

Paper 112-30

Rounding After Multiple Imputation With Non-binary Categorical Covariates

Christopher F. Ake, VA Healthcare System, San Diego, CA

Keywords: rounding, multiple imputation, categorical variables

ABSTRACT

At some point in their careers many SAS® users will confront a problem of missing data. A variety of statistical approaches have been developed to handle this problem. One of the most promising of these to emerge in the last several decades is multiple imputation. In SAS multiple imputation can now be performed in many contexts with the use of the MI and MIANALYZE procedures. To use these procedures, however, unless your data have a special structure known as monotone, they must be approximately multivariate normal, or transformable into multivariate normal. Methods have been proposed to accommodate data that include binary variables or categorical variables with more than two levels with missing values. These methods involve replacing non-binary categorical variables with sets of binary dummies, treating the binaries like normal variables in the imputation step, and then doing some kind of rounding of the imputed values. But such rounding may lead to bias in estimates derived from the resulting “filled-in” data. This paper provides estimates of such bias in the case of a sample mean in a real-life setting with extensive data, and also considers the alternative of retaining the unrounded imputed values for your analyses instead of rounding them.

INTRODUCTION

A variety of approaches have been proposed for doing statistical analyses with missing data, and from these a small number have become established in usage.

Perhaps the most widely used is listwise deletion, also referred to as complete case analysis, in which any observation with a missing value for any of its variables is deleted from the analysis. While this method is simple to use and applicable for any kind of statistical analysis, it can easily result in biased estimates, especially as the percentage of observations that are deleted becomes larger. And this method makes no use of whatever information is contained in the non-missing values among the deleted cases.

Pairwise deletion improves upon listwise deletion in the use of available information by using all available cases for a given analysis. For example, when computing a correlation coefficient, PROC CORR will use all those cases having non-missing values for the two variables in question, even though some of those cases may have missing values for some of their other variables. On the other hand, when using pairwise deletion in analyzing a single data set, different analyses may be based on different subsets of the data. And results may still be seriously biased.

Each of these deletion methods loses available information in your data. Imputation methods, on the other hand, retain all non-missing values. As an approach to handling missing data, imputation refers to any of a number of methods that involve in some way “filling in” missing values with a reasonable substitute and then analyzing the resulting data set consisting of the non-missing values together with the imputed values.

Among imputation methods the simplest may be mean imputation, in which, for any given variable with missing data, the mean of the non-missing values for that variable is substituted for each missing value. Since mean imputation only substitutes the same value for every missing value, it always underestimates uncertainty and for non-linear models usually produces biased results. In general, it can seriously underestimate any quantity that is not a linear function of the data, such as percentiles or measures of shape of the distribution (Little and Rubin, 2002). An extension of mean imputation that imputes mean values within subclasses of the data still suffers from the same problem of underestimation, although to a lesser degree.

In hot deck imputation, missing values are replaced by non-missing values taken from other observations that are classified as being sufficiently similar to the observation with the missing value. Again, with the same values being assigned to observations whose true values may differ, variances can be underestimated.

MULTIPLE IMPUTATION

Beginning with Version 8, SAS has included two procedures, MI and MIANALYZE, which allow you to generate multiple “completed” data sets from incomplete data by repeated processes of imputation. Each “completed” data set has missing values replaced by imputed values. These are imputed in such a way that in many situations you can then analyze each of these “completed” data sets with the procedures you would otherwise use for such data in the case of no missing values, and then combine the results of these separate analyses. When properly done, the results of these combined analyses not only yield unbiased estimators for parameters, but adequately incorporate the uncertainty involved because of the missing data, i.e., produce valid estimates of the variances of these parameter estimates.

There are several requirements for using MI and MIANALYZE to produce valid results. One is that your missing data are missing completely at random (MCAR) or at least missing at random (MAR), as defined by Rubin (1976). Roughly, your data are MCAR if the following condition is satisfied: whether the values in your data set that are missing are indeed missing or not is independent of what the values of any of your data are, i.e., both the observed values and the missing values. Your data are MAR if whether the missing values are missing or not is independent of what those missing values are, once you have conditioned on the observed values in your data. In a survey with an income question, for example, if people with higher incomes tended not to respond to a question about their income more than other people, the resulting survey data would not be MCAR. They could still be MAR, however, if, say, none of the other variables in your survey data had missing values, and it turned out that within subgroups defined by occupation and educational level whether the income question was answered was independent of income.

Ordinarily you would not know if your data satisfy the MAR assumption. It is impossible to verify empirically from your observed data alone—to do so you would need to know what the missing values are. But it is one that can often be assessed for its degree of plausibility in a given setting, depending on such things as your knowledge of the process that generated the data values, both observed and missing.

Here, however, we want to focus on a different kind of requirement for the use of MI and MIANALYZE. In order to use them you must specify an imputation model. You must provide MI with an input data set and provide MI with information as to how the variables in this data set are related to one another, e.g., by specifying that the distribution of all the data is multivariate normal, or else that the data have a certain other kind of structure known as “monotone.”

The imputation model that you use when running MI does not need to be the same as the model you will use in your subsequent analyses with your “completed” data sets, whatever those analyses might be, e.g., with PROC REG or PROC MIXED. When considering what variables to include in your imputation data set it is advisable to be conservative and include any variables that might contain information about the missing values or the missingness of the variables for which you will be imputing values. But in selecting the variables for the data you input to MI, whichever of its methods options you use, you must be mindful of the structural requirements for your data.

USING MI

METHODS AVAILABLE IN MI

PROC MI provides three different methods to perform multiple imputation.

Two of these require that your data exhibit a pattern of missingness known as monotone. For your data to have this structure, you must be able, working in the usual rectangular framework of each of n rows representing an observation and each of m columns representing a variable, to order your variables in such a way that any observation you choose then has all its missing values on the right hand side. In other words, for any given observation, if the i th variable has a missing value then so do all the later variables for that observation: the $(i+1)$ th variable, the $(i+2)$ th, all they way out to the m th variable. If you can set your data up in this way, then each variable has at least as many non-missing values as the one following it.

The first of these two methods in MI, specified with the METHOD=REGRESSION statement, then involves running a series of linear regressions, in which each of these regressions is used to produce imputed values for a given variable using as regressors all the variables to the left of it. In SAS 9 MI allows you the flexibility of specifying a particular set of covariates, not necessarily all those that precede it in the order given in the VAR statement, to be used in the imputation for any given variable, via options in the MONOTONE statement.

If you use the METHOD=PROPENSITY option, a propensity score is assigned for every observation of each variable with missing values, where the propensity score can be considered the conditional probability, given other covariate

values for that observation, of that variable having a missing value. In this method a logistic regression is used to calculate a propensity score, and as in the REGRESSION method, these logistic regressions are run for a given variable using as regressors all of the variables to the left, i.e., earlier in the ordering. Then for each variable all the observations are divided into groups based on their propensity scores and a bootstrap method is used to impute values within each group.

Beginning with Version 8.2 of PROC MI, if your data are not monotone but are close enough to being so, you can use the IMPUTE=MONOTONE option in the MCMC statement to fill in enough of a subset of the missing values to make the resulting imputed data sets have a monotone pattern.

If your data are not monotone or nearly so, multiple imputation can still be conducted by MI, but only on the basis of assuming your data are multivariate normal. Based on this assumption the Markov Chain Monte Carlo (METHOD=MCMC) method utilizes Bayesian inference to simulate draws from the posterior distribution of the missing data given the observed data. This method can accommodate only special departures from the assumption of multivariate normality. (Schafer 1997, pp. 147-148, 211-218).

Also in Version 8.2 MI features a TRANSFORM statement, which can be used to convert certain non-normal variables to normality, or approximate normality, via Box-Cox, exponential, logarithmic, logit, or power transformations prior to imputation, and then reverse-transform them afterward. This approach may work well, for example, if you need to transform highly skewed but continuous variables into roughly normal variables. But you cannot use it to transform non-continuous variables such as categorical variables into normal variables.

ROUNDING

What do you do if your data are not monotone so that you thereby want to use the MCMC option in PROC MI which assumes your data are multivariate normal, but your data are not and cannot be transformed into multivariate normal, or even nearly so?

This problem may be less serious in some cases than in others. If your only non-normal data are binary, an approach that has been proposed (Schafer, 1997, p. 148) as effective is to treat a binary variable as continuous in the imputation process and subsequently round each imputed value to the nearest category, i.e., to a '0' or a '1'.

But what if you want to impute values for a categorical variable with more than two levels? A straightforward extension of the approach used for binaries would be to replace each categorical variable by a corresponding set of binary dummy variables with 0-1 values, and then treat each of these dummies as an individual normal variate for imputation purposes, so that MI will impute a continuous value to it. You would need to follow this with a routine to assign a 1 to one of the dummies and 0s to the other dummies based on their imputed values in that iteration. In the last stage of each round of the imputation process, the set of dummy variables can be replaced by the original categorical variable, which will receive as a value the level whose dummy received the 1 in that round.

Does such an approach take care of your problem satisfactorily? Little work has been published on the adequacy of rounding the values imputed to binary variables. Horton et al., (2003), however, have shown that in the case of imputing missing values in a data set consisting of a single binary variable, rounding from an imputed value produced by treating the variable as normal, and imputing values via a regression relation based on the non-missing values, can induce a biased estimate of the mean. In contrast, they found that simply using the imputed nonrounded values, i.e., not rounding the value assigned by the imputation process at all, does not lead to bias for this estimate. They conjecture that "use of rounding in a more complicated setting may also yield nontrivial bias."

Horton et al. also speculated that the kind of rounding they examined would be less of a problem for discrete outcomes with many levels, on the grounds that the effect of rounding relative to the range of levels would be small, thereby reducing the relative bias, but they did not attempt to provide any estimate of how much less of a problem it would actually be.

IMPUTED VS. PLAUSIBLE VALUES

So we seem to face a conflict between two objectives you might have when you use multiple imputation. On the one hand, when you use any imputation procedure, you want the imputed values that result to be "plausible" in the sense of being sufficiently like those that would have been there if the data had not been missing, since you are going to use them as your "filled-in" data. Thus you would want the imputed values to be out of the range of the possible non-missing values, you would not want them to be continuous if the non-missing values are discrete, etc. On the other hand, when using the imputed data for various analyses, you will want the results of these analyses to be unbiased.

But if only unrounded imputes lead to unbiased results, and unrounded values are not “plausible” you cannot realize both objectives simultaneously.

If you have any expectation of running analyses where unbiased results will be sought, such conflict suggests you think carefully about exactly in what way you need your imputed data to have plausible values.

To assess the value of having plausible values in a completed data set you create by imputation, it might be useful to recall that multiple imputation was developed in a context strongly influenced by the need to provide end-users with public use data sets based on data with missing values. (Rubin, 1996) Because the person or persons creating the database might have access to confidential or other information not accessible to the end users, the imputation step, which might benefit from making use of this additional information, would be separated from any analysis step. In this context it would not necessarily be known at all who the actual end-users would be, or what they would want to do with the completed data they would receive. But the goal was to provide them with data sets of the same nature they would otherwise be working with, which rendered imputed values that were plausible especially desirable.

But if you are conducting the imputation process on some data with missing values and you will also be performing the data analyses, or know what they will be, then you may be in a position to decide for yourself how important it is that all values that are imputed need to be plausible in the sense of being the same as observed values (or values that could have been observed for the variable in question). A decision may depend on the cost involved in not rounding and thereby incurring the risk of biased estimates in your analyses. A small bias might be tolerable, for example, while a larger one would not be.

METHODS

In what follows we will consider results we have obtained on the extent of bias arising from not rounding in the simple case of estimating a mean. The extent of such bias in a somewhat complicated real-life setting is based on simulations using data from an ongoing VA study. The simulations first looked at estimates for a four-level categorical variable specially constructed from these data. In doing so it became apparent that we should examine how rounding affects a simple binary also. The following is a description of the data and imputation model our research group at the VA has been using. Simulation methods, results for binary variables, and then for our four-level categorical variable are then presented.

DATA USED IN SIMULATIONS

In the first round of an ongoing observational study, our research group has used proportional hazards models to examine whether the 36,766 patients in the Veterans Administration Healthcare system who have received VA care for HIV/AIDS between January 1, 1993, and June 30, 2001, show a higher risk for certain adverse effects if they have received some form of highly-active antiretroviral therapy (HAART) compared to those under VA care for HIV/AIDS who have not received any such treatment. We have been examining a variety of adverse events, including all-cause mortality, mortality due to cardiovascular or cerebrovascular disease, as well as inpatient admissions for such diseases. For each such outcome event, we examine the effect of exposure to nucleoside analogs, non-nucleoside reverse transcriptase inhibitors, and protease inhibitors singly and in combination, both in the form of being ever exposed and of amount of cumulative exposure. In all models we adjust for certain patient characteristics with demographic and other covariates.

OUR MISSING DATA

Among the covariates in all our models are four with missing values: age at first VA care for HIV (4 cases out of 36,766 or 0.01%), baseline severity of illness (181 cases or 0.49%), race (2684 cases or 7.30%), and risk group at first VA care for HIV (4632 cases or 12.60%). These total to slightly less than 1% of all entries in the various 36,766 x 22 to 36,766 x 26 data matrices of the time-independent covariates we use in our modeling. In our analyses we have handled the age variable as continuous, the severity variable as ordinal, and the race and risk variables as strictly categorical.

OUR IMPUTATION MODEL

The 19 variables used in the original imputation modeling were: race, HIV risk category, severity of illness at entry into VA care for HIV, age at first VA care for HIV, lower bound on severity of illness, year of initial VA care for HIV, facility switcher (whether person is prone to switching facilities or not), number of months in VA care prior to initial VA care for HIV, and a set of binary indicators for exposure to ARV treatment within 3 months of first VA care for HIV, treatment before or at initiation of VA care for HIV for -AIDS-defining illness, for a cardiovascular diagnosis or

procedure, for a cerebrovascular diagnosis or procedure, for other vascular disease, for diabetes, for smoking, for hypertension, for drug abuse, for hyperlipidemia, or for a certain set of other risk factors.

For our imputations, we have used the MCMC option in PROC MI version 8.1 with single long chains and Jeffrey's non-informative prior with the "INITIAL=EM" option. We used 400 burn-in iterations and 200 iterations between imputations in a single chain. We used one macro to create the binary indicators, with the data sets it created passed to a separate macro to run the imputations. Other details of our modeling are discussed in Ake and Carpenter (2002).

BINARY VARIABLES: METHOD

We selected two binary variables from our data set for our examination of bias due to non-rounding. The first (PREVSMOF) is an indicator that records for each patient whether there is any record of VA treatment for various smoking-related conditions before that person began their VA treatment for HIV. The second (PREVAVASF) is an indicator that similarly records whether a patient has received certain kinds of vascular care in the same time period. We chose these two for two reasons. First, their "p" values, i.e., the proportion of 1s vs. 0s in the data, are close to one another: 0.04123 for PREVSMOF, and 0.03998 for PREVAVASF. Results for these two might thereby provide a rough idea of how much variability to expect among very similar values. Second, each has a p value that the theoretical results of Horton et al. (2003) suggested might make them susceptible to nearly as high a degree of absolute bias, when estimating their mean based on imputation of missing values for them, as any binary variables might be.

In our data neither of these variables in fact has any missing values. So we set values of each to missing in simple (pseudo-)random samples of our observations comprising 1%, 5%, 10%, 20% and 40% of our total observations. We generated these samples using the PROC SURVEYSELECT with the METHOD=SRS option and the SAMPRATE option set to the various percentages of interest.

We then evaluated three different methods of handling those missing values, by comparing how close the estimated mean based on the remaining non-missing values together with the values imputed by that method came to the true mean of all the original non-missing values. To do so we used the same data set for our imputation modeling as we have when imputing value for the age, severity, race, and risk variables. For each of these two binaries and each of the variously sized random samples, we ran our multiple imputation process five times with $m=5$ completed data sets, imputing values for the missing binary as well as our other four variables with missing values.


In the first of these methods, listed below as ROUND1, we treat the binary as a continuous normal variable by using the MCMC option of MI. After the values are imputed we round by assigning a "1" if the imputed value is greater than or equal to .5 and assigning a "0" otherwise. In the second method, ROUND2, we simply include the ROUND=1 option in the PROC MI statement. When this option is included values are rounded to the unit specified, so this method should produce results like the first (although the second method does not prohibit an integer value like 2 or -1 being assigned, however unlikely this might be; this can be avoided by specifying the further options MIN=0 and MAX=1, which we did not do). In the third method, NOROUND, we simply leave the imputed value unrounded.

In each of these three methods, after running PROC MI and then either rounding the values or leaving them unrounded, we follow this by using PROC CORR to generate mean values for each of the $m=5$ completed data sets for each run, and then use PROC MIANALYZE, with the BY _IMPUTATION_ statement, to combine the mean estimates across the imputation runs. Finally we compare the resulting estimates for the mean value of each of the dummies across the three methods.


BINARY VARIABLES: RESULTS

Figures A1 and A2 below show results for these two binaries in terms of relative bias, i.e., $(p' - p)/p$, where p' is the estimate of the mean and p is the actual value. As expected, results are very similar for the two variables, even with only five rounds of multiple imputation (the same sample and imputation seeds were used for both variables). Note that the horizontal scale is neither linear nor logarithmic. Clearly the relative bias increases with sample size for ROUND1 and ROUND2, reaching a value in the range of -0.25 to -0.30 when the missingness reaches a proportion of 0.40. And clearly NOROUND does much better—in fact, these results suggest that it may indeed represent an unbiased estimator. ROUND1 seems to perform slightly better than ROUND2; perhaps this is because for ROUND2 we used the ROUND=1 option without the MIN and MAX options.

A1. Relative Bias of Binary with $p=0.04123$



A2. Relative Bias of Binary with $p = 0.03998$



CATEGORICAL VARIABLES: METHOD

To examine how such bias might work in the case of a categorical variable with more than two levels, we decided to use the cross product of two of the binary variables in our data set, because the percentage of 1s occurring at each of the four resulting levels was spread out from 4.7%, through 9.7 and 23.0, to 62.7%. Thus we could check the relative bias for smaller as well as larger percentages.

We ran PROC MI with binary dummies included in the imputation model to represent the various levels of this four-way categorical variable. In each method we treated each such binary variable as a continuous normal variable. As

in the case of examining a single binary variable, in each method we use with categorical variables we first selected a random sample of observations for which we then set these dummy variables to missing.

For categorical variables we used four methods to assign values. In method ROUND1 we include four binary dummies, one for each level of our categorical variable, and after running PROC MI assign a '1' to whichever of the four levels' dummy had the highest imputed value, and a '0' to each of the other levels.

ROUND2 is a slight variant of ROUND1. As described in Allison (2001, pp. 39-40), you can, in the imputation process, use n-1 instead of n dummies for a variable with n categories. After you run PROC MI you assign a value to the nth category, in our case the fourth, by subtracting the sum of the other n-1 imputed values from 1. You then, as in the first method, assign a '1' to the category whose dummy got the highest value, and a '0' to the other categories.

NOROUND1 starts out the same as ROUND1, by including a dummy for each level of the categorical variable you are imputing values for, but then, rather than assigning 1s and 0s, simply leaves each imputed value as is, i.e., unrounded.


Finally, in method NOROUND2 you start out as you would in ROUND2, i.e., you put n-1 dummies in the imputation model corresponding to an n-level categorical variable. After you run PROC MI you assign a value to the nth level by subtracting the sum of the imputed values for the other n-1 levels from 1. You then leave all these n values unrounded.

In each of these four methods, after running PROC MI and then either rounding the values or leaving them unrounded, we follow this, as we did with the binaries, by using PROC CORR to generate mean values for each of the m=5 completed data sets for each run, and then use PROC MIANALYZE, with the BY `_IMPUTATION_` statement, to combine the mean estimates across the imputation runs. Finally we compare the resulting estimates for the mean value of each of the dummies across the four methods as indicated in the tables below. Each estimate is based on 25 rounds of multiple imputation with m=5 completed data sets generated in each round. Seeds for the sample selection process in PROC SURVEYSELECT as well as the seeds for imputation process in PROC MI were varied across these 25 runs.


CATEGORICAL VARIABLES: RESULTS



Figures B1 through B4 below show the relative bias for the estimated mean for each of the four levels of our categorical variable. Again, the horizontal scale is not linear in percent missing. In addition, the vertical scales vary from one figure to another, given the differing ranges of relative biases. In figure B1 results for the level whose 'p', i.e., true proportion of 1s in our data, is 0.0466, show a close similarity to those for the two for our two binaries above, whose p values are 0.03998 and 0.04123. As with the two binaries, ROUND1 and ROUND2 here result in relative bias in the -0.25 range with 40% missing values. NOROUND1 and NOROUND2 produce nearly identical results that again suggest they represent unbiased estimators. In figures B2 and B3 ROUND1 and ROUND2 again produce extremely similar estimates, but the relative bias is now positive, in the range of 0.02 in B2 and 0.045 in B3 for the 40% missing samples. Again NOROUND1 and NOROUND2 produce estimates suggesting unbiasedness, save for larger relative bias for the 20% sample in B2. Finally, in B4 the estimates show more variability, with ROUND1 and ROUND2 diverging more from each other in relative bias, and the results for NOROUND1 and NOROUND2 displaying much more departure from the 0 line than in the other levels. And in B4 ROUND1 and ROUND2 exhibit negative relative bias, as in B1, although here of much smaller magnitude, in the -0.002 range for the 40% sample.

B1. Relative Bias of Dummy for Level with $p=0.0466$



B2. Relative Bias of Dummy for Level with $p = 0.0966$



B3. Relative Bias of Dummy for Level with $p=0.2295$ B4. Relative Bias of Dummy for Level with $p= 0.6273$ 

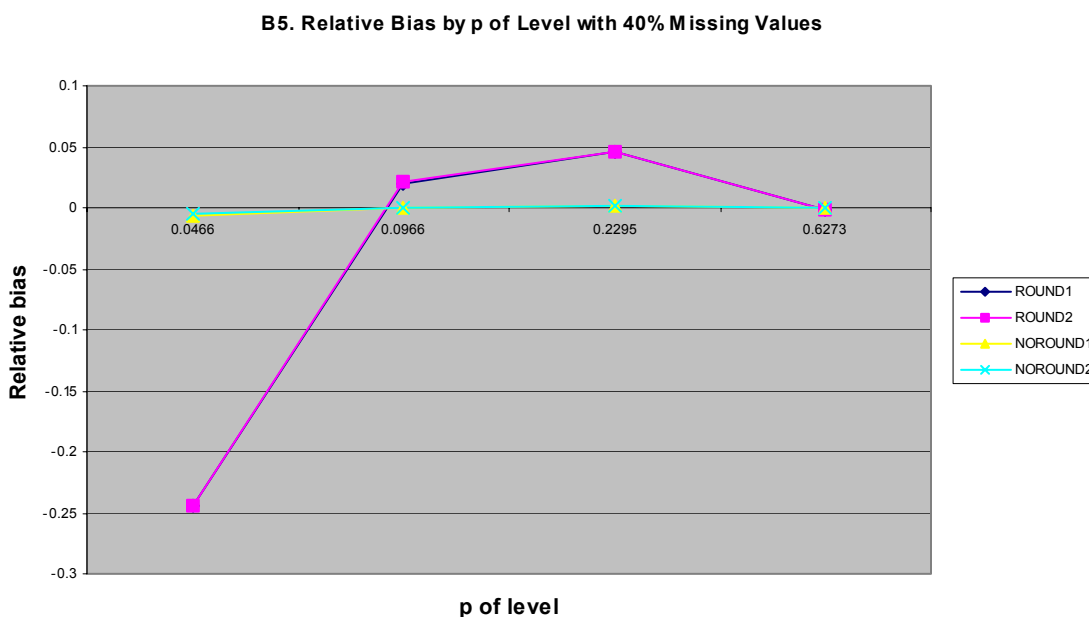
DISCUSSION

Our project here has been to evaluate various procedures that could be used with PROC MI, some involving rounding and some not, by how well each of these procedures performs in one particular real-data setting with respect to one component, relative bias, of one kind of statistical analysis, estimation of a mean.

These results are based on a limited number of runs of MI and MIANALYZE for each method and percentage of missing values, five in the case of the binary variable and 25 for the four level categorical variable. They are also based on using one specific data set and imputation model with two individual binaries from that data set, and one particular combination of two binaries chosen to create a four level categorical variable. Our results might differ, for

example, if we had selected a different binary which turned out to be more correlated with the other variables in the data set than the one we chose—perhaps this would ameliorate the bias in the results.

A clearer picture of the extent of bias in various settings may require further simulation, given how complicated the algorithms underlying the multiple imputation process can be. And the larger question remains of what the effect of rounding on various other parameter estimates than the mean would be. It is noteworthy, nevertheless, that the results for rounding error shown above in Figures A1-A2 and B1-B4 all appear consistent in an interesting way with the theoretical result of Horton et al. (2003, Figure 2) for a single univariate binary. They found that relative bias from rounding should move from roughly -1 to 0 as the p value went from 0 to roughly 0.1 , then assume smaller positive values until $p=0.5$, then presumably take positive but smaller values until roughly $p=0.9$. Our results accord with this even in the case of a four-level categorical, as indicated by Figure B5 below, which compares the relative bias for all levels for the 40% missing value case.



CONCLUSION

In SAS 9 PROC MI has an experimental CLASS statement to accommodate categorical variables, but this only works in the case of monotone data. Therefore analysts who use MI and MIANALYZE to handle problems of missing data should continue to use a special purpose procedure for handling categorical variables. For analyses involving such variables, at least with respect to bias, leaving imputed values unrounded appears preferable to rounding to the nearest observed or plausible individual value.

TRADEMARK CITATION

SAS is a registered trademark of the SAS Institute in the USA and other countries. ® indicates USA registration.

REFERENCES

Ake, Christopher F. and Arthur L. Carpenter, (2002), "Survival Analysis with PHREG: Using MI and MIANALYZE to Accommodate Missing Data," Proceedings of the 10th Annual Western Users of SAS Software Regional Users Group Conference (2002), pp. 102-107, Cary, NC: SAS Institute, Inc.

Allison, P. D. (2001), *Missing Data*, Sage University Papers Series on Quantitative Applications in the Social Sciences, 07-136. Thousand Oaks, CA: Sage.

Horton, N.J., Lipsitz, S.R., and Parzen, M. (2003), "A Potential for Bias When Rounding in Multiple Imputation," *American Statistician*, 57, 229-232.

Little, R.J.A. and Rubin, D.R., (2002), *Statistical Analysis with Missing Data*, 2nd ed., John Wiley & Sons. Hoboken, NJ.

Rubin, Donald (1976), "Inference and Missing Data," *Biometrika* 63: 581-592.

Rubin, Donald (1996), "Multiple Imputation After 18+ Years," *Journal of the American Statistical Association* 91: 473-489.

Schafer, J.L. (1997), *Analysis of Incomplete Multivariate Data*, London: Chapman and Hall.

ACKNOWLEDGEMENTS

Dr. Jacinte Jean was instrumental in initial design of data analyses. Dr. Samuel A. Bozzette has directed the entire research project. Dr. Thomas A. Louis of the Johns Hopkins Bloomberg School of Public Health has served as a statistical consultant for the project. Art Carpenter and Rick Smith have created almost all the macros the project uses.

AUTHOR CONTACT INFORMATION

Chris Ake
Health Services Research & Development
Veterans Affairs San Diego Healthcare System
3350 La Jolla Drive, 111N-1
San Diego, CA 92161
(858) 552-8585 x 5492
Christopher.Ake@med.va.gov