STATISTICAL ISSUES IN A RECORD CHECK STUDY OF CHILDHOOD IMMUNIZATION

J. Michael Brick, Graham Kalton, Mary Nixon, Westat;
Jimmie Givens, Trena Ezzati-Rice, NCHS

Abstract: The National Immunization Provider Record Check Study is used to assess the quality of household reporting of vaccination levels for children aged 19 to 35 months. Provider nonresponse is handled by a statistical adjustment and estimates of reporting quality are produced. The findings show that both response bias and variance should be considered in this evaluation.

1. Introduction

The National Immunization Provider Record Check Study (NIPRCS) is conducted in conjunction with the Immunization Supplement of the National Health Interview Survey (NHIS) to assess and improve estimates of vaccination levels of children between 19 and 35 months of age. In the Immunization Supplement, household respondents report on the numbers of vaccinations the children in the sampled households have received. The providers of eligible children are contacted in the NIPRCS and asked to report on the children's vaccinations. The provider and household data are compared and "best values" for the numbers of vaccinations are assigned for as many children as possible. Children for whom provider data cannot be obtained have missing best values. Statistical procedures are used to adjust for missing data. The analysis includes an examination of the response bias and variance in the Immunization Supplement household reports and the development of improved estimates of number of vaccinations.

This article outlines each step in the NIPRCS. The next section presents the overall study design. Section 3 summarizes the outcomes of each of the NIPRCS data collection components for the period 1994-1996. The assignment of best values is discussed in Section 4. The procedures used to produce the estimates and compute their sampling errors, including the methods used to adjust for the missing best values, are covered in Section 5. Sections 6 and 7 describe the analytic methods used to estimate the response bias and variance and apply these methods to the NIPRCS data from 1994 to 1996. The last section contains some concluding remarks.

2. Study Design

The NHIS is a large-scale, on-going household survey of the civilian noninstitutionalized population of the U.S. conducted by the National Center for Health Statistics (NCHS). The main objective of the NHIS is to monitor the health of the U.S. population. The data collection instrument consists of a core questionnaire and several specialized supplemental questionnaires, one of which is the Immunization Supplement which gathers vaccination data about sample persons under the age of six. Massey *et al*. (1989) describe the NHIS sample design.

The NIPRCS focuses on the quality of reporting for each of five recommended vaccines for children aged 19 to 35 months. The five vaccines are: diphtheria, tetanus toxoids, and pertussis vaccine (DTP); poliovirus vaccine (polio); measles, mumps, and rubella vaccine (MMR); *Haemophilus influenzae* type b vaccine (Hib); and hepatitis B vaccine (Hep B). Particular attention is paid to whether the child has received the recommended number of doses for each of these. The recommended number of doses for DTP is 4 or more, for polio 3 or more doses, for MMR 1 or more doses, for Hib 3 or more doses, and for Hep B 3 or more doses. For the purposes of this study, a child who received the recommended number of doses of a particular vaccine is considered up-to-date with respect to that vaccine. In addition to the reporting of individual vaccinations, three combinations of vaccines are considered: at least 4 doses of DTP, 3 of Polio and 1 of MMR (431); at least 4 doses of DTP, 3 of Polio, 1 of MMR, and 3 of Hib (4313); and at least 4 doses of DTP, 3 of Polio, 1 of MMR, 3 of Hib, and 3 of Hep B (43133).

An important issue in estimating measurement error from a record check study is the definition of error in the survey reports. Defining error as the difference between the survey and the record value assumes that the record value is the "true" value. As with other administrative systems, provider records are often error-prone, with missing records, records that are in the system erroneously, and errors in the values recorded. In particular, a provider's record for a child may be incomplete if some of the child's vaccinations were given by other providers. The definition of "best values" rather than true values was developed with this issue in mind.

A second issue is that the comparison of the values from the survey and the provider can only be made if the records from the two sources can be matched or linked together. Errors in matching can result in distorted estimates of measurement error. Care was, therefore, taken to avoid mismatches, including matching on date of birth as well as name and comparing the dates of vaccinations, where available, for consistency.

3. Data Collection

The NIPRCS data collection process starts with the NHIS interviews. In addition to providing data about a child's vaccinations, the household respondent was asked to give the names and addresses for up to three of the child's vaccination providers and to sign a permission form permitting the Centers for Disease Control (CDC) to request vaccination information from these providers. Nonrespondents to the NHIS Immunization Supplement were not included in the NIPRCS (except in 1994). Weighting adjustments were made to account for this nonresponse.

The three data collection efforts undertaken in the NIPRCS in 1994 to 1996 were:

- *Original Provider Survey.* This survey contacted health care providers and requested vaccination information for sampled children. The survey was restricted to providers of children for whom the permission form was signed and to providers where addresses, as given by the household respondents, were adequate for mailing the questionnaire. About half the children had providers surveyed in the Original Provider Survey, leaving the other half to be included in the Followup Study.

- *Followup Study.* This study comprised two components, the Household Followup Survey and the Provider Followup Survey. The households of children from the NHIS with no locatable providers, no named providers, or without signed permission forms were included in the Household Followup Survey. Households of children with no reported vaccinations were included in this survey and asked to identify general medical providers for the children. All providers identified during these household contacts who were locatable and for whom written or verbal contact permission could be obtained from the household were included in the Provider Followup Survey.

- *Reconciliation Contacts.* Telephone calls were made to providers responding to the Original Provider Survey and/or to the corresponding NHIS household respondents to resolve discrepancies in the reported numbers of doses between the provider and household respondents, and sometimes to check that the data provided related to the correct child. Names and addresses of additional providers were also requested from households. Newly identified providers were then contacted for vaccination information.

The Original Provider Survey and the Provider Followup Survey were conducted initially by mail with telephone followup calls being made to providers who did not respond to the mailings. The Household Followup Survey and virtually all of the Reconciliation contacts were conducted by telephone. The Followup Study and the Reconciliation contacts could not be conducted until 18 to 30 months after the original NHIS interview when the NHIS data were available.

Table 1 summarizes overall response rates among households responding to the NHIS for the three survey years. The table does not show the response rate to the NHIS core interview which is greater than 90 percent of eligible households. Row C of Table 1 shows that Immunization Supplements were completed for about 92 to 94 percent of the eligible 19- to 35-month-old children in responding households in the NHIS. Children are classified as having completed Immunization Supplements if the supplement was administered, even if the household respondent did not know whether vaccinations had been given. Between 47 and 53 percent of the children with Immunization Supplements had signed permission forms that were used in the Original Provider Survey. Over the three years, useable permission forms were obtained from the Followup Study for between 41 percent and 67 percent of the children surveyed.

The last rows of Table 1 show the number and percent of children with useable provider data after completion of all provider data collection activities. Between 65 and 78 percent of the children with completed NHIS Immunization Supplements had provider vaccination information after these efforts. For most of the children with provider data, the data came from the Original Provider Survey (from 55% to 72% of children with provider data over the years). The Followup Study was also productive, accounting for between 27 and 43 percent of all the children with provider data. The Reconciliation Contacts were less productive in these terms, accounting for only 1 to 2 percent of the children with provider data.

4. Best Value Determination

Since provider records are not error-free, the vaccinations reported by a child's provider (or the combined report from several providers) were not accepted as the true values. Similarly, even household reports from shot records cannot be assumed to be correct and completely current. Since neither source on its own gives the true values for a child, best vaccination values were developed for the NIPRCS to provide the best estimates of true values.

The best values for the numbers of doses of the various vaccines that a child had received by the time of the NHIS interview were determined by examining both household and vaccination provider reports. Best values could not be determined for children for whom no provider reports were received, with one exception. If a child was reported to be 4313 up-to-date from records and attempts to contact providers were unsuccessful, the household report was taken as the best value (there were 22 such cases in 1996). The process of assigning best values was complex and is described more fully by Ezzati-Rice *et al*. (1996). Below, we outline a few of the key steps in the process.

Table 1. Response rates for the 1994 - 1996 NIPRCS

	Description	Response rate[1]	1994		1995		1996	
			No.	%	No.	%	No.	%
A	Eligible children 19-35 months in the NHIS	---	2,651		2,255		1,386	
B	Children with completed Immunization Supplement	B/A	2,439	92	2,089	93	1,296	94
C	Children without completed Immunization Supplement	C/A	212	8	166	7	90	6
D	Total children with permission to contact provider	D/B	1,365	56	1,175	56	778	60
E	Total children with useable permission forms[2]	E/B	1,303	53	1,069	51	609	47
F	Eligible for Followup Study		1,342 [3]	---	1,020	---	687	---
G	Children with useable permission forms from Followup Study	G/F	733	55	415	41	461	67
H	Total number of children with useable provider data:	H/B	1,798	74	1,352	65	1,005	78
I	From Original Provider Survey with Immunization Supplement	I/H	1,192	66	969	72	553	55
J	From Reconciliation Contacts	J/H	28	2	18	1	20	2
K	From Followup Study with Immunization Supplement	K/H	578	32	365	27	432	43

[1]The letters in parentheses indicate how response rates were calculated for each category.
[2]Useable permissions were those remaining after removing cases with missing provider name and/or incomplete address information and/or out-of-country address and/or permission from respondent other than parent/guardian.
[3]Includes 206 children for whom the Immunization Supplement was not completed for 1994.

The initial process of matching the provider-reported data against the household-reported values began by classifying the degree of difference in the reports. For most of the cases with identical or nearly identical reports, best values corresponded to both the household and provider-reported values. If the provider reported more doses than the household and the provider reported the child was up-to-date on the 4313 combination, the best values were set to the values reported by the provider. These general rules were not followed if the data for the individual child appeared suspect. For example, no best values were assigned if the pattern of doses was not plausible. Similarly, if the household reported more doses or there was a complex pattern of differences in the number of doses, then the reconciliation was attempted.

The percentage of children with best values varied over the three NIPRCS survey years. In 1994, best values were assigned for 74 percent of children with an Immunization Supplement. In 1995 and 1996, the percentages were 66 percent and 79 percent respectively.

5. Estimation

Since the NIPRCS is a followup of the NHIS Immunization Supplement, the estimation strategy starts with the Immunization Supplement estimation weights and then compensates for missing best values. The NHIS weighting process has several steps that result in national person weights for each sampled person. The first step is the inverse of the probability of selection. The second step is an adjustment for household nonresponse. The final steps in developing these weights are ratio adjustments to census totals in 60 classes defined by age, gender, and race. For the Immunization Supplement, a further adjustment is made to compensate for Immunization Supplement nonresponse among households responding to the NHIS. The resultant Immunization Supplement nonresponse adjusted weight was the initial weight used in the NIPRCS.

The inability to assign best values for all the responding children in the NHIS Immunization Supplement sample presents the problem of how these should be handled in the analysis. Two general-purpose methods of handling such missing data in survey analysis are weighting adjustments and imputation. In the NIPRCS both options are feasible because the best values are either all missing or all present (i.e., if the best value is assigned for any vaccination it is assigned for all). Both the weighting and imputation methods were applied to all three years of the NIPRCS, but weighting was the preferred approach and is the one described here.

Weighting was applied for missing best values by restricting the analysis file to the children with best values and assigning nonresponse adjustment weights to these children to compensate for the children without best values. These adjustments were computed separately within classes defined by a set of auxiliary variables. For each of the three NIPRCS survey years, the auxiliary variables were chosen based on exploratory analyses conducted to find good predictor variables for the propensity of the children to have best values (response propensities) and for a characteristic of prime interest; chosen to be up-to-date for the 4313 combination. These two dependent variables were chosen because using powerful predictors of either of them for forming weighting classes reduces the large sample bias of the estimate.

Exploratory analyses were performed using the NIPRCS data for each year. A categorical search algorithm was employed to predict the two dichotomous target variables (having a best value and being up-to-date on the 4313 combination). The weighting classes were then based on auxiliary variables that were important predictors in these analyses. The auxiliary variables chosen were not the same for all three years, but always included the household up-to-date status for the 4313 combination (up-to-date, not up-to-date, and unknown), and whether or not records were consulted (records and recall) for the child's NHIS immunization report. Other predictors included poverty, race, age of child, and parent's education.

The nonresponse adjustment for a weighting class was computed as the ratio of the sum of the weights for all children to the sum of the weights for the children with best values. The adjusted weight is then the nonresponse adjustment multiplied by the Immunization Supplement weight. Estimates derived using this weight are called *weighted best value estimates*.

Variance Estimation

Since the NIPRCS is a probability sample, its estimates are subject to sampling variability. The jackknife replication method of variance estimation was used for estimating the precision of NIPRCS estimates. A replicate is created by deleting a PSU (or variance unit) from a variance stratum and reweighting the other PSUs in the stratum to compensate for the deleted unit. Each variance stratum in the NIPRCS has exactly two PSUs. By separately computing the weighting adjustments for each replicate, the effects of weight adjustments can be reflected in the estimates of variance. Since it was not possible to compute replicate base weights for the NHIS from the information available, the replication process had to start from the weights that included the adjustments for nonresponse to the Immunization Supplement. Thus, the variance estimates obtained do not fully reflect the effects of the weight adjustments.

Replicates for all three NIPRCS survey years were created based on NHIS public use file design information. The NHIS documentation discusses two approximations to the NHIS design, termed Method 1 and Method 2. For the 1994 and 1995 NIPRCS, Method 1 was used to form replicates for variance estimation purposes. This approach resulted in 62 strata for the 1994 survey, and 187 strata for the 1995 survey.

For the 1996 NIPRCS, the methods suggested in the NHIS documentation would have resulted in a large number of replicates and increased computing time for variance estimation. An alternative method similar to an approach examined for the 1995 NHIS was used (Nixon *et al*., 1998). The alternative combines the NHIS PSUs into 150 pseudo-PSUs, which are assumed to be sampled with replacement from 75 variance strata, each with exactly 2 sampled pseudo-PSUs.

Weighting adjustments and imputation compensate for missing data but must be treated carefully in variance estimation to obtain approximately unbiased estimates of standard errors. Rao and Shao (1992) and Shao *et al.* (1998) developed replication variance estimation techniques for computing standard errors for sample means and estimates of population totals when the missing data have been imputed using hot deck methods.

An evaluation of several methods of variance estimation for the NIPRCS was performed with preliminary Quarter 1 and Quarter 2 1994 NIPRCS data (Nixon *et al*., 1996) using the methods referenced above for imputation. The evaluation compared the weighting and imputation approaches for compensating for missing best values and examined different techniques for replicated variance estimation for both approaches.

When the weighting classes are the same as the imputation classes, the estimates produced by the two procedures will be equal in expectation and the variances of the weighted estimates will be smaller than those of the imputed estimates. The empirical findings comparing the weighted and imputed best value estimates are in line with these theoretical results. There is very little difference in the point estimates. However, the standard errors of the weighted estimates are smaller than those of the imputed estimates. Moreover, the estimation of standard errors for the weighted estimates is simpler and can be performed using existing variance estimation software. The standard error computations for imputed data requires special programming. In view of these considerations, weighted best value estimates were adopted for the main NIPRCS analyses.

6. Analysis Methods

The analytic methods for evaluating the quality of the household reports of vaccination status are now outlined. The key indices used in analyzing the data available from the NIPRCS are the net and gross difference rates, as typically defined in reinterview studies. The gross difference (*gdr*) rate is the percentage of erroneous household reports (treating the best value as the truth). Erroneous household reports are of two types, understating and overstating vaccinations. The net difference rate (*ndr*) is the nonoffsetting part of the gross difference rate.

Consistent estimators for the *gdr* and *ndr* when weighting adjustments are used to compensate for missing best values are:

$$gdr = \frac{\sum w_i d_i^2}{\sum w_i} \tag{1}$$

$$ndr = \frac{\sum w_i d_i}{\sum w_i} \tag{2}$$

where w_i are the household weights modified to account for Immunization Supplement nonresponse and missing best values, and $d_i = 0$ if the household and best values agree, $d_i = 1$ if the household report is up-to-date and the best value is not, and $d_i = -1$ if the household report is not up-to-date and the best value is up-to-date for child *i*. In computing the *gdr* and *ndr*, the children with don't know (DK) household responses for a particular vaccination are excluded.

The above formula can also be applied when imputation is used to compensate for missing best values, but in this case the weights w_i are only adjusted to compensate for Immunization Supplement nonresponse. It is well known that imputation attenuates the associations between variables (Brick and Kalton, 1996). Thus, cases with imputed values should be excluded for the computation of the *gdr* and *ndr*. The household records with DK responses in the Immunization Supplement are also excluded.

7. Findings

The household and weighted best value estimates of the percentage of children who are up-to-date for the individual vaccinations and the three combinations of vaccinations are given in Table 2 for 1994, 1995, and 1996. The results show that while the household and best value estimates have the same general trend over time, the household estimates are markedly different from the best value estimates. The magnitude of the differences demonstrates the need for provider data to improve the estimates of vaccination levels.

Previous studies such as those by Goldstein *et al*. (1993) have suggested that parents tend to underestimate the number of doses received for multiple-dose vaccines and to overestimate coverage for single-dose vaccines. The net difference rates for the vaccinations given in Table 3 shed some light on this measurement error phenomenon for the NHIS Immunization Supplement. The net difference rates for most of the estimates are relatively large and negative, indicating significant household underreporting. For MMR, the only single-dose vaccination studied, the net difference rates are not statistically significantly different from zero. (All the significant tests were t-tests of the difference between the household and weighted best values estimates for a given year, incorporating Bonferroni adjustments for multiple comparisons.)

Table 2. Percentage of children up-to-date for individual vaccinations and combinations, by household and weighted best values for 1994, 1995 and 1996

	Household			Weighted best values		
	1994	1995	1996	1994	1995	1996
Individual vaccination						
DTP	70.1 (1.1)	69.8 (1.3)	76.5 (1.5)	76.8 (1.3)	80.5 (1.1)	83.4 (1.3)
Polio	79.2 (0.9)	81.6 (0.9)	88.8 (1.2)	85.6 (1.1)	87.8 (0.8)	93.4 (1.0)
MMR	90.4 (0.7)	88.5 (0.9)	91.4 (1.0)	89.0 (0.8)	91.9 (0.8)	92.4 (0.9)
Hib	75.0 (1.1)	74.8 (1.3)	77.3 (1.7)	91.1 (0.7)	91.9 (0.7)	94.0 (0.8)
Hep B	34.4 (1.4)	60.8 (1.6)	74.0 (1.6)	29.4 (1.1)	66.0 (1.6)	83.3 (1.4)
Combinations						
431	67.5 (1.1)	67.3 (1.3)	73.3 (1.7)	75.2 (1.2)	78.2 (1.2)	81.8 (1.3)
4313	59.8 (1.2)	59.4 (1.4)	65.3 (2.0)	73.5 (1.1)	76.4 (1.2)	80.4 (1.3)
43133	27.4 (1.3)	44.2 (1.5)	56.8 (2.0)	24.2 (1.0)	55.7 (1.6)	71.9 (1.6)

Note: Percentages are based on weighted values. DK responses in the Immunization Supplement are dropped from the household estimates for that vaccine or a combination that includes that vaccine. Standard errors are in parentheses.

Table 3 also shows the gross difference rates for the vaccinations for the three years. The size of the gross difference rates indicates that the household respondents do not report the vaccination levels very accurately. The level of misclassification is consistent over time. For example, about 25 percent of the children are misclassified for the 431 combination each year.

Table 3. Net and gross difference rates for individual vaccinations and combinations for 1994, 1995 and 1996

	Net difference rate			Gross difference rate		
	1994	1995	1996	1994	1995	1996
Individual vaccination						
DTP	-8.0 (1.3)	-11.6 (1.6)	-8.8 (1.8)	21.9 (1.2)	24.8 (1.5)	22.5 (1.3)
Polio	-6.9 (1.2)	-5.9 (1.4)	-5.0 (1.4)	19.1 (1.2)	17.2 (1.4)	11.9 (1.4)
MMR	0.5 (0.9)	-2.8 (1.2)	-1.9 (1.4)	10.6 (0.8)	11.9 (1.1)	11.9 (1.2)
Hib	-16.2 (1.2)	-17.8 (1.4)	-16.7 (1.8)	20.3 (1.2)	20.5 (1.3)	21.1 (1.7)
Hep B	4.1 (1.3)	-7.0 (1.7)	-10.4 (2.0)	19.6 (1.2)	21.0 (1.4)	21.7 (1.4)
Combinations						
431	-9.4 (1.3)	-12.5 (1.5)	-10.7 (2.0)	23.7 (1.3)	26.9 (1.4)	25.1 (1.6)
4313	-15.7 (1.5)	-20.1 (1.6)	-17.5 (2.1)	28.0 (1.4)	30.9 (1.4)	31.2 (1.9)
43133	1.4 (1.4)	-14.6 (2.1)	-18.3 (2.3)	21.5 (1.3)	31.3 (1.5)	34.0 (1.9)

Note: Percentages are based on weighted values. Standard errors are in parentheses.

The net and gross difference rates were also examined for a variety of different subgroups covering a range of social, economic, and demographic characteristics of the children and their households. For almost all of the subgroups, the rates did not vary significantly from the corresponding rates for the full population. One implication of this finding is that the household response errors are not highly related to characteristics that could be exploited to improve the estimates from household reports, without the use of record check data.

One subgroup analysis that did show important differences was according to whether the household respondent used a shot card (record) in reporting the child's vaccinations or not (recall). As an example, consider DTP. Table 3 shows that the net difference rates for all children for this vaccination are negative for all three years. Table 4 which presents net and gross difference rates for DTP by record or recall for 1994-1996 shows that the net difference rates are appreciably larger for those reported from records than those reported from recall. When records were used, the households more consistently underestimated the up-to-date status of children. This is because the predominate error in a household record is a failure to include a vaccination, rather than including one that did not occur. Thus, the errors do not cancel and underreporting biases result.

Table 4. Net and gross difference rates for DTP, by household use of records for 1994-1996

	Net difference rate		Gross difference rate	
	Recall	Records	Recall	Records
1994	-3.4	-11.6	30.9	14.9
1995	-8.7	-14.1	34.0	17.0
1996	-4.6	-13.0	28.0	17.0

Another important finding in Table 4 is that the gross difference rates are substantially smaller for the children for whom records were used. Since the estimates have significant response bias (net differences not equal to zero), the gross difference rates estimate mean square errors rather than simple response variances. Thus, the household reports based on records have lower mean square errors than those based on recall.

These results point out the difficulty in attempting to measure the quality of the vaccination data by a single measure. The estimates where the household used records are subject to lower mean square error than those where the household used recall. On the other hand, the former estimates are more biased than the latter. Depending on whether bias or mean square error is used as the criterion of quality, the assessment of which subgroup reported “better” is switched.

8. Concluding Comments

The NIPRCS provides a mechanism to better understand errors in the reporting of vaccinations of children from the NHIS Immunization Supplement. In addition it is a source for better estimates of vaccination levels for children. The provider record check involved an extensive effort to obtain data for as many of the children in the NHIS as possible. Furthermore, a procedure for assigning best values was developed that considered both the household and the provider reports of vaccinations rather than simply treating the provider reports as true values. This procedure reduced the dependence on the assumption that provider reports were always correct and improved the assessment of household reporting errors.

The efforts made to obtain a high provider response rate were essential because children without provider data must be handled by statistical adjustment methods. Although these adjustments aim to reduce the nonresponse bias in the estimates, they are never fully able to do so. For the NIPRCS, the weighting adjustment method was determined to be simpler and more efficient than imputation. A replicate variance estimation method that accounted for both the complex sample design and the missing data adjustments was implemented. The weighting approach had a distinct advantage over the imputation method in terms of ease of variance estimation.

The findings from the NIPRCS for 1994 to 1996 reveal that the households systematically underreported the number of vaccinations for children between 19 and 35 months of age. They also show that households do not report accurately, with up to 35 percent of the children being misclassified with respect to their up-to-date vaccination status. The types of errors made in household reporting depend on whether the household reported using shot records or recall. An interesting finding relates to the widely-used data collection technique of asking respondents to use records in responding to survey questions. In the NIPRCS, the vaccination coverage estimates for children for whom shot records were used in responding to the Immunization Supplement had more response bias than those for children for whom shot records were not used. This result came about because of the presence of systematic errors in the household records. On the other hand, respondents using records reported vaccinations more consistently than those who relied on recall. These findings point out the importance of examining both response bias and variance in this type of study.

9. References

Brick, J.M., and Kalton, G. (1996). Handling missing data in survey research. *Statistical Methods in Medical Research*, **5**, 215-230.

Ezzati-Rice, T.M., Zell, E.R., Massey, J.T., and Nixon, M. (1996). Improving the assessment of vaccination coverage rates with the use of both household and medical provider data. *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 335-339.

Goldstein, K.P., Kviz, F.J., Daum, R.S. (1993). Accuracy of immunization histories provided by adults accompanying preschool children to a pediatric emergency department. *Journal of the American Medical Association*, **270**, 2190-2194.

Massey, J.T., Moore, T.F., Parsons, V.L., and Tadros, W. (1989). *Design and estimation for the National Health Interview Survey, 1985-94.* Vital and Health Statistics, Series 2, No. 110, National Center for Health Statistics, Hyattsville, MD.

Nixon, M., Brick, J.M., Kalton, G., and Lee, H. (1998). Alternative variance estimation methods for the NHIS. *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 326-331.

Nixon, M., Kalton, G., Brick, J.M., Ezzati-Rice, T., and Zell, E. (1996). Variance estimation with missing best values in the NIPRCS. *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 347-352.

Rao, J.N.K., and Shao, J. (1992). Jackknife variance estimation with survey data under hot deck imputation. *Biometrika*, **79**, 811-822.

Shao, J., Chen, Y., and Chen, Y. (1998). Balanced repeated replication for stratified multistage survey data under imputation. *Journal of the American Statistical Association,* **93**, 819-831.