Journal of Official Statistics, Vol. 17, No. 3, 2001, pp. 391–406

Estimation of the Rates and Composition of Employment in Norwegian Municipalities

Nicholas T. Longford[1]

A multivariate shrinkage method is applied to estimate the municipality-level statistics about the employment status of adult Norwegian residents in 1990. It can be motivated as borrowing strength across municipalities and components of the vector of estimated percentages. The gains in precision of the method over sample (naive) estimation are discussed and related to the municipality-level heterogeneity. A general approach to validation of the method is illustrated and a more comprehensive application of the method outlined.

Key words: Between-area variation; employment status; municipalities in Norway; shrinkage estimation; small area estimation.

1. Introduction

As part of the 1990 decennial population census in Norway, statistics about the employment status of the adult (registered) residents of the country's 448 municipalities (*kommune*) were collected. For each municipality, two kinds of summaries are of interest:

- rates (percentages) of employment within categories defined by sex and age;
- composition of the employed by the industrial sector and sex-by-age categories.

Each adult person who has in the past been employed is assigned to one of nine industrial sectors. Five age groups are defined for each sex, delimited by 20, 25, 60, and 66 years of age. By composition we mean the (multinomial) percentages of employed in each sector or sex-by-age category. The counts, or the corresponding percentages, are required by the local (municipal) and national administrative authorities for planning, and are an important resource for labour market research. The municipalities maintain registers of all employed residents, with personal details and type of employment (including the industrial sector), but these are often not complete or up-to-date. A survey embedded in the 1990 Census provides information with less bias, but with sampling variation.

Within the 1990 Census, information about employment was collected from all adult residents (aged 16 or over on the day of the Census, 3rd November 1990) in municipalities with (estimated) population of less than 6,000 (275 municipalities), and from independent

[1] Department of Medical Statistics, De Montfort University, The Gateway, Leicester LE1 9BH, England. Email: ntl@dmu.ac.uk

Acknowledgments: A substantial part of the work on this article was carried out while the author was a visitor at the Norwegian Social Science Data Services (NSD) in Bergen, Norway, under the Training and Mobility of Researchers Programme funded by the European Union. Hospitality of NSD is acknowledged. Astrid Nilsen assisted me by retrieving the analysed data and identifying the relevant background literature.

Table 1. Limits for reporting of cell counts

	Sampling fraction			
	1:12	1:10	1:7	1:5
Not reported	< 124	< 101	< 67	< 45
In parentheses	124–276	101–224	67–147	45–98
	For estimates combined with the registers			
Not reported	< 41	< 33	< 22	< 15
In parentheses	41–92	33–75	22–49	15–33

Note: Reproduced from Statistisk Sentralbryå (1994a).

random samples without replacement in the other municipalities: with 20 per cent sampling fraction in 41 municipalities with population in the range 6,000–7,999, 14.3 per cent (1:7) in 26 municipalities with population in the range 8,000–9,999, 10 per cent in 91 municipalities with population in the range 10,000–49,999, and 8.3 per cent (1:12) in the eight most populous (urban) municipalities which were estimated to have (in 1989) more than 50,000 inhabitants each. Further, in two and five municipalities with respective populations in the ranges 6,000–7,999 and 10,000–49,999, complete enumeration was conducted. (The additional costs were covered by the municipal authorities concerned.) The employment status, recorded as a binary variable, is declared positive if the resident was engaged in paid employment for at least 100 hours in the period 3rd November 1989–2nd November 1990, and negative otherwise. A similar binary variable is recorded regarding employment in the week immediately preceding the Census date of 3rd November 1990.

The 1990 Census in Norway yielded the population count 4,247,600, with 2,124,000 employed (*yrkesaktive*) persons (50.0 per cent), of which 1,163,000 (54.7 per cent) were men and 962,000 were women. In the survey, the sampling fractions were set so as to obtain most sample estimates of the counts of employed residents with sufficient precision while staying within the allocated budget. The employment survey involved 595,300 subjects, 415,300 (69.8 per cent) of them from municipalities with full enumeration. In the tables distributed by the Norwegian Social Science Data Services (NSD), the estimates are reported without standard errors, but when the relative standard error (RSE) is larger than 30 per cent the estimate is not given, and is given in parentheses when RSE is in the range 20–30 per cent. For simplicity, these thresholds are expressed in terms of minimal cell counts (Statistics Norway 1994a); see Table 1.

For estimating the composition of employment by sector, the method of Lillegård (1993), combining the survey information with the imperfect information from the registers, was also applied. In this method, the survey data are compared with the register, and the observed rates of the two kinds of error are incorporated in the population estimate by an approach similar to poststratification. The resulting estimates, evaluated separately for each municipality, are much more precise and the minimum estimated counts that are reported are about three times smaller; see bottom part of Table 1. Nevertheless, 100 cells in the tables of municipality-level counts crossclassified by industrial sector, involving three sectors and 89 municipalities, are not given because of insufficient precision. Note, however, that for 78 municipalities for which only one count is not reported this count (in some cases zero) can be recovered from the municipality's total number of employed residents.

This article illustrates how the estimation of the various counts, or percentages, can be improved by (multivariate) shrinkage methods (Longford 1999a), and outlines how the sampling strategy could be adjusted if shrinkage estimators were applied in future. Its focus is on estimating the municipality-level composition of the employed by sex and age (age groups 16–19, 20–24, 25–59, 60–66, and over 66 years), where the shrinkage estimators are particularly effective, but estimation of percentages (counts) within the industrial sectors is also discussed.

Shrinkage has been firmly established as an important element of small-area estimation by Fay and Herriot (1979). Fay (1987) and Battese, Harter, and Fuller (1988) have developed extensions for multivariate outcomes and for regression adjustment. Their focus has been on continuous (normally distributed) outcomes, applying empirical Bayes or hierarchical models (Gelman et al. 1995). Although these models have extensions to the exponential family, fitting them by iterative maximum likelihood algorithms entails several problems; see Elston, Koch, and Weissert (1991) and Thomas, Longford, and Rolph (1994) for applications. Longford (1999a) makes no distributional assumptions, and his algorithm is noniterative, based on minimizing the expected mean squared error of a linear combination of estimators. Assuming normality, the best linear unbiased predictor (BLUP), obtained by maximum likelihood, has similar properties, and they are not eroded substantially with moderate departures from normality. In multivariate shrinkage, covariate information, usually incorporated by a regression model, can be exploited by the device of auxiliary variables. All approaches mentioned pursue the same goal of borrowing strength from auxiliary information. For comprehensive reviews of small-area estimation, see Ghosh and Rao (1994) and Marker (1999). Singh, Stukel, and Pfeffermann (1998) compared empirically several Bayesian approaches.

The next section gives a technical overview of shrinkage estimation as applicable to composition and sets of rates. The application to the sex-by-age categories and industrial sectors is presented in Section 3 and the results are compared with the (standard) sample estimates. The conclusions are supported by a validation exercise involving sampling from the municipalities with full enumeration. Section 4 summarizes the results and outlines the full potential of shrinkage estimation in the context of municipality-level statistics.

2. Shrinkage estimation

Let $\mathbf{p}_i$ be the vector of proportions associated with a partitioning of the population of adult residents (aged 16 years or over) in municipality $i = 1, \ldots, L = 448$, and let $\hat{\mathbf{p}}_i$ be the sample (or another unbiased) estimator of $\mathbf{p}_i$; $\mathrm{E}_s(\hat{\mathbf{p}}_i | \mathbf{p}_i) = \mathbf{p}_i$. We use subscripts s and i to indicate that the moments (means E or variances var) are associated with sampling and municipalities, respectively. In the formulae below, $\mathbf{0}$ and $\mathbf{I}$ denote the zero and the identity matrices of dimensions implied by the context.

Denote by $\mathbf{U}_i$ the sampling variance matrix of $\hat{\mathbf{p}}_i$; $\mathbf{U}_i = \mathrm{var}_s(\hat{\mathbf{p}}_i | \mathbf{p}_i)$. We assume that the municipality population sizes N_i are precise, and are large enough, so that the sampling variation of each of the sample sizes n_i can be ignored. In fact, the surveys employed a systematic sampling design with a sorting of the adult residents in such a manner as to ensure representativeness and randomness (using the randomly assigned digits of the

Personal Number issued to each Norwegian resident at birth or when becoming a resident). The national population proportions $\mathbf{p} = N^{-1}(N_1\mathbf{p}_1 + N_2\mathbf{p}_2 + \cdots + N_L\,\mathbf{p}_L)$ are estimated by $\hat{\mathbf{p}} = N^{-1}(N_1\hat{\mathbf{p}}_1 + N_2\hat{\mathbf{p}}_2 + \cdots + N_L\hat{\mathbf{p}}_L)$, where $N = N_1 + N_2 + \cdots + N_L$ is the national count of adults; $n = n_1 + n_2 + \cdots + n_L$ is the overall (national) sample size. Let $f_i = n_i/N_i$ be the sampling fraction used in the survey of municipality i.

The key quantity in shrinkage estimation is the between-municipality variance matrix $\boldsymbol{\Sigma} = \mathrm{var}_i(\,\mathbf{p}_i)$. If the municipalities had identical vectors of proportions, $\mathbf{p}_i \equiv \mathbf{p}$, $\hat{\mathbf{p}}$ would be a far superior estimator of $\mathbf{p}_i$ than $\hat{\mathbf{p}}_i$. Although it is unlikely that $\boldsymbol{\Sigma} = \mathbf{0}$, $\boldsymbol{\Sigma}$ may be quite small, and $\hat{\mathbf{p}}$ may be more efficient than $\hat{\mathbf{p}}_i$ for *some* municipalities. The shrinkage estimator combines the two alternatives, $\hat{\mathbf{p}}_i$ and $\hat{\mathbf{p}}$,

$$\tilde{\mathbf{p}}_i = (\mathbf{I} - \mathbf{b}_i)\hat{\mathbf{p}}_i + \mathbf{b}_i\hat{\mathbf{p}} \tag{1}$$

with the matrix of coefficients $\mathbf{b}_i$ for which the expected mean squared error (EMSE) of the linear combination $\tilde{\mathbf{p}}_i^{\top}\mathbf{w}$ is minimised. The EMSE is defined as

$$\mathrm{EMSE}(\mathbf{b}_i) = \mathrm{E}_i[\mathrm{E}_s\{\mathbf{w}^{\top}(\tilde{\mathbf{p}}_i - \mathbf{p}_i)(\tilde{\mathbf{p}}_i - \mathbf{p}_i)^{\top}\mathbf{w}|\,\mathbf{p}_i\}]$$

Note that $\mathbf{w} = (0, \ldots, 0, 1, 0, \ldots, 0)^{\top}$ corresponds to estimating a component of $\mathbf{p}_i$. The expectation over the municipalities (E_i) removes the dependence on the unknown $\mathbf{p}_i$. The estimator $\tilde{\mathbf{p}}$ in (1) is conditionally biased, given $\mathbf{p}_i$. Minimizing the variance or its expectation would ignore this bias. Therefore, it is essential to consider the (expected) mean squared error.

For composition data, the minimum EMSE is attained for

$$\mathbf{b}_i^* = \left(1 - \frac{n_i}{n}\right)\mathbf{D}_i^{-1}\mathbf{U}_i \tag{2}$$

where $\mathbf{D}_i = \boldsymbol{\Sigma} + \mathrm{var}(\hat{\mathbf{p}}) + (1 - 2n_i/n)\mathbf{U}_i$. Since the national sample size is very large, the contribution of $\mathrm{var}(\hat{\mathbf{p}})$ to $\mathbf{D}_i$ can usually be ignored. The ratio n_i/n is often also small enough to be ignored. For rates, the minimum EMSE is attained for

$$\mathbf{b}_i^* = (\mathbf{I} - \mathbf{Q}_i)\mathbf{D}_i^{-1}\mathbf{U}_i$$

where $\mathbf{Q}_i$ is the diagonal matrix of the composition fractions:

$$\mathbf{Q}_i = \mathrm{diag}\left(\frac{n_{i1}}{n_i}, \frac{n_{i2}}{n_i}, \ldots, \frac{n_{iH}}{n_i}\right)$$

n_{ih} is the count of survey subjects in municipality i and category h, and $\mathbf{D}_i = \boldsymbol{\Sigma} + \mathrm{var}(\hat{\mathbf{p}}) + (\mathbf{I} - 2\mathbf{Q}_i)\mathbf{U}_i$. In either case (rates or composition), the solution $\mathbf{b}_i^*$ is common to all vectors $\mathbf{w}$.

The estimator $\tilde{\mathbf{p}}_i$ can be interpreted as pulling (shrinking) the sample estimator $\hat{\mathbf{p}}_i$ closer to the overall (national) estimator $\hat{\mathbf{p}}$. This can be motivated by realizing that even if the underlying vectors $\mathbf{p}_i$ were identical, the sample estimates $\hat{\mathbf{p}}_i$ would not be, but would vary as much as is implied by the sampling variance matrices $\mathrm{var}_s(\hat{\mathbf{p}}_i)$. Since the sampling errors $\hat{\mathbf{p}}_i - \mathbf{p}_i$ and municipality-level deviations $\mathbf{p}_i - \mathbf{p}$ are uncorrelated, their contributions to the (observed) variation of the $\hat{\mathbf{p}}_i$ can be decomposed. The sampling variation of $\hat{\mathbf{p}}_i$ is estimated from the sampling design, and the between-municipality variation as the component of the *observed* between-municipality variation that is in excess of the sampling variation:

$$\hat{\boldsymbol{\Sigma}} = \mathrm{var}_i(\hat{\mathbf{p}}_i) - S\{\mathrm{var}_s(\hat{\mathbf{p}}_i),\ i = 1, \ldots, L\}$$

where S is a suitable function. When the between-municipality variance matrix $\boldsymbol{\Sigma}$ vanishes, $\hat{\mathbf{p}}$ in (2) is given almost full weight ($\mathbf{b}^* \doteq \mathbf{I}$); when $\boldsymbol{\Sigma}$ is large relative to $\mathbf{U}_i$, $\hat{\mathbf{p}}_i$ is accorded almost full weight ($\mathbf{b}^* \doteq \mathbf{0}$) because $\hat{\mathbf{p}}$ is not very useful for estimating $\mathbf{p}_i$. For intermediate cases, the weights accorded to $\hat{\mathbf{p}}_i$ and $\hat{\mathbf{p}}$ reflect the relative precisions of $\hat{\mathbf{p}}_i$ and $\hat{\mathbf{p}}$ as estimators of $\mathbf{p}_i$. For small sample sizes n_i, the national proportions $\hat{\mathbf{p}}$ may have much smaller sampling variances than $\hat{\mathbf{p}}_i$, especially when $\boldsymbol{\Sigma}$ is small. In general, shrinkage estimators assign larger weight to the national proportions for smaller sample sizes and when $\boldsymbol{\Sigma}$ is small.

For composition, in simple random sampling without replacement,

$$\mathbf{U}_i = \frac{1 - f_i}{n_i}\{\mathrm{diag}(\mathbf{p}_i) - \mathbf{p}_i\mathbf{p}_i^\top\} \tag{3}$$

and for rates,

$$\mathbf{U}_i = (1 - f_i)\mathrm{diag}_h\left\{\frac{p_{ih}(1 - p_{ih})}{n_{ih}}\right\}$$

When $\boldsymbol{\Sigma}$ is small relative to $\mathbf{U}_i$, and especially when n_i is small, $\mathbf{U}_i$ is estimated more efficiently by substituting $\hat{\mathbf{p}}$, *not* $\hat{\mathbf{p}}_i$, for $\mathbf{p}_i$ in (3), because the entries of the naively estimated variance matrix, $\mathrm{diag}(\hat{\mathbf{p}}_i) - \hat{\mathbf{p}}_i\hat{\mathbf{p}}_i^\top$, have large variances; see Longford (1999a) for a discussion. We denote

$$\mathbf{R} = \mathrm{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\top \quad \text{and} \quad \mathbf{V}_i = n_i^{-1}(1 - f_i)\mathbf{R}$$

for composition, and

$$\mathbf{R} = \mathrm{diag}\{p_h(1 - p_h)\} \quad \text{and} \quad \mathbf{V}_i = (1 - f_i)\mathrm{diag}\{p_h(1 - p_h)/n_{ih}\}$$

for rates. When information from the municipality registers is incorporated, $\mathbf{V}_i$ is approximated as $(1 - K_i)\mathbf{U}_i$ or $(1 - K_i)\mathbf{V}_i$, where the factor K_i is estimated from the cross-tabulation of survey responses and register records; see Lillegård (1993) and Swensen (1988).

Without conditioning on $\mathbf{p}_i$, we have

$$\mathrm{var}(\hat{\mathbf{p}}_i) = \mathbf{V}_i + \frac{n_i - g_i}{n_i}\boldsymbol{\Sigma}$$

where $g_i = (1 - f_i)(1 - K_i)$ when the information from the register is incorporated and $g_i = (1 - f_i)$ otherwise. Since the sampling processes in the L municipalities are independent of one another,

$$\begin{aligned}\mathrm{var}(\hat{\mathbf{p}}) &= \frac{1}{N^2}\sum_{i=1}^{L} N_i^2\left(\mathbf{V}_i + \frac{n_i - g_i}{n_i}\boldsymbol{\Sigma}\right)\\ &= \frac{1}{N^2}\sum_i \frac{N_i}{f_i}\{g_i\mathbf{R} + (n_i - g_i)\boldsymbol{\Sigma}\}\end{aligned} \tag{4}$$

The variance matrix $\boldsymbol{\Sigma}$ is estimated by moment matching applied to the sum-of-squares statistic

$$\mathbf{S}_B = \sum_i n_i(\hat{\mathbf{p}}_i - \hat{\mathbf{p}})(\hat{\mathbf{p}}_i - \hat{\mathbf{p}})^\top$$

Since $\text{cov}_s(\hat{\mathbf{p}}_i, \hat{\mathbf{p}}) = N^{-1}N_i\text{var}(\hat{\mathbf{p}}_i)$,

$$
\begin{aligned}
\mathrm{E}_{si}(\mathbf{S}_B) &= \sum_i n_i\left(1 - 2\frac{N_i}{N}\right)\mathbf{V}_i + n\ \text{var}(\hat{\mathbf{p}}) \\
&= \sum_i \left(1 - 2\frac{N_i}{N}\right)g_i\mathbf{R} + \frac{n}{N^2}\sum_i \frac{N_i g_i}{f_i}\mathbf{R} \\
&\quad + \sum_i \left(1 - 2\frac{N_i}{N}\right)(n_i - g_i)\boldsymbol{\Sigma} + \frac{n}{N^2}\sum_i \frac{n_i}{f_i}(n_i - g_i)\boldsymbol{\Sigma} \\
&= A\mathbf{R} + B\boldsymbol{\Sigma}
\end{aligned}
$$

for implicitly defined constants A and B. Hence,

$$
\hat{\boldsymbol{\Sigma}} = \frac{1}{B}(\mathbf{S}_B - A\mathbf{R})
$$

is an unbiased estimator of $\boldsymbol{\Sigma}$. Its sampling variance matrix is quite small when the number of municipalities L is large, and when many of them are fully enumerated ($f_i = 1$). The information about $\boldsymbol{\Sigma}$ can be effectively summarized by the conventional degrees of freedom. For m fully enumerated municipalities we have m degrees of freedom. Each municipality in which only a sample is available contributes only with a fraction of a degree. See Longford (1996) for a related discussion. Also, as with estimating univariate variances, $\text{var}_s(\boldsymbol{\Sigma})$ is smaller the smaller the actual value of $\boldsymbol{\Sigma}$.

We define the shrinkage estimator $\tilde{\mathbf{p}}_i$ by (1) with the matrix coefficients $\mathbf{b} = \mathbf{b}^*$ given by (2) with estimates (and approximations, where applicable) in place of the matrices $\boldsymbol{\Sigma}$, $\mathbf{V}_i$ and $\text{var}(\hat{\mathbf{p}})$. All our evaluations of the sampling variation ignore the uncertainty in their estimation. Estimation of $\mathbf{V}_i$ is usually associated with the largest sampling variation. Underestimating (the variances in) $\mathbf{V}_i$ reduces the efficiency of $\tilde{\mathbf{p}}_i$, but at the same time it reduces the chances that $\tilde{\mathbf{p}}_i$ is conditionally less efficient than $\hat{\mathbf{p}}_i$, given the value of $\mathbf{p}_i$. See Longford (1999a) for a detailed discussion.

When $\boldsymbol{\Sigma}$, $\text{var}_i(\hat{\mathbf{p}})$, and $\mathbf{V}_i$ are known, the EMSE of $\tilde{\mathbf{p}}$ is $\mathbf{V}_i - \mathbf{b}_i^{*\top}\mathbf{D}_i\mathbf{b}^*$, so the gain in precision of $\tilde{\mathbf{p}}_i$ over $\hat{\mathbf{p}}_i$ is $\mathbf{b}_i^{*\top}\mathbf{D}_i\mathbf{b}_i^*$. The effect on $\widehat{\text{var}}(\hat{\mathbf{p}}_i)$ of uncertainty about the variance matrices can be explored by bootstrap. Longford (1999a) found, in a setting with fewer small areas and sparser sampling (1:50), that this can be ignored; by regarding the variance matrices as known, the standard errors are underestimated by not more than two per cent. In the analyses in Section 3, the standard errors are underestimated by less than one per cent. The singularity of the H components of a multinomial vector of proportions is resolved by handling only the first $H - 1$ components.

3. Application

In Section 3.1, we apply shrinkage to estimate the rates of employment in the (2×5) sex-by-age categories of the residents of each municipalities. In Sections 3.2 and 3.3, the composition of the employed population in each municipality is estimated, by sex and age (10 categories) and by industrial sector (9 categories), respectively. Shrinkage is very effective for the rates and composition with regard to sex-age groups because their distributions are fairly homogeneous across the municipalities. On the other hand, shrinkage is not very effective in the analysis of the industrial sectors because the distribution of employment across the sectors varies a great deal among the municipalities. Some improvement might be attained by classifying the municipalities into a small number of relatively homogeneous categories, and then applying the shrinkage within each category separately. However, such categories are difficult to identify. Past data are not sufficient because they are sparse, both in sampling and in time, and the composition of employment is gradually changing in response to disparate economic stimuli.

3.1. Estimating the rates

Although we are interested in estimating the numbers of employed residents in only 166 municipalities in which simple random samples were drawn, the (complete) information from the other 282 municipalities contributes to estimating the national proportions $\mathbf{p}$ and the between-municipality variance matrix $\boldsymbol{\Sigma}$. To avoid reporting very small numbers, we convert all proportions to percentages. Throughout, we use two-letter symbols for the categories given by age and sex, comprising T (teenage, 16–19 years of age), Y (young, 20–24 years), M (middle-aged, 25–59 years), V (veteran, 60–66 years), and R (retirement age, 67 years and above), and m and w for men and women, respectively. For instance, Yw stands for women aged 20–24 years.

The nation-level summary of the estimated percentages is given in Table 2. The relative homogeneity, indicated by the small values of the standard deviations in $\hat{\boldsymbol{\Sigma}}$, suggests that shrinkage is quite effective – by shrinkage, a lot of strength is borrowed across the municipalities. For illustration, a municipality with 24,000 residents, 12,000 of whom are employed, will be represented in the sample by 1,000 subjects, of whom about 25 are 16–19 year-old women (Tw). If the underlying rate of employment in Tw is 45 per cent, the standard error of the sample percentage is 9.95 per cent. If we estimate the rate by the national rate, 43.72 per cent, we fare not much worse; the associated standard error is $\sqrt{9.91^2 + 1.17^2} = 9.98$ per cent. But by combining these two estimators, with (effectively) equal weights, the root-EMSE is reduced to 7.05 per cent, a substantial gain in precision.

The estimated between-municipality correlation matrix is

$$\begin{array}{cc} \begin{pmatrix} 1.00 & 0.50 & 0.49 & 0.48 & 0.44 & 0.46 & 0.55 & 0.44 & 0.49 & 0.50 \\ 0.50 & 1.00 & 0.42 & 0.56 & 0.33 & 0.53 & 0.37 & 0.49 & 0.32 & 0.45 \\ 0.49 & 0.42 & 1.00 & 0.46 & 0.60 & 0.21 & 0.56 & 0.41 & 0.38 & 0.28 \\ 0.48 & 0.56 & 0.46 & 1.00 & 0.41 & 0.58 & 0.39 & 0.46 & 0.24 & 0.40 \\ 0.44 & 0.33 & 0.60 & 0.41 & 1.00 & 0.49 & 0.77 & 0.60 & 0.39 & 0.40 \\ 0.46 & 0.53 & 0.21 & 0.58 & 0.49 & 1.00 & 0.58 & 0.65 & 0.35 & 0.49 \\ 0.55 & 0.37 & 0.56 & 0.39 & 0.77 & 0.58 & 1.00 & 0.75 & 0.55 & 0.61 \\ 0.44 & 0.49 & 0.41 & 0.46 & 0.60 & 0.65 & 0.75 & 1.00 & 0.51 & 0.62 \\ 0.49 & 0.32 & 0.38 & 0.24 & 0.39 & 0.35 & 0.55 & 0.51 & 1.00 & 0.53 \\ 0.50 & 0.45 & 0.28 & 0.40 & 0.40 & 0.49 & 0.61 & 0.62 & 0.53 & 1.00 \end{pmatrix} & \begin{array}{c} \mathrm{Tm} \\ \mathrm{Tw} \\ \mathrm{Ym} \\ \mathrm{Yw} \\ \mathrm{Mm} \\ \mathrm{Mw} \\ \mathrm{Vm} \\ \mathrm{Vw} \\ \mathrm{Rm} \\ \mathrm{Rw} \end{array} \end{array}$$

All the correlations are positive. This is not surprising; when the rate of employment in a municipality is higher for one category it tends to be higher also for the other categories.

Table 2. Estimated rates of employment for the sex-age groups

	Tm	Tw	Ym	Yw	Mm	Mw	Vm	Vw	Rm	Rw
Estimate	46.00	43.72	80.32	74.31	88.56	76.47	57.72	39.83	9.88	3.81
St. error	1.08	1.17	0.75	1.02	0.57	0.79	1.44	1.44	0.58	0.33
$\sqrt{\mathrm{diag}(\hat{\boldsymbol{\Sigma}})}$	9.31	9.91	5.45	6.57	3.87	5.19	10.30	9.41	4.08	1.92

Multivariate shrinkage exploits these associations, and improves the estimation for the less populous categories, the youngest (T and Y) and oldest (V and R), by borrowing strength from the most populous categories (Mm and Mw), in addition to borrowing strength across the municipalities.

For instance, in municipality 101 (Halden in county Østfold, population 25,900), the sample rates of employment among teenagers are 45.4 per cent (men, sample size 76) and 46.3 per cent (women, sample size 75), both with standard errors 5.4 per cent. Multivariate shrinkage adjusts the rates to 43.8 per cent (men) and 43.6 per cent (women), and the standard errors are reduced to 4.3 per cent and 4.4 per cent, respectively. Note that the shrinkage moves these estimates in directions different from the componentwise overall percentages. This is a consequence of incorporating information (borrowing strength) from the municipality's other eight sex-by-age groups. Since the correlations in $\boldsymbol{\Sigma}$ are high and the other groups contain many subjects, a lot of strength is borrowed. Multivariate shrinkage differs from univariate (componentwise) shrinkage substantially. The two ways of shrinkage would coincide if the estimated between-municipality variance matrix $\hat{\boldsymbol{\Sigma}}$ were diagonal; then no strength would be borrowed across groups.

Among those aged 67 or older (188 men and 271 women in the sample for Halden), the sample rates are 6.9 per cent for men and 0.0 per cent for women. The corresponding standard errors are estimated from the national rates, $(1 - f_i)\hat{p}_k(1 - \hat{p}_k)/n_{ik}$: 2.1 per cent (men) and 1.1 per cent (women). Had they been estimated as $(1 - f_i)\hat{p}_{ik}(1 - \hat{p}_{ik})/n_{ik}$, the sample rate $p_{ik} = 0.0$ for Rw would be falsely indicated as precise and no shrinkage should take place. However, in view of the small national percentage, the absence of any employed women in the sample is not in contradiction with a positive (though small) employment rate. The shrinkage estimate for men is 6.8 per cent (standard error is 1.7 per cent), and for women 1.1 per cent (0.9 per cent).

In general, the shrinkage estimates are more closely aligned with the national rates, especially for the less populous municipalities for which more shrinkage takes place. The estimates are difficult to present graphically because they involve 166×10 quantities. In any case, they should be used only for reporting to the municipality concerned. For ranking of municipalities, or for exploration of between-municipality variation, different estimators should be used; see Shen and Louis (1998). Table 3 summarizes the estimated gains in precision. The absolute gains are the differences $\sqrt{\widehat{\mathrm{var}}_s(\hat{\mathbf{p}}_i)} - \sqrt{\widehat{\mathrm{var}}_s(\tilde{\mathbf{p}}_i)}$ and the relative gains are defined as $1 - \sqrt{\widehat{\mathrm{var}}_s(\tilde{\mathbf{p}}_i)/\widehat{\mathrm{var}}_s(\hat{\mathbf{p}}_i)}$. The median, smallest, and largest gain are given for each combination of sex-by-age category (row) and municipality size (column).

The largest relative gains are attained for categories Y, V, and Rw, and the smallest for categories M. The gains are smallest for the most populous municipalities (1:12). The

Table 3. Absolute and relative gains due to shrinkage; reductions of the estimated standard errors of the rates of employment by size of the municipality (sampling fraction) and sex-by-age categories.
For each category and municipality size, the median gain is given, followed by the range of the gains in parentheses

	6,000–7,999 (1:5)		8,000–9,999 (1:7)		10,000–49,999 (1:10)		50,000+ (1:12)	
Absolute gains								
Men								
Tm	1.8	(1.1–2.3)	2.1	(1.5–2.6)	1.7	(0.5–3.1)	0.3	(0.1–0.7)
Ym	1.6	(0.9–2.1)	1.9	(1.4–2.3)	1.6	(0.5–2.5)	0.3	(0.0–0.6)
Mm	0.2	(0.2–0.3)	0.3	(0.3–0.4)	0.2	(0.1–0.4)	0.0	(0.0–0.1)
Vm	2.3	(1.8–4.2)	3.0	(2.0–4.6)	2.5	(0.8–4.7)	0.5	(0.1–1.2)
Rm	0.7	(0.4–1.6)	0.9	(0.5–1.6)	0.8	(0.2–1.8)	0.1	(0.0–0.4)
Women								
Tw	1.6	(1.2–2.1)	1.9	(1.5–2.4)	1.6	(0.5–2.9)	0.3	(0.0–0.6)
Yw	1.8	(1.3–2.2)	2.1	(1.6–2.6)	1.7	(0.5–2.9)	0.3	(0.0–0.5)
Mw	0.3	(0.3–0.5)	0.4	(0.3–0.5)	0.3	(0.1–0.6)	0.0	(0.0–0.1)
Vw	2.1	(1.5–3.7)	2.5	(1.9–4.2)	2.1	(0.5–4.3)	0.3	(0.0–0.9)
Rw	0.5	(0.3–1.0)	0.6	(0.5–1.3)	0.6	(0.1–1.2)	0.1	(0.0–0.3)
Relative gains (per cent)								
Men								
Tm	27	(20–31)	29	(24–33)	26	(13–36)	10	(3–15)
Ym	35	(25–39)	38	(32–41)	35	(19–43)	14	(4–21)
Mm	15	(13–18)	17	(15–19)	14	(7–20)	5	(0–7)
Vm	36	(32–47)	40	(33–49)	37	(21–49)	16	(4–27)
Rm	23	(18–37)	28	(21–37)	25	(10–40)	8	(2–18)
Women								
Tw	24	(20–28)	27	(23–30)	24	(12–33)	8	(2–13)
Yw	33	(27–36)	36	(30–39)	31	(15–41)	11	(2–16)
Mw	15	(13–18)	17	(15–19)	14	(6–20)	4	(0–7)
Vw	33	(27–43)	36	(31–45)	33	(15–46)	11	(2–21)
Rw	33	(26–46)	36	(30–50)	34	(15–49)	13	(4–24)

gains for categories T are not as large as for Y, even though their sample sizes are comparable in most municipalities. This is so because the rates of employment in T vary much more across the municipalities than the rates for Y (see Table 2). For the rates of employment in categories R, the gains are not as large as what the associated sample sizes would suggest. The reason for this is that their rates are very low, and so the sample percentages are quite precise, limiting the scope for improvement. The absolute gains are largely in correspondence with the relative gains. They are largest for the age group 60–66 (V). A wider range of improvement is observed for the sampling ratio 1:10 because it contains the widest range of population sizes.

Without applying shrinkage, 25 per cent reduction of the standard error is equivalent to $1/0.75^2 \doteq 1.75$-fold increase of the sample size. By applying shrinkage, the sample sizes required could be reduced 1.75 times. Further, for smaller sample sizes more shrinkage and larger gains in precision take place, so the 1.75-fold reduction of the sample sizes is associated with a somewhat smaller than $100/75 = 1.33$-fold increase of standard errors.

3.2. Composition of the employed population by sex and age

The nation-level summary of the composition of the employed population by sex and age groups is given in Table 4. The estimated between-municipality correlation matrix derived from $\hat{\boldsymbol{\Sigma}}$ is

$$\begin{array}{cc} \begin{pmatrix} 1.00 & 0.42 & -0.15 & 0.31 & 0.46 & \mathbf{0.52} & -0.31 & \mathbf{-0.77} & -0.07 & -0.16 \\ 0.42 & 1.00 & 0.15 & -0.05 & 0.07 & 0.23 & 0.03 & \mathbf{-0.62} & -0.29 & 0.21 \\ -0.15 & 0.15 & 1.00 & -0.53 & -0.52 & -0.37 & -0.17 & 0.05 & -0.59 & 0.56 \\ 0.31 & -0.05 & -0.53 & 1.00 & 0.64 & 0.09 & -0.42 & -0.50 & 0.63 & -0.50 \\ 0.46 & 0.07 & -0.52 & 0.64 & 1.00 & 0.26 & -0.39 & -0.63 & 0.49 & -0.62 \\ \mathbf{0.52} & 0.23 & -0.37 & 0.09 & 0.26 & 1.00 & -0.07 & -0.38 & -0.08 & -0.15 \\ -0.31 & 0.03 & -0.17 & -0.42 & -0.39 & -0.07 & 1.00 & 0.37 & -0.20 & 0.26 \\ \mathbf{-0.77} & \mathbf{-0.62} & 0.05 & -0.50 & -0.63 & -0.38 & 0.37 & 1.00 & -0.15 & 0.27 \\ -0.07 & -0.29 & -0.59 & 0.63 & 0.49 & -0.08 & -0.20 & -0.15 & 1.00 & -0.50 \\ -0.16 & 0.21 & 0.56 & -0.50 & -0.62 & -0.15 & 0.26 & 0.27 & -0.50 & 1.00 \end{pmatrix} & \begin{array}{c} \text{Tm} \\ \text{Ym} \\ \text{Mm} \\ \text{Vm} \\ \text{Rm} \\ \text{Tw} \\ \text{Yw} \\ \text{Mw} \\ \text{Vw} \\ \text{Rw} \end{array} \end{array}$$

The correlations do not display any simple pattern that would be easy to interpret, but the large (absolute) size of some of them indicates that some information is borrowed across the categories. For instance, the percentages for Tm and Tw are positively correlated (0.52), and Tm and Ym are negatively correlated with Mw (−0.77 and −0.62); these figures are highlighted in the correlation matrix. The relative gains in precision are mainly due to the small variances in $\boldsymbol{\Sigma}$; the sex-age composition of the employed is very homogeneous across the municipalities. Except for the categories Mm and Mw, the composition varies somewhat more for men than for women.

The estimated percentages are in a very wide range, and so the scale on which the changes (sample versus shrinkage) are assessed should be considered carefully. For instance, the odds or the log-odds scales are in some contexts suitable. Since the percentages are large only for Mm and Mw categories, the extent of shrinkage does not reflect the size of the between-municipality standard deviation (the third row in Table 4). The gains in precision are also much larger for these categories than what the sample sizes would suggest. Nontrivial gains are realized even for the largest municipalities. In fact, the variation relative to the overall percentage is much smaller for Mm and Mw than for the other (sparser) categories. In general, the changes are largest for the least populous municipalities.

In the largest municipalities (cities) relatively lower percentages of the employed are

Table 4. Estimated national composition for the sex-by-age categories

	Men					Women				
	Tm	Ym	Mm	Vm	Rm	Tw	Yw	Mw	Vw	Rw
Estimate	2.81	6.52	40.58	3.64	1.18	2.56	5.74	33.63	2.69	0.66
St. error	0.12	0.13	0.27	0.13	0.11	0.09	0.11	0.33	0.10	0.06
$\sqrt{\text{diag}(\hat{\boldsymbol{\Sigma}})}$	0.85	0.87	1.75	0.90	0.79	0.59	0.68	2.23	0.67	0.40

young people (Tm, Ym, and Tw), compared to the smallest municipalities (up to 8,000 residents). The gains in precision can be summarized in the same format as in Table 3; for brevity, details are omitted. A coarse but functional summary of the gains in precision is provided by the number of items which have relative standard errors larger than 30 per cent (that are not reported) or are in the range 20–30 per cent (reported in parentheses). Without applying poststratification, among the 10×166 estimated items (the municipalities with full enumeration are not considered), 271 items (16 per cent) are not reported and 326 items (20 per cent) are reported in parentheses. With shrinkage estimation, only 171 items (10 per cent) would not be reported and 127 (8 per cent) would be reported in parentheses. Among the ten categories, the most striking reduction occurs for teenagers (Tm and Tw, 332 items); nine items are not reported and 117 are reported in parentheses. With shrinkage estimation, every item would be reported and only 30 items would be in parentheses. For categories Vw, Rm, and Rw, most items are not reported. Shrinkage estimates would be reported for many more items in Vw, but the gains are much more modest for Rm and Rw. In the other categories, Y, M, and Vm, all items are reported and only a few of them in parentheses (for the sample estimates only).

3.3. Industrial sectors

The following nine primary industrial sectors are considered: Agriculture, hunting, forestry and fishing (AF); Mining and quarrying including oil production (MQ); Manufacturing (MN); Electricity, gas and water supply (EW); Construction (CB); Wholesale and retail trade, restaurants and hotels (WR); Transport, storage and communications (TC); Finance, insurance, real estate and business services (FB); and Social and personal services (SP). The abbreviations given in the parentheses are used in the tables summarising the results. Each employed resident is associated with one of these nine sectors. Of interest is the composition of the employed population with respect to these sectors in the municipalities.

The estimated national percentages, their standard errors and the associated between-municipality variances are given in Table 5. Despite the substantial sample size, the national percentages are estimated with only modest precision. This is a consequence of the substantial between-municipality variation in the percentages. For instance, the estimated percentage for Agriculture, hunting, forestry and fishing (AF) is 6.09, with standard error 1.50 per cent, and the estimated standard deviation of the municipality-level percentages is 10.56 per cent. This should come as no surprise; in some municipalities, AF is the principal industrial sector, in others its employment in less than 1 per cent (e.g., in some noncoastal urban municipalities). For sectors that are much more evenly distributed across the country, EW, CB,

Table 5. Estimated national percentages for the primary industrial sectors

	Industrial sector								
	AF	MQ	MN	EW	CB	WR	TC	FB	SP
Estimate	6.09	1.17	15.85	1.07	7.53	18.10	7.75	7.69	34.75
St. error	1.50	0.32	1.04	0.20	0.38	0.75	0.30	0.61	0.87
$\sqrt{\mathrm{diag}(\hat{\boldsymbol{\Sigma}})}$	10.56	2.24	7.31	1.42	2.63	5.27	2.10	4.28	6.06

and TC in particular, the standard errors are much smaller, even when taking into account the dependence of the standard error on the estimated percentage.

The large standard deviations might indicate that shrinkage is not very effective; the national percentages are not very useful for estimating municipality-level percentages. However, this would not be the case if the correlations in $\mathbf{\Sigma}$ were high because information could then be ‘borrowed’ across the sectors. The correlation matrix corresponding to $\hat{\mathbf{\Sigma}}$ is

$$\begin{array}{cl}\left(\begin{array}{rrrrrrrrr}
1.00 & -0.07 & -0.20 & 0.19 & 0.28 & \mathbf{-0.74} & -0.17 & \mathbf{-0.76} & 0.40 \\
-0.07 & 1.00 & -0.07 & -0.02 & -0.10 & -0.02 & -0.09 & 0.06 & 0.10 \\
-0.20 & -0.07 & 1.00 & -0.12 & -0.13 & -0.12 & -0.20 & -0.10 & 0.50 \\
0.19 & -0.02 & -0.12 & 1.00 & 0.39 & -0.30 & -0.14 & -0.34 & 0.03 \\
0.28 & -0.10 & -0.13 & 0.39 & 1.00 & -0.28 & -0.13 & -0.51 & 0.16 \\
\mathbf{-0.74} & -0.02 & -0.12 & -0.30 & -0.28 & 1.00 & 0.10 & \mathbf{0.75} & -0.20 \\
-0.17 & -0.09 & -0.20 & -0.14 & -0.13 & 0.10 & 1.00 & 0.19 & -0.10 \\
\mathbf{-0.76} & 0.06 & -0.10 & -0.34 & -0.51 & \mathbf{0.75} & 0.19 & 1.00 & 0.32 \\
0.40 & 0.10 & 0.50 & 0.03 & 0.16 & -0.20 & -0.10 & -0.32 & 1.00
\end{array}\right) & \begin{array}{l}\mathrm{AF}\\ \mathrm{MQ}\\ \mathrm{MN}\\ \mathrm{EW}\\ \mathrm{CB}\\ \mathrm{WR}\\ \mathrm{TC}\\ \mathrm{FB}\\ \mathrm{SP}\end{array}\end{array}$$

Most of the (estimated) correlations are small. Notable exceptions are the high negative correlations of AF with WR and FB (−0.74 and −0.76, respectively), and the positive correlation of WR with FB. They highlight a particular aspect of the differences in the structure of economic activity in rural areas (AF) and in urban areas (WR and FB). For most percentages, little shrinkage takes place because little information is borrowed from other municipalities or across the sectors.

The relative gains in precision of the municipality-level percentages are summarized in Table 6. The absolute gains are smaller than 0.3 for all estimated percentages. The gains are very modest for the first three industrial sectors (AF, MN, and MQ), but they are not trivial for the other sectors. The gains are related to the relative sizes of the national percentages and the between-municipality standard deviations. Of course, the gains are smallest for the most populous municipalities.

Table 6. Relative gains due to shrinkage; reduction of the estimated standard errors of the composition of employment, by size of the municipality (sampling fraction) and industrial sector.
For each sector and municipality size, the median gain is given, followed by the range of gains in parentheses

	Relative gains							
	6,000–7,999 (1:5)		8,000–9,999 (1:7)		10,000–49,999 (1:10)		50,000+ (1:12)	
AF	0.8	(0.0–1.6)	0.8	(0.3–1.5)	0.5	(0.0–0.5)	0.1	(0.0–0.2)
MQ	0.9	(0.0–3.3)	1.3	(0.7–7.9)	0.9	(0.2–4.8)	0.3	(0.0–1.1)
MN	1.6	(1.0–2.5)	2.0	(1.2–2.6)	1.4	(0.4–3.2)	0.3	(0.0–0.6)
EW	3.5	(1.9–8.2)	4.2	(2.3–6.8)	3.3	(0.9–11.9)	0.6	(0.1–1.3)
CB	8.1	(5.3–11.9)	8.9	(6.8–11.3)	7.4	(2.3–13.4)	1.4	(0.2–3.1)
WR	7.9	(6.5–9.9)	8.9	(7.8–10.8)	8.1	(3.1–12.8)	1.9	(0.4–3.7)
TC	8.5	(6.5–13.7)	10.4	(7.5–14.0)	8.5	(2.8–14.6)	1.8	(0.3–3.8)
FB	6.2	(4.8–7.9)	7.6	(5.9–10.7)	7.1	(2.5–13.0)	2.3	(0.4–3.8)
SP	5.1	(3.9–6.3)	5.8	(5.2–7.0)	5.2	(1.8–8.4)	1.1	(0.2–2.1)

3.4. Validation

The most complete form of validation of the method is by estimating population quantities which are known with precision. In our case, we can apply the sample and shrinkage estimation to samples drawn from the data for the municipalities which were fully enumerated. For rates of employment, we drew a random sample (without replacement) with sampling fraction 1:3 from each of these 282 municipalities, all but seven of them with fewer than 6,000 residents each. The between-municipality variance matrix $\mathbf{\Sigma}$ was estimated from these samples and the (original) samples drawn in the larger municipalities.

The shrinkage estimates were closer to the 'true' percentages for 58.3 per cent of the items, ranging from 52.4 per cent for Mw to 63.5 per cent for Vw. A complementary summary is obtained by comparing the means of the 'discrepancy' quantities $(\hat{p}_{ih} - p_{ih})^2/p_{ih}$ and $(\tilde{p}_{ih} - p_{ih})^2/p_{ih}$ for the sample and shrinkage estimates, respectively. The overall mean for the sample estimates is 1.03, whereas for the shrinkage estimates it is 0.65 (percentage-squared). Within the categories, the means of these discrepancy quantities are more than twice as large for the sample versions for Ym and Yw, while the differences are small for Mm and Mw. The shrinkage summaries are smaller for each category.

Naturally, such validation is not complete because it is conducted on an opportunistically selected subsample of municipalities. However, they are 63 per cent of all the municipalities and their validation sample sizes, 72–2,000, and seven municipalities exceeding 2,000, overlap with the sample sizes of the remaining municipalities (1,208–4,787 and eight municipalities with larger sample sizes). Only 30 municipalities have fewer than 1,200 residents.

A similar validation exercise was conducted for the composition data; although comparisons are less impressive for the industrial sectors, shrinkage is shown to be superior in all comparisons.

3.5. Stratification and poststratification

The shrinkage method is not effective for the composition of employment with respect to the industrial sectors because the municipalities are very heterogeneous. Stratification to a few more homogeneous subsets of municipalities might improve the performance of shrinkage. The municipalities are *a priori* classified by their industrial structure, population density, and centrality (proximity to urban centres). However, these groupings, with various aggregations, are not well suited for stratification because they do not reduce the between-municipality variation a great deal.

For instance, Industrial Link is a classification of the municipalities into 22 categories (Statistics Norway 1994b). Category 20, by far the largest, contains 135 municipalities. For this stratum, the between-municipality standard deviations, $\sqrt{\mathrm{diag}(\hat{\mathbf{\Sigma}})}$, are over twice as small as for the whole of Norway for AF (5.04 per cent versus 10.56 per cent) and FB (2.06 per cent versus 4.28 per cent), but the reductions for the other sectors are much more modest. Further, for the remaining 313 municipalities the standard deviations are only slightly smaller than for the entire country, and it is difficult to find for them a suitable partitioning to a small number of strata. For Density and Centrality, two ordinal classifications defined for the municipalities, natural strata are easy to define, but they are much more heterogeneous with respect to the industrial sectors than what their definitions might suggest.

For the industrial sectors, the poststratification estimator of Lillegård (1993) is much more effective and yields much larger gains in precision than the shrinkage estimator. With the information from the registers enabling poststratification at the municipality level, the shrinkage is likely to be less effective than it is without using the register, and it could even become redundant because the within-municipality information would overwhelm information from outside. Nevertheless, poststratification could be combined with shrinkage for the sex-by-age categories. (The register data from 1990 are no longer available.) Also, information could be pooled across the 2×2 municipality-level tables crossclassifying the survey records with the register. Details are given in Longford (1999b).

3.6. *Planning future studies*

If small area estimation is employed in the analysis of a survey of employment, the sample size considerations are somewhat more complex than for the sample estimates. Of course, this could be circumvented by being extremely conservative and anticipating no gains in efficiency of the shrinkage over sample estimation. A more constructive approach is to assume that the between-municipality variance matrix $\mathbf{\Sigma}$, as well as the populations *and* subpopulations of the municipalities, has changed little from the previous Census. Conservatism can be incorporated by inflating the considered variances in $\mathbf{\Sigma}$ (e.g., based on the previous survey). Also, the gains due to multivariateness of the shrinkage can be discarded by assuming that $\mathbf{\Sigma}$ is diagonal. In the planning, the focus should be on the items based on the smallest sample sizes: smallest municipalities, rarest industrial sectors, and teenagers. For municipality-level sample size n_i, the sampling variance of a sample proportion is $u_i = p_i(1 - p_i)/n_i$, and $u_i \doteq p(1 - p)/n_i$, where p is the national proportion. An approximate upper bound for the EMSE of the shrinkage estimator of the proportion is $u_i\sigma^2/(u_i + \sigma^2)$, where σ^2 is the between-municipality variance (an element of $\mathbf{\Sigma}$). This approximation enables a simple discussion of the association of the sample size n_i and EMSE of $\tilde{p}_i$. Information from past studies can be incorporated by defining suitable auxiliary variables.

4. Discussion

Small area methods are an indispensable tool for efficient estimation of quantities defined at the level of an administrative or geographical division of a country. Historically, the product of a survey collecting such information has been tables of sample estimates (means, proportions, or percentages) for various subpopulations and variables, with an implied argument that secondary users could improve on these estimates. The constructor of such tables is much better suited for this task because they are more likely to have the statistical and computing expertise and facilities as well as easier access to background information. Also, they can avoid the duplication of effort that would be expended by the secondary users.

The algorithms for the multivariate shrinkage method applied in Section 3 involve no iterations and are computationally undemanding. The most complex operations involved are matrix inversion or solving linear systems of equations (near singularity can arise in them only in some esoteric settings). Concerns about the validity are pertinent only when the method is applied to a moderate or small number of small areas, because then the variance matrices required are estimated with little precision. This concern does not arise in our analyses. The method can be adjusted to be conservative simply by inflating the estimated between-area variance matrix. In this way, more weight is assigned to the local information and, as a consequence, a less biased estimator is obtained. The chances that such an estimator would be conditionally less efficient than the sample estimator are reduced, in exchange for less optimality on average. Of course, extreme conservatism leads to the sample estimator which is rather inefficient for small areas.

The method can be validated directly by subsampling from the survey data, and can be combined with poststratification, or applied separately to each of a small number of strata. The software implementing the multivariate shrinkage method, developed in Splus, can be obtained from the author on request.

5. References

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, 28–36.

Elston, J.M., Koch, G.G., and Weissert, W.G. (1991). Regression-adjusted Small Area Estimates of Functional Dependency in the Noninstitutionalized American Population Age 65 and Over. American Journal of Public Health, 81, 335–343.

Fay, R.E. (1987). Application of Multivariate Regression to Small Domain Estimation. In Small Area Statistics. An International Symposium, eds. R. Platek, J.N.K. Rao, C.-E. Särndal, and M.P. Singh, New York: John Wiley and Sons, 91–102.

Fay, R.E. and Herriot, R. (1979). Estimates of Income for Small Places: An Application of James-Stein Procedures to Census Data. Journal of the American Statistical Association, 74, 269–277.

Gelman, A., Carlin, J.B., Stern, H.S., and Rubin, D.B. (1995). Bayesian Data Analysis. London: Chapman and Hall.

Ghosh, M. and Rao, J.N.K. (1994). Small-area Estimation: An Appraisal. Statistical Science, 9, 55–93.

Lillegård, M. (1993). Folke- og boligtelling 1990. Dokumentasjon av de statistiske metodene. (Census of the Population and Dwellings 1990. Documentation of the Statistical Methods.) Rapporter 93/4. Statistisk sentralbyrå (Statistics Norway), Oslo, Norway.

Longford, N.T. (1996). Model-based Variance Estimation in Surveys with Stratified Clustered Design. The Australian Journal of Statistics, 38, 333–352.

Longford, N.T. (1999a). Multivariate Shrinkage Estimation of Small-area Means and Proportions. Journal of the Royal Statistical Society, Ser. A, 162, 227–245.

Longford, N.T. (1999b). Combining Information About Small Areas from Surveys and Imperfect Registers. Unpublished manuscript.

Marker, D. A. (1999). Organization of Small Area Estimators Using a Generalized Linear Regression Framework. Journal of Official Statistics, 15, 1–24.

Shen, W. and Louis, T. A. (1998). Triple-goal Estimates in Two-stage Hierarchical Models. Journal of the Royal Statistical Society, Ser. B, 60, 455–471.

Singh, A.C., Stukel, D.M., and Pfeffermann, D. (1998). Bayesian Versus Frequentist Measures of Error in Small Area Estimation. Journal of the Royal Statistical Society, Ser. B, 60, 377–396.

Statistisk sentralbyrå (Statistics Norway). (1994a). Folke- og boligtelling 1990 – Evaluering. Instilling fra et utvalg somme har vurdert tellingen. (Report of the committee evaluating the census methodology.) Notater 1994/3. Oslo, Norway.

Statistisk sentralbyrå (Statistics Norway). (1994b). Standard for kommuneklassifisering 1993. (Standard for Classification of Municipalities 1993.) Official Statistics of Norway, C 185. Oslo, Norway.

Swensen, A.R. (1988). Estimating Change in a Proportion by Combining Measurements from a True and a Fallible Classifier. Scandinavian Journal of Statistics, 15, 139–145.

Thomas, N., Longford, N.T., and Rolph, J.E. (1994). Empirical Bayes Methods for Estimating Hospital Specific Mortality Rates. Statistics in Medicine, 13, 889–903.

Received October 1999
Revised January 2001