panel offers the following recommendations for facilitating NASS access to data collected via precision agriculture measurements.

Recommendation 3-3. NASS should collaborate with farmer cooperatives to ensure that it is one of the government agencies with which farmers can choose to share their relevant precision agriculture data.

Recommendation 3-4. NASS should collaborate with current vendors of software that may be used to report relevant farm-level data directly to FSA or indirectly to RMA through AIPs/insurance agents to ensure that the software allows the option of reporting relevant data to NASS.

Recommendation 3-5. NASS should develop a precision agriculture reporting option for the County Agricultural Production Survey (CAPS)/Acreage, Production, and Stocks (APS) survey system. Farmers who reported relevant precision agriculture data would either not receive an additional survey form or receive one that was simplified and easy to use.

Recommendation 3-6. NASS should work to educate farmers that use precision agriculture concerning the advantages of using that technology to report to NASS. A convincing argument might be that having NASS collect the relevant data with a promise of privacy and making summary data available to all farmers at the same time would keep the playing field level.

According to Johansson and colleagues (2017):

> The fact that USDA reports its estimates freely means that both buyers and sellers can have equal information about the supply and demand of a crop. Such information may come directly through USDA's own reports, but often reaches users through news and information sources that depend on USDA reports to inform their clients and customers. In a market without this free information, large firms might well be able to invest in market intelligence that small firms and farms would not have available.

DATA OBSERVED BY SATELLITE AND OTHER REMOTE SENSING TOOLS

An increasing number of satellites, aircraft, drones, flux towers, and weather stations collect geospatially referenced data that may be useful for monitoring crop-growing conditions. These data may be available from other government agencies or for purchase from private companies.[7] The current NASS program based on satellite remote sensing and other external sources of information provides indications of planted acres and yield, but only for some crops and some regions.

A number of private organizations are now producing satellite-based estimates of yield and publishing them as alternatives to the official NASS monthly estimates of state-level yields during the growing season. Many of these organizations are publishing estimates frequently (daily or weekly) and at a fine level of detail (county). These estimates are direct competitors with the NASS state-level in-season yield forecasts that are not within the purview of this panel. In time, however, these methods may lead to estimates of county-level final yield that are viewed as competitors with those produced by NASS. Of course, one of the disadvantages of any county-level estimates prepared by these companies would be that the methodology used to derive them would be proprietary and not publicly released. At this stage of development, these companies judge the quality of their estimates by comparison with end-of season estimates from NASS.

The Agricultural Research Service (ARS) has a system of energy flux towers in the Midwest that provide ground-based remote measurements of a number of variables potentially relevant in the preparation of yield estimates over time. Each flux tower has a GIS location and collects a vast array of data, depending on its sensors. ARS uses these data, in conjunction with satellite data, to quantify carbon, measure production, and understand ecosystem dynamics. It would be valuable for NASS to consider how it might use the data from these flux towers as input to improve its yield models, especially those that link ground-based farm-specific reports with location information. ARS also is conducting research on the use of the OCO-2 satellite[8] data in conjunction with the flux tower data (Wood et al., 2017), especially to monitor carbon emissions and photosynthetic production, commonly known as gross primary productivity (GPP). NASS would be well advised to monitor the progress of this research.

In documentation provided to the panel at its November 2015 meeting, NASS included a list of weather-related resources. These include the National Oceanic and Atmospheric Administration's (NOAA) Climate Divisional Database, which contains a variety of weather and precipitation data;[9] Oregon State University's Program in Statistics and Methodology (PRISM),[10] which provides station-level temperature and precipitation data that may map to CLUs; the Applied Climate Information System, operated and maintained by NOAA Regional Climate Centers;[11] and the United States Drought Monitor.[12]

[7]For example, for data collected by drones, see http://agribotix.com [July 12, 2017].
[8]The OCO-2 satellite was launched July 2, 2014. It carries an imaging sensor that can detect chlorophyll fluorescence and hence provide estimates of crop growth.
[9]See https://catalog.data.gov/dataset/noaas-climate-divisional-database-nclimdiv [July 12, 2017].
[10]See http://www.prism.oregonstate.edu [July 12, 2017].
[11]See http://www.rcc-acis.org/index.html [July 12, 2017].
[12]See http://www.droughtmonitor.unl.edu [July 12, 2017].

Multiple Sources of Remote Sensing Data, Sources of Uncertainty, and Models

Each of the sources being considered as inputs to modeling has footprints on the earth's surface:

- **Satellite remote sensing.** Each pixel in an earth image covers a square area on a flattened earth. The size of the square varies depending on the sensor, from less than 1 m to 1 km in linear measure. The projection used to flatten the earth also varies. Merging data from different sensors and merging with other sources of data thus requires a complex process that takes into account the position, orientation, and geometry of each pixel on the flattened earth and the projection used for flattening. The nature of the response recorded in each pixel also must be addressed. This is a convolution of what is actually on the ground and the optics of the sensor; in the simplest case, it is an average over the area of the pixel. When the pixel's geometry is overlaid on another data source, the unknown heterogeneity of the pixel becomes important, although in most cases it is simply assumed away.
- **Digitized county boundaries.** These boundaries are commonly available, created from published maps, and have a positional accuracy of approximately 0.5 mm at the scale of the map. For example, boundaries digitized from 1:100,000 mapping will have a positional accuracy of roughly 50 m. Some sources may be more accurate: the Topologically Integrated Geographic Encoding and Referencing (TIGER) data used by the Census Bureau have been upgraded to roughly 5 m accuracy. Boundaries of standard reporting units are likely expressed in a common coordinate system, such as Universal Transverse Mercator (UTM), state plane, or latitude/longitude.
- **CLUs.** The CLUs used by the U.S. Department of Agriculture (USDA) have been digitized from digital orthophotos. The standard digital ortho quarter quad (DOQQ) program has a stated and tested positional accuracy of 6 m, although the orthophotos used by USDA in some areas may be more or less accurate. Unless CLUs have been snapped to county boundaries during the digitizing process, they will not match, and instead, small slivers will exist between the county boundaries and those CLU boundaries that follow county lines.

In summary, each source of geospatial data comes with its own geometry, expressed in one of a number of coordinate systems. Combining sources as input to models requires a complex process of addressing spatial mismatch that takes into account map projections and coordinate systems. Suppose, for example, that the modeling effort seeks to predict yield y over a small area, say, a CLU, using a number of input variables x_i, $i = 1, n$ derived from geospatial layers. Estimating each x for the small area from data whose spatial units do not precisely match the small area introduces a degree of uncertainty. The degree of uncertainty depends on the size of the original spatial unit relative to the size of the area to be estimated. For satellite remote sensing, the natural heterogeneity of pixel contents should be addressed, but likely will be assumed away.

NASS Satellite Remote Sensing Indications

The methodology used for producing NASS satellite remote sensing indications of planted acres was first described by Graham and Iwig (1996). The methodology for planted acres has not changed appreciably since that time, although the newer Landsat satellites have a smaller pixel size, and FSA CLUs now provide training data for classification, whereas in the 1990s, the June Area Survey was used as the source of training data for classification.

Classification

The Cropland Data Layer (CDL) is a georeferenced, crop-specific data layer of land cover created annually by NASS for the continental United States using Landsat 8's moderate-resolution (30 m) satellite imagery, Deimos-1 and UK-DMC2 moderate-resolution (22 m) satellite imagery, and agricultural ground truth. It offers nationwide coverage of pixels classified by land use.[13] The CDL is created using 70 percent of the FSA CLU and Form FSA-578 data as ground truth along with satellite data to classify pixels by crop in a decision tree. The remaining 30 percent of CLUs are used in a pixel-by-pixel comparison to assess the quality of the classification. NASS also uses the U.S. Geological Survey (USGS) Land Cover Data Base as a training and validation dataset for nonagricultural categories. Software used for classification includes ERDAS Imagine 2011, Rulequest See5/C5, and SAS.

The accuracies associated with crop-specific covers for major crops in selected states are shown in Table 3-2. The table counts only "buffered" pixels, that is, pixels that are wholly contained within a field, and ignores pixels that straddle a field edge and are therefore planted only partially to the crop. Errors of omission occur when a pixel that does indeed contain the crop is nevertheless not classified as such from remote sensing; errors of commission occur when a pixel is classified incorrectly as containing the crop. Remote sensing clearly has the greatest difficulty identifying winter wheat in Minnesota, followed by cotton in Oklahoma. But for other crops, notably corn and soybeans, the imagery clearly provides high levels of accuracy, at least for these buffered pixels.

Estimation of Planted Acres

NASS uses the classified pixels from the CDL, June Area Survey data, and the regression described by Battese and Fuller (1981) adjusted according to Walker and Sigman (1982) to estimate planted acres. The CDL pixel count is an auxiliary data source in the unit-level regression model, the unit in this case being a June Area Survey segment. The CDL pixels can be matched to the segment boundaries from the June Area Survey. County estimates are constructed by using the coefficients estimated for the state along with the constant term adjusted based on June Area Survey segments in the county. In 2014, satellite remote sensing estimates for planted acres were available for corn in 37 states, for soybeans and wheat in 27 states, for alfalfa in 14 states, and for cotton in 11 states. Estimates for planted acres of barley, beets, canola, dry beans, peanuts, potatoes, rice, sorghum, sugar cane, sunflowers, tobacco, durum wheat, and spring wheat were each available in fewer than 7 states.

A regression model such as the Battese-Fuller (1981) model provides estimates of uncertainty. However, the variance is likely to be understated because when estimates are rolled

[13]See https://nassgeodata.gmu.edu/CropScape [July 12, 2017].

up for larger areas (estimates from pixels rolled up to CLUs, for example), spatial dependence becomes an issue. The estimates being rolled up are not independent, and the result is that variance is likely to be underestimated.

Estimation of Yield

The NASS CDL classification is used to establish in-season corn and soybean fields for estimating yield. Currently, yield estimates are developed only for corn and soybeans, although yield estimates for other crops are under consideration.

The current methodology for estimating yield was implemented in 2013 and is described by Johnson (2014). The Normalized Difference Vegetation Index (NDVI) and Daytime Land Surface Temperature (DLST) have been demonstrated to be correlated with plant yield during the growing season for a variety of commodities in a number of references, including Johnson (2016) and Hatfield (1983). Johnson (2016) states that the peak correlation between yield and NDVI during the growing season for corn and soybeans was .8 and .7 respectively, and that correlation with DLST was negative .6 and negative .5, respectively. The seasonal NDVI profile describes the crop growth and development; surface temperature provides additional information on potential crop stress conditions.

NASS computes an 8-day composite of the NDVI and daytime land surface temperature from the Moderate Resolution Imaging Spectroradiometer (MODIS) on two earth science research–oriented satellites, Terra and Aqua, operated by the National Aeronautics and Space Administration. MODIS provides NDVI measurements with 250 m resolution and daytime and nighttime land temperature measurements with 1 km resolution.

PREPUBLICATION COPY, UNCORRECTED PROOFS

TABLE 3-2 Examples of Classification Accuracy

Commodity	State	Production	Buffered Pixels	Pixels Classified as Commodity	Pixels Correctly Classified as Commodity	Actual Pixels of Commodity	Percent Pixels Incorrectly Classified as Commodity	Percent Pixels Incorrectly not Classified as Commodity
Corn	Iowa	2740.5 million bushels	944,966	397,529	391,624	396,284	1.2	1.5
Corn	Ohio	524.7 million bushels	928,450	226,117	220,008	225,195	2.3	2.7
Soybeans	Illinois	560.9 million bushels	939,373	280,544	273,966	279,843	2.1	2.3
Soybeans	Arkansas	146.6 million bushels	949,805	272,025	264,779	280,643	5.7	2.7
Cotton	Texas	7436 thousand bales	1,964,758	222,419	209,102	233,831	10.6	6.0
Cotton	Oklahoma	350 thousand bales	943,356	14,729	12,807	15,541	17.6	13.0
Winter Wheat	Kansas	4674 thousand bushels	962,256	205,127	196,774	206,806	4.9	4.1
Winter Wheat	Minnesota	75 thousand bushels	969,727	276	94	121	22.3	65.9

SOURCE: Prepared by the panel based on data from https://www.nass.usda.gov/Research_and_Science/Cropland/sarsfaqs2.php [August 2017]. Data from links are in answer to question 12.

According to Johnson (2014), a data mining software product, Rulequest's Cubist,[14] is used to develop a predictive model for yield. A grid-based reference frame, the 250 m sinusoidal MODIS-based equal-area map projection, has been selected to prepare the data for input to the model. With this reference frame, a 250 m resolution raster-based mask file is created for each crop and state/county to identify the classified 30 m Landsat pixels within the state/county border. Boundary definitions are derived from Esri's Data and Maps Detailed Counties shapefile. A majority-area rule is used to assign pixels split by two or more states/counties. Each final state mask is aligned spatially with the 250 m MODIS reference grid. Counts of classified pixels of the crop of interest (corn or soy) within each 250 m reference grid are determined. Johnson (2014) reports that 16–49 pixels are assigned to each grid for a crop. The total number of classified pixels of the crop of interest in each MODIS grid cell are divided by the maximum possible to derive an areal proportion. Finally, that proportion is compared with .9 to determine whether that grid cell contained enough of the crop of interest to be used in estimation. County-level means of NDVI and land surface temperature are computed for each of the 32 time periods from mid-February to mid-October. Inputs to the model include the 32 values collected for NDVI and land surface temperature during the season; the areal proportion of pixels of the crop of interest; and the NASS-estimated county yields for a 6-year period, 2006–2011. Johnson (2014) reports that midseason values of NDVI and land surface temperature were most important in the predictive model, and that a county-level within-sample R^2 of .93 was found with both crops.

The model and its estimated parameters are used with current-year satellite measurements of NDVI and land surface temperature in projecting yield for the current year to provide out-of-sample forecasts. Johnson (2014) reports that results compared reasonably well with published official NASS estimates, with an R^2 of .77 for corn and .71 for soybeans. The root mean square errors (RMSEs) were 1.26 and .42 metric tons per hectare. Johnson also "hindcast" years 2006 through 2011 showing similar results to those for 2012 (e.g., a hindcast of 2009 would remove the data from 2009, refit the model, and use the model to predict results for 2009). The resulting R^2 for corn ranged from .66 to .77, with an RMSE of .96 to 1.26 metric tons per hectare, and for soy ranged from .47 to .71, with an RMSE of .38 to .47 metric tons per hectare.

At present, NASS does not provide measures of uncertainty for its remote sensing estimates of acreage or yield, even though both are derived from models. ASB considers indications for planted acres and yield for selected regions and crops when preparing estimates; however, these indications receive little weight in computation of the composite indication. The panel concludes that NASS does not yet view the current remote sensing indications as being as accurate as other sources of information.

The Statistics Canada/Agriculture and Agri-Food Canada Approach to In-Season Estimates of Yield

At its January 2017 meeting, the panel heard from Gordon Reichert about Statistics Canada's Integrated Canadian Crop Yield Forecaster (ICCYF) (see also Reichart et al., 2016). The ICCYF is a crop yield monitoring model that uses remote sensing and agroclimatic data along with survey data to estimate yield during the growing season. (See Chipanshi et al. [2015] for an evaluation of the model, and Newlands et al. [2014] for a more detailed description of the

[14]Described on the Rulequest website as a machine learning tool that automatically determines the best-fitting piecewise linear model to predict a continuous outcome variable.

model.) Based on the results of the modeling effort, Statistics Canada decided in 2016 that it could replace one of its growing season surveys, the September survey, with the remote sensing estimate, thereby reducing the burden on farmers, reducing costs, and enhancing timeliness.

The goal of the model is to predict final crop yield for the current year. Hence the model's dependent variable is the crop yield estimate from the November Farm Survey. During development of the model, 80 potential explanatory variables were considered, including same-year yield estimates from the July survey, NDVI (1 km resolution) from the NOAA series of satellites, and many weather and precipitation variables obtained from Agriculture and Agri-Food Canada's (AAFC) Versatile Soil Moisture Budget (VSMB) model. These weather variables were obtained from 416 climate stations throughout the agriculture region of interest and were prepared using information about soil type and crop phenology. The variables ultimately selected for the model include daily measures of the cumulative (through the growing season) growing degree days (GDDs), cumulative precipitation, and crop moisture stress. Also included are the standard deviation of daily GDD and the standard deviation of daily precipitation.

The ICCYF uses a stepwise selection process—the least absolute shrinkage and selection operator (LASSO)—to determine the variables to use as independent variables. Previous work (Bédard and Reichert, 2013) had established that five of the climate variables should be included.

The ICCYF raises several points of potential relevance to NASS. First, it makes use of agroclimatic data not from remote sensing but from a large number of weather stations. Second, Statistics Canada uses the VSMB model to integrate the agroclimatic data, along with soil type and crop phenology information, to derive a variety of relevant county-level input variables. Two of the key input variables are measures of uncertainty: the standard deviation of the daily stress index and the standard deviation of daily GDD. Chipanshi and colleagues (2015) provide an evaluation of the ICCYF model for spring wheat, barley and canola. Their paper describes how the authors selected three model comparison measures: the coefficient of determination (R^2), the mean absolute percent error (MAPE), and the Nash-Sutcliffe efficiency index (Krause et al., 2005; Szulczewski et al., 2012) (one minus the sum of squared differences between predicted and observed divided by the sum of squared differences between observed and the mean of observed values). Reichart and colleagues (2016) further evaluated the model for 15 crops in total. The authors chose to examine these measures using a leave one out validation approach. They concluded that the model results and September survey results at the national level were generally comparable for the 15 products considered. Their conclusions led to the adoption of the ICCYF as a replacement for the September survey in 2016.

A model such as the VSMB could prove very helpful to NASS in identifying and preparing relevant independent variables that could be used to improve both remote sensing models and area-level models for yield. Most likely such a model could be obtained, used, and/or adapted from work done by other agencies within USDA, such as ARS. Having such a model also would help NASS make better use of data from ground-based weather stations, and possibly even from flux towers.

Near-Term Enhancements to NASS Remote Sensing Indications

Estimating Acreage

Currently, NASS uses 70 percent of FSA data for "training" to classify pixels and withholds 30 percent for assessing the accuracy of the classification. NASS uses the June Area

Survey data by segment in the Battese-Fuller regression model (Battesse and Fuller, 1981) for estimating acreage by crop. NASS could consider using the 30 percent of FSA CLUs withheld for accuracy assessment as ground truth for the Battese-Fuller regression, either as a replacement for or in addition to the June Area Survey segments. NASS also could consider using the Battese-Harter-Fuller (Battese et al., 1988) regression model instead of the Battese-Fuller model since there is no stratification to account for with the use of FSA CLUs. Additionally, NASS could consider using the uncertainty measures from these regression estimates in developing its uncertainty measures for remote sensing estimates of acreage.

Recommendation 3-7. NASS should consider using the 30 percent of CLUs withheld to assess the accuracy of classification in the Battese-Fuller regression either as a replacement for or in addition to data from the June Area Survey segments.

Estimating Yield

RMA data can currently be used to identify insured acres as a proxy for planted acres as well as failed acres by CLU/RLU and production for a farm by crop. RMA has required reporting by CLU/RLU for all farms since 2016, but collected these data for many farms in preceding years. Data on acreage are reported to insurance agents by July 15 of the crop year. Farm-level production is reported to insurance agents by about May 20. There are delays in the reporting of data to RMA. This dataset would provide ground truth information for use in building a model that would use the model and fitted parameters from the previous year along with satellite remote sensing data from the current year to estimate yield. The RMA data can provide yield information for a specific GIS location on the ground. Location information can be linked to satellite measurements, as well as to information about land productivity, such as the NCCPI, and to weather data.[15]

Additional Sources of Remote Sensing Data

If NASS does not use the most up-to-date satellite data when they are pertinent, private organizations using these data may produce better estimates. The increasing number of satellites and the complex types of information they provide make it an ongoing challenge to incorporate the latest satellite data into estimates in a timely fashion. As NASS states in its Estimation Manual,[16] “Having the flexibility to use multiple sensors and adapt to new ones is essential in the NASS operational model.” Many agencies within USDA are considering how to develop improved production/yield estimates, especially using satellite data and growth models. In addition to NASS, the Foreign Agricultural Service and ARS are examples. NASS needs to be a full partner in these efforts.

[15]NASS has identified the following potential sources of weather data: (1) NOAA National Climactic Data Center—Climate Divisional Database nCLIMDIV (http://catalog.data.gov/dataset/noaas-climate-divisional-database); (2) Oregon State University’s PRISM (http://www.prism.oregonstate.edu); (3) Applied Climate Information System (http://www.rcc-acis.org/index.html); (4) United States Drought Monitor (http://www.droughtmonitor.unl.edu); and (5) NOAA/National Weather System (NWS) Advanced Hydrologic Prediction Service (http://www.nws.noaa.gov/ohd/ahps).

[16]NASS Estimation Manual, Volume 2, Section 6.2.3. Provided to the panel.

The JASON Deliberative Study (2016) entailed a comprehensive review of the field of remote sensing and the many new satellites and sensors that have been launched in recent years, and this review might provide useful data for NASS. While NASS is already using Landsat data (30 m resolution), the JASON report also draws attention to similarly free data from the European Space Agency (ESA)-2 satellites (10 m resolution). It further suggests possible augmentation with data purchased from Rapideye (from the Satellite Imaging Corporation, with five satellites and 5 m resolution) and the Disaster Monitoring Constellation (a number of remote sensing satellites constructed by Surrey Satellite Technology, operated by the Disaster Monitoring Constellation for International Imaging [DMCii], and designed to be comparable in resolution to Landsat). NASS reported to the panel that it has access to the second-generation DMC satellites Deimos-1 and UK-DMC2, with moderate resolution (22 m). In its report, JASON also mentions ResourceSat-2 (launched in December 2016), with sensors that include the Advanced Wide-Field Sensor (AWiFS) with 56 m resolution, the Linear Imaging Self-Scanning Sensor (LISS-III) with 23.5 m resolution, and the LISS-IV Camera with 5.8 m resolution. Each of these sensors has its own repeat interval (the number of hours or days between successive imaging of a given point), its own spectral characteristics (the portion or portions of the electromagnetic spectrum sensed by the imager), and its own acquisition details (cost and length of time between the capturing of an image and its availability for analysis by an agency such as NASS).

In addition to the sensors discussed in the JASON report, many other newer sensors are being developed for deployment on satellites, aircraft, drones, towers, and agricultural machinery. Mention has already been made of OCO-2, a satellite designed to detect emissions of special relevance to plant growth. Radar, light detection and ranging (LIDAR), and microwave are examples of types of active sensors that both emit and receive radiation, in contrast to the passive basis of traditional remote sensing.

The days when only the federal government possessed the resources to invest in remote sensing are long gone, and the private sector and farmers themselves are increasingly important sources of data for crop monitoring and estimation. The number of sensors that might be useful to NASS is clearly growing rapidly, and at an increasing rate. To NASS this should be a simple matter of costs and benefits: How do the costs (including the programmatic costs of switching to a new source of imagery) compare with the benefits? Finer spatial resolution will require more powerful processing, and is unlikely to yield more accurate identification of crops except in extremely small fields. Estimation of total crop area, on the other hand, will be improved with smaller pixels because of the issues at field boundaries. If each pixel is length b on a side and assuming a square field of B on a side, the number of edge pixels (pixels rejected as compromised by the boundary) will be approximately $4B/b$, covering an area of $4Bb$, or $4b/B$ as a proportion of the total. Thus the error due to edge pixels is proportional to b, and reducing pixel size from 30 m to 10 m can be expected to reduce error in estimating field size and thus yield by a factor of 3.

Recommendation 3-8. NASS should explore collaboration with other USDA agencies that are actively involved in remote sensing applications to obtain access to data with finer spatial resolution, and possibly also to share in the costs of processing those data.

NASS should continue to focus on producing timely and accurate estimates. In light of changing markets and emerging data sources such as precision agriculture measurements and remote sensing measurements, it should consider whether purchasing private data would support this effort.

There are disadvantages to private data. Transparency of source data and methods may be difficult for proprietary data vendors. Data vendors may change their products over time. They may change their marketing strategy. They may go out of business. However, they may also offer data for purchase that are not easily produced by NASS (or other government surveys). NASS should evaluate the costs and benefits of external data, considering the following questions: Do the data fill a gap? How closely do the data fit with NASS concepts and definitions? What is the time frame of the data? How representative are the data? How consistent are the data, and how likely is it that consistency will be maintained? How transparent are the data?

Recommendation 3-9. NASS should keep abreast of emerging data sources; how they are used; and how they might be used to improve county estimates, especially of yield. Based on a careful evaluation, NASS might consider purchasing data.

4

Sources of Data for Cash Rents

Because land is the primary input to agricultural production, its value and the rent it garners are key indicators of the financial health of the farm sector. According to the 2012 Census of Agriculture, almost 40 percent of agricultural land is rented.[1] NASS estimates thus provide reference points for transactions in a significant market for farmland. Accordingly, NASS has regularly sought to improve the quality of its cash rent estimates and, as with crop yields, has sought additional sources of data to bolster its predictions.

The current NASS Cash Rents Survey is sponsored by the Farm Services Agency (FSA), which uses the results in determining payments to farmers under the Conservation Reserve Program. In its surveys, NASS asks respondents to report acres rented and either the cash rental rate or the total dollars of rent paid. NASS computes cash rental rates as the ratio of total dollars paid in rent within a geographic area to total acres rented in that area based on responses to the Cash Rents Survey (used for county-level estimates) or the annual June Area Survey (used for state-level estimates). The methodology for these surveys is described in Appendix A. The Cash Rents Survey is sent to about 275,000 farms every other year, while the June Area Survey results in about 35,000 interviews.[2]

This chapter summarizes the auxiliary data available for improving the estimation of county-level cash rents, describes the model for cash rents that currently provides input to the Agricultural Statistics Board (ASB), and summarizes how NASS might transition the ASB process for estimation of cash rents to one that is more transparent and reproducible.

AUXILIARY DATA FOR CASH RENTS

Auxiliary data for cash rents are potentially valuable for improving estimates derived through models. NASS has noted that spatial and temporal relationships of cash rents are quite stable and that cash rent values may be related to land values, soil productivity, yield, and climate. Cash rents reflect the value of the farmland asset in crop production, a relationship that can be modeled statistically. A number of sources of auxiliary data for use in a cash rents model are available.

First, the U.S. Department of Agriculture's (UDSA) Economic Research Service (ERS) and NASS jointly sponsor the Agricultural Resource Management Survey (ARMS), which, among other things, collects land value information from farms annually. The sample size of the

[1]See https://www.agcensus.usda.gov/Publications/2012/Online_Resources/Highlights/Farms_and_Farmland/Highlights_Farms_and_Farmland.pdf [July 12, 2017].

[2]See https://www.nass.usda.gov/Surveys/Guide_to_NASS_Surveys/June_Area [July 12, 2017].

ARMS is about 30,000 farms and ranch operations each year. The ARMS Tenure, Ownership, and Transition of Agricultural Land (TOTAL) survey collected cash rental rates and farmland values in 2014 from landlord owners of agricultural land, including nonfarm operators of agricultural land, as a follow-on component of the Census of Agriculture program. Data are available only for the U.S. total and 25 states. The U.S. sample size was 41,205.[3] A similar survey, the Agricultural Economics and Land Ownership Survey, was conducted in 1999. Farm real estate values (land and buildings) are collected in the Census of Agriculture.[4]

Second, county tax offices may have tax information on land values and real estate transactions. However, the definitions and formats they use are inconsistent. Commercial vendors have compiled tax information and made it publicly available for purchase. Examples include the following:

- Core Logic has data on land values and geospatial parcels.
- Boundary Solutions Inc., another parcel data company (a competitor of Core Logic, it says on its home page that it is the custodian of the national Parcelmap Data Portal, which also includes extensive tax roll attributes).
- Acrevalue by Granular advertises on its website that users can “view GIS [geographic information systems] maps that compile agricultural data, including farmland values, soil productivity ratings, crop mix, and parcel ownership information.”

Third, many, if not all, of the Federal Reserve Banks, including the Chicago, Kansas City, Dallas, and Minneapolis Federal Reserve Banks, conduct agricultural credit surveys. Additionally, land grant universities such as Iowa State and Michigan State collect information on and prepare estimates of cash rents and land values. The Purdue farmland value survey collects information from rural appraisers, commercial bank and agricultural loan officers, FSA personnel, farm managers, and farmers. These private surveys tend to be collected during the spring of the year and published in August, timing similar to that of the NASS Cash Rents Survey.

Finally, land values are a function of characteristics of the land itself, geographic location, market prices, and eligibility for government payments. Nickerson and colleagues (2012) document a trend and correlation analysis of farmland values and the related macroeconomic and parcel-specific factors. For this analysis, they used estimates for cropland and pastureland values from the NASS June Area Survey. In particular, because the June Area Survey is an area survey, they were able to use its GIS identification of farm parcels to develop linkages to such parcel-specific (spatial) factors as soil quality, access to market terminals, government payments, proximity to development potential, distance to population concentrations, access to roads, and the land’s amenity value. Data available from the June Area Survey include cash rental rates, land values, irrigation status, and planted acres by crop. Among the variables found by the authors to affect farmland values were proximity to urban areas, soil

[3]See https://www.agcensus.usda.gov/Publications/2012/Online_Resources/Highlights/TOTAL/TOTAL_Highlights.pdf [July 12, 2017].

[4]See https://www.agcensus.usda.gov/Publications/2012/Online_Resources/Highlights/Farm_Land_and_Buildings/Highlights.pdf [July 12, 2017].

quality, and irrigation status. Of note, they found that relationships between land values and the factors they considered were complex and varied in importance across location and time. For example, they found that the relationship between land value and soil quality was fairly consistent for farms that were at least 25 miles away from a city of at least 50,000, and was even stronger for parcels at least 40 miles from a population center. They also observed that patterns varied substantially across regions, with positive correlation between soil quality and cropland values in the Corn Belt, Lake States, and North Dakota (in the Northern Plains region), but a negative or insignificant correlation in the Appalachian region.

While the analysis of Nickerson and colleagues (2012) was intended to illuminate land values, a similar approach could be used to evaluate the factors associated with cash rents. In particular, the authors' Appendix A provides a detailed description of the preparation of the input variables and indices used in the study. The examination and development of parcel-specific variables is of particular importance since the addition of Common Land Units (CLUs) to the NASS list frame would make the development of unit-level models for cash rents feasible in the future. Currently, as described in the next section, NASS makes use of a cash rents model that considers previous-year rents but also a limited set of covariates that relate mainly to yield and soil productivity. Recruiting additional sources of data for the modeling effort and tying estimates to specific parcels of land could potentially improve prediction results.

Recommendation 4-1. NASS should work with the Economic Research Service to extend the Nickerson et al. (2012) analysis of parcel-specific variables that influence farm values to parcel-specific variables that influence cash rental rates. The results of this effort would illuminate the potential for additional modeling of cash rental rates once geospatial identifiers are available.

MODEL-BASED INDICATIONS FOR CASH RENTS

Berg and colleagues (2014) developed a model for producing county-level estimates for cash rents based on the (then) annual Cash Rents Survey, and illustrated the approach by applying it to estimation of cash rents for irrigated cropland, nonirrigated cropland, and pastureland in six diverse states. A primary motivation for investigating model-based estimates of cash rental rates was to guide ASB with a transparent and reproducible method for developing county-level estimates for cash rents in the above three categories, along with mean squared error estimates of those rents.

NASS first incorporated the Berg and colleagues (2014) model for cash rental rates into its official estimation process in 2013. In addition to direct estimates obtained from the Cash Rents Survey, which serve as the primary basis for setting official estimates, the Berg and colleagues (2014) model-based estimates were supplied as additional indications for ASB's consideration in 2013, 2014, and 2016 (because, as described in Chapter 2, the survey was not conducted in 2015). See Bellow et al. (2017) for a summary of the impacts of the skipped year on the performance of the cash rents model. The Cash Rents Survey will be conducted in 2017, indicating that NASS will likely need to maintain two versions of the model, one for the situation when surveys are conducted in adjacent years, and one for the situation when there is a 2-year gap between surveys.

The model described by Berg and colleagues (2014) consists of two univariate Fay-Herriot models—one for the average of the two yearly means and the other for the difference between the two yearly means—each with an auxiliary input index, as described below. The estimate of cash rents for the current year is then the sum of the average cash rents from the first model and half of the average cash rents difference from the second model. The procedure recognizes the correlation between cash rents in the two consecutive years but assumes equality of variances in these two years.

The auxiliary input index is based on the county-level total dollar value of agricultural production from the 2007 Census of Agriculture; NASS's published county yields for 2004–2009 to reflect the quality of the land in the county; one of four yield indices described below, depending on the county; and National Commodity Crop Productivity Indices (NCCPIs) for corn, wheat, and cotton from the National Resources Conservation Service. The yield indices include one based on combined yield from irrigated and nonirrigated croplands in states where these yields are not split, two indices for states where yields are split by irrigated and nonirrigated, and a separate hay yield index used to exclude hay crops from the irrigated and nonirrigated croplands.

Although this model meets many of the salient challenges that arise in actual operations, the simplifying assumption of equality of variances in the two years may not always hold and merits further investigation. The assumption is less likely to hold when surveys are conducted every other year rather than annually. A panel member discussed with NASS staff a bivariate hierarchical Bayesian model that circumvents this difficulty and warrants evaluation. Porter and colleagues (2015) provide another possibility, as do Bradley and colleagues (2015a). An evaluation of the model is described below.

Recommendation 4-2. NASS should develop and evaluate a bivariate hierarchical Bayesian model for cash rents that does not rely on the assumption of equal variances in two survey years. The intent would be to improve model performance, especially when surveys are conducted 2 years apart.

Benchmarking

The model of Berg and colleagues (2014) implements two-stage benchmarking that enables benchmarking the Agricultural Statistics District (ASD) estimates to the corresponding aggregated state estimates, and the county-level estimates to the corresponding benchmarked ASD estimates. The proposed method benchmarks the county cash rental rates to the average cash rental rate for the state, with weights based on the fraction of rented acres in the county, a procedure more sophisticated than the ratio adjustment currently used by NASS. The latter adjustment suffers from the drawback that all units are equally benchmarked to the corresponding aggregated unit and that the variation introduced by this process is not included in measures of uncertainty.

The current NASS procedure constrains the model's predictions because they are benchmarked to survey results. Such benchmarking might appear logical in the context of crop production estimates, when county and state estimates must sum to a preset, finite total of national output. But it is not obvious that the cash rents model should be treated similarly, as there is no analogous constraint on what could be spent on cash rents. NASS could develop unconstrained estimates using models that include the auxiliary data on land characteristics, geographic location, and subsidy eligibility discussed earlier. The performance of the two modeling approaches could then be compared. However, several references describe potential advantages to benchmarking. Pfeffermann (2013), for example, notes that benchmarking to higher levels when design-based estimates are deemed reliable may provide some protection against cases in which the model is misspecified. This benefit is especially important for time series models since they are "slow to adapt to abrupt changes." Additionally, as observed by Berg and colleagues (2014, p. 23), benchmarking can reduce bias that may result from the modification of outliers, as noted by Gershunskaya and Lahiri (2010).

NASS Evaluation of Model-Based Estimates

At its January 2017 meeting, the panel heard a presentation by Cruze, Erciulescu, and Bellow, that has since been published in Bellow et al. (2017), assessing the suitability of the model of Berg and colleagues (2014) adjusted to accommodate survey data collected every 2 years. They described a careful evaluation of the model and its performance, expressing some frustration that the model did not always appear to result in improvements. They observed that, according to the 2012 Census of Agriculture, there were 2,764 counties with 20,000 or more acres of combined cropland and pasture. Cash rents are intended to be published for these counties. A cash rental rate for at least one of the categories (nonirrigated, irrigated, pasture) has been published for 93–95 percent of these counties since 2009.

The model is a clever way to gain strength by combining survey estimates for 2 years. However, it relies on the difference between individual farm-level survey responses in the 2 years as part of computing the estimate for the difference between the 2 years. This approach has obvious benefits except in counties for which the sample is thin, in which case there may be no or few such matched observations. Berg and colleagues (2014) note that the average number of respondents per year was approximately the same when the Cash Rents Survey was conducted annually. However, the average response rate in both years was 25–50 percent. The average number of farms responding in both survey years is a key indicator of the quality of the resulting model-based indications. Table 1 in Berg et al. (2014) illustrates that irrigated and nonirrigated cropland and pastureland have very different error characteristics. For example, that table shows that there was an average of less than one matched respondent for irrigated cropland in four states, fewer than two in another state, and fewer than three in another state. If this characteristic persists, it may be unreasonable to expect to make much use of model-based indications for irrigated cropland. Results were somewhat better for pastureland, with only two states having an average of fewer than three matched respondents. Results were most promising for nonirrigated cropland. Because of the low match rate between respondents in the two survey years, the revised model using data from 2014 and 2016 could not produce an estimate for all counties for which ASB produced estimates.

The analysis presented to the panel displayed model results versus ASB estimates (indicating those suppressed for confidentiality purposes). The straight-line plots looked promising, but there were outliers, some of the most egregious of which were suppressed in ASB estimates. The model indication looked most like the Board estimate (adherence to a straight line) for nonirrigated cropland and irrigated cropland (except for suppressed values). Pastureland followed a line with outliers and variance clearly increasing with cash rent per acre.

A final table presented to the panel showed the number of times the ASB estimate fell within the 95 percent confidence interval of the model-based estimate. This table showed that 92–95 percent of the ASB estimates fell within the model's 95 percent confidence interval.

NASS needs to determine how the model can help in providing high-quality estimates, and in which counties (or for which the characteristics of those counties) it will help the most. The model might be useful to NASS in two key ways. First, by combining 2 years of survey data and using covariates, the model may provide publishable estimates for some of the counties that are not currently publishable. Second, in some counties that pass the current publication standard, the model-based indications may have lower mean squared error, leading one to think that they would provide higher-quality estimates. This model is unlikely to support estimates of cash rents for irrigated cropland in counties where the sample is already thin.

The panel proposes that NASS prepare a summary illustrating the root mean square error or other summary uncertainty measure of the model-based indication versus the design-based indication for several categories of counties (a county has fewer than 30 respondents but accounts for more than 25% of acreage, a county has fewer than 30 respondents and accounts for less than 25% of acreage, and a county has 30 or more respondents). Within the latter two groups, some classification based on the number of matched respondents might also be useful to help in determining the characteristics associated with the best model performance. Summary statistics for the uncertainty measures might include the maximum, 75th percentile, median, 25th percentile, and minimum, as well as the average number of matched responses in the county. Matched respondent groups might be 10 to 19, 20 to 29, 29 to 50, and more than 50. Ideally, this table would indicate where the model provides the greatest (or least) improvement. A similar table using counties published only because respondents cover 25 percent of acreage might indicate how well the model does for this set of counties.

This analysis would represent only a starting point. The intent would be to identify where the model works best and hence where and how it should be used, as well as where the model does not work well. Note that this model requires matched responses in two subsequent years to produce useful estimates, and cannot be expected to perform well in all counties for all land use categories. It may, however, provide improved estimates in some counties, and its ability to do so needs to be ascertained.

Recommendation 4-3. NASS should work with ASB and stakeholders to determine the circumstances in which the model of Berg and colleagues (2014) performs best and develop guidance on how the model indications should be used by ASB. In particular, tables illustrating uncertainty estimates for counties not currently publishable, those publishable only because respondents cover 25 percent of acreage, and those currently meeting publication standards would be most valuable.

THE AGRICULTURAL STATISTICS BOARD'S PROCESS FOR REVIEWING ESTIMATES FOR CASH RENTS

The remaining challenge is that, although model-based indications for cash rents are now available to ASB, there is no public information on how (or whether) they are used. NASS's estimates are derived through its ASB process. The panel was told that the primary input to the determination of cash rental rates for irrigated and nonirrigated cropland and pastureland is direct indications from the Cash Rents Survey. Indications of total rented acres by category for ASDs are ratio benchmarked to add to the state totals used to prepare published estimates of cash rents at the state level. Next, county-level indications are ratio benchmarked to sum to ASD totals.

The panel was told that ASB strives for consistency in its estimates and makes use of all available indications (direct survey indications, past-year cash rents estimates, and model-based indications). Edit guidelines are described in the NASS Estimation Manual. The guidelines state, for example, that the current-year estimate should be between or on the current survey indication and the previous-year estimate. If the direct survey estimate has a sufficient number of reports and satisfies this inequality, it is likely to be used. The panel is unaware of guidance on specifically what is to be done if the inequality does not hold. The panel was told that large counties with good reports are typically not adjusted by the ASB review process. Adjustments are made more often in smaller counties to make them consistent and fill in gaps so as to match the state total. Year-to-year outlier changes are verified with reports, historical knowledge, and field staff notes. Field offices are instructed to comment on, document, and justify changes to the survey-based indications. In 2014, 34,452 county-level estimates were developed with 939 (8.2%) ASB changes.

While the process is not particularly transparent and reproducible, the panel's interpretation of the ASB changes is that the direct survey estimate was used for more than 90 percent of the count-level estimates. There is also evidence that NASS is keeping good records of decisions made as part of the process of selecting estimates. ASB's cash rents estimation process would likely be the easiest to use as an example for converting to a transparent and reproducible process.

NASS already has established some rules for review of estimates that it uses to assist in ASB's deliberation process. This experience can be used in creating formal and well-documented guidelines that would lead to a more transparent review process.

Recommendation 4-4. NASS should use its experience with the cash rents model results and other model-based approaches as a starting point for establishing a set of clear guidelines and rules for review of estimates before they are released for publication.

5

Implementing the Vision and Beyond

Chapter 2 lays out a vision for NASS in 2025 that includes evolving the role of the Agricultural Statistics Board (ASB) from integrating multiple data sources through a process that appears subjective to one of evaluating estimates prepared through a statistical model-based integration of these alternative information sources (Recommendation 2-1), preparing its county estimates using a transparent and well-documented process (Recommendation 2-2), and developing and publishing measures of uncertainty along with point estimates (Recommendations 2-3 and 2-4). Taking these steps will bring NASS into compliance with Standard 4.1 of the statistical standards promulgated by the U.S. Office of Management and Budget (2006). To adhere to these standards, NASS will need to develop, evaluate, and use statistical models; formulate a plan for ongoing evaluations to assess model and survey methodologies; and provide documentation on its website concerning its methodologies for developing county-level estimates (Recommendation 2-6).

Foundational for achieving this vision is for NASS to have a georeferenced list frame that can be kept up to date using available administrative information, will support improved use of administrative data in NASS's survey operations, and will provide the geospatial information needed to support farm- and unit-level modeling. Development of a georeferenced list frame will be accomplished most expeditiously if NASS adopts the Common Land Unit (CLU), the geospatial convention already in use by the Farm Services Agency (FSA) and the Risk Management Agency (RMA), as its basic spatial unit (Recommendation 2-8). NASS also will need to be prepared to maintain alternative geospatial field boundary data, such as those from precision agriculture or the resource land units (RLUs) approved by RMA, in its databases (Recommendation 2-9). NASS will need to identify the CLUs that make up each NASS farm, which will likely require changes to the structure of the list frame to accommodate the georeferenced CLUs (Recommendation 2-10). NASS also will need to determine how to collect or identify CLU (or equivalent) data for farms that are not on FSA or RMA lists.

Recommendation 5-1: NASS should undertake a staged, systematic effort to implement the vision presented in Chapter 2 of this report.

Achieving this vision will take many years and a focus on achieving results. Senior NASS leadership will need to adopt the vision and clarify the overall goal. They will need to identify one or more champions who can identify others to be part of the process and can promote the importance of evolving to achieve the vision.

This chapter breaks down the vision and the panel's recommendations into projects that could be accomplished by different groups of people within two stages of effort. This detail is provided to assist senior leadership in their planning and staff in better understanding the panel's recommendations.

First-stage projects can start now. They include enhancing liaisons both within and outside the U.S. Department of Agriculture (USDA): within to leverage each agency's unique areas of expertise, and outside to gain access to precision agriculture data in a way that reduces the burden on farmers and enhances accuracy (Recommendations 3-3, 3-4, 3-5, and 3-6). First-stage projects further include working to enhance, document, and use the current cash rents model (Recommendations 4-2 and 4-3); developing useful small-area models for acreage and yield; planning for the inclusion of CLUs and equivalent geospatial information (such as RLUs and boundary information from precision agriculture) in the Enhanced List Management Operations (ELMO) database; and linking FSA and RMA data to the NASS list frame (Recommendation 2-10).

Second-stage projects can be accomplished once sufficient progress has been achieved in linking FSA and RMA farms with the farms in the NASS list frame. These projects include adding a new Internet reporting option for linked FSA and RMA farms to take advantage of administrative data and reduce reporting burden (Recommendation 2-11), enhancing surveys to make use of administrative data in imputation (Recommendation 2-12) and estimation (Recommendation 2-13), and expanding the modeling effort to consider farm-level or unit-level models.

The goal of this chapter is to help NASS see how the panel's vision can be realized. Moving to make this vision a reality will present many challenges for NASS, but the adoption of new ways of doing business is necessary to maintain the agency's credibility and promote efficiency in its operations.

FIRST-STAGE PROJECTS: PROJECTS THAT CAN BEGIN NOW

Project 1: Collaboration

NASS is a statistical agency within USDA, a department including many other agencies with expertise, data, and skills that complement those of NASS. NASS staff who collaborate on intra-USDA agency projects can identify ideas and models (broadly construed) that ultimately might improve NASS methods. It is likely, for example, that NASS cannot accomplish all the remote sensing work on its own and would benefit from an exchange of ideas with other geospatial experts within USDA. NASS modelers would similarly benefit from consultation and collaboration with modelers from sister agencies. NASS needs to collaborate broadly, especially with other USDA agencies that have substantial expertise in remote sensing and modeling, such as the Economic Research Service (ERS), the Agricultural Research Service (ARS), the Foreign Agricultural Service, RMA, FSA, and the National Resources Conservation Service. Specific steps to this end include the following:

- ERS and NASS administrators now meet monthly, and they could add approaches to achieving the vision to their agenda. ERS has extensive modeling and geographic information systems (GIS) capability that could be invaluable to NASS, while improvements in NASS data and estimates would benefit ERS, other agencies in USDA, and other data users.

- NASS leadership could encourage ERS to propose and establish joint projects that might lead to improved matching of the NASS list frame to FSA and RMA farms.
- NASS leadership could encourage ERS to extend its analysis of farm values (Nickerson et al., 2012) to include similar research on the impact of parcel-specific variables on cash rents so as to identify variables that might improve models (Recommendation 4-1).
- NASS leadership could establish working relationships between modelers within NASS and ERS modelers to facilitate skills enhancement.

* RMA and NASS could collaborate on sharing of data (Recommendation 3-1), on sharing of ideas and methodology for modeling, and on NASS's collection of data from precision agriculture.
  - RMA and NASS could consult concerning access to individually identifiable data and how NASS will protect confidential data (Recommendation 2-5).
  - NASS could collaborate with RMA actuaries and statisticians to learn how they have accounted for heteroscedasticity and incorporated soil variability and geographic differences into their estimates of yield. The intent would be to see whether those methods could improve NASS's approaches to modeling.
  - NASS could consult with RMA and the Approved Insurance Providers (AIPs) on the development of an option for reporting of planted acres, failed acres, and production through use of precision agriculture.
* FSA and NASS could collaborate on sharing data and CLU identification.
  - FSA and NASS could agree on how FSA data will be protected (Recommendation 2-5).
  - NASS could consult with FSA on the determination of CLUs for farms that are not in the FSA data system.
  - NASS could collaborate with FSA concerning definitions of CLUs, both to understand current definitions and procedures and potentially to influence FSA to adapt CLUs to NASS's needs.
  - NASS could consult with FSA on the development of an option for reporting of planted acres, failed acres, and production through use of precision agriculture.

ARS, ERS, and NASS could collaborate on the development and use of a crop model that can provide the most relevant variables for predicting yield of a given crop in a given county (based on soil, precipitation, growing degree days, etc.). The panel views this model as being similar to the Versatile Soil Moisture Budget (VSMB) used by Statistics Canada to prepare agroclimatic data for use as input to Statistics Canada's Integrated Canadian Crop Yield Forecaster (ICCYF) (see Chapter 3). The VSMB is a yield model developed by Agriculture and Agri-Food Canada and modified within Statistics Canada. The model synthesizes soil, weather, and precipitation data from many ground weather stations in Canada. It is used in preparing estimates for crop- and Canadian Census Agriculture Region–specific variables such as growing degree day (GDD) accumulation, growing season precipitation accumulation, water stress indices, and other variables scientifically recognized as being potentially useful for explaining crop growth through the growing season. A model such as the VSMB could prove very helpful to NASS in identifying and preparing relevant independent variables that could be used to improve both remote sensing models and area-level models for yield. NASS likely could obtain, use, and/or adapt such a model from work done by other agencies within USDA, such as ARS.

The panel is not aware of NASS's use of data from ground-based weather stations and ARS flux towers, which represent one approach for acquiring the information in a form that would be most useful as direct input to NASS modeling efforts, both for remote sensing yield indications (where augmenting the current methodology for yield estimation with new input variables may be quite simple) and for including weather- and soil-specific variables in area models.

NASS also will need to collaborate outside of USDA to keep abreast of emerging data sources (Recommendation 3-9)—possibly with Statistics Canada and/or other international agencies on the development of a crop-specific model, described above as a collaboration with ARS, and with independent software vendors—to achieve the goal of developing a precision agriculture reporting option (Recommendations 3-4 and 3-5). NASS also could collaborate with farmer cooperatives to develop new approaches to obtaining data in ways least burdensome to farmers (Recommendation 3-3). More generally, NASS will need to collaborate with farmers as NASS pursues these enhancements (Recommendation 3-6).

Project 2: Cash Rents Model

NASS needs to enhance, document, and develop guidance on how best to use the cash rents model–based estimates in developing official NASS county estimates. The cash rents model developed by Berg and colleagues (2014) integrates survey data from 2 years and was originally developed for use when the Cash Rents Survey was conducted annually. The model-based estimates were provided as input to the ASB process in 2013 and 2014. The 2014 farm bill stated that the survey should be conducted no less frequently than every other year. The survey was not conducted in 2015, but was conducted in 2016 and is planned to be conducted again in 2017. The model was adapted to account for the change in survey frequency, and results were provided to ASB in 2016. Especially for the case in which there is a 2-year gap between surveys, the model would be enhanced by using Bayesian approaches that would support relaxation of the assumption that the survey variances are the same in both years (Recommendation 4-2). NASS modelers also could collaborate with ASB and FSA (users of the estimates) to develop convincing summaries of survey and model uncertainties that would inform decisions about how model outputs (estimates) can best be used by ASB (Recommendation 4-3). NASS also will need to prepare model documentation and provide it on its website (Recommendation 2-6).

NASS will also need to decide how to improve its benchmarking to previously published totals. As noted in Chapter 4, the panel favors including benchmarking as constraints in the modeling process.

Project 3: Development of Models for Crop Statistics

NASS needs to continue working on the development of area models that can be used to synthesize data and prepare accurate estimates of planted acres, harvested acres, and production or yield. This project has two key aspects: developing the models, and building the technical capacity for model development and improvement.

Developing Models

As described in Chapter 2, further development is needed before NASS will have models that can be used to provide county-level estimates suitable for publication after careful review.

As it is now, development of such models would be a focus of the Research Division. Once modelers had identified candidate models, they could work with ASB (the user of the models and their results, and responsible for maintaining the quality of official NASS estimates) and with RMA and FSA (key users of NASS county-level estimates) to reach agreement on the utility of model-based estimates. (Success with the cash rents project described above could provide some guidance on how this can be accomplished.)

The Fay-Herriot type of subarea model currently under consideration by NASS represents an excellent start at model development. However, this type of model considers direct survey estimates to be primary and does not account for the error structure of other potential input variables. Use of a more complete measurement error structure might evolve the current approach into one that is capable of integrating multiple data sources. For more detail, see Appendix C, which also addresses unit-level models, as well as extending models to include space and time.

As summarized in Chapter 2, NASS already has experience with subarea-level modeling (Cruze et al., 2016; Erciulescu at al., 2016) using Bayesian approaches. The panel encourages NASS to continue exploring and extending Bayesian models to incorporate the full measurement error structure, especially for the key data sources (those that provide alternative estimates of one of the key variables). NASS needs to explore new ways of addressing such current issues as skewness and multicolinearity. It could consider including alternative independent variables such as the Normalized Difference Vegetative Index (NDVI), as it is already computed for many counties in support of the current NASS remote sensing estimate of yield. Note that ecological bias can be avoided (under certain assumptions) by including the within-area variance of a variable, as well as its mean, in the model (Wakefield, 2008). Adding spatial dependence is conceptually straightforward using Besag or Leroux spatial formulations (see Appendix C). Spatial random effects can be incorporated to leverage spatial similarity due to unmeasured variables. Stratification can be accounted for by including fixed effects in the model, and cluster sampling by including random effects. (Scott and Smith [1969] is an early reference; relevant references for the non-Gaussian setting and ecological bias are Bradley et al. [2016, 2017]). Computation with the integrated nested Laplace approximation (INLA) (Rue et al., 2009) approach is fast (compared with Markov chain Monte Carol [MCMC], which NASS has been using) and accurate. There is a reliable R implementation of the INLA method, although it is not a standard package.

The Bayesian approach to modeling naturally leads to intuitive measures of uncertainty. The fundamental output of a Bayesian analysis is a multivariate posterior distribution over all unknown quantities in the model. This distribution is typically of high dimension, so summarization is required. In particular, summaries of univariate posterior distributions of quantities of interest may be reported. For example, the posterior median (or posterior mean if the posterior on the quantity of interest is symmetric) may be quoted along with quantiles such as the 2.5 percent and 97.5 percent to give a 95 percent interval. In a Bayesian analysis, the posterior variance is a standard measure of uncertainty when the posterior distribution is normal.

Panel members discussed the importance of including soil productivity, and possibly weather, in models. As described to the panel, the experience of NASS modelers to date has been that such indices when included as independent variables are of marginal significance. Following the lead of RMA, NASS could develop better groupings of counties than Agricultural Statistics Districts (ASDs) and use them to borrow strength in model estimation. Sometimes a county on the border of a state is more like the adjacent county in the next state than it is like the other counties in its ASD. The section in Chapter 3 on RMA's yield modeling addresses some of these points. NASS needs to publish county-level data, and there are legitimate reasons for the aggregation of county estimates within a state to add to previously published state estimates. If alternative grouping of counties leads to measurably more accurate results relative to ASDs and fewer outliers, NASS may choose to benchmark county estimates directly to state totals, computing ASD totals as the sum of benchmarked county estimates.

NASS has been investigating benchmarking options as part of its modeling efforts. The panel agrees that including benchmarking as a constraint in models will ultimately lead to more defensible estimates.

Building Technical Capacity

The second challenge for NASS is maintaining and expanding the technical expertise of NASS staff. In addition to the points discussed below, many of the collaborative efforts described earlier under Project 1 will support this goal by helping staff learn new approaches and incorporate them into NASS processes:

- NASS modelers need to work to understand the ASB process and incorporate relevant features and data in their model development efforts. This would be facilitated by assigning a modeler to be a member of the ASB (Recommendation 2-7).
- NASS needs to enhance its statistical, economic, and geospatial modeling expertise with new hires over time. (The panel recognizes that NASS has been attempting to do this and that it is not easy to identify and recruit highly qualified staff.)
- NASS needs to enhance the capabilities of current staff through ongoing training opportunities in such areas as new modeling approaches, software, and incorporation of model-based approaches that use auxiliary data into traditional survey methods. NASS could consider bringing in experts to provide this training for staff.
- NASS could arrange for ad hoc expert review committees and/or (non–Federal Advisory Committee Act [FACA]) advisory committees such as the American Statistical Association's (ASA) Committee on Energy Statistics, sponsored by the Energy Information Administration, to bring in external expertise and advice.
- NASS technical staff need access to appropriate hardware and software. With regard to hardware, the methods most useful for addressing geospatial issues and modeling require significant storage and speed for computation. Staff also need to have access to the most up-to-date software, especially in a test environment. Ways in which software products that enhance NASS capabilities can be moved into operational environments are needed as well.

Project 4: Geospatial Database

Linking FSA and RMA data (CLUs and RLUs) to the NASS list frame is a high-priority first-stage project. Achieving complete linkage with these administrative sources will be time-consuming, so NASS's work on this effort needs to begin right away, starting with the development of an approach that will build to success. One first step will entail planning how to accommodate the CLU, RLU, and field boundaries derived from precision agriculture, as well as FSA and RMA identifiers (IDs) in ELMO (or in files easily accessed by ELMO). Second, NASS has pursued studies in the past that involved identifying FSA farms that match NASS farms. If information from those studies concerning farms that match is still available, it may be able to serve as a starting point for a matched list. For example, for several years the June Area Survey has made use of FSA CLUs, and remnants of the 2010 record linkage effort in Nebraska may be available. Third, NASS's traditional matching approaches using owner/operator names can provide linkage for the least complicated farms—those for which one NASS farm corresponds to one FSA farm. The Nebraska matching experiment revealed that an "easy" one-to-one exact match existed for almost 50 percent of farms. NASS could develop revised survey forms for matched farms that take advantage of the administrative data to simplify reporting and reduce respondent burden (Recommendation 2-11). With the promise of a simple reporting form, perhaps NASS can work with all respondents to encourage them to identify the FSA IDs with which they are associated. Completing the match though manual reviews is possible, but time-consuming. Expensive, manual efforts at matching will best be dedicated to achieving matches for the largest farms.

While the linkage effort may be complex, maintaining the linked database should be relatively simple. FSA and RMA update their data annually. The panel learned that FSA introduces some changes in CLU definitions over time—possibly because of changes in ownership or land use. At its second meeting, the panel heard from ERS analysts who have accomplished linkages between FSA data and other USDA data and maintained those linkages over time. They estimated that no more than 10 percent of CLUs are changed from one year to the next. NASS needs to plan how to incorporate these changes in CLUs into its databases, preferably using automated methods.

Project 5: Evolving the Role of the ASB and Updating Documentation

Ultimately, multiple data sources will be integrated through the use of formal statistical models that will provide the estimates to be reviewed by ASB. ASB will exercise judgment in deciding when model estimates fail to reflect the impact of unforeseen events (e.g., droughts or hurricanes) or of systemic changes (e.g., rapid farm consolidation) that are not well captured in the models. ASB will make the final determination as to whether any county estimate should not be published. As a means of providing quality control, ASB will drive the feedback loop with analysts by suggesting modifications to improve model performance and interpretation. Documentation of models and reasons for changes from the documented methodology will be available to the public so that increased transparency will bolster confidence in the robust nature of the process.

The panel believes the ASB process for cash rents can begin to evolve toward a more transparent approach now, as summarized above under Project 2 and in Chapter 4. The crop estimates program is more complex at present, and also requires that a candidate model be developed and adopted before serious evolution can begin.

For now, NASS could take the following steps:

- Establish templates for posting estimation methodology on the NASS website. Examples include the Bureau of Labor Statistics, the Bureau of Economic Analysis, and the Census Bureau.
- Develop detail on how outliers are identified and evaluated for whether they can be explained, and why decisions are made to revise them.

- Develop a list of key reasons for making changes, and use it to keep a record of changes made.
- Formulate new publication standards based on Recommendation 2-3, and determine how to apply them now and as models are adopted in the future.
- Determine how the software product DICE—on which ASB currently relies to support its review and evaluation process, including benchmarking of district and county estimates to previously published state totals—should be revised to support new processes.

SECOND-STAGE PROJECTS: PROJECTS OCCURRING ONCE FSA, RMA, AND NASS FARMS HAVE BEEN LINKED AND INCLUDED IN THE LIST FRAME OR OTHER ACCESSIBLE DATABASE

Projects in the second stage of implementation of the panel's vision will depend on progress on first-stage projects described above. Collaborative efforts under first-stage Project 1 will lay the foundation by obtaining access to administrative and auxiliary data and facilitating the development and use of new data sources. Having the CLU as the basic spatial unit and incorporating CLUs for (most) farms in the NASS list frame will have many advantages for NASS, including the following:

- FSA will provide annual updates of owner, operator, and CLU data for all FSA farms in the database, supporting NASS's efforts to keep the list frame up to date.
- Current-year FSA and RMA data on planted and failed acres will be available to improve imputation and estimation for the County Agricultural Production Survey (CAPS) and Acreage, Production, and Stocks (APS) survey system.
- Additional types of models that require use of the geospatial information associated with CLUs or farms (groups of CLUs) can be explored.

By the end of the first-stage phase projects, the cash rents model will have been refined and built into a new and more transparent ASB process, and models that successfully integrate multiple data sources for crop estimates will have been developed and adopted. Second-stage projects include improving surveys forms, improving survey imputation and estimation, developing unit-level models, and potentially redesigning NASS surveys to take advantage of newly available information.

Project 1: Improving Survey Forms

With the availability of linked FSA data, NASS could be working to develop a new Internet survey form for CAPS/APS that would ask matched respondents to provide only the current-year data items that are not available from administrative data. For example, since the crops planted are known, each survey form would specifically request production or yield information only for relevant crops. The intent would be to reduce the burden of responding to NASS surveys (Recommendation 2-11). It is anticipated that collaboration with farmers, farmer cooperatives, FSA, RMA, and software developers that translate precision agriculture data into variables required by NASS will have been fruitful, and that NASS will be able to make arrangements for collecting precision agriculture data as an alternative response mode in CAPS/APS (Recommendation 3-5).

Project 2: Improving Survey Imputation and Estimation

As described in Chapter 2, once a reasonable number of FSA farms and their CLUs have been linked with the NASS list frame, the use of FSA current-year data on planted acres by crop for imputation will be feasible.[1] As time passes and more FSA farms are linked, NASS might realize significant improvements in accuracy due to the inclusion of current-year data on crops planted. As nonsurvey precision agriculture data become available, they may also prove valuable in imputation for nonresponse. The current sample-based data collection and estimation method could be used, but nonresponse on planted and harvested acreage items could be imputed with data from the actual nonresponding farm or a "similar" farm (Recommendation 2-11).

FSA and RMA data both define subpopulations of all farms, and NASS may be able to access data from precision agriculture databases on an ongoing basis for a substantial fraction of the total population. These auxiliary data for these well-defined subpopulations could be used to provide more efficient estimators relative to the current NASS sampling and estimation methods.

Data from precision agriculture databases, FSA program participants, and RMA purchasers of insurance clearly are not probability samples and are not likely to be representative of the population. Thus the data could not be treated as a direct replacement for or even an augmentation of the probability sample. However, methods for making use of data from nonprobability samples can be used in ways that are familiar to those accustomed to probability sampling methodology. All of these methods require that it be possible to link the data from the nonprobability samples to the list frame. With the new linkages established in first-stage projects, the methodologies described in Chapter 2 could be used to improve estimation in CAPS/APS and possibly even the Cash Rents Survey (Recommendation 2-13).

Project 3: Developing Unit-Level Models

Appendix C summarizes modeling strategies, including both area-level models (for counties) in NASS applications and unit-level models (for farms or CLUs). The Battese-Fuller model (Battese and Fuller, 1981), used to develop remote sensing estimates of acreage, and the Battese-Harter-Fuller model (Battese et al., 1988) are examples of unit-level models. NASS could consider models for its basic spatial unit (CLUs) or for farms (collections of CLUs). One disadvantage of farm-level modeling is that farms can be of vastly different sizes. One advantage of unit-level modeling is that it would support incorporation of parcel-specific variables and observations from remote sensing and precision agriculture with the potential to enhance model performance.

Project 4: Redesigning NASS Surveys

Potentially the most important project in the second stage could occur only after substantial linkage had been accomplished, and either production or yield data by crop were

[1] As noted in Chapter 2, the panel was told that crop switching causes considerable concern during imputation. Imputation relies on past-year data for a missing farm to select crops for imputation. However, past-year data may not be related to the crops a farmer plants during the current year.

available from an alternative source (either from RMA in time for NASS to use the data in preparing county-level estimates, or from a sufficiently accurate yield model that had been developed, documented, and accepted and did not rely on current-year survey data). To the extent that NASS can acquire information already provided by farmers to USDA, it should be possible to greatly reduce respondent burden in surveys. Greater reliance on administrative data could allow redirection of scarce survey resources to farms not in FSA or RMA records or to areas where survey response is low. In this situation, NASS could consider revamping its survey program to take advantage of available administrative data and models and to refocus survey efforts on areas or commodities not well covered by alternative sources.

CONCLUSION

The panel's vision for the future of NASS is for 2025. The panel debated a shorter time frame, concluding that a time frame with the greatest feasibility would be preferable. Many of the projects described in this chapter will take time and dedication to complete, and in the interim NASS will necessarily continue with its current ongoing schedule and workload. Roughly speaking, each stage might be viewed as taking approximately 3 years to complete, although individual projects might be completed sooner.

As noted above, attaining this vision will require a focus on achieving results. Most organizations that have tried to implement change have found that support by the most senior leaders is critical to success. Therefore, senior NASS leadership will need to adopt the panel's vision and clarify overall goals, and they will need to establish and monitor projects. The panel proposes that senior leaders who potentially could serve as key champions include the head of ASB, to oversee the Board's evolution to a more transparent and reproducible mode of operation; the director of research, to serve as the champion for model development activities; and the head of the list frame organization within NASS, to oversee the development of a georeferenced list frame. These champions will need to identify others to be part of the process and will need to promote the importance of the changes. They will need to plan to provide resources for continuing operations while simultaneously working to identify, develop, and implement these changes. They will need to lead in the collaboration with other USDA agencies, and identify and involve supportive midlevel managers and champions. They will need to prepare a timeline, scoping the component projects essential to implementing the vision, identifying necessary steps, and monitoring progress. And they will need to celebrate success.[2]

[2]Celebrating success is an important aspect of leading change. See http://www.brendabence.com/media-room/articles/The-Top-10-Reasons.pdf for example.

REFERENCES

Adrian, D.W. (2012). *A Model-Based Approach to Forecasting Corn and Soybean Yields*. Paper presented at the Fourth International Conference on Establishment Surveys, Montréal, Québec, Canada, June 11–14. Available: http://www.amstat.org/meetings/ices/2012/papers/302190.pdf [September 16, 2017].

Atwood, J.A., Shaik, S., and Watts, M.J. (2002). Can normality of yields be assumed for crop insurance? *Canadian Journal of Agricultural Economics*, *50*(2), 177–184.

Atwood, J.A., Shaik, S., and Watts, M.J. (2003). Are crop yields normally distributed? A reexamination. *American Journal of Agricultural Economics*, *85*(4), 888–901.

Babcock, B.A. (2008). Corn belt contributions to the crop insurance industry. *Iowa Ag Review*, *14*(2), 1–3.

Bailey, J.T., and Kott, P.S. (1997). An application of multiple list frame sampling for multi-purpose surveys. In *Proceedings of the Section on Survey Research Methods*. Washington, DC: American Statistical Association. http://ww2.amstat.org/sections/srms/Proceedings/papers/1997_084.pdf [September 16, 2017].

Baker, R., Brick, M., Bates, N., Battaglia, M., Couper, M., Dever, J., Gile, K., and Tourangeau, R. (2017). *Report of the AAPOR Task Force on Non-Probability Sampling*. Available: https://www.aapor.org/AAPOR_Main/media/MainSiteFiles/NPS_TF_Report_Final_7_revised_FNL_6_22_13.pdf [August 29, 2017].

Banerjee, S., Gelfand, A., Finley, A.O., and Sang, H. (2008). Gaussian predictive process models for large spatial data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *70*(4), 825–848.

Barboza, W.J., and Iwig, B. (2010). *Assessing an Administrative Data Source as a Sampling Frame*. Available: http://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/reports/conferences/JSM-2010/jsm2010-paper-Barboza-Iwig.pdf [September 16, 2017].

Battese, G.E., and Fuller, W.A. (1981). Prediction of county crop areas using survey and satellite data. In *Proceedings of the Section on Survey Research Methods*. Washington, DC: American Statistical Association.

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. *Journal of the American Statistical Association*, *83*(401), 28–26.

Bédard, F., and Reichert, G. (2013). *Integrated Crop Yield and Production Forecasting Using Remote Sensing and Agri-Climatic Data*. Analytical Projects Initiatives Final Report. Ottawa, Ontario: Statistics Canada.

Bell, J., and Barboza, W. (2012). *Evaluation of Using CVs as a Publication Standard*. Paper presented at the Fourth International Conference on Establishment Surveys, Montréal, Québec, Canada, June 11–14.

Bellow, M.E. (2007). *Comparison of Methods for Estimating Crop Yield at the County Level*. RDD Research Report RDD-07-05. Washington, DC: U.S. Department of Agriculture, National Agricultural Statistics Service. Available: http://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/reports/Crop_Yield_Research_Report_MEB.pdf [September 16, 2017].

Bellow, M.E., and Lahiri, P. (2012). *Evaluation of Methods for County Level Estimation of Crop Harvested Area That Employ Mixed Models*. Paper presented at the Fourth International Conference on Establishment Surveys, Montréal, Québec, Canada, June 11–14.

Bellow, M.E., Cruze, N., and Erciulescu, A. (2017). Developments in model-based county-level estimation of agricultural cash rental rates." Proceedings of the Section on Survey Research Methods. Joint Statistical Meetings. American Statistical Association.

Berg, E., Cecere, W., and Ghosh, M. (2014). Small area estimation for county-level farmland cash rental rates. *Journal of Survey Statistics and Methodology*, *2*(1), 1–37.

Berrocal, V.J., Gelfand, A.E., and Holland, D.M. (2010). A spatio-temporal downscaler for output from numerical models. *Journal of Agricultural, Biological, and Environmental Statistics*, *15*(2), 176–197.

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B*, *36*(2), 192–236.

Besag, J., York, J., and Mollié, A. (1991). Bayesian image restoration with two applications in spatial statistics. *Annals of the Institute of Statistics and Mathematics*, *43*(1), 1–59.

Bradley, J.R., Holan, S.H., and Wikle, C.K. (2015a). Multivariate spatio-temporal models for high-dimensional areal data with application to longitudinal employer–household dynamics. *Annals of Applied Statistics*, *9*(4), 1761–1791.

Bradley, J.R., Wikle, C.K., and Holan, S.H. (2015b). Spatio-temporal change of support with application to American Community Survey multi-year period estimates. *Stat*, *4*(1), 255–270.

Bradley, J.R., Cressie, N., and Shi, T. (2016). A comparison of spatial predictors when datasets could be very large. *Statistics Surveys*, *10*, 100–131. doi: 10.1214/16-SS115.

Bradley, J.R., Wikle, C.K., and Holan, S.H. (2016). Bayesian spatial change of support for count-valued survey data with application to the American Community Survey. *Journal of the American Statistical Association*, *111*(514), 472–487.

Bradley, J.R., Wikle, C.K., and Holan, S.H. (2017). Regionalization of multiscale spatial processes using a criterion for spatial aggregation error. *Journal of the Royal Statistical Society, Series B*, *79*(3), 815–832.

Bunn, D., and Wright, G. (1991). Interaction of judgmental and statistical forecasting methods: Issues and analysis. *Management Science*, *37*(5), 501–518.

Busselberg, S. (2011). *Bridging Livestock Survey Results to Published Estimates through State–Space Models: A Time Series Approach*. Available: https://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/reports/conferences/JSM-2011/JSM-2011-Busselberg.pdf [September 16, 2017].

Cecere, W., Berg, E., and Ghosh, M. (2012). *A Problem in Small Area Estimation: Cash Rental Rates.* Paper presented at the Fourth International Conference on Establishment Surveys, Montréal, Québec, Canada, June 11–14. Available: www.amstat.org/meetings/ices/2012/papers/301998.pdf [September 16, 2017].

Chipanshi, A., Zhang, Y., Kouadio, L., Newlands, N., Davidson, A., Hill, H., Warren, R., Qian, B., Daneshfar, B., Bedard, F., and Reichert, G. (2015). Evaluation of the Integrated

Canadian Crop Yield Forecaster (ICCYF) model for in-season prediction of crop yield across the Canadian agricultural landscape. *Agricultural and Forest Meteorology*, *206*, 137–150. doi: 10.1016/j.agrformet.2015.03.007.

Coble, K.H., Knight, T.O., Goodwin, B.K., Miller, M.F., and Rejesus, R.M. (2010). *A Comprehensive Review of the RMA APH and COMBO Rating Methodology*. Available: http://www.rma.usda.gov/pubs/2009/comprehensivereview.pdf [March 15, 2010].

Congressional Research Service. (2017). *NASS and U.S. Crop Production Forecasts: Methods and Issues*. Report 7-5700. Available: https://fas.org/sgp/crs/misc/R44814.pdf [July 11, 2017].

Cressie, N.A.C. (1993). *Statistics for Spatial Data*. New York: Wiley.

Cressie, N.A.C., and Johannesson, G. (2008). Fixed rank kriging for very large spatial data sets. *Journal of the Royal Statistical Society, Series B (Statistical Methodology), 70*(1), 209–226. Available: http://www.jstor.org/stable/20203819 [September 16, 2017].

Cressie, N.A.C., and Wikle, C.K. (2011). *Statistics for Spatio-Temporal Data*. Hoboken, NJ: John Wiley & Sons, Inc.

Crouse, C. (2000). *Evaluation of the Use of Spatial Modeling to Improve County Yield Estimates*. Research Report RD-00-05. Washington, DC: U.S. Department of Agriculture, National Agricultural Statistics Service.

Cruze, N.B. (2015a). *A Bayesian Hierarchical Model for Combining Several Crop Yield Indications*. Briefing presented at Federal Committee on Statistical Methodology Research Conference, Washington, DC, December 1. Available: http://fcsm.sites.usa.gov/files/2016/03/C3_Cruze_2015FCSM.pdf [September 16, 2017].

Cruze, N.B. (2015b). *Integrating Survey Data with Auxiliary Sources of Information to Estimate Crop Yields*. Proceedings of the Survey Research Methods Section of the Joint Statistical Meetings, Seattle, WA, August 13. Available: https://ww2.amstat.org/meetings/jsm/2015/onlineprogram/AbstractDetails.cfm?abstractid=314604 [September 16, 2017].

Cruze, N.B., Erciulesscu, A.L., Nandram, B., Barboza, W.J., and Young, L.J. (2016). *Developments in Model-Based Estimation of County-Level Agricultural Estimates*. Paper presented at the Fifth International Conference on Establishment Surveys, Geneva, Switzerland, June 20–23. Available: http://ww2.amstat.org/meetings/ices/2016/proceedings/131_ices15Final00229.pdf [September 16, 2017].

Diggle, P.J., Moraga, P., Rowlingson, B., and Taylor, B.M. (2013). Spatial and spatio-temporal log-Gaussian Cox processes: Extending the geostatistical paradigm. *Statistical Science*, *28*(4), 542–563.

Dorfman, A. (2017). *Towards a Routine External Evaluation Protocol for Small Area Estimation*. Submitted to International Statistical Review.

Dow AgroSciences. (2016). Dow AgroSciences puts growers at the center of a new precision agronomy program. *Dow AgroSciences Newsroom*, March 3. Available: http://www.dowagro.com/en-US/usag/News%20and%20Resources/NewsRoom/2016/March/03/Dow%20AgroSciences%20Puts%20Growers%20at%20the%20Center%20of%20a%20New%20Precision%20Agronomy%20Program [July 12, 2017].

Erciulescu, A.L., Cruze, N.B., and Nandram, B. (2016). Model-based county-level crop estimates incorporating auxiliary sources of information. In *Proceedings of the Survey*

Research Methods Section of the Joint Statistical Meetings. Alexandria, VA: American Statistical Association.

Fay, R.E., and Herriot, R.A. (1979). An application of James–Stein procedures to census data. *Journal of the American Statistical Association*, *74*(366), 269–277.

Franco, C., and Bell, W.R. (2016). Borrowing information over time in binomial/logit normal models for small area estimation. In M. Hidiroglou and W. Okrasa (Eds.), *Joint Issue of Statistics in Transition New Series and Survey Methodology Small Area Estimation 2014*, *16*(4), 563–584. Available: http://stat.gov.pl/en/sit-en/joint-issue-part-i-sae-poznan-2014 [September 16, 2017].

Fuentes, M., and Raftery, A.E. (2005). Model evaluation and spatial interpolation by Bayesian combination of observations with outputs from numerical models. *Biometrics*, *61*(1), 36–45.

Fuller, W.A., and Goyeneche, J.J. (1998, June). *Estimation of the State Variance Component*. Working Paper. Ames, IA: Iowa State University, Statistical Laboratory.

Gallagher, P. (1987). U.S. soybean yields: Estimation and forecasting with nonsymmetric disturbances. *American Journal of Agricultural Economics*, *69*(4), 796–803.

Gelfand, A.E. (2010). Misaligned spatial data. In A.E. Gelfand, P.J. Diggle, M. Fuentes, and P. Guttorp (Eds.), *Handbook of Spatial Statistics* (pp. 517–539). Boca Raton, FL: CRC Press.

Gershunskaya, J. (2012). Estimation for detailed publication levels in the current employment statistics survey. In *Proceedings of the American Statistical Association, Survey Research Methods Section*. Alexandria, VA: American Statistical Association.

Gershunskaya, J., and Lahiri, P. (2010). Robust small area estimation using a mixture model. In *Proceedings of the American Statistical Association, Survey Research Methods Section*. Alexandria, VA: American Statistical Association. Available: http://www.amstat.org/sections/srms/proceedings/y2010/files/307425_58536.pdf [September 16, 2017].

Glauber, J.W. (2004). Crop insurance reconsidered. *American Journal of Agricultural Economics*, *86*(5), 1179–1195.

Good, D.L., and Irwin, S.H. (2016). Using FSA acreage data to project NASS January planted acreage estimates for corn and soybeans. *Farmdoc Daily*, *6*, 17. Available: http://farmdocdaily.illinois.edu/2016/01/using-fsa-acreage-data-project-nass-january-planted.html [July 12, 2017].

Goodwin, B.K., and Ker, A.P. (1998). Nonparametric estimation of crop yield distributions: Implications for rating group-risk crop insurance contracts. *American Journal of Agricultural Economics*, *80*(1), 139–153.

Gotway, C., and Young, L. (2002). Combining incompatible spatial data. *Journal of the American Statistical Association*, *97*(458), 632–648.

Graham, M., and Iwig, W. (1996). County estimation of crop acreage using satellite data. In W. Schaible (Ed.), *Indirect Estimators in U.S. Federal Programs* (Chapter 6). New York: Springer. Available: https://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/GIS_Reports/The%20National%20Agricultural%20Statistics%20Service%20County%20Estimates%20Program.pdf [July 11, 2017].

Haining, R.P. (2003). *Spatial Data Analysis: Theory and Practice*. New York: Cambridge University Press.

Harri, A., Erdem, C., Coble, K.H., and Knight, T.O. (2009). Crop yield distributions: A reconciliation of previous research and statistical tests for normality. *Review of Agricultural Economics*, *31*(1), 163–182.

Hatfield, J. (1983). Remote sensing estimators of potential and actual crop yield. *Remote Sensing of Environment*, *13*(4), 301–311.

Hazell, P.B.R. (1984). Sources of increased instability in Indian and U.S. cereal production. *American Journal of Agricultural Economics*, *66*(3), 302–311. doi: 10.2307/1240797.

Hodges, J.S., and Reich, B.J. (2010). Adding spatially-correlated errors can mess up the fixed effect you love. *The American Statistician*, *64*(4), 325–334.

Hurst, B. (2015). *Big Data and Agriculture: Innovation and Implications*. Statement of the American Farm Bureau Federation to the House Committee on Agriculture. Available: https://agriculture.house.gov/uploadedfiles/10.28.15_hurst_testimony.pdf [July 12, 2017].

Iwig, W. (1996). The National Agricultural Statistics Service County Estimates Program. In W. Schaible (Ed.), *Indirect Estimators in U.S. Federal Programs* (Chapter 7) (pp. 129–144). New York: Springer. Available: https://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/GIS_Reports/The%20National%20Agricultural%20Statistics%20Service%20County%20Estimates%20Program.pdf [July 11, 2017].

JASON Deliberative Study. (2016). *New Techniques for Evaluation of Crop Production*. JST-16-Task-006. McLean, VA: The MITRE Corporation.

Johansson, R., Effland, A., and Coble, K. (2017). Falling response rates to USDA crop surveys: Why it matters. *Farmdoc Daily,* January 19. Available: http://farmdocdaily.illinois.edu/2017/01/falling-response-rates-to-usda-crop-surveys.html [July 11, 2017].

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 Sensing of Environment 141*, 116–128. doi: 10.1016/j.rse.2013.10.027.

Johnson, D.M. (2016). A comprehensive assessment of the correlations between field crop yields and commonly used MODIS products. *International Journal of Applied Earth Observation and Geoinformation, 52*, 65–81. doi: 10.1016/j.jag.2016.05.010.

Kalton, G. (2002). Models in the practice of survey sampling (revisited.). *Journal of Official Statistics, 18*(2), 129–154.

Kott, P.S. (1989). Assessing linearization variance estimators. In *Proceedings of the American Statistical Association, Survey Research Methods Section*. Alexandria, VA: American Statistical Association. Available: http://ww2.amstat.org/sections/SRMS/Proceedings/papers/1989_030.pdf [September 16, 2017].

Krause, P., Boyle, D.P., and Base, F. (2005). Comparison of different efficiency criteria for hydrological model assessment. Adv. Geosci. 5, 89-97.

Leroux, B.G., Lei, X., and Breslow, N. (1999). Estimation of disease rates in small areas: A new mixed model for spatial dependence. In M.E. Halloran and D.A. Berry (Eds.), *Statistical Models in Epidemiology, the Environment and Clinical Trials* (pp. 179–192). New York: Springer.

Li, Y., Brown, P., Gesink, D.C., and Rue, H. (2012). Log Gaussian Cox processes and spatially aggregated disease incidence data. *Statistical Methods in Medical Research*, *21*(5), 479–507.

Lindgren, F., Rue, H., and Linstrom, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic differential equation approach (with discussion). *Journal of the Royal Statistical Society, Series B*, *73*(4), 423–498.

Liu, B., Stokes, S.L., Topping, T., and Stunz, G. (2017). Estimation of total from a population of unknown size and application to estimating recreational red snapper catch in Texas. *Journal of Survey Statistics and Methodology*, *5*(3), 350–371. doi: 10.1093/jssam/smx006.

Lokupitiya, E., Breidt, F.J., Williams, S., and Keith, P. (2007). Deriving comprehensive county-level crop yield and area data for U.S. cropland. *Agronomy Journal, 99*, 673–681. doi: 10.2134/agronj2006.0143.

McCarl, B.A., Villavicencio, X., and Wu, X. (2008). Climate change and future analysis: Is stationarity dying? *American Journal of Agricultural Economics*, *90*(5), 1241–1247.

Mercer, L., Wakefield, J., Pantazis, A., Lutambi, A., Masanja, H., and Clark, S. (2015). Small area estimation of child mortality in the absence of vital registration. *The Annals of Applied Statistics*, *9*(4), 1889–1905.

Nandram, B., Berg, E., and Barboza, W. (2014). A hierarchical Bayesian model for forecasting state-level yield. *Environmental and Ecological Statistics*, *21*(3), 507–530.

National Academy of Sciences, Engineering, and Medicine. (2017). *Innovations in Federal Statistics: Combining Data Sources While Protecting Privacy*. Panel on Improving Federal Statistics for Policy and Social Science Research Using Multiple Data Sources and State of the Art Estimation Methods. R.M. Groves and B.A. Harris-Kojetin (Eds.). Committee on National Statistics, Division of Behavioral and Social Sciences and Education. Washington, DC: The National Academies Press.

National Research Council. (1980). *Estimating Population and Income of Small Areas*. Panel on Small-Area Estimates of Population and Income. Committee on National Statistics, Assembly of Behavioral and Social Sciences. Washington, DC: National Academy Press.

National Research Council. (1991a). *Improving Information for Social Policy Decisions—The Uses of Microsimulation Modeling: Volume I, Review and Recommendations*. Panel to Evaluate Microsimulation Models for Social Welfare Programs. C.F. Citro and E.A. Hanuschek (Eds.). Committee on National Statistics, Commission on Behavioral and Social Sciences and Education. Washington, DC: National Academy Press.

National Research Council. (1991b). *Improving Information for Social Policy Decisions—The Uses of Microsimulation Modeling: Volume II, Technical Papers*. Panel to Evaluate Microsimulation Models for Social Welfare Programs. C.F. Citro and E.A. Hanuschek (Eds.). Committee on National Statistics, Commission on Behavioral and Social Sciences and Education. Washington, DC: National Academy Press.

National Research Council. (1997). *Assessing Policies for Retirement Income: Needs for Data, Research and Models*. Panel on Retirement Income Modeling. C.F. Citro and E.A. Hanuschek (Eds.). Committee on National Statistics, Commission on Behavioral and Social Sciences and Education. Washington, DC: National Academy Press.

National Research Council. (2000a). *Small-Area Estimates of School-Age Children in Poverty: Evaluation of the Current Methodology*. Panel on Estimates of Poverty for Small Geographic Areas. C.F. Citro and G. Kalton (Eds.). Committee on National Statistics, Commission on Behavioral and Social Sciences and Education. Washington, DC: National Academy Press.

National Research Council. (2000b). *Small-Area Estimates of School-Age Children in Poverty: Priorities for 2000 and Beyond*. Panel on Estimates of Poverty for Small Geographic Areas. C.F. Citro and G. Kalton (Eds.). Committee on National Statistics, Commission on Behavioral and Social Sciences and Education. Washington, DC: National Academy Press.

National Research Council. (2008). *Understanding American Agriculture: Challenges for the Agricultural Resource Management Survey*. Panel to Review USDA's Agricultural Resource Management Survey, Committee on National Statistics, Division of Behavioral and Social Sciences and Education. Washington, DC: The National Academies Press.

National Research Council. (2013). *Nonresponse in Social Science Surveys: A Research Agenda.* Panel on a Research Agenda for the Future of Social Science Data Collection. R. Tourangeau and T.J. Plewes (Eds.). Committee on National Statistics, Division of Behavioral and Social Sciences and Education. Washington, DC: The National Academies Press.

Neill, J. (2011). FSA programs explained: Reconstitution. *Tri-State Livestock News*, June 3. Available: http://www.tsln.com/news/fsa-programs-explained-reconstitution [September 16, 2017].

Newlands, N., Zamar, D., Kouadio, L., Zhang, Y., Chipanshi, A., Potgeiter, A., Toure, S., and Hill, H. (2014). An integrated, probabilistic model for improved seasonal forecasting of agricultural crop yield under environmental uncertainty. *Frontiers in Environmental Science*, *2*(17). doi: 10.3389/fenvs.2014.00017.

Nickerson, C., Morehart, M., Kuethe, T., Beckman, J., Ifft, J., and Williams, R. (2012). *Trends in U.S. Farmland Values and Ownership.* Economic Information Bulletin 92. Washington, DC: U.S. Department of Agriculture, Economic Research Service. Available: https://www.ers.usda.gov/webdocs/publications/44656/16748_eib92_2_.pdf?v=41055 [July 2017].

Nychka, D., Bandyopadhyay, S., Hammerling, D., Lindgren, F., and Sain, S. (2015). A multiresolution Gaussian process model for the analysis of large spatial datasets. *Journal of Computational and Graphical Statistics*, *24*(2), 579–599.

Pfeffermann, D. (2013). New important developments in small area estimation. *Statistical Science*, *28*(1), 40–68.

Porter, A.T., Wikle, C.K., and Holan, S.H. (2015). Small area estimation via multivariate Fay-Herriot models with latent spatial dependence. *Australian & New Zealand Journal of Statistics*, *57*, 15–29.

Reichart, G., Bédard, F., Mohl, C., Benjamin, W., Dongmo Jiongo, V., Chipanshi, A., and Zhang, Y. (2016). *Crop Yield Modelling Using Remote Sensing, Agroclimatic Data and Statistical Survey Data*. Briefing presented at Seventh International Conference on Agricultural Statistics, Rome, Italy, October 26–28. Slides Available: https://www.dropbox.com/s/v5j50nqenuikciy/CROP_YIELD_MODELLING_USING_REMOTE_SENSING_AGROCLIMATIC_DATA_AND_STATISTICAL_SURVEY_DATA_Reichert_Statistics_Canada_ICAS_VII.pptx?dl=0 . [Sept 26, 2017}.

Riebler, A., Sorbye, S., Simpson, D., and Rue, H. (2016). An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical Methods in Medical Research*, *25*(4), 1145–1165.

Rue, H., and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. Boca Raton, FL: Chapman and Hall/CRC Press.

Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society, Series B*, *71*(2), 319–392.

Schaible, W.L. (Ed.). (1996). *Indirect Estimators in U.S. Federal Programs*. New York: Springer.

Schimmelpfennig, D. (2016). *Precision Agriculture Technologies and Factors Affecting Their Adoption*. Washington, DC: U.S. Department of Agriculture, Economic Research Service. Available: https://www.ers.usda.gov/amber-waves/2016/december/precision-agriculture-technologies-and-factors-affecting-their-adoption [June 2016].

Schleusener, M. (2016). NASS surveys have direct impact on critical farm programs. *Farmdoc Daily*, *6*, 222. Available: http://farmdocdaily.illinois.edu/2016/11/nass-surveys-impact-critical-farm-programs.html [July 11, 2017].

Scott, A., and Smith, T.M. (1969). Estimation in multi-stage surveys. *Journal of the American Statistical Association*, *64*(327), 830–840.

Stein, M.L. (1999), Interpolation of Spatial Data: Some Theory for Kriging, Springer.

Szulczewski, W., Zyromski, A., Biniak-Pierog, M. (2012). New approach in modeling spring wheat yielding based on dry periods. *Agric. Water Manage*, 103, 105–113.

Tang, M., Slud, E., and Pfeiffer, R. (2014). Goodness of fit tests for linear mixed models. *Journal of Multivariate Analysis*, *130*, 176–193. doi: 10.1016/j.jmva.2014.03.012.

Taylor, B.M., Davies, T.M., Rowlingson, B., and Diggle, P.J. (2015). An R package for inference with spatial and spatio-temporal log-Gaussian Cox processes. *Journal of Statistical Software*, *63*, 1–29.

Torabi, M., and Rao, J.N.K. (2014), On Small Area Estimation under a Sub-Area Model, Journal of Multivariate Analysis, 127, 36–55

Traxler, G., Falck-Zepeda, J., Ortiz-Monasterio, J.I., and Sayre, K. (1995). Production risk and the evolution of varietal technology. *American Journal of Agricultural Economics*, *77*(1), 1–7.

U.S. Department of Agriculture. (2016). *Requirement to Report Field Location on the Acreage Report*. Bulletin No. MGR-16-005. Available: https://www.rma.usda.gov/bulletins/managers/2016/mgr-16-005.pdf [July 12, 2017].

U.S. Department of Agriculture-National Agricultural Statistics Service. (2007). *Safeguarding America's Agricultural Statistics: A Century of Successful and Secure Procedures, 1905–2005*. Available: https://www.nass.usda.gov/About_NASS/pdf/asb_historical.pdf [July 11, 2017].

U.S. Department of Agriculture-Natural Resources Conservation Service. (2012). *User Guide for the National Commodity Crop Productivity Index (NCCPI)*. Version 2.0. Washington, DC: Author.

U.S. Department of Agriculture-National Agricultural Statistics Service. (2014). *Cash Rents Methodology and Quality Measures*. Available: https://www.nass.usda.gov/Publications/Methodology_and_Data_Quality/Cash_Rents/08_2014/rentqm14.pdf [August 26 2017].

U.S. Department of Agriculture-National Agricultural Statistics Service. (2016). *Cash Rents Methodology and Quality Measures.* Available: https://www.nass.usda.gov/Publications/Methodology_and_Data_Quality/Cash_Rents/08_2016/rentqm16.pdf [August 26 2017].

U.S. Office of Management and Budget. (2006). *Standards and Guidelines for Statistical Surveys*. Available: https://obamawhitehouse.archives.gov/sites/default/files/omb/inforeg/statpolicy/standards_stat_surveys.pdf [July 11, 2017].

Wakefield, J. (2008). Ecologic studies revisited. *Annual Review of Public Health*, *29*, 75–90.

Walker, G., and Sigman, R. (1982). *The Use of LANDSAT for County Estimates of Crop Area: Evaluation of the Huddleston–Ray and Battese–Fuller Estimators*. Available: https://www.nass.usda.gov/Education_and_Outreach/Reports,_Presentations_and_Conferences/GIS_Reports/The%20Use%20of%20LANDSAT%20for%20County%20Estimates%20of%20Crop%20Areas%20Evaluation.pdf [September 16, 2017].

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. *Journal of Agricultural, Biological, and Environmental Statistics*, *17*(1), 84–106.

Williams, M. (2013). *Small Area Modeling of County Estimates for Corn and Soybean Yields in the U.S.* Paper presented at the Federal Committee on Statistical Methodology (FCSM) Research Conference, Washington, DC, November 4–6. Available: http://fcsm.sites.usa.gov/files/2014/05/C2_Williams_2013FCSM.pdf [September 16, 2017].

Wood, J.D., Griffis, T.J., Baker, J.M., Frankenburg, C., Verma, M., and Yuen, K. (2017). Multiscale analyses of solar-induced florescence and gross primary production. *Geophysical Research Letters*, *44*(1), 533–541. Available: https://www.ars.usda.gov/research/publications/publication/?seqNo115=331937 [July 12, 2017].

Yang, S.-R., Koo, W.W., and Wilson, W.W. (1992). Heteroskedasticity in crop yield models. *Journal of Agriculture and Resource Economics*, *17*(1), 103–109.

You, Y., and Zhou, Q.M. (2011). Hierarchical Bayes small area estimation under a spatial model with application to health survey data. *Survey Methodology, 37*, 25–37.

Zhu, Y., Goodwin, B.K., and Ghosh, S. (2008). *Time-Varying Yield Distributions and the Implications for Crop Insurance Pricing*. Paper presented at the NC State Agricultural Economics Workshop, Raleigh, NC, May 6.

Appendix A

NASS Country-Level Survey Programs

COUNTY CROP ESTIMATES[1]

The quarterly crops Acreage, Production, and Stocks (APS) surveys collect data on crop acreage, yield, and production by commodity and quantities of grain and oilseeds stored on farms used to develop state- and national-level estimates. Since 2011, the fully implemented annual County Agricultural Production Survey (CAPS) has collected data on crop acreage, yield, and production by commodity from a supplemental probability sample of farms; these data are designed to be used with the September and December APS list samples to produce county-level estimates.

The annual APS cycle begins with the June Area Survey (JAS), an area sample described in more detail in the next section. List samples are conducted in June, September, December, and March, with the survey content differing each quarter to capture the seasonality of agriculture: the June survey collects data on planted acres for spring-planted crops and acres harvested and to be harvested for spring crops and winter wheat; the September survey collects data on final harvested acres and production (or yield) of small grains; the December survey collects data on seeded acres of winter wheat (new crop) and final harvested acres and production (or yield) for row crops; and the March survey collects data on winter wheat acres to be harvested as grain, in addition to planting intentions for the spring-planted crops. Final state-level estimates of planted and harvested area, production, and yield for small grains are published in the *Small Grains Annual Summary* each September; final national- and state-level estimates of acreage, production, and yield for major row crops are published in the *Crop Production Annual Summary* each January.

NASS's county crop estimates program is defined jointly by NASS, the U.S. Department of Agriculture's (USDA) Risk Management Agency (RMA), and USDA's Farm Service Agency (FSA). Forty-three states are partners in the process and may add commodities to the program to cover special needs of local cooperators provided external funding. Currently, small grain county estimates are published for 37 states, while row crop county-level estimates are published for 43 states. As of 2016, the list of federal program commodities included

- barley;
- dry edible beans;
- corn—for grain and for silage;
- cotton—upland and pima;
- flaxseed;

[1]This appendix draws heavily on information provided on the NASS website as well as information provided to the panel at its meetings by NASS staff. This appendix has been fact-checked by NASS.

- hay—alfalfa and other;
- oats;
- peanuts;
- potatoes;
- rice;
- sorghum;
- soybeans;
- sugarcane for sugar;
- sugarbeets;
- sunflower;
- tobacco; and
- wheat—Durum, other spring, and winter.

Survey Timeline

The data collection period for each APS survey lasts roughly 15 days, beginning near the first of the month in which the survey is conducted. The first of the month serves as a reference date for all questions regarding stocks; acreage and production data are reported based on the date of the interview. Thus, interviews for the September and December APS surveys capture end-of-season acreage and production data for small grains and row crops, respectively.

Like the APS surveys, CAPS collects data shortly after harvest so that final production and yield are known, yet close enough to harvest so that memory bias is not an issue. Each CAPS survey (small grains and row crops) is administered across two tiers or groups of states, with data for the first tier being collected earlier than data for the second tier. The first-tier states are southern states whose harvest period is earlier than that of northern states. Data collection for Small Grains CAPS is completed by mid-October (subsequent to the publication of state-level estimates in the *Small Grains Annual Summary*). Row Crops CAPS also is split into two tiers based on the same rationale. Data collection for Row Crops CAPS is completed by mid-January, but only after the release of state-level estimates in the *Crop Production Annual Summary*. County-level estimates of acreage, yields, and production are released annually. County-level estimates of acreage and production of small grains are released in mid-December. Federal row crop estimates are published in late February through October, depending on the commodity.

Sampling

The target population for the Row Crops APS is all agricultural operators with cropland and/or storage capacity. NASS uses the dual-frame approach noted above, consisting of four APS list frame samples conducted in June, March, September, and December and an area frame sample of operators that are not on the list (NOL) of operators eligible to be selected from the APS sampling population. The target population for CAPS—a list-only survey—is operators with cropland and/or storage capacity in the 43 collaborating states. The NASS list frame includes all known agricultural establishments. The list frame for CAPS consists of those NASS list frame records with positive planted acres or storage capacity of the desired commodities in the previous year.

The Row Crops APS and CAPS list frame samples are selected using a multivariate probability proportional to size (MPPS) sampling scheme in which the measure of “size” is determined by more than one item (see Bailey and Kott, 1997). The MPPS design allows target sample sizes for the commodities of interest to be set at the county level. The probability of selection is determined by taking the maximum (over all targeted commodities) of products involving

- the targeted sample size for each commodity, and
- the sampling unit's size with respect to each targeted commodity—a proportion based on that unit's planted area relative to the county's planted area for each commodity.

Calibration of sampling weights to list frame totals is performed.

The sampled records from the list portion of the Row Crops APS are used in the production of state- and national-level estimates, and they are used again in conjunction with CAPS records to produce county-level estimates. Hence each sample record has two different MPPS weights. The first is used in the estimation process for the Row Crops APS surveys at the state and national levels. The second is used to produce estimates for the combined Row Crops APS and CAPS surveys at the county level (both in composite estimators that use the JAS data); effectively, the sampled APS and CAPS records from the list frame are reweighted to represent a single probability sample.

The target coefficient of variation (CV) for acreage is 1–4 percent at the national level and 5–10 percent at the state level. As of 2016, the list-only sample sizes were approximately 59,000 for the September APS survey and 75,000 for the December APS survey. The area NOL sample provided an additional 6,000 records in both September and December to account for undercoverage in the state- and national-level estimates. The sample size was 80,000 for Small Grains CAPS and 171,000 for Row Crops CAPS. When the respective APS and CAPS list samples are pooled, county-level estimates for small grains are based on a sample size of 139,000, and county-level estimates for row crops are based on a sample size of 246,000.

Data Collection

Sampled farms receive a presurvey letter explaining the survey and indicating that respondents will be contacted for survey purposes only. The letter provides the questions that will be asked to allow respondents to prepare in advance, and also provides a pass code they can use to complete the survey online. All modes of data collection are utilized for the September and December Row Crops APS surveys as well as the Small Grains CAPS and Row Crops CAPS. Regional Field Offices (RFOs) have the option of conducting a mail-out/mail-back phase. Most of the data are collected through computer-assisted telephone interviews (CATIs) administered by individual RFOs and Data Collection Centers. Limited personal interviewing is done, generally for large operations or those with special handling arrangements.

Before CAPS data collection is complete, an adaptive data collection strategy is used to prioritize counties for nonresponse follow-up to increase commodity coverage and meet publication standards. There are four priority levels based on individual county crop profiles. The first priority is a county with a crop that is close to meeting publication standards. Second is a county for which the publication standard has not been met for at least one crop but for which estimates were published the previous year. The third priority is remaining counties that do not meet publication standards for at least one crop. Finally, the fourth priority is all other counties that meet publication standards. For this last group, the county is still available for nonresponse follow-up, but the remaining nonresponding farms are at the bottom of the call scheduler's list.

Nonresponse Adjustment

Reweighting is used to account for nonresponse. Nonresponse weighting groups are based on operation size and type as well as Agricultural Statistics District (ASD), a geographically defined grouping of neighboring counties within a state. The nonresponse weighting groups for the Row Crops APS and CAPS samples are based on the control data items for total cropland; on-farm grain storage capacity; rice acreage (rice-producing states); potato acreage (potato-producing states); and some rare or specialty crops, which depend on the state. These nonresponse weighting groups ensure that operation size and location are taken into account during reweighting. In each nonresponse weighting group, the adjustment is calculated by summing the weights for all sampled records and dividing by the sum of the weights from the completed records. This ratio is applied to the weights of the completed records and assumes that the data for the nonrespondents are similar to the data for the respondents. No item-level imputation is performed.

Estimators

The list-only APS surveys and NOL component obtained from the JAS are used together to estimate national- and state-level acreage, production, and yield. Data from the list-only portions of the APS surveys and CAPS are combined to produce ASD- and county-level estimates. The small grains (wheat, barley, and oats) county-level estimates are direct expansions utilizing the combined Small Grains CAPS and September Row Crops APS list records. Similarly, the county-level estimates for row crops (e.g., corn, soybeans, and cotton) are direct expansions of the combined Row Crops CAPS and December Row Crops APS list records.

Explicit formulae used by NASS may be found in Kott (1989). Two kinds of estimators are used for county-level indications: direct expansions and ratio estimators. Direct expansions, or Horwitz-Thompson estimators, are used to estimate such totals as planted and harvested acres and production. Reweighted direct expansions are calculated by summing the reported commodity values multiplied by the nonresponse-adjusted sample weights in each nonresponse weighting group. All reweighted direct expansions are computed at the state, district, and county levels.

The yield ratio estimator takes the form of a ratio of two reweighted direct expansion estimates—production and harvested acres—computed as described above. Other survey ratio indications are commodity planted acres to total cropland acres and commodity harvested acres to planted acres. In addition, ratio estimators are used for all across-survey ratios (e.g., current-year to previous-year planted acreages). For the survey-to-survey ratios, both the current and previous survey data must be reported (positive or valid zero) to be included in the ratio. If either of these components is not reported, the sampling unit is excluded from the estimate, and the weights of the completed records are adjusted accordingly. All ratio indications are calculated at the state, district, and county levels.

Variances for all total indications are calculated using the delete-a-group jackknife method with 15 replicate groups. Variances for ratio indications are constructed using a second-order Taylor series expansion for the ratio. The variance–covariance structures for the numerator and denominator (totals) are constructed using the delete-a-group jackknife with 15 replicated weights (Kott, 1989).

The Agricultural Statistics Board (ASB) Process for County Estimates

RFO statisticians collect survey data, perform data editing, and conduct further data analysis. Estimates established by RFO staff and appointed headquarters staff are submitted for review to ASB. Official county-level estimates are set by incorporating a variety of input sources, including direct survey estimates of totals and ratios, FSA data on planted acres, a variant of the famed Battese-Harter-Fuller small-area model of planted acreage (Battese et al., 1988), RMA data on failed acres (to improve harvested area totals), and remotely sensed estimates for production and yield for corn and soybeans in selected states. Past official estimates also are reviewed. From these inputs, NASS constructs composite estimates for its planted area, harvested area, and production totals, which serve as the starting point for setting official estimates. Prescribed initial weights are provided by ASB, although these weights may be adjusted as necessary. The composite estimate at the ASD level is ratio benchmarked to the previously published state totals. NASS rounding rules are enforced, generating official ASD estimates. Subsequently, the composite estimates at the county level are ratio adjusted to the rounded ASD estimates. Enforcement of NASS rounding rules at the county level generates the official county-level totals. All official estimates of yield are derived as the ratio of corresponding final estimates of the production and harvested area totals.

ASB reviews estimates established by RFO staff and appointed headquarters staff. The RFO provides justification to headquarters when recommended estimates deviate from survey results. ASB members review all recommended state-, district-, and county-level estimates for accuracy and consistency across state boundaries and verify that proper procedures were followed. Through this review process, ASB has final approval of all official crop estimates nationwide.

NASS verifies that estimates meet publication standards that must be met before any official estimate can be published. First, if the combined sample sizes realized for the CAPS and APS surveys include at least 30 positive reports of crop production, the county-level (or ASD-level) official estimate is suitable for publication. In counties with fewer than 30 but at least 3 positive reports of crop production, a 25 percent coverage condition is verified; these county estimates may still be published provided the corresponding *unweighted* harvested acreage reports account for at least 25 percent of the official estimate of harvested area set by ASB. Counties with fewer than 3 positive reports of crop production are automatically suppressed. Complementary suppressions may be necessary to avoid disclosing suppressed county estimates. These standards are verified independently for every commodity.

Official estimates are open to revision on a preannounced schedule only if new information becomes available. If changes are made to the state-level official estimates during the normal annual revision period (timing varies by commodity), the county-level data are revised to ensure that county- and ASD-level estimates continue to sum to state-level estimates. These previous-year revisions are released at the same time that the data for the current year are published.

The panel was told that small grain county estimates were published for about 17–53 percent of eligible counties in 2014. The percentage of U.S. total production covered by these published individual county-level estimates of small grains ranged from 52 percent to 82 percent, depending on the commodity. County-level estimates of row crops were published for 29–68 percent of all eligible counties. The percentage of U.S. total production covered by these published individual county-level estimates ranged from 73 percent to 96 percent, depending on the commodity.

CASH RENTS SURVEY PROGRAM

Since 2008, state-level estimates for cash rents have been set based on a combination of the Cash Rents Survey and JAS in the years in which the Cash Rent Survey is conducted. In years when the Cash Rents Survey is not conducted, state-level estimates are set using the JAS. Estimates are published in the first week of August for all states except Alaska and Hawaii. District- and county-level cash rents estimates are published the second week of September, no less frequently than every other year as mandated by the 2014 farm bill.

The June Area Survey

The JAS includes questions on cash rental rates in addition to questions on other agricultural topics. The JAS sample is drawn from the NASS area frame, which covers all land in the United States except for Alaska and Hawaii. The JAS sample design is a stratified random design, with strata based on percent cultivation. Within each stratum, primary sampling units (PSUs) of 2 to 8 square miles are selected. The selected PSUs are divided into segments of approximately 1 square mile, and a segment is randomly selected from within the PSU. Once selected, a segment remains in the sample for 5 years. Each year approximately 20 percent of the segments enter the sample, and 20 percent rotate out. Each segment is divided into tracts, each representing a unique operation. The national sample consists of about 11,000 segments comprising some 38,000 tracts.

Data collection for the JAS is conducted by personal interview. Enumerators must account for all operations and land contained in their assigned segments. Enumerators conduct interviews and collect responses from the end of May through mid-June. Survey questionnaires are returned to the RFOs, where they are reviewed visually and entered manually into a database.

For the JAS cash rents, item-level nonresponse is accounted for by imputing data where values are missing. Imputed values are calculated through an automated imputation algorithm that requires a minimum of five complete reports within the imputation group. When a group lacks a sufficient number of responses, groups are collapsed according to a defined hierarchy, preserving as much of the homogeneity as possible, until five complete reports have been identified. The first imputation group is the segment, the second is similar strata within the county, the third is the county, the fourth is the district, and the fifth is the state.[2]

The JAS uses direct expansion estimates of cash rent items and their respective variances. Because cash rent items pertain to the entire farm, including those portions lying outside the sampled tract, the survey weight for a record is the product of the original segment sampling weight from the area frame and the proportion of the farm residing within the segment boundaries (called the farm-to-tract ratio). Rent per acre is computed as the ratio of the weighted estimate of dollars of rent paid to the weighted estimate of acres rented.

---

[2]Drawn from USDA, NASS ISSN 2167-129X, “Cash Rents Methodology and Quality Measures.” See https://www.nass.usda.gov/Publications/Methodology_and_Data_Quality/Cash_Rents/08_2016/rentqm16.pdf p. 2.

The Cash Rents Survey

The target population for the Cash Rents Survey is all farms and ranches with $1,000 or more in agricultural sales (or potential sales) that operate land rented from others on a cash basis in counties with at least 20,000 acres in cropland or pastureland (excluding Alaska) based on the most recent Census of Agriculture. The Cash Rents Survey collects data on acres rented and cash rental rates or total dollars paid in rent for three land use categories: irrigated cropland, nonirrigated cropland, and permanent pastureland. The data collection period for the survey is from late February through the end of June.[3]

The Cash Rents Survey sample is selected from the NASS list frame. The sample design for the survey is a stratified systematic sample at the county level, stratified by irrigated acres, nonirrigated acres, and pastureland. Within strata, farms are sorted by rent expenses. The certainty strata are those with a sampling interval of 1 in Table A-1, usually the strata that include farms making up more than 10 percent of the county total in certain categories. Most of the noncertainty strata use a 1 in 2 sampling rate. The noncertainty stratum for rent paid for all land and buildings uses a 1 in 3 sampling rate. The category of land rented or leased from others uses a sampling interval of 1 in 40. In 2014, the population was 745,489, farms, and the sample size was 241,176.

TABLE A-1 General Stratified Sample Design for Each County

Stratum	Description	Sampling Interval
98	Any farms with irrigated cropland acres >10% of total	1
96	Any farms with pastureland acres >10% of total	1
94	Any farms with nonirrigated cropland acres >10% of total	1
90	Irrigated cropland acres	1
80	Pastureland acres	2
70	Nonirrigated cropland acres	2
60	Rent paid for all land and buildings	3
50	Unknown cash rent expenses	1
40	Land rented or leased from others	40

The Cash Rents Survey sample is drawn in November, and the questionnaire is finalized in January. The survey data are collected early in the year, when farmers know their rental rates. The initial mailing starts in February; calling begins in March; and nonresponse mailing begins in April. Data collection by personal visit,[4] to either large operations or those with special handling arrangements, concludes in July. There is also an electronic reporting option. State-level estimates are published in August and county-level estimates in September. Data collection is coordinated with other surveys during this time period to minimize respondent burden. Data are collected on acres (by irrigated cropland, nonirrigated cropland, and pastureland) by acres rented to and acres rented from. Data on rent in dollars are collected for the same breakdown.

[3]To help with survey coordination in the month of March, the Cash Rents Survey questions are included on the March Agricultural Survey (see http://www.nass.usda.gov/Surveys/Guide_to_NASS_Surveys/Crops_Stocks [October 2015]) questionnaire. And during June, cash rents data are collected on the JAS in all states except Alaska and Hawaii.

[4]Data are collected either on a paper form or by a computer-assisted personal interviewing instrument.

Data on the total cash rents paid for all land and buildings also are collected. Family rents and cash rents are included in the survey, while share, flex, hybrid, and bonus rents are excluded.

The overall response rate to the Cash Rents Survey was about 76 percent in 2014. Once a county reaches the threshold required by the publication standard, attention to improving response is turned to other counties near the threshold, although responses are still accepted. The goal is to publish estimates for as many counties as possible. With respect to the mode by which completed surveys are received, the breakdown is 1 percent by fax; 2.7 percent by Internet; 39.4 percent by mail; 27.6 percent by telephone; and the remainder, about 5.3 percent, by field enumeration. NASS is considering approaches to encouraging response and cognitively testing the form. It hopes to have a more attractive Internet reporting option in the future.

All nonresponse is accounted for by reweighting. Weighting groups are based on ASD and design strata. Variances of totals are calculated using delete-a-group jackknife. Variances of ratios are calculated based on a second-order Taylor series expansion for ratios (see Kott [1989] for details). No item-level imputation is performed.

Cecere and colleagues (2012) and Berg and colleagues (2014) describe an evolution of the NASS cash rents model. The Berg, Cecere, and Ghosh (2014) model has been used to provide indications to ASB since 2013. The model is univariate, summing results from two separate univariate Fay-Herriott models that estimate the average rental rate (dollars/acre) in 2 years and the difference between rental rates (dollars/acre) in the same 2 years. Both models also include as covariates the National Commodity Crop Productivity Index from the National Resources Conservation Service (NRCS), a yield index from NASS historical estimates, and the total value of production from the Census of Agriculture. A two-stage benchmarking process applied to the modeled county-level estimates generates ASD-level estimates and satisfies relationships with the published state-level cash rental rates.

As is common with most NASS surveys, RFOs collect the data, resolve record-level errors identified during the editing process, and conduct further analysis within and across records before computing survey estimates. Final estimates are set through an ASB process that involves appointed NASS RFO and Headquarters staff. For cash rents, the primary basis for official estimates is the direct survey estimate. The Cecere, Berg, and Ghosh modeled estimates provide an additional source of information for ASB’s consideration, as do previously reported survey and frame data.

The NASS publication standard for the Cash Rents Survey is that there must be at least 30 positive reports of rent in the county, or if there are fewer responses, the sum of their corresponding unweighted acreage must account for at least 25 percent of rented acreage estimated. Estimates based on samples of fewer than 3 reports are automatically suppressed. While NASS computes estimates for total acres rented and total rent expense in a county, it does not publish these separate series; it publishes only the dollar per acre rental rate. The publication standard is verified independently for each of the three land use categories: irrigated cropland, nonirrigated cropland, and permanent pastureland.

Appendix B

Routine External Evaluation Protocol

Dorfman (2017) proposes that the Routine External Evaluation Protocol (REEP) be used as part of external evaluation.[1] He argues that, as the demand for small-area estimation is ever-increasing, statisticians have the responsibility to develop a protocol aimed at setting criteria and enabling detection of "a tipping point" at which estimates produced using small-area models cannot be regarded as satisfactory. This is particularly difficult to determine in groups of areas where there is little or no sample and that may behave differently from larger areas.

For NASS, REEP would be based on comparing county model estimates with the corresponding direct sample-based estimates where the latter are based on a *large enough* sample. The procedure would be implemented by selecting a random sample of counties within which a supplementary sample of farms would be selected and used to prepare direct estimates.

Step 1. Select *at random* a certain number of counties from the set of all counties that are expected to be estimated using models because the regular sample size does not meet publication standards.

Step 2. From each county selected at Step 1, take a *supplementary sample*. This sample should be large enough so that even with nonresponse, the direct sample-based estimates meet publication standards. These direct estimates and their variances will serve to test the quality of the model estimates.

> **Remark 1:** For the County Agricultural Production Survey (CAPS) , the current sample design already attempts to sample enough units from every county (the target sample size for each county is 70 for each commodity, with the goal of having 30 positive reports or at least 25% coverage within the county). The adaptive design approach used by NASS to target counties for survey nonresponse follow-up may need to be modified to give priority to the counties selected for external evaluation.

[1] If NASS decides to research the feasibility of the above REEP, a number of questions will have to be answered before the implementation:

- the procedure for selecting the random sample at Step 1, the number of counties needed, and the size of the sample selected within each county;
- specific procedure to amend the current adaptive design for data collection;
- specific method of evaluation and diagnostics, including developing formal criteria;
- estimated additional cost;
- estimated benefits from having the additional sample (beyond the goal of the model evaluation);
- possibility of evaluation (simulations?) based on historical data;
- whether it is possible to construct a procedure in such a way as to be able to modify the set of publishable counties based on the test results, and in the case where a model fails, whether it can still be applied to a subset of counties.

Remark 2: In the presence of high nonresponse, this supplementary sample may also be potentially useful in assessing the nonresponse bias. Compare the estimates in selected counties with all the other counties. Since the counties at Step 1 were selected at random, the average estimate that includes the supplementary sample can be compared with the rest of the sample; if there is no bias due to nonresponse, the estimates should be close.

Step 3. Use the "main sample" to compute model-based estimates and the full sample for evaluation and diagnostics. Note: The advantage of this "external" test is that it is *formal (objective)* and *contemporaneous.*

Step 4. Once the goal of model evaluation is accomplished, the supplementary sample can be combined with the main sample to produce better direct survey estimates and model-based estimates. Thus, even in the situation where model-based estimates fail, the estimates for counties selected at Step 1 still will be published.

Appendix C

Small-Area Modeling in Space and Time with Multiple Data Sources

Small-area estimation (SAE) methods refer to a wide range of modeling techniques generally devised to create improved estimates for domains where direct sample-based estimates are not reliable because of small sample sizes. Sample surveys are usually designed to produce reliable estimates for large domains (such as at the state or national level). However, there has been widespread growing interest in developing estimates for progressively finer domains, which has led to rapid development of a multitude of SAE methods in recent years. In SAE, the term "small area" refers to any domain that does not have enough samples for reliable direct sample-based estimates. In various surveys, examples include small areas defined by the intersection of detailed industry, geography, demography, etc. County levels in NASS surveys are typical examples of "small areas."

The underlying concept of SAE methods is to "borrow strength" by integrating information from multiple data sources, including survey data, or across time and space to improve estimates for small areas. Application of SAE methods at NASS has become increasingly important because of the demand for crop estimates at detailed geographic levels. Using traditional survey approaches to produce reliable estimates at these levels would require a much larger sample; this may be prohibitively expensive and impractical in terms of data collection and processing time. Consequently, reliance solely on sample-based estimates may result in many estimates being withheld from publication because of their poor quality. This type of situation is seen in the current NASS processing. Small-area modeling has been found to provide useful estimates for areas that are otherwise unpublishable; for areas that are considered publishable, it can also lead to improved efficiency over direct survey estimates. Two key success stories of SAE are the Small Area Income and Poverty Estimates program and the Small Area Health Insurance Estimates program of the U.S. Census Bureau, described elsewhere in this report.

This appendix gives an overview of modeling approaches and suggests those that NASS might pursue in the near and longer terms. It goes on to describe uncertainty measures, characteristics of area and unit models, extension to time, and related problems.

MODELS

The universe of methods that may be considered for SAE to combine information from multiple sources to improve county-level crop estimates is large and varied. Some of the most promising/interesting include geospatial methods, machine learning, and Bayesian methods. These are not mutually exclusive. For example, Bayesian methods are used in geospatial analysis and time domain analysis, as well as with incorporation of multiple types of measurements.

When used in a geospatial context, it can be helpful to include explicit spatial effects in order to broaden or borrow spatial support and thus to reduce uncertainty, especially when modeling estimates for small areas. A large number of model types have been proposed and demonstrated in applications, and can be found in texts on spatial statistics, spatial econometrics, and spatial modeling, notably those described in Cressie and Wikle (2011) and Haining (2003). Many of these models are implemented and readily accessible in software packages such as R.

In the classic SAE literature, there is a distinction between unit-level models (e.g., a model for a farm or for a Common Land Unit [CLU]) and area-level models (e.g., a model for a county). Both can be cast as linear mixed models with spatial random effects written as a spatial basis expansion (as there are now areal basis functions). If spatial models are used, then, roughly speaking, unit-level models correspond to the covariance (point-level) modeling approach, while area-level models correspond to the precision (polygon-level) approach. The majority of models available for point-level spatial and spatiotemporal modeling were developed in a noncomplex design setting, and often for Gaussian data.

The panel believes that the Bayesian approach holds great promise as recent developments have allowed combining design-based estimates with space–time smoothing models. For example, Mercer and colleagues (2015) effectively use a spatial Fay-Herriot (1979) model in the context of modeling childhood mortality based on complex survey data. The basic idea is to assume a hierarchical model in which the first stage is taken as the asymptotic distribution of the direct (design-based) estimator. Porter and colleagues (2015) use a similar model with an intrinsic conditional autoregressive (ICAR) spatial model, and emphasize covariate modeling. You and Zhou (2011) discuss both ICAR and Leroux specifications for the spatial model. Researchers who have worked on multivariate extensions to area-level models include Ghosh and Datta and colleagues and Bell and colleagues (e.g., Franco and Bell, 2016). Multivariate space–time models have been developed by Bradley and colleagues (2015a).

Near-Term Approaches

The panel suggests starting with area-level models. It is straightforward to add covariates to such models. The covariates may be added via a simple linear model or via a more flexible form, such as those used in the machine learning literature; it would be best to begin with simple, interpretable models. In an area-level model, satellite data can be included by taking within-area averages (for example). Note that ecological bias can be avoided (under certain assumptions) by including the within-area variance of the variable in the model, as well as the mean of the variable (Wakefield, 2008).

NASS already has experience with area-level modeling (Cruze et al., 2016; Erciulescu et al., 2016). So far, the use of spatial random effects has not been extensive. The panel suggests that NASS begin by exploring county-level models using the area-level spatial Fay-Herriot model to describe survey measurements. Each alternative data source could be given its own data model, linked to the larger model in a hierarchical Bayes framework. The model could use Besag or Leroux spatial formulations. See Hodges and Reich (2010) for a discussion of possible spatial confounding. Computation with the integrated nested Laplace approximation (INLA) (Rue et al., 2009) approach is fast (as compared with Markov chain Monte Carol [MCMC], which NASS has been using) and accurate. There is a reliable R implementation of the INLA method, though it is not a standard package.

Longer-Term Approaches

Modeling at the unit level may be more difficult than at the area level but potentially could lead to improved estimates, especially if auxiliary information is available at the unit level. To avoid estimation bias, the sample design should be properly accounted for in the model (for example, stratification could be accounted for by including fixed effects in the model, and cluster sampling by including random effects).

Point-level Fay-Herriot models also are possible. If a spatial Fay-Herriot unit model were used, then the data model could correspond to the asymptotic distribution of the direct estimate, with the spatial model appearing in the process model. Point-level models are somewhat troubling here as farms are not points, but can be very large. It may be that once CLU information is available, the area-level modeling could be extended to farms, again with a caveat that farms vary significantly in size.

EXPRESSING UNCERTAINTY

The Bayesian approach to modeling naturally leads to intuitive measures of uncertainty. The fundamental output of a Bayesian analysis is a multivariate posterior distribution over all unknown quantities in the model. This distribution is typically of high dimension, and so summarization is required. In particular, summaries of the univariate posterior distribution of quantities of interest may be reported. For example, the posterior median (or posterior mean, if the posterior on the quantity of interest is symmetric) may be quoted along with such quantiles as the 2.5 percent and 97.5 percent, to give a 95 percent interval.

Maps of posterior summaries may be produced. For example, a map of posterior medians may be accompanied by a map of the width of an interval of some percentage coverage (for example 95%). Hatching may also indicate uncertainty. For example, one may map the posterior median at the county level, but hatch with increased hatching as the associated uncertainty (interval estimate, for example) increases.

CHARACTERISTICS OF POINT-LEVEL AND AREA-LEVEL MODELS

There are two main approaches to modeling spatial data:

Point-Level (or Unit-Level) Modeling

- This modeling conceptually treats space as continuous.
- Spatial modeling concentrates on specification of the covariance matrix (e.g., Stein, 1999).
- Intuitive isotropic correlation models based on distance lead to dense matrices, i.e., matrices with few zeroes.
- Unfortunately, if n (the number of units) is large, the fitting of the model is computationally expensive because one must carry out operations (determinants and inverses) on n x n matrices (Rue and Held, 2005).
- This approach is also often referred to as *Gaussian random field (GRF)* or *geostatistical* modeling.
- This approach is suited to unit-level (point) modeling.

- Much of the literature (particularly in the environmental sciences) splits the modeling into a data model and a process model.
- Numerous approaches have been suggested to modeling the continuous surface with efficient computation (to give a variance–covariance matrix that is guiding the development).
  - Fixed-rank kriging (Cressie and Johanneson, 2008) and R implementation via the FRK package may be used.
  - Lattice kriging (Nychka et al., 2015) and R implementation via the LatticeKrig package may be used.
  - Stochastic partial differential equations (SPDE) (Lindgren et al., 2011) and R implementation via the INLA package may be used.
  - Predictive processes (Banerjee et al., 2008) and R implementation in the spBayes package may be used.
  - Bradley, Cressie, and Shi (2016) provide a review of approaches, with an emphasis on big data.

Polygon-Level (Area-Level) Modeling

- Modeling focuses on the precision (inverse covariance) matrix.
- This is also referred to as a Markovian or conditional modeling approach. Besag (1974) is an early influential reference. More specifically, the approach is also referred to as *Gaussian Markov random field (GMRF)* modeling.
- The key idea is to model local structure, which leads to sparse matrices (with zeros in the precision matrix corresponding to conditional independencies). There is less intuition on the implied covariances.
- Computation is very efficient with either MCMC or INLA (Rue et al., 2009).
- The approach discretizes space, usually based on administrative regions. This can lead to somewhat ad hoc neighborhood definitions.
- These are also called area (or aggregate) models. There has been much experience in different fields of county-level modeling, particularly with health and census data.
- Common approaches to area-level modeling include:
  - intrinsic conditional autoregressive (ICAR) models (Besag et al., 1991);
  - Leroux et al. (1999); and
  - variants as used in INLA (Riebler et al., 2016).
- Neighbors need to be defined, and the most common method is based on sharing a common boundary.

EXTENSION TO TIME

Extension to time is conceptually straightforward, but joint space–time correlation models require care. Including time may not be important for county crop estimates, even though there is substantial correlation from year to year for some variables. Alternative data sources from the current year may provide stronger predictors than previous-year crops data. There is a literature in spatial and spatiotemporal modeling of specifying first-order models rather than second-order models, as this is often easier. Conditional space–time models may also be simpler.

See Cressie and Wikle (2011) for further information. See also work on multivariate spatiotemporal mixed-effects models by Bradley and colleagues (2015a).

RELATED PROBLEMS

The change of support problem (COSP) has seen much interest (Cressie, 1993; Gotway and Young, 2002; Cressie and Wikle, 2011; Gelfand, 2010; Bradley, Wikle, and Holan, 2016). This problem occurs when one would like to make an inference at a particular spatial resolution, but the data are available at another resolution. Much of this work focuses on normal data and kriging-type approaches, in which block kriging is used. For example, Fuentes and Raftery (2005) combine point and aggregate pollution data, with the latter consisting of outputs from numerical models produced over a gridded surface using MCMC, and evaluate the block kriging integrals on a grid. Berrocal and colleagues (2010) considered the same class of problem, but added a time dimension and used a regression model with coefficients that varied spatially to relate the observed data to the modeled output. Bradley, Wikle, and Holan (2015b, 2016) describe non-Gaussian spatial change of support (COS) and space-time COS.

Diggle and colleagues (2013) take a different approach for discrete data and model various applications using log-Gaussian Cox point processes, including the reconstruction of a continuous spatial surface from aggregate data. Their approach is based on MCMC and follows Li and colleagues (2012) in simulating random locations of cases within areas, which is a computationally expensive step. Software to implement this approach is described in Taylor et al. (2015).

Appendix D

Biosketches of Panel Members and Staff

Mary Ellen Bock (*Chair*) is professor emerita of statistics at Purdue University. Her current research interests are in massive data sets and data mining and bioinformatics. She is a fellow of the American Association for the Advancement of Science, the Institute of Mathematical Statistics, and the American Statistical Association. She is a past president of the American Statistical Association. She is a member of the Committee on National Statistics of the National Academy of Sciences. She holds a Ph.D. in mathematics from the University of Illinois.

Julie Gershunskaya is a mathematical statistician with the Statistical Methods Staff of the Office of Employment and Unemployment Statistics at the U.S. Bureau of Labor Statistics. Her research has involved small-area estimation and treatment of influential observations, with special application to the U.S. Current Employment Statistics Program. She holds an M.S. in mathematics from Moscow State University and a Ph.D. in survey methodology from the University of Maryland.

Malay Ghosh is distinguished professor of statistics at the University of Florida. He has been a leader in research on and applications of small-area estimation methods. In addition to small-area estimation, he is interested in Bayesian inference, decision theory, and likelihoods. His current research includes, in particular, Bayesian analysis of case-control data, multiple hypothesis testing, Bayesian variable selection, empirical likelihood, and composite likelihood. He is a fellow of the American Statistical Association and the Institute of Mathematical Statistics. He holds a Ph.D. in statistics from the University of North Carolina at Chapel Hill.

Michael F. Goodchild (NAS) is emeritus professor of geography at the University of California, Santa Barbara. He is a member of the National Academy of Sciences and the American Academy of Arts and Sciences and a foreign fellow of the Royal Society of Canada. Dr. Goodchild's research achievements center on the measurement, description, and analysis of phenomena on the surface of the earth. He has explored using digital information gathered by remote sensing satellites to create spatial and environmental models of the planet, construct maps, and create digital libraries of geographic information that can be widely accessed electronically. He has also developed mathematical models to help quantify the difference between these geographic measurements and the real world. Dr. Goodchild holds a Ph.D. in geography from McMaster University.

Chad Hart is associate professor of economics, crops markets specialist, and extension economist at Iowa State University. Prior to that he was a U.S. policy and insurance analyst with the Food and Agricultural Policy Research Institute (FAPRI) and a scientist with the Center for

Agricultural and Rural Development at Iowa State University. For FAPRI, he was responsible for directing econometric and modeling efforts for the crop insurance component of the FAPRI modeling system, and he has engaged in research examining the interaction between the agricultural commitments within the World Trade Organization (WTO) and the agricultural policies and programs of WTO members. Dr. Hart received his Ph.D. in economics and statistics from Iowa State University.

Olga Isengildina Massa has been associate professor in the Department of Agricultural and Applied Economics at Virginia Tech University since fall 2015. She worked as associate clinical professor in the Economics Department of the College of Business at the University of Texas at Arlington, and previously worked at the University of Texas, Clemson University, and the Office for Futures and Options Research at the University of Illinois. Her research focuses on forecast analysis and evaluation, price risk management, value of information, agricultural marketing, and agribusiness. She holds a Ph.D. in agricultural economics from Mississippi State University.

Susan E. Offutt was chief economist for the U.S. Government Accountability Office (GAO) until her retirement in May 2015. Before joining GAO, she served as administrator of the Economic Research Service at the U.S. Department of Agriculture (USDA) for 10 years, and prior to that tour was executive director of the National Research Council's Board on Agriculture and Natural Resources, which conducts studies on a range of topics in agricultural science. She was chief of the Agriculture Branch at the Office of Management and Budget (OMB). During her tenure at OMB, she coordinated budget and policy analyses of the farm bill and trade negotiations. She began her career on the faculty at the University of Illinois Urbana-Champaign, where she taught econometrics and public policy in the agricultural economics department. She received an M.S. and a Ph.D. from Cornell.

S. Lynne Stokes is a professor in the Department of Statistical Science at Southern Methodist University. Her current research interests include sampling and nonsampling error modeling, psychometrics, and capture–recapture methodology. She is an elected fellow of the American Statistical Association (ASA) and a recipient of ASA's Founder's Award. For the past 15 years, she has been a member of the Technical Advisory Committee for the U.S. Department of Education's National Assessment of Educational Progress. She received her Ph.D. in mathematical statistics from the University of North Carolina at Chapel Hill.

Jonathan Wakefield is a professor in the Department of Statistics and the Department of Biostatistics at the University of Washington. His research interests include spatial epidemiology, ecological inference, genetic epidemiology, genome-wide association studies, the analysis of next-generation RNA sequence data, space–time models for infectious disease data, small-area estimates, hierarchical models for survey data, and the links between Bayes and frequentist procedures. He is a fellow of the American Statistical Association and a recipient of the Guy Medal in Bronze from the Royal Statistical Society. He holds a Ph.D. in mathematics from the University of Nottingham.

Robert E. Young is chief economist and deputy executive director of public policy at the American Farm Bureau Federation. He directs the organization's Economic Analysis Team, which conducts and coordinates economic research to support Farm Bureau public policy positions on such topics as farm policy, agricultural trade, regulatory costs, labor, and taxation. Dr. Young previously served as co-director for the Food and Agricultural Policy Research Institute at the University of Missouri; as a research associate professor in agricultural economics at the University of Missouri; and as chief economist for the U.S. Senate Committee on Agriculture, Nutrition and Forestry. He holds a Ph.D. in agricultural economics from the University of Missouri-Columbia.

Nancy J. Kirkendall (*Study Director*) is a senior program officer for the Committee on National Statistics. Previously, she served as director of the Statistics and Methods Group of the Energy Information Administration (EIA) and a member of EIA's senior staff. She also served as senior mathematical statistician in the Statistical Policy Branch of the Office of Information and Regulatory Affairs of the U.S. Office of Management and Budget, serving as the desk officer for the U.S. Census Bureau and chair of the Federal Committee on Statistical Methodology. She is a fellow and past vice president of the American Statistical Association and a past president of the Washington Statistical Society. She is a recipient of the American Statistical Association's Roger Herriot Award for Innovation in Federal Statistics and its Founder's Award. She holds a Ph.D. in mathematical statistics from The George Washington University.

Glenn D. White, Jr. *(Senior Program Officer)* is a senior program officer of the Committee on National Statistics. Previously, he was a senior manager at Ernst &Young's Quantitative Economic and Statistics Practice, where he established and directed a quantitative survey practice. His primary responsibilities as practice leader included general statistical consulting, with an emphasis on survey and web data collection. He also previously was a senior mathematical statistician at the Internal Revenue Service and a supervisory mathematical statistician at the U.S. Bureau of the Census. He holds a B.A. from the University of San Diego and an M.S. from the University of Vermont.