LARGE NUMBERS OF ESTIMATES FOR SMALL AREAS

Allen L. Schirm, Mathematica Policy Research; Alan M. Zaslavsky, Harvard University; and John L. Czajka, Mathematica Policy Research

1. Introduction

National surveys and even national samples of administrative records can typically support only imprecise state or substate estimates because state sample sizes are small. Borrowing strength with an empirical Bayes or similar indirect estimator is a common solution to the problem of imprecise direct estimates, and has been used successfully in many applications. For example, an indirect estimator has been used for several years to derive state estimates for allocating federal funds under the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC) (Schirm and Long 1995). Similar estimators have also been used to obtain state and county estimates of poor school-aged children for allocating federal Title I funds for compensatory education in elementary and secondary schools (National Research Council 1998).

The estimators used in these applications are suitable for deriving a single estimate or a few closely related estimates for each geographic area. However, they are not suitable for deriving large numbers of estimates--for example, filling in a large table--for each area. The problem is that the modeling undertaken for empirical Bayes or similar estimation is specific to the estimates being produced, and it would not be practical to develop a model for each cell of a large table.

To address this problem, Schirm and Zaslavsky (1997) proposed a method for reweighting a survey or administrative records database to borrow strength and improve precision. A Poisson regression model is fitted to obtain an estimated prevalence in each state (or other small area) of every household type in the database, where types are defined by households' characteristics. This model is specified to control important aggregates at the state level, and the prevalences are expressed as a matrix of weights, with each household having a weight for every state. After this Poisson regression model is estimated, no further modeling is required. Any estimates sought for a state are obtained using all the households in the database, not just the households actually in that state. By applying the appropriate weight for each household, the database is weighted to look like the state, rather than the whole country. This reweighting method can be combined with empirical Bayes methods by using empirical Bayes estimates as control totals in the reweighting.

This paper describes two applications of this reweighting approach. The first application is the derivation of state estimates of the numbers and characteristics of children without health insurance. Such estimates are currently being prepared to assist state and federal efforts in designing, implementing, and evaluating state Children's Health Insurance Program (CHIP) initiatives. The need to borrow strength combined with the need to develop several fairly large tables of estimates for each state makes the reweighting method an appropriate approach.

The second application discussed in this paper is microsimulation of proposed food stamp and welfare program reforms. Microsimulation has been used extensively for many years to assess the national effects of reforms to these programs. However, welfare reform has heightened the importance of state-level analyses. An approach for borrowing strength is needed because state samples in all of the relevant databases are small. Because any one execution of a microsimulation model produces many tables of estimates, the reweighting approach is appropriate for this application as well. In fact, this second application is substantially similar to the first, although microsimulation generally entails more complex manipulations of the microdata and the derivation of many more estimates. An important difference between the two applications arises because of the nature of the interaction between microsimulation and the policymaking process. While the tables to be produced by a microsimulation may be well specified before the database is reweighted, the analyst who will carry out the reweighting may not know very well which of a wide range of potential program reforms will be proposed by policymakers and, thus, will need to be simulated. In contrast, in an application like the derivation of estimates of uninsured children, the policies under consideration are more limited and better known if not already finalized.

Before describing further these two applications of the reweighting method, we describe the method itself. Then, we conclude by discussing very briefly the questions addressed in our ongoing evaluation of the reweighting method. In that evaluation, we are estimating the biases and variances of the model-based estimates from a reweighted database and evaluating the model-based estimates relative to direct sample estimates for key estimands. Results from the evaluation will be presented in the future.

2. Reweighting to Borrow Strength and Improve Precision

Basic Ideas. Schirm and Zaslavsky (1997) proposed a method for reweighting sample observations in a database to borrow strength and improve precision. The basic idea of the reweighting approach is to use households from many states to borrow strength and improve precision when deriving estimates for any one state. How reweighting can be used to borrow strength is illustrated by comparing (1) the direct estimator that uses the original sample weights and does not borrow strength with (2) the indirect estimator that uses reweighted data and does borrow strength.

To calculate an estimate for Virginia, for example, the direct estimator uses only the sample households for Virginia and their original sample weights. Observations for other states are ignored. This is equivalent to using all the observations in the database weighted by "Virginia weights" that equal the original sample weights for households in Virginia but are zero for households in all other states.

In contrast, for indirect estimation, nonzero Virginia weights would be assigned to households in not only Virginia but also other states. For example, suppose that the only two states in our database are Virginia and Maryland. Suppose also that each state has a sample household with similar income and other characteristics, but the only child in the Virginia household is age 4, while the only child in the Maryland household is age 5. If the age distributions in the two states are similar, each of the households (or ones like them) could about equally well have appeared in the other state. In other words, the presence of the household with the 4 year old in Virginia rather than Maryland reflects sampling variability. Thus, if a count of the number of households with children under age 5 is needed or if some program provision being simulated with a microsimulation model is triggered by the presence of a child under age 5, better estimates could be obtained for both states by giving each state a copy of each of the two households but with half as much weight. On the other hand, if we have evidence that the age distribution of children in Maryland is shifted upward relative to the age distribution in Virginia, we might want to give the household with the older child slightly more weight in Maryland and slightly less weight in Virginia.

Although giving weight to out-of-state households introduces some bias, using many more observations that are similar except for state of residence should substantially improve precision. The objective of reweighting and, more generally, indirect estimation is to enhance accuracy as measured by a standard like mean squared error (MSE) that reflects the tradeoff between bias and variance.

Under our proposed approach, we derive a matrix of state weights. Every household in the database gets as many new weights as there are states (51 counting the District of Columbia as a state). For every state, there is a weight for each household in the database, although some weights may be small or (by design) zero. To derive estimates for any one state, we use all households in the database--regardless of actual state of residence--and apply the appropriate set of weights. Virginia weights are used to derive estimates for Virginia, Maryland weights are used to derive estimates for Maryland, and so forth. Thus, Virginia borrows strength from other states that have households with nonzero Virginia weights.

Using a Poisson regression model, our reweighting method assigns a Virginia weight to a household according to how prevalent that "type" of household is in Virginia. The more prevalent it is, the more Virginia weight it gets. A household's type is defined by all the characteristics in the database, some of which are measured directly while others are calculated or simulated. A household's prevalence is determined, under the model, by a set of household characteristics that (1) are policy-relevant (e.g., characteristics determining program eligibility and benefits), (2) capture the key dimensions along which households in different states are different, and (3) have about the same meanings across states. The problem that arises when this last condition is seriously violated for a particular characteristic is discussed by Schirm and Zaslavsky (1998).

The variables included in the reweighting model serve as control variables, and households are reweighted so that weighted sums (indirect estimates) equal specified control totals. These totals can be direct sample estimates, indirect estimates smoothed using empirical Bayes methods, or administrative totals. For example, if the number of people in the household is a control variable, the total state population is a control total. If household income is a control variable, total personal income in the state is a control total.

With the original (national) weights, the database looks like the entire United States. With Virginia weights, the database looks like Virginia in terms of some key aggregates (the control totals). We then conjecture that the reweighted database resembles Virginia in terms of many other relevant aggregates for which we cannot control, including, for example, the main estimands of a microsimulation model. Our evaluation will address the extent to which this is accomplished.

The Formal Model. The reweighting model is:

$$w_{hs} = \gamma_{hs}\, exp(\beta_s' x_h + \delta_h),$$

where w_{hs} is the expected number of households of type h in (the population of) state s. A type is, practically speaking, unique on the database because no two households are exactly alike. Therefore, each household in the database represents its own type, and w_{hs} is the weight that will be given to household h when deriving estimates for state s. γ_{hs} is an indicator set by the modeler to one if state s is allowed to borrow from the state in which household h actually resides, and zero otherwise. Although we will generally assume that each state borrows from every other state, Schirm and Zaslavsky (1998) describe an application in which the extent of borrowing is restricted. x_h is a column vector of I control variables, that is, household characteristics for household h. β_s is a vector of I unknown parameters to be estimated for each state. δ_h is an unknown parameter to be estimated for each household. The first term in the exponent on the right side of the model ($\beta_s' x_h$) reflects the "general" prevalence in state s of households like household h, that is, households with similar characteristics. The second term (δ_h) reflects the "specific" prevalence of household h.

The β_s and δ_h parameters are estimated by maximum likelihood and satisfy two constraints:

$$\text{Constraint 1: } \sum_s w_{hs} = W_h \text{ for each } h,$$

where W_h is the control weight, that is, the original sample weight or national weight of household h, and

$$\text{Constraint 2: } \sum_h w_{hs}\, x_{hi} = X_{si} \text{ for each } s \text{ and } i,$$

where X_{si} is the control total for control variable i in state s. According to the first constraint, reweighting does not change the total weight given to a household across all states, that is, at the national level, ensuring that the household contributes the same to a national estimate after reweighting as before. According to the second constraint, all control totals are satisfied for every state. Schirm and Zaslavsky (1997) describe an iterative two-step procedure for estimating the parameters of the reweighting model. A simple numerical example for which state weights are derived is given by Schirm and Zaslavsky (1998).

The general philosophy underlying this approach to modeling can be understood best by thinking of the entire database as a high-dimensional contingency table with dimensions that are defined by the many household characteristics included in the database, including state of residence. If we describe this table by a log-linear model, our modeling assumption is that some margins and low-order interactions of the household characteristics--the ones that appear in the vector x_h--are interacted with state, while the high-order interactions of the household characteristics are not interacted with state and, therefore, are the same in each state. In other words, the model assumes that the ways some household characteristics interact with each other are similar across states. Fitting a model that includes low-order interactions and excludes high-order interactions is a standard approach to smoothing a contingency table. The estimated table that is fit under the model is smoother, that is, less affected by sampling variability than the sample table that is obtained by tabulating every state separately with the original sample weights, which corresponds to the model in which every interaction among household characteristics is also interacted with state so that no interactions are excluded. By using weights fitted under the model with only low-order interactions, rather than the original sample weights, we reduce the variances for the entries in the high-dimensional table that is our database and, hence, for estimates calculated from the table.

3. An Application: Estimating the Numbers and Characteristics of Uninsured Children

Although there are prominent examples--such as those cited in the opening paragraph of this paper--where a single estimate may be all that is required for each state or substate area, there are other occasions when policymakers want to know how a particular subpopulation is distributed by two or more characteristics. What may be required in this case is a moderate to large table for each small area. For example, the following table shows the number of uninsured children in a state by several categories of age and poverty level:

Numbers of Uninsured Children

	Age In Years				
Poverty Level	< 1	1 to 5	6 to 13	14 to 18	Total
< 50% FPL					
50% to < 100%					
100% to < 150%					
150% to < 200%					
200% to < 350%					
350% or Greater					
Total					

We are currently deriving estimates of such tables to assist state and federal policymakers in the design, implementation, and evaluation of state CHIP initiatives. To develop the estimates, we are using both the empirical Bayes and reweighting methods for borrowing strength. Specifically, we develop empirical Bayes estimates of the total number of uninsured children in each state and the total number of children in each of several family income categories (measured relative to the poverty level). Taking these empirical Bayes estimates and other state-level administrative and sample estimates as control totals, we apply the reweighting method to derive state weights that can be used to produce the table shown above for each of the 51 states. With the reweighted data, we can produce other tables relevant to uninsured or low-income children. We could produce the same set of tables for all 51 states, or we could generate tabulations specific to individual states. For example, we might want to specify income levels that match the varying CHIP eligibility thresholds across the states, or we might want to develop in-depth analyses for particular states. With the reweighted CPS database, these alternative applications are straightforward to carry out.

4. An Application: Microsimulation of Food Stamp and Welfare Program Reforms

What Is Microsimulation? A microsimulation model simulates how proposed changes to a government program affect the program and its participants. The model has two elements: (1) a micro database and (2) a computer program. The database is constructed from administrative or survey data with information on households in the population targeted by the government program. The model's computer program codes the rules of the government program under both the "baseline" policy, which is typically the current policy, and a "reform" policy, which is an alternative under consideration. The computer program also simulates what a caseworker does--that is, it determines whether a household is eligible for the government program and the benefits for which the household would qualify. In addition, the computer program simulates a household's behavioral response, determining whether the household will participate in the program. Processing all the households in the database, the model counts participants to estimate the caseload of the government program and adds up their benefits to estimate costs. By performing these operations under both baseline and reform policies and comparing the results, the model estimates the cost and caseload effects of the proposed reform. The model can also estimate the distributional effects of the reform, identifying the population subgroups that gain and lose benefits.

We have focused our work on the Micro Analysis of Transfers to Households (MATH®) family of models. These models, developed by Mathematica Policy Research beginning in 1974, have been used extensively to simulate reforms to the Food Stamp Program (FSP), the Aid to Families with Dependent Children (AFDC) and Temporary Assistance to Needy Families (TANF) programs, and the Supplemental Security Income (SSI) program. The database for the current MATH SIPP model was constructed by combining data for January 1994 from Waves 7 and 4 of the 1992 and 1993 Panels of the Survey of Income and Program Participation (SIPP). The database for the QC Minimodel--another member of the MATH family--is constructed from the most recently available fiscal year Integrated Quality Control System (IQCS) sample, an administrative records database.

Microsimulation for States. Although the MATH models have been used in recent years to estimate the national effects of hundreds of potential reforms to national programs (mainly the FSP), welfare devolution and the resulting need for state analyses and estimates create a new challenge. Because a database developed from a national survey like the SIPP or even a national administrative database like the IQCS sample has state samples that are small for general purposes or, at least, for some important applications, direct microsimulation estimates are typically imprecise. Thus, there is substantial uncertainty about the likely impacts of proposed reforms and little guidance for policymakers.

Although it is not practically feasible to develop an empirical Bayes model for each of the many estimands of a microsimulation model, we can use the reweighting approach to improve the precision of microsimulation estimates. Schirm and Zaslavsky (1998) describe preliminary reweightings of the MATH SIPP and QC Minimodel databases. Alternative reweightings of the MATH SIPP database are examined in our evaluation of the reweighting method.

5. An Evaluation of the Reweighting Method

Schirm and Zaslavsky (1997) report very encouraging findings from a preliminary evaluation of the reweighting method for borrowing strength. The objective of our current research is to conduct a more thorough evaluation of the method. The evaluation addresses the following questions:

- Are model-based estimates from a reweighted database more accurate than direct estimates? That is, does the reduction in variance from reweighting more than offset the biases introduced? If the model-based estimates are more accurate, how much more accurate are they?

- What are the gains in accuracy, that is, what are the bias-variance tradeoffs from adding control variables to the reweighting model? In other words, when does a model with more control variables produce more accurate estimates than a model with fewer control variables?

- What are the gains in accuracy from using administrative or empirical Bayes estimates as control totals instead of direct sample estimates when reweighting a database?

Results from our evaluation of the reweighting method will be presented in the future.

References

National Research Council, Committee on National Statistics, Panel on Estimates of Poverty for Small Geographic Areas. *Small-Area Estimates of School-Age Children in Poverty, Interim Report 2, Evaluation of Revised 1993 County Estimates for Title I Allocations*, edited by Constance F. Citro, Michael L. Cohen, and Graham Kalton. Washington, DC: National Academy Press, 1998.

Schirm, Allen L. and Cindy Long. "Fund Allocation and Small Area Estimation in the WIC Program." *1995 Proceedings of the Section on Survey Research Methods*. Alexandria, VA: American Statistical Association, 1995.

Schirm, Allen L. and Alan M. Zaslavsky. "Model-Based Microsimulation Estimates for States When State Programs Vary." *1998 Proceedings of the Section on Survey Research Methods*. Alexandria, VA: American Statistical Association, 1998.

Schirm, Allen L. and Alan M. Zaslavsky. "Reweighting Households to Develop Microsimulation Estimates for States." *1997 Proceedings of the Section on Survey Research Methods*. Alexandria, VA: American Statistical Association, 1997.