when λ is unknown. The point estimates are identical or very similar depending on the choice of $\hat{\lambda}_1$ for each approach but the associated measures of uncertainty could be quite different. In addition to the above modifications which rely on the δ-method, Singh, Stukel and Pfeffermann (1993) also obtain a modification of the asymptotic Bayes method of Hamilton (1986) which uses Monte Carlo integration (MCI) for evaluating the two terms of the posterior variance given by (11) thus avoiding computation of partial derivatives. The MCI simply entails generating λ_1-values from the approximate posterior distribution of λ_1 which is given by $N(\hat{\lambda}_1, \overline{V}(\hat{\lambda}_1))$. It is not difficult to show that the order of the neglected terms in the Hamilton (H) approximation is $O(m^{-1})$ and not $o(m^{-1})$. However, if the posterior distribution of λ_1 is approximated by $N(\hat{\hat{\lambda}}_1, \overline{V}(\hat{\lambda}_1))$, then the modified Hamilton (MH) approximation is of the desired order. Singh, Stukel and Pfeffermann (1993) report results of a Monte Carlo study on the frequentist properties of various approximations. Empirically, it is found that the KS-I approximation is biased downward, but KS-II* adds a positive term (similar to PR) and tends to be conservative. The behaviour of the MH approximation is quite similar to KS-II*, but H tends to be more biased downward than KS-I. The performance of the PR approximation is found to be best overall with respect to the frequentist properties, although other approximations provide useful alternatives. In particular, Bayesian approximations KS-II* and MH have the distinct advantage of having a dual interpretation in both frequentist and Bayesian contexts.

It may be noted that if m is quite large, then there will be hardly any difference between various approximations.

4. REMARKS

It is evident from the paper of Ghosh and Rao that great advances have been made in the field of small area estimation by both Bayesians and frequentists. It is also evident from the present discussion that there may be quite a bit of agreement between the two approaches. However, these advanced tools are not in widespread use, especially by statistical agencies conducting large scale complex surveys who face probably the greatest demand for small area statistics. Perhaps, the reason for this is the practitioner's skepticism in modelling complex survey data. Indeed, for complex surveys there is very little by way of model validation and more so for element-level modelling because of possible selection bias [see section 4 of Ghosh and Rao and a recent review by Pfeffermann (1993)]. There is no doubt that the area of model validation for complex survey data needs more research. This is also recognized by Ghosh and Rao and I would like to emphasize by noting that further work in this direction will be a very valuable contribution.

ACKNOWLEDGMENT

This research was partially supported by a grant from Natural Sciences and Engineering Research Council of Canada.

Comment

Elizabeth A. Stasny

Ghosh and Rao are to be congratulated for their timely paper reviewing methods for small-area estimation. My main complaint is that a paper such as this was not available five years ago when I began working on small-area estimation problems. I particularly enjoyed the historical perspective offered in the demographics methods section of the paper; I was sorry that section was so short since much of the material described in that section is not readily available to statisticians outside of the government agencies.

As the authors noted, there is a growing demand for small-area estimates and a corresponding interest in research on procedures for producing such estimates. The widely publicized debate on adjusting the U.S. population census for the undercount to produce adjusted counts for states and large cities has made many researchers focus on small area estimation problems related to the population census. There are, however, other long-standing small-area estimation programs. One of these is the USDA's program of county-level estimation of crop and live-

Elizabeth A. Stasny is Associate Professor, Department of Statistics, 1958 Neil Avenue, 148D Cockins Hall, Ohio State University, Columbus, Ohio 43210.

stock production. Because most of my work on small-area estimation has been on the problem of producing county-level estimates of crop production, I would like to add a brief discussion of this program to Ghosh and Rao's list of examples. A more detailed description of the National Agricultural Statistics Service (NASS) county estimation program is provided by Iwig (1993).

The USDA's NASS, in cooperation with state governments, has published county crop estimates for every state in every year since 1917. These estimates include acreage, yield and production data for many crops, both common (for example, wheat and corn) and rare (for example, rice and peanuts). For example, the 1990 Annual Report of the Ohio Agricultural Statistics Service includes county-level estimates of number of farms, acres in farms; acres harvested, yield (in bushels per acre), and production (in bushels) for a number of common crops including corn, soybeans, wheat, oats, hay and, for counties producing them, for less common crops such as tomatoes and sugar beets; and grain storage capacity.

Funding for extra data collection and the production of county-level estimates is provided by individual states; the USDA's national Quarterly Agricultural Surveys (QAS) are used to produce estimates only at the state and national level. Although a recent effort has been made to standardize county estimation programs (see the task force report by Bass et al., 1989), the sampling and estimation procedures used in county estimation programs differ from state to state. The task force recommendations for sampling procedures are that each state stratify farms by commodity or group of commodities and by size of operation, choose samples from within these strata, combine these samples into a single sample and delete farms that were already sampled for the QAS (since information for those farms is already available). Within this basic sampling plan, however, individual states have a considerable amount of flexibility. States often choose to sample a high proportion of large farms and of farms that responded to the survey in the previous year (possibly 100% in both cases). In addition, since information on farm operations is required to maintain the control data for the sampling frame, states often sample all farms that have not responded to a survey within a certain number of years.

Typically, 15,000 to 20,000 farms are sampled within a state; the response rate is about 30% with no follow-up of nonrespondents. There is no attempt to obtain sample-based weights for the responding cases since they may have been sampled from one or more commodity lists or from the QAS. Thus, small-area estimation methods that require known sampling weights cannot be used.

Methods of county estimation of crop production vary from state to state. A typical procedure involves initially obtaining the direct county estimates from the data available within each county. (Note that there may be few or no observations within a county for a certain crop, particularly if the county is largely urban or if the crop is relatively rare.) Then one or more experts review the estimates and adjust them in light of their personal knowledge of the farms in the sample, weather conditions, production in the county in previous years and other factors. The experts then look at the implications of the adjustments on the estimated production for the state. These last two steps may be repeated several times until the experts are satisfied with the estimates.

The county estimation program provides an example of the constrained estimation problem described in Section 7.2 of Ghosh and Rao's paper. The QAS data, along with other information such as historical data, administrative data (for example, on land set-aside programs) and weather data, is used by NASS to set the state-level estimates of crop production. Because the QAS is a large, probability sample of farms and the estimates produced using QAS data are believed to be fairly accurate, county estimates are typically constrained to agree with NASS's state estimates. Stasny, Goel and Rumsey (1991) consider the problem of how to scale wheat production estimates to agree with the NASS state total. They consider 1) a constant scaling factor, 2) scaling factors that minimize the sum of squared differences between initial and adjusted estimates and 3) scaling factors that minimize the sum of squared relative differences between initial and adjusted estimates. They found that Method 2) was clearly inferior to the other methods, but there was not much difference between estimates scaled using Methods 1) and 3).

While county estimates obtained following a procedure such as that described above have been used successfully for many years, there are many problems with the procedure. The lack of a formal statistical methodology makes it impossible to repeat the estimation process, to compare estimates from different states and to obtain estimates of the uncertainty. On the other hand, the willingness to base county estimates on expert opinion and on data from many sources (current surveys, historical data and administrative records) makes this an exciting area for continued research.

One area for research is in using information from other crops and from neighboring counties or states to improve the county estimates. Pawel and Fesco (1988) explored historical estimates of crop yields and found that, as expected, there are high positive correlations between yields in neighboring states and between agronomically related crops that are grown in overlapping regions. This research, however, was conducted at the state rather than county level; it is still an open question whether similar relations will be useful at the county level.

Another area for research is in using the historical data on crop production in current county estimates. A natural way to use this information would be in a Bayesian setting such as the hierarchical Bayes estimates described in Section 5.3 of Ghosh and Rao's paper. Indeed, it seems surprising that a noninformative prior would be used in small-area estimation problems involving census data or data from continuing surveys; there is certainly a wealth of information on which to base an informative prior.

Finally, I would like to mention a success story in research in the production of county estimates. Ghosh and Rao describe the experimental research of Battese, Harter and Fuller (1988) on county estimation of crop production using satellite data. This year, for the first time, Arkansas is using satellite data to aid in production of crop acreage estimates as part of their county estimates program. Over the next few years, other states are expected to begin using such data to aid in the production of their crop acreage estimates.

ACKNOWLEDGMENT

This research was supported in part by the USDA/NASS under Cooperative Agreements 58-3AEU-9-80040 and 43-3AEU-3-80083. The author takes sole responsibility for the work.

Comment

Ib Thomsen

It takes talent and hard work to provide an overview and evaluation of a rapidly evolving subject like small area estimation. In my opinion the authors have succeeded in doing this, and I want to congratulate them with a very useful review. In many statistical offices, substantial methodological work is being done to find suitable estimators for small areas. People involved in such work will be grateful to Ghosh and Rao for their present contribution.

Below I shall communicate some experiences gained when developing and using small area estimates within Statistics Norway. But first a few comments to the example given in Section 6 of the paper. In this example a synthetic population is constructed by fitting a nested error regression model to a business population. For this synthetic population, the EBLUB (or EB) and the HB estimators are shown to produce small area estimators which are superior to the ratio-synthetic and a sample-size dependent estimator. As pointed out by the authors, this demonstrates the advantages of using EBLUB or HB estimators when the model fits the data well. A question remains concerning the robustness of these estimators as compared to the simpler sample-size dependent estimator. A column in Table 3 showing the small area means of the real business population could have thrown some light on the robustness of the estimators studied in the paper.

At Statistics Norway, small area estimators have been used for some years now (Laake, 1978). In the beginning we concentrated on synthetic estimators, but more recently composite estimators are being used. In what follows some of our experiences concerning the feasibility of the EB estimator are presented.

I shall look at a very simple situation in which θ_i, $(i = 1, \ldots, T)$ is a small area parameter, and $\overline{X}_i$, $(i = 1, \ldots, T)$ is a direct estimator such that

$$E(\overline{X}_i|\theta_i) = \theta_i \quad i = 1, \ldots, T.$$

The parameters $\theta_1, \theta_2, \ldots, \theta_T$ are considered realizations of a random variable with unknown distribution $G(\cdot)$. The mean μ and variance σ^2 are assumed to be known or that estimates are available. For a set of small areas, unbiased estimators $\overline{X}_1, \ldots, \overline{X}_T$ are available with conditional distributions equal to the binomial.

When $G(\theta)$ is unknown, empirical Bayes estimators generally employ $(\overline{X}_1, \ldots, \overline{X}_T)$ to estimate $E(\theta|\overline{X}_1, \ldots, \overline{X}_T)$. However, for many distribution, $E(\theta|\overline{X}_1, \ldots, \overline{X}_T)$ cannot be consistently estimated un-

Ib Thomsen is Director of Research, Statistics Norway, and Professor of Statistics, University of Oslo, P.B. 8131 Dep, 0033 Oslo, Norway.