STATISTICS IN TRANSITION-new series, October 2011
Vol. 12, No. 2, pp. 345—356

PERCENTILE-ADJUSTED ESTIMATION OF POVERTY INDICATORS FOR DOMAINS UNDER OUTLIER CONTAMINATION

Ari Veijanen[1], **Risto Lehtonen**[2]

ABSTRACT

Traditional estimation of poverty and inequality indicators, such as the Gini coefficient, for regions does not currently use auxiliary information or models fitted to income survey data. A predictor-type estimator constructed from ordinary mixed model predictions is not necessarily useful, as the predictions have too small spread for estimation of income statistics. Ordinary bias corrections are aimed at correcting the expectation of predictions, but poverty indicators would not be affected at all by a correction involving multiplication of predictions. We need a method improving the shape of the distribution of predictions, as poverty indicators describe differences of income between people. We therefore introduce a transformation bringing the percentiles of transformed predictions closer to the percentiles of sample values. The experiments show that the transformation results in smaller MSE of a predictor. If unit-level data from population are not available, the marginal domain frequencies of qualitative auxiliary variables can be successfully incorporated into a new calibration-based predictor-type estimator. The results are based on design-based simulation experiments where we use a population generated from an EU-wide income survey. The study is a part of the AMELI project funded by the European Union under the Seventh Framework Programme for research and technological development (FP7).

Key words: small area estimation; poverty indicator; income data; bias correction; auxiliary information; mixed model; prediction.

1. Introduction

Regional income statistics have received a lot of attention in recent years. Thresholds for a poverty indicator have been used in regional allocation of resources (e.g. Zaslavsky and Schirm, 2002). The European Union conducts regular income surveys (Statistics on income and living conditions, SILC, see e.g.

1 Statistics Finland. E-mail: ari.veijanen@stat.fi
2 University of Helsinki. E-mail: risto.lehtonen@helsinki.fi

Clemenceau, Museux and Bauer, 2006) yielding information about income and social attributes of households. In the AMELI project (Advanced Methodology for European Laeken Indicators), we studied the estimation of indicators on poverty and social exclusion (the so-called Laeken indicators) including poverty gap, quintile share ratio, and the Gini coefficient. These statistics are nonlinear, so methods designed for estimation of domain totals, such as GREG (Särndal, Swensson and Wretman, 1992; Lehtonen and Veijanen, 2009) or EBLUP (empirical best linear unbiased predictor) cannot be applied. In this paper, we introduce new methods for estimation of poverty indicators in regions (small areas) and other domains.

The equivalized income constitutes the key variable underlying the monetary poverty indicators. It is defined as a household's total disposable income divided by its "equivalent size", to take account of the composition of the household (European Commission, 2003). Equivalization is made on the basis of the OECD modified scale, which assigns weight 1.0 for the first adult, 0.5 for every additional person aged 14 or over, and 0.3 for every child under 14. The equivalized income is attributed to each household member including children, and the poverty indicators are defined for unit-level data composed of persons.

We compare currently widely used “default” estimators of poverty indicators with corresponding predictor-type estimators. A design-based default estimator incorporates design weights and sample values but does not use auxiliary data or modelling. It is a so-called direct estimator, as observations from other domains do not contribute to a domain estimate (Lehtonen and Veijanen, 2009, p. 223). The direct estimator probably has small design bias, but its variance can be large in domains with a small sample size. In an indirect predictor-type estimator based on unit-level auxiliary information, predictions are plugged into the formula of the poverty indicator defined at the population level. The estimator is expected to have small variance but it may be seriously design biased.

As the equivalized income is approximately distributed as lognormal, a model is often fitted to log-transformed data and the fitted values are back-transformed to the original scale. In estimation of domain totals, this should be followed by a bias correction, such as RAST (Ratio Adjusted by Sample Total; Chambers and Dorfman, 2003; Fabrizi, Ferrante and Pacei, 2007). Chandra and Chambers (2011) discuss non-linear transformations and introduce more accurate bias corrections. However, the quintile share especially is determined by the shape of the income distribution, not its average. As the distribution of predictions is concentrated around the average, they yield too large quintile shares or too small Gini coefficients and poverty gaps. We define a transformation that brings the percentiles of transformed predictions closer to the percentiles of sample values. Our technique is model-based but it aims to correct for design bias with the use of design weights, so it is comparable to design consistent pseudo synthetic, pseudo EBLUP and pseudo EBP (empirical best predictor) type approaches (Rao 2003; You and Rao, 2002; Jiang and Lahiri, 2006).

Predictions can be avoided altogether by simulating unknown observations from their conditional distribution given the sample values and auxiliary data (Molina and Rao, 2010). However, in our simulations deviations from the assumed model resulted in large bias. Estimating quantiles of the income distribution by M-quantile regression (Chambers and Tzavidis, 2006) might also provide an alternative basis for small area estimation of poverty indicators.

This article is organized as follows. Section 2 contains notation and basic definitions. Section 3 introduces the technique of transforming the predictions. As unit-level auxiliary data are not always available, we introduce in Section 4 new frequency-calibrated estimators incorporating known marginal totals of auxiliary variables. In Section 5, we evaluate the design bias and accuracy of the estimators by Monte Carlo experiments using a synthetic Amelia population, which is based on the EU-wide SILC survey (Alfons et al., 2011). Short discussion is in Section 6.

2. Definitions

2.1. Domain models

The fixed and finite population of interest is denoted $U = \{1, 2, ..., k, ..., N\}$, where k refers to a unit's label. The population is divided into subsets called *domains* by region, for example. Each domain U_d, indexed by d, has N_d elements. The sample s is composed of corresponding subsets s_d of size n_d. The domains are called unplanned unless the sampling design is stratified by domains (Lehtonen and Veijanen, 2009, p. 222). Design weights, inverses of inclusion probabilities, are denoted by a_k.

To account for differences between domains, a linear mixed model incorporates domain-specific random effects $u_d \sim N(0, \sigma_u^2)$. The model is given by

$$Y_k = \mathbf{x}_k'\boldsymbol{\beta} + u_d + \varepsilon_k, k \in U_d, \varepsilon_k \sim N(0, \sigma^2).$$

The random effects may also be associated with aggregates of domains. The parameters $\boldsymbol{\beta}$, σ_u^2 and σ^2 are first estimated from the data by using ML or REML methods, and the values of the random effects are then predicted. This yields predictions $\hat{y}_k = \mathbf{x}_k'\hat{\boldsymbol{\beta}} + \hat{u}_d, k \in U_d$. Typically, design weights are not incorporated in the estimation of the model.

2.2. Poverty indicators

The *S20/S80 ratio,* or *quintile share ratio*, compares the average equivalized incomes in the poorest and the richest quintile. Each quintile contains 20 % of people; in the design-based case accounting for 20 % of design weights. The *default* (direct) *estimators* of the first (S20) and the fifth quintile average (S80) are Hájek estimators, that is, weighted domain averages involving design weights.

The direct quintile share estimate is the ratio of S20 to S80. The predictor-type estimator of quintile share in a domain is the ratio of averages of predictions in the first and fifth quintiles.

Consider ordered equivalized incomes *y* in a domain U_d. The *Lorenz curve* L_d describes how the first *k* ($k = 1, 2, ..., N_d$) persons' proportion of the total income depends on their numerical proportion. The curve is approximated by a line between points

$$L_d\left(\frac{k}{N_d}\right) = \frac{\sum_{i \le k; i \in U_d} y_{(i)}}{\sum_{t \in U_d} y_t} \quad (y_{(1)} \le y_{(2)} \le ... \le y_{(N_d)}) \tag{1}$$

The *Gini coefficient* G_d is defined as

$$G_d = 1 - 2\int_0^1 L_d(x)dx . \tag{2}$$

A direct estimator of the Lorenz curve is defined in ordered sample: for unit *k*,

$$L_{HT;d}\left(\frac{\sum_{i \le k; i \in s_d} a_i^s}{\sum_{t \in s_d} a_t}\right) = \frac{\sum_{i \le k; i \in s_d} a_i^s y_{(i)}}{\sum_{t \in s_d} a_t y_t} \quad (y_{(1)} \le y_{(2)} \le ... \le y_{(n_d)}),$$

where a_i^s denotes the design weight of the *i*th element in the ordered sample. The predictor-type estimator of the Lorenz curve is simply obtained by plugging the ordered predicted incomes into (1). The *default* (direct) estimator and the predictor of the Gini coefficient are then defined as integrals similar to (2).

Poverty gap in a domain describes the difference between the at-risk-of-poverty threshold *r* (typically 60 % of the median income in the whole country) and the median income $m_d(r) = Md\{y_k; y_k \le r; k \in U_d\}$ below the threshold:

$$g_d = \left(r - m_d(r)\right)/r . \tag{3}$$

The *default* (direct) *estimator* of the poverty gap is calculated as a similar ratio involving estimated threshold $\hat{r}$ and the design-based estimate of median income below the threshold obtained from estimated distribution function

$$\hat{F}_{HT}(y;\hat{r}) = \frac{\sum_{k \in p_d} a_k I\{y_k \le y\}}{\sum_{k \in p_d} a_k}; \; p_d = s_d \cap \{k : y_k \le \hat{r}\} .$$

The predictor contains $\hat{r}$ and predictions' median $\hat{m}_{pred;d}(\hat{r}) = Md\{\hat{y}_k; \hat{y}_k \le \hat{r}; k \in U_d\}$.

3. Percentile-adjusted predictors

The distribution of the equivalized income y is skewed with a long right tail. A candidate for a distributional assumption is the lognormal distribution, but it is often not realistic. In any case, after logarithmic transformation a model fits the data better. For pragmatic reasons, a model is fitted to $z_k = \log(y_k + 1)$, and the fitted values $\hat{z}_k$ are back-transformed to $\hat{y}_k = \exp(\hat{z}_k) - 1$. To correct for bias in estimates of domain totals, Chambers and Dorfman (2003) and Fabrizi, Ferrante and Pacei (2007) calculate a RAST multiplier c_R ensuring that the weighted sample sum of $c_R \hat{y}_k$ equals the corresponding sum of original incomes. However, the poverty indicators are not affected at all by multiplication amounting to a change in currency. Correction for shape of distribution is also required. We correct for bias and spread of predictions $\hat{y}_k = \exp(\hat{z}_k) - 1$ ($k \in U_d$) by a nonlinear transformation that brings the distribution of predictions closer to the distribution of observed values y_k ($k \in s_d$) in terms of percentiles, denoted by $\hat{p}_{cd}$ and p_{cd}, respectively. The percentiles p_{cd} of sample values are obtained from the estimated cumulative distribution function

$$\hat{F}_{HT;d}(y) = \frac{1}{\hat{N}_d} \sum_{k \in s_d} a_k I\{y_k \le y\};\ \hat{N}_d = \sum_{k \in s_d} a_k .$$

Our goal is to obtain transformed predictions $\tilde{y}_k = e^{\alpha_d} \hat{y}_k^{\gamma}$ whose percentiles, denoted $\tilde{p}_{cd}$, are close to p_{cd} on logarithmic scale. The using of logarithms reduces the influence of large percentiles. To avoid unstable estimates in the smallest domains, we pooled the percentile data from all domains and minimized

$$\sum_d \sum_{c=1}^{C} \left(\log\left(\tilde{p}_{cd}\right) - \log\left(p_{cd}\right)\right)^2 = \sum_d \sum_{c=1}^{C} \left(\alpha_d + \gamma \log\left(\hat{p}_{cd}\right) - \log\left(p_{cd}\right)\right)^2 ,$$

where C=50 for poverty gap and C=99 for other indicators. The *percentile-adjusted*, or *p-adjusted*, predictions involve OLS estimates of parameters α_d and γ:

$$\log\left(\tilde{y}_k\right) = \hat{\alpha}_d + \hat{\gamma} \log\left(\hat{y}_k\right) (k \in U_d) . \qquad (4)$$

We experimented also with a mixed model fitted to the percentile data but it did not yield significantly different results.

About 1.5 % of the people in the samples of our experiments had zero equivalized income. To incorporate zero predictions in the predictors we replaced a roughly identical proportion of the smallest predictions in each domain by zero. The transformation (4) was applied only to the positive predictions, with percentiles $\hat{p}_{cd}$ and p_{cd} calculated from positive predictions and sample values. Occurrence of zeroes could also be described by a model that includes a logistic component (Karlberg, 2000).

4. Frequency-calibrated predictors

We develop here a new method that may be feasible in situations where only aggregate-level qualitative auxiliary data are available. Suppose only the domain sizes and domain totals of qualitative auxiliary variables are known. Then we cannot access the values of $\mathbf{x}_k = \left(x_{1k}, x_{2k}, ..., x_{pk}\right)$ for every unit. However, for the calculation of a predictor, it is actually enough to know the population frequencies of distinct values of $\mathbf{x}_k$, that is, the frequencies of cells in the crosstabulation of the **x**-variables. We propose a method of estimating these frequencies.

Consider domain *d*. Let us denote the set of distinct values of $\mathbf{x}_k$, $k \in s_d$, by $X_d = \{z_1, z_2, ..., z_m\}$. A direct estimate of the domain frequency of $z \in X_d$ is

$$\hat{n}_z = \sum_{k \in s_d} a_k I\{\mathbf{x}_k = z\} .$$

These do not, in general, sum up to the known marginal totals t_d. This requirement is formulated as a calibration equation

$$\sum_{k \in U_d} \mathbf{x}_k = \sum_{z \in X_d} n_z z = t_d . \qquad (5)$$

Calibration (e.g. Singh and Mohl, 1996; Särndal, 2007) yields new frequencies $\tilde{n}_z$ that are close to the $\hat{n}_z$ and also satisfy the calibration equations. We measure the distance of $\tilde{n} = (\tilde{n}_z; z \in X_d)$ to $\hat{n} = (\hat{n}_z; z \in X_d)$ by chi-squared distance

$$\sum_{z \in X_d} \frac{1}{\hat{n}_z} \left(\hat{n}_z - \tilde{n}_z\right)^2 .$$

This distance is minimized subject to the calibration equations (5) by

$$\tilde{n}_z = \hat{n}_z \left(1 + \lambda_d' z\right), \qquad (6)$$

where the Lagrange multiplier λ_d is

$$\lambda_d = \left(t_d - \sum_{z \in X_d} \hat{n}_z z \right) \left(\sum_{z \in X_d} \hat{n}_z z z' \right)^{-1} .$$

To avoid singular matrices, we excluded from each $z \in X_d$ the indicator variables of classes that did not appear in the sample domain. Moreover, if two variables had identical values in the domain, the latter variable was removed. Corresponding modifications were made in the vector t_d. If the algorithm still failed due to linear dependencies, for example, we used the initial estimates $\hat{n}_z$. This occurred rarely. Unfortunately, about 10 % of the $\tilde{n}_z$ were negative in our simulations. We replaced them by zero. After this, the calibration equations do not necessarily hold. Negative estimates might be avoided by the approach of

Singh and Mohl (1996) in taking into account range restrictions of calibrated weights.

The vector of predictions in the domain is finally obtained by repeating $\tilde{n}_z$ times the fitted value associated with each $z \in X_d$ (using rounded frequency estimates). Transformation (4) may then be applied. We call the resulting predictor a *frequency-calibrated*, or an *n-calibrated* predictor.

5. Monte carlo simulation experiments

5.1. Introduction

We present numerical results from design-based Monte Carlo simulation experiments with synthetic Amelia data set constructed by Alfons et al. (2011) to mimic the regional and demographic variation of income statistics in European Union. It contains about ten million persons. We applied SRSWOR (n = 2000). As domains we used 40 regions (variable DIS) or, in the case of poverty gap, 110 demographic population subsets defined by age class, gender and NUTS2 region. The domains were classified to minor, medium and major domains by expected sample size with class boundaries at 45 and 55 units for the DIS regions and at 20 and 30 units for the demographic domains. Our models fitted to equivalized income variable EDI2 incorporated age class and gender with interactions, attained education level (ISCED), activity (working, unemployed, retired, or otherwise inactive) and degree of urbanisation of residence (three classes). The mixed models, which had random intercepts associated with districts (DIS) or NUTS2 regions (in the case of poverty gap), were fitted by R package nlme using ML without design weights.

In contamination experiments, outliers were created in each sample without modifying the population. One percent of sampled persons were declared as outliers, chosen completely at random, and a normally distributed value from $N(500000, 10000^2)$ was added to the personal cash or near-cash income of an outlier. We evaluated the estimates by absolute relative bias (ARB, absolute average error divided by the true value) and the relative root mean squared error

$$RRMSE = \frac{1}{\theta_d}\sqrt{\frac{1}{K}\sum_{k=1}^{K}(\hat{\theta}_{dk} - \theta_d)^2} \ .$$

5.2. Results

The results for quintile share, Gini coefficient and poverty gap are in tables 1, 2 and 3, respectively. Domain size is the most important factor affecting accuracy of estimation in a domain. In general, ARB and RRMSE were largest in small domains. With a direct estimator and small samples, the estimates vary greatly, and show too large disparities between domains.

Table 1. Results with quintile share in regions (DIS).

Estimator	*ARB (%)*				*RRMSE (%)*			
	Expected domain sample size				*Expected domain sample size*			
	minor	*medium*	*major*	*all*	*minor*	*medium*	*major*	*all*
No contamination								
Direct	4.9	4.6	3.4	4.4	43.5	41.7	38.5	41.3
Ordinary predictor	110.2	125.8	129.6	122.2	113.7	129.4	133.2	125.7
p-adjusted predictor	12.3	8.6	5.7	8.9	16.0	13.6	11.4	13.7
n-calibrated predictor	11.1	13.3	10.6	11.9	31.3	29.6	25.9	29.1
Contaminated								
Direct	7.9	9.1	10.8	9.2	43.8	41.8	39.3	41.7
p-adjusted predictor	14.3	8.5	5.7	9.5	18.1	14.2	12.2	14.8
n-calibrated predictor	10.9	10.3	7.0	9.6	30.6	27.7	23.7	27.5

Table 2. Results for Gini coefficient in regions (DIS).

Estimator	*ARB (%)*				*RRMSE (%)*			
	Expected domain sample size				*Expected domain sample size*			
	minor	*medium*	*major*	*all*	*minor*	*medium*	*major*	*all*
No contamination								
Direct	2.6	2.1	1.6	2.1	12.9	11.6	10.4	11.7
Ordinary predictor	21.7	23.3	23.7	22.9	22.6	24.0	24.4	23.7
p-adjusted predictor	10.7	8.4	7.7	8.9	11.5	9.3	8.7	9.8
n-calibrated predictor	6.2	4.3	3.7	4.7	11.3	9.1	7.8	9.4
Contaminated								
Direct	7.8	8.7	9.8	8.7	21.5	20.6	19.9	20.7
p-adjusted predictor	12.7	10.3	9.6	10.8	13.4	11.1	10.5	11.6
n-calibrated predictor	7.8	6.0	5.5	6.4	12.5	10.1	9.0	10.5

Table 3. Results of poverty gap in domains defined by age class, gender and NUTS2 region.

Estimator	*ARB (%)*				*RRMSE (%)*			
	Expected domain sample size				*Expected domain sample size*			
	minor	*medium*	*major*	*all*	*minor*	*medium*	*major*	*all*
No contamination								
Direct	6.5	3.3	2.4	5.5	51.8	43.8	38.5	48.8
Ordinary predictor	36.9	42.5	39.4	38.2	46.7	46.8	43.2	46.4
p-adjusted predictor	18.1	23.8	22.0	19.6	44.4	37.8	30.7	41.6
n-calibrated predictor	14.1	18.4	20.8	15.7	63.0	51.2	41.4	58.4
Contaminated								
Direct	6.2	3.0	2.4	5.2	51.6	43.7	38.4	48.7
p-adjusted predictor	17.1	23.0	21.3	18.7	44.4	37.6	30.4	41.6
n-calibrated predictor	14.3	17.9	20.1	15.7	63.3	51.3	41.5	58.6

The ordinary predictor was usually so biased that it had even larger RRMSE than the direct estimator. The percentile-adjusted predictor and the n-calibrated predictor benefitted greatly from the transformation (4) bringing the distribution of predictions closer to the distribution of observations. Both bias and RRMSE decreased. The percentile-adjusted predictor had smaller RRMSE than the direct estimator in every domain size class. Moreover, these predictors were more robust to contamination than the direct estimator, although in the case of poverty gap all estimators were robust as the median is not affected by outliers.

The frequency-calibrated estimator (Eq. 6) was not usually as accurate as the p-adjusted predictor. This was expected, as the frequency-calibrated predictor has no access to unit-level information. The n-calibrated predictor was better than the direct estimator of quintile share and Gini coefficient, but not better than the direct estimator of poverty gap. The estimator appears to have similar robustness properties as the percentile-adjusted predictor, since the transformation (4) was applied.

According to our experiments, the method of Molina and Rao (2010) seems to be sensitive to deviations from assumed lognormal distribution of income (results not shown here). In our data, the right tail of the income distribution was “thinner” than expected under the distributional assumption. The domain maxima of the simulated income were often at least ten times larger than in the data and the estimated quintile shares were about 10 % of true values.

6. Discussion

The percentile-adjusted predictors incorporating transformed predictions yielded substantial improvements over current default method and ordinary domain predictor. When the transformation incorporates percentiles of observations up to the 99th percentile, rare outliers occurring with frequency smaller than one percent do not affect the percentile-adjusted predictor too much. The breakdown point of the estimator can probably be adjusted by changing the range of percentage points used in the transformation. The frequency-calibrated estimator (Eq. 6) performed surprisingly well. The estimation of mean square error with bootstrap appears feasible, although time-demanding, as it is necessary to fit the mixed model to each bootstrap sample.

Acknowledgments

This work was completed as a part of the AMELI project, which was supported by European Commission funding from the Seventh Framework Programme for research and technological development. Statistics Finland supported the work of the first author.

REFERENCES

ALFONS, A., TEMPL, M., FILZMOSER, P., KRAFT, S., HULLIGER, B, KOLB, J.-P. and MUNNICH, R., 2011. *Report on outcome of the simulation study.* (Deliverable 6.2 of WP6) [online] Trier: Project AMELI. Available at: <http://svn.uni-trier.de/AMELI>.

CHAMBERS, R. L. and DORFMAN, A. H., 2003. Transformed variables in survey sampling, Working paper M03/21, Southampton Statistical Sciences Research Institute.

CHAMBERS, R. and TZAVIDIS, N., 2006. M-quantile models for small area estimation. *Biometrika*, 93, pp. 255-268.

CHANDRA, H. and CHAMBERS, R., 2011. Small area estimation under transformation to linearity. *Survey Methodology*, 37, pp. 39-51.

CLEMENCEAU, A., J., MUSEUX M. and BAUER, M., 2006. *EU - SILC (Community statistics on Income and Living Conditions): issues and challenges.* [online]. Available at:
<http://www.stat.fi/eusilc/clemenceau_museux_bauer_rev2.pdf>.

European Commission (EUROSTAT Working Group statistics on income, poverty and social exclusion), 2003. *Laeken indicators. Detailed calculation methodology* (DOC. E2/IPSE/2003). Luxembourg.

FABRIZI, E., FERRANTE, M. R. and PACEI, S., 2007. Comparing alternative distributional assumptions in mixed models used for small area estimation of income parameters. *Statistics in Transition*, 8, pp. 423-439.

JIANG, J. and LAHIRI, P., 2006. Mixed model prediction and small area estimation. *Sociedad de Estadistica e Investigacion Operative Test*, 15, pp. 1-96.

KARLBERG, F., 2000. Survey estimation for highly skewed populations in the presence of zeroes. *Journal of Official Statistics*, 16, pp. 229-241.

LEHTONEN, R. and VEIJANEN, A., 2009. Design-based methods of estimation for domains and small areas. In: C. R. Rao and Pfeffermann, D., eds. *Handbook of statistics, vol. 29(B). Sample surveys: theory, methods and inference*, pp. 219-249. Oxford: Elsevier.

MOLINA, I. and RAO, J. N. K., 2010. Small area estimation of poverty indicators. *The Canadian Journal of Statistics*, 38, pp. 369-385.

RAO, J. N. K., 2003. *Small area estimation*. New York: John Wiley & Sons.

SINGH, A. C. and MOHL, C. A., 1996. Understanding calibration estimators in survey sampling. *Survey Methodology*, 22, pp. 107-115.

SÄRNDAL, C.-E., 2007. The calibration approach in survey theory and practice. *Survey Methodology*, 33, pp. 99-119.

SÄRNDAL, C.-E., SWENSSON, B. and WRETMAN, J., 1992. *Model assisted survey sampling*. New York: Springer-Verlag.

YOU, Y. and RAO, J. N. K., 2002. A Pseudo-Empirical Best Linear Unbiased Prediction Approach to Small Area Estimation Using Survey Weights. *The Canadian Journal of Statistics,* 30, pp. 431-439.

ZASLAVSKY, A. M. and SCHIRM, A. L., 2002. Interactions between survey estimates and federal funding formulas. *Journal of Official Statistics*, 18, pp. 371-391.