

Improving Small Area Estimates of Disability: Combining the American Community Survey with the Survey of Income and Program Participation*

Jerry J. Maples[†] and Matthew Brault[‡]

Abstract

The Survey of Income and Program Participation (SIPP) is designed to make national level estimates of changes in income, eligibility for and participation in transfer programs, household and family composition, labor force behavior, and other associated events. Used cross-sectionally, the SIPP is the source for commonly accepted estimates of disability prevalence, having been cited in the findings clause of the Americans with Disability Act. Because of its sample size, SIPP is not designed to produce highly reliable estimates for individual states. The American Community Survey (ACS) is a large sample survey which is designed to support estimates of characteristics at the state and county level, however, the questions about disability in the ACS are not as comprehensive and detailed as in SIPP. We propose combining the information from the SIPP and ACS surveys to improve, i.e. lower variances of, state estimates of disability (as defined by SIPP).

1. Introduction

The Survey of Income and Program Participation (SIPP) is a nationally representative sample to estimate the sources and amounts of income, labor force information, program participation and eligibility data, and general demographic characteristics. The SIPP survey is a longitudinal survey which interviews the sampled respondents every 4 months, called a wave. The survey contains a set of core questions (asked at every wave) and sets of topical modules which are only asked during specific waves, e.g. one the topics of child care, wealth, disability, and school enrollment. The disability module has detailed questions (the same questions as in the standard Activities of Daily Living and Instrumented Activities of Daily Living Battery) about many different aspects of disability. Questions address physical and mental conditions affecting the persons in the household, including the use of mobility aids, vision and hearing impairments, speech difficulties, lifting and aerobic difficulties, and the ability to function independently within the home (Brault 2012). The disability estimates from the SIPP have been referenced in the findings clause of the Americans with Disability Act.

The American Community Survey (ACS) is a large scale general survey designed to make estimates at various geographic levels, e.g. states and counties. The ACS is a cross sectional survey drawing a new sample each year. Starting with the 2008 sample, the ACS added six general questions on disability, covering: hearing, vision, cognitive, ambulatory, self-care and independent living difficulties (Brault 2008). Obviously, many of the national level estimates of disability from the ACS will differ from the SIPP due to the different set of questions being asked. However, the two responses (from the same person) should be highly correlated. The 2010 sample for the ACS had approximately 1.9 million responding households and 3.7 million people 15 years or older, a much larger scale than SIPP. In contrast, 2008 SIPP had a sample of approximately 37,000 households and 70,000 people (at least 15 years old) in wave 6, the 2010 wave which included the disability module.

*Disclaimer: This presentation is released to inform interested parties of ongoing research and to encourage discussion of work in progress. The views expressed on statistical and methodological issues are those of the authors and not necessarily those of the U.S. Census Bureau.

[†]U.S. Census Bureau, CSRMS, 4600 Silver Hill Road, Washington, D.C. 20233

[‡]U.S. Census Bureau, SEHSD, 4600 Silver Hill Road, Washington, D.C. 20233

The methodological question is whether the two surveys can be coupled together to bring the power of the large ACS survey to improve estimates of the smaller, but more focused topical estimates from the SIPP. In this paper, we will explore state level estimates for total disability, as defined by the SIPP.

2. Regression Projection Estimator

The regression projection method will follow the work of Kim and Rao (2012). For this method, there are two surveys. Survey 1 contains a large sample, A_1 , but only collects information on a set of variables \mathbf{x} . Survey 2 is a much smaller (relative to survey 1) sample, A_2 , and contains data on \mathbf{x} but also on y which is the variable of interest. Note that the y is not observed from survey 1. It is assumed that the \mathbf{x} 's between the two surveys are comparable. Ideally, they should be the same survey question obtained by the same response mode, but the main requirement is that the questions prompt the same response from a hypothetical survey respondent. The goal is to create a synthetic dataset with proxy values, y_i 's, for the unobserved for survey 1. With these synthetic values, traditional design-based estimates of totals can be created.

From survey 2, a working model is created to model the relationship between y_i and \mathbf{x}_i , then used to predict the y_i 's with the \mathbf{x}_i 's from survey 1. Let w_{i1} be the sampling weight associated with unit i in sample 1. If y_i was observed in survey 1, then the unbiased estimate of the total would be $\hat{Y}_1 = \sum_{i \in A_1} w_{i1} y_i$. Instead, we use the synthetic value \tilde{y}_i and estimate

$$\hat{Y}_p = \sum_{i \in A_1} w_{i1} \tilde{y}_i. \quad (1)$$

Let $\hat{m}_i = m(\mathbf{x}_i; \hat{\beta})$ be the model predicted synthetic estimate for y_i given \mathbf{x}_i in survey 1, where $\hat{\beta}$ is obtained by fitting a model on the $\{(y_i, \mathbf{x}_i) : i \in A_2\}$ data from survey 2. If y_i is a binary outcome, then the synthetic prediction will provide two results: $\tilde{y}_i = 1$ with weight $w_{i1} \hat{m}_i$ and $\tilde{y}_i = 0$ with weight $w_{i1}(1 - \hat{m}_i)$ respectively. Asymptotic bias of the estimated total (at the national level) can be eliminated if $\hat{\beta}$ is estimated such that $\sum_{i \in A_2} w_{i2}(y_i - m(\mathbf{x}_i; \hat{\beta})) = 0$ is satisfied. Therefore, \hat{Y}_p based only on the synthetic values \tilde{y}_i is asymptotically design unbiased if this condition on $\hat{\beta}$ holds (Kim and Rao, 2012).

Suppose the working model is $E(y_i | \mathbf{x}_i) = m(\mathbf{x}_i; \beta) = m_i$ and that $Var(y_i | \mathbf{x}_i) = \sigma^2 a(m_i)$ for some known function $a(\cdot)$. Then $\hat{\beta}$ is obtained as a solution to

$$\sum_{i \in A_2} \frac{w_{i2}}{a(m_i)} \frac{\partial m_i}{\partial \beta} (y_i - m_i) = \sum_{i \in A_2} w_{i2} h_i (y_i - m_i) = 0 \quad (2)$$

where $h_i = (\partial m_i / \partial \beta) / a(m_i)$. This estimating equation can be thought of as finding the solution β that minimizes a weighted sum of squared residuals. To satisfy unbiasedness, one column of \mathbf{x}_i should be equal to unity. For a continuous variable, $m(\mathbf{x}_i; \beta) = \mathbf{x}_i' \beta$ and $a(m_i) = 1$, we have that $h_i = \mathbf{x}_i$. For a binary outcome, such as presence of any disability, a logistic regression $\text{logit}(m(\mathbf{x}_i; \beta)) = \mathbf{x}_i' \beta$ and $a(m_i) = m_i(1 - m_i)$, we have that $h_i = \mathbf{x}_i$. Note that for the binary outcome with the logistic regression link for m_i and $a(m_i) = m_i(1 - m_i)$, the same estimating equation is obtained from the weighted score equations of the log likelihood. Through asymptotic arguments (see Kim and Rao, 2012 for details) the projection estimator is shown to have approximate variance

$$Var(\hat{Y}_p) \approx Var \left\{ \sum_{i \in A_1} w_{i1} m(\mathbf{x}_i; \beta) \right\} + Var \left\{ \sum_{i \in A_2} w_{i2} (y_i - m(\mathbf{x}_i; \beta)) \right\} \quad (3)$$

The approximate variance of \hat{Y}_p is composed of two parts from (3). The first component is the survey variance as if the model prediction was an actual survey response on survey 1, the large sample survey. The usual design-based methods used to estimate sampling variance can be used to estimate it. The second component is the contribution to the variance from the prediction error from survey 2. Ignoring the second component typically underestimates the variance because it does not account for the uncertainty in estimating m_i . If the model is completely uninformative about the outcome of interest, then the prediction will just be the survey weighted mean. The first component will estimate no variance and the second component will be the full design-based variance from survey 2. The other extreme is when the model has perfect prediction. The first component then contributes all of the variance because the second component has all of the residuals as 0. In practice, both parts will contribute to the variance.

2.1 Domain Estimators

The projection estimator in (1) can be used to make estimates for areas and domains. Let $\delta_i(d) = 1$ if unit i belongs to domain d and 0 otherwise. The domain projection estimate $\hat{Y}_{d,p}$ is

$$\hat{Y}_{d,p} = \sum_{i \in A_1} w_{1i} \delta_i(d) \tilde{y}_i \quad (4)$$

In general, (4) is asymptotically design biased. The bias can be eliminated if the vector of domain indicators are included in the set of predictor variables \mathbf{x}_i . The approximate variance of the domain projection estimator is

$$Var(\hat{Y}_{d,p}) \approx Var \left\{ \sum_{i \in A_1} w_{i1} \delta_i(d) m(\mathbf{x}_i; \beta) \right\} + Var \left\{ \sum_{i \in A_2} w_{i2} \delta_i(d) (y_i - m(\mathbf{x}_i; \beta)) \right\} \quad (5)$$

The estimation of the second term in (5) may be problematic when there is little or no sample in domain d . The details of variance estimation will be discussed in Section 3.2.

3. State Level Disability Estimates

The goal is to improve state level estimates of disability, i.e. having any disability, using the SIPP definition of disability as defined by the disability module in wave 6. Many of the basic demographic variables are collected in both the SIPP and ACS such as age, race, sex, hispanic origin, employment status, education level, marital status, citizenship and veteran status. In wave 7, the SIPP module asked the 6 ACS disability questions. Ignoring the complications that arise due to wave 6 and wave 7 being collected 4 months apart with difference reference periods, this adds to the pool of x 's that can be used to model the probability of having a disability. Of these variables, a priori we expect age and the ACS disability questions to play a major role in building the predictive model.

One issue that must be addressed in the SIPP is that with some of the x 's being collected in a later wave, there will be some attrition of the sample between wave 6 and wave 7. In wave 6 of the 2008 panel, there was a sample of 70,306 people aged 15 years or older. In wave 7, only 62,865 of these people remained in the sample. Additionally, only 56,541 records had complete data on the ACS disability questions. Imputations made in wave 7 do not take into account data collected in wave 6 such as disability status. Any relationship between the SIPP disability module in wave 6 and the ACS style disability questions in wave 7 could potentially be destroyed in the imputed cases. Therefore, we will only base the model fit on the non-imputed cases from wave 7. To account for the cases

Table 1: Comparison of sample sizes, weights, CVs and effective sample sizes due to attrition of respondents between wave 6 and 7, additionally due to dropping the imputed cases from wave 7 and adjusting the weights to reflect the original wave 6 survey sample.

Data	disability rate	n	sum of weights	CV	effective n
Wave 6	21.02%	70,306	241,816,071	0.99%	37,989
Wave 7	21.21%	62,865	216,312,873	0.99%	37,337
Wave 7 (no imp)	21.38%	56,541	195,265,525	1.05%	33,524
W7 (no imp) adjusted	21.02%	56,541	241,816,071	0.99%	37,989

dropped for attrition and imputation, we can do a weight adjustment as if we were doing a non-response adjustment. We do not assume a non-ignorable non-response mechanism, therefore we post-stratify by SIPP disability status for adjustment cells. The main survey weight and each of the replicates weights are given this adjustment which preserves the overall disability rate with the wave 6 full SIPP data. Table 1 shows the difference in unweighted sample size, sum of weights, CVs on total disability estimate (national) and the effective sample size between the wave 6 and wave 7 datasets in SIPP. The effective sample size, n^* can be computed by solving for the n^* that reproduces the estimated CV under a simple random sampling design:

$$CV = 100 \sqrt{\frac{\text{Var}}{\text{Mean}^2}} = 100 \times \sqrt{\frac{p(1-p)N^2/n^*}{p^2N^2}} \quad (6)$$

$$n^* = \left(\frac{100}{CV}\right)^2 \frac{1-p}{p} \quad (7)$$


The estimated design effects are 1.85 for wave 6 and 1.68 (no imputations) for wave 7. Notice that wave 6 has a slightly lower disability rate. The cases in wave 6 but not in wave 7 have a disability rate of 19.43%. The weight adjustment for attribution and imputation make the wave 7 data, which will be used for model fitting, look similar to the full wave 6 data.

3.1 Model Fitting

Weighted logistic regression score equations are identical to (2) and allows us to use standard software to fit various predictive models. The covariates of interest are: age, race, sex, Hispanic origin and 6 ACS disability questions. These will be extended to include: veteran status, citizenship status, employment status and education attainment. To control the statistical power in testing the contribution of different covariates, the weights are ratio adjusted to sum up to the effective sample size, 37,989. This will not change any parameter estimates since the relative weights are still the same, but it will affect the standard errors which determine which covariates are kept in the model.

The first approach to building a model was to look at the marginal relationships between rate of having disability (defined by SIPP) and demographic variables age, race, Hispanic origin and sex. There are also some differences by sex and race/origin. Females had a higher disability rate than males (25% to 20%). Race and Hispanic origin were combined into a 5 level factor: Hispanic, Non Hispanic White (alone), Non Hispanic Black (alone), Non Hispanic Asian (alone), Non Hispanic Residual (anything not in previous categories). These will be referred to as Hispanic, White, Black, Asian and Other and they had disability


Figure 1: SIPP disability rates by Age with least squares 4th order polynomial regression line



rates of 23.8%, 28.2%, 14.9%, 30.2% and 16.4% respectively. Age clearly shows a pattern of increasing rates of disability as a person ages (see Figure 1). The relationship between the 6 ACS questions of disability and the presence of any disability (defined by SIPP) required some exploratory analysis. The rates of SIPP disability by different breakdowns of ACS disability questions are given in Table 2. Using regression partition trees, the best predictor was a three level variable (ACS-DIS-GROUP) for 0 ACS disabilities, 1 ACS disability and 2 or more ACS disabilities. The specific disability did not matter, only the number of disability questions answered yes. Note that the SIPP disability rate for people answering all 6 ACS questions yes was lower than for those answering 2 through 5 questions yes. This could be evidence of response error. In models, an indicator for answering all 6 ACS yes questions may be useful. The age by ACS-DIS-GROUP interactions showed different age trends by number of ACS disabilities (Figure 2). These suggest interaction structures to test in the model. A least squares fit line, using a 4th order polynomial, was added to the plot for each group. Two measures will be used to assess the predictive power of the model: the reduction in total squared error r^2 using the estimated \hat{m}_i , and a 'model-free' R^2 of Tjur (2009) which measures the difference in mean predicted probabilities for cases with disability and those without, $R_T^2 = \text{Mean}(\hat{m}_i|y = 1) - \text{Mean}(\hat{m}_i|y = 0)$. To avoid overfitting the data we will use the BIC. The log-likelihood used for the BIC is the sample weighted (with normalized weights to sum up to effective sample size) log-likelihood based on the Bernoulli distribution. This log-likelihood gives score equations that are the same as (2). We choose BIC over AIC because due to the larger penalty given for extra parameters in larger datasets where it is much easier to find spurious effects that are statistically significant. In our application, we will use the effective sample size (which is also the sum of the adjusted sample weights) 37,989 which give a penalty term to the log likelihood of 10.54 times the number of parameters in the model.

In order to not have to deal with a bias correction, the state effects are always kept in

Figure 2: SIPP disability rates by Age for 0, 1 or 2+ ACS disability questions answered yes with least squares 4th order polynomial regression lines



the model. As a set of 50 parameters (including D.C.), their contribution is statistically significant. However, most of the individual state parameters are not statistically different than the baseline (California - state with the largest sample size).

The following models will be evaluated for fit:

M0 : intercept only

M1 : State effects only

M2 : State + indicator for any ACS disability

M3 : State + 3 level ACS-DIS-GROUP variable

M4 : State + 4th order polynomial for Age (Age-P4)

M5 : State + 4th Age-p4 interacted with ACS-DIS-GROUP

M6 : State + 2nd order polynomial for Age (Age-P2) interacted with ACS-DIS-GROUP + indicator for all 6 ACS disability questions answered yes.

M6ns : M6 without the state effects

M7 : M6 + Age-P2 interacted with Sex + Age-P2 interacted with 5 level Race-origin

M8 : M7 + receiving SSI interacted with Age-P2, receiving SSI interacted with receiving Social Security, and receiving Social Security interacted with Sex

M8ns : M8 without the state effects

Table 2: SIPP disability rates by ACS disability survey questions

ACS Disability	SIPP Disability		Num of ACS Dis	Disability
	With	Without		
Hearing	66.0%	19.3%	0	11.4%
Sight	74.0%	20.0%	1	65.1%
Memory	82.4%	18.0%	2	86.0%
Physical Movement	84.9%	14.7%	3	89.8%
Dressing Self	88.6%	19.6%	4	94.0%
Going out of Home	89.6%	17.6%	5	95.2%
			6	64.5%
			1+	76.3%
			2+	87.9%

The above list of models is not represent an exhaustive search to find the best model. Model M8 is the current ‘best’ model and results from trying many different interactions. The second order polynomial in age was sufficient to explain the trend in age, once interacted with the disability count group variable. Model M6 is a reasonable model with less interactions compared to the more complex model M8, which pushed to find as many significant interactions as possible. The models M6ns and M8ns are the models M6 and M8 respectively without the state level effects. These models are purely synthetic and do not take into account any local area effects and will be used for comparison purposes.

Figure 3 shows the state level estimates of disability rates under the SIPP survey and the models M6, M8 and M8ns. The estimates from M6 and M8 are very similar to each other and both track the SIPP direct estimates fairly well. One noticeable difference is that when the SIPP estimates an extremely high disability rate (more than 30%), the model-assisted predictions are less. The synthetic models M6ns and M8ns are also very similar to each other in their state level predictions. However, model M8ns shows much less variation between states in the estimated disability rate. Even though we are using small area techniques to improve the state level disability estimates, the estimate from SIPP are design-consistent and if the model is approximately correct, then the model-based estimates should fall into the design-based confidence interval. The estimates from M6 and M8 fall into the 95% confidence interval of the design-based estimate 49 out of 51 times (96%) with one model-based estimate being higher and one being lower. The synthetic model M8ns fell into the 95% confidence interval only 35 out of 51 times (68.6%) with 5 model-based estimates larger (on the state with a lower estimated disability rate from SIPP) and 11 estimates lower (all on state with a higher disability estimate from SIPP). This is evidence that the synthetic assumption, i.e. no the state effects, is wrong and that there are non-trivial state level effects that should be accounted for in the model.

3.2 Variance estimation


Two different methods are given by Kim and Rao (2012) to estimate the variances from (5). Recall the variance approximation:

$$\begin{aligned}
 Var(\hat{Y}_{d,p}) &\approx Var \left\{ \sum_{i \in A_1} w_{i1} \delta_i(d) m(\mathbf{x}_i; \beta) \right\} + Var \left\{ \sum_{i \in A_2} w_{i2} \delta_i(d) (y_i - m(\mathbf{x}_i; \beta)) \right\} \\
 &= V_{1,d} + V_{2,d}.
 \end{aligned}$$

Table 3: Model fits

Model	Num Params	r^2 (MSE)	R_T^2	BIC
M0	1	0	0	39,063.0
M1	51	.009	.009	39,297.6
M2	52	.327	.330	29,402.9
M3	53	.339	.343	28,974.4
M4	55	.150	.152	34,542.8
M5	65	.375	.379	27,365.5
M6	60	.376	.380	27,273.0
M6ns	10	.372	.377	26,897.7
M7	75	.381	.385	27,219.5
M8	80	.402	.408	26,408.8
M8ns	30	.398	.405	26,034.4

Figure 3: Comparison of estimates of state level disability rates



The two parts could be estimated using standard design-based methods from their respective surveys. Both the SIPP and ACS use a replicate weight variance method. The SIPP uses 160 replicates, while the ACS uses 80 replicates. The first component computes the design-based variance using the model predictions on the ACS as if they were real survey data. The second component computes the variance of the sum of residuals from the SIPP. Since the ACS has a much larger sample size, we expect $V_{1,d}$ to be smaller than $V_{2,d}$. However, if we have a good model for $E(y_i|x_i)$, then that should help reduce the size of $V_{2,d}$. Another consideration is that the quality of the estimate for $V_{2,d}$ is dependent on the sample size of the domain in the SIPP survey. This method is problematic when there is little or no sample data in the second survey for a domain of interest. For example, only about half of the counties have data in the SIPP disability dataset. Thus, this variance method would not be useful for counties with no SIPP data, but there needs to be a way to incorporate the model prediction uncertainty or we would risk severe underestimation of the variance.

A second method is proposed in Kim and Rao (2012) to estimate the variance. First, they assume that both surveys use replicate weights (either by design or the weights can be constructed). They also assume that both surveys use the same number of replicates, for this application $R = 80$. For the SIPP survey, using the first half, 80, of the replicates produces valid variance estimates. This method computes the $\hat{\beta}$'s under the original survey weights and each set of R replicate weights for survey 2 (SIPP). The $R+1$ (81) predictions, including, \hat{Y}_p from (1), are appended to survey 1 (ACS). Then the replicate weight formula for survey 1 is modified as follows:

$$Var(\hat{Y}_{d,p}) = c_R \sum_{k=1}^R (\hat{Y}_p^k - \hat{Y}_p)^2 \quad (8)$$

where \hat{Y}_p^k is the prediction using the k^{th} set of $\hat{\beta}$'s and c_R is the factor associated with the replicate weight method for survey 1. This formulation includes both the variability of the design of survey 1 and the uncertainty of model prediction from survey 2. Additionally, it does not exclusively use data only from domain d for survey 2 to incorporate uncertainty about the model. The correctness of this variance method may depend implicitly on the model being correct and must be studied further. Another benefit from this second method to estimate variances is that it does not require access to the data from survey 2 once the $R+1$ sets of $\hat{\beta}$'s have been computed. This may be useful when the data from survey 1 can be made available, but the data from survey 2 can not.

The two variance methods give similar results for the state level estimates of disability rates when the working model is reasonable. Figures 4 and 5 show the reduction of the variance for regression projection estimates compared to the design-based variances of the direct survey estimates using the two variance methods. The two variance estimation methods have a 98% correlation for M6 and M8. There is a big difference between the two variance methods for model M6ns. Method 2 estimates a much lower variance for the synthetic model. If the model was correct (or even approximately correct), this would be a great application of the regression projection method. However, when combined with the results in Figure 3 it is clear that the synthetic model is not even approximately correct and the second method of estimating the variance is adversely affected by model misspecification.

For the national estimate of total disability, the variance reduction using the regression projection method was about 20%. For the state level estimates, the results were mixed with seven states showing higher variability after modeling compared to the direct survey estimates. These seven states showed a range of sample sizes, some on the smaller end and some in the middle. However, 12 states saw a variance reduction of 40% or more.

Figure 4: Variance estimates - method 1





Figure 5: Variance estimates - method 2



4. Discussion

The regression projection method is a method to take the relationship between a set of variables and an outcome of interest from one survey and impute through model prediction onto another survey that contains all of the same variables except the outcome of interest. This model works best when there are common survey variables that are highly related to the outcome of interest. This is a model-assisted approach and is robust to model misspecification provided indicators for all domains of interest are included in the model. Two different methods to compute the variance of the regression projection estimator are given. The application of this method to state level disability estimates with the SIPP and ACS surveys showed some of the limitations with the methodology. By including the states (domains) as fixed effects, the variability of the those domain effects are governed by the potential small sample sizes in survey 2 for that domain. It is also a problem when some domains may not have any data, which would occur in SIPP for county level estimates.

This methodology does have some promising features in that once the regression projection estimates are appended to survey 1, any aggregate can be computed (unplanned domains). Method 2 for the variance estimation does not even require going back to the data from survey 2 in order to calculate the variances for unplanned domains. However, method 2 does not seem to be robust in the estimate of variance when the model is misspecified. Additional research for improving state level estimates of disability is planned. First, a switch to random effects for domains instead of fixed effects using as much of the framework from Kim and Rao (2012) as possible will be investigated. For comparison, a bivariate Fay-Herriot model will be developed to model the state and county level disability rates. Another extension will be to broaden the outcome variable of interest. Currently, the focus has been on whether a person has any disability. We could also consider rates of a very specific disabilities. These will occur at much smaller rates, which may present new problems.

REFERENCES

- Brault, M. (2009), "Review of Changes to the Measurement of Disability in the 2008 American Community Survey", U.S. Census Bureau, Washington, DC, http://www.census.gov/people/disability/files/2008ACS_disability.pdf.
- Brault, M. (2012) "Americans with Disabilities: 2010", Current Population Reports P70-131, U.S. Census Bureau, Washington, DC.
- Kim, J.K. and Rao, J.N.K. (2012). "Combing data from two independent surveys: a model-assisted approach", *Biometrika*, **99**, 85-100.
- Tjur, T. (2009) "Coefficients of determination in logistic regression models, a new proposal: The coefficient of discrimination." *The American Statistician*, **63**, 366-372.