Survey Methodology, June 1986
Vol. 12, No. 1, pp. 37–47
Statistics Canada

Basic Ideas of Multiple Imputation for Nonresponse

DONALD B. RUBIN[1]

ABSTRACT

Multiple imputation is a technique for handling survey nonresponse that replaces each missing value created by nonresponse by a vector of possible values that reflect uncertainty about which values to impute. A simple example and brief overview of the underlying theory are used to introduce the general procedure.

KEY WORDS: Survey nonresponse; Proper imputation methods; Multiple imputation.

1. INTRODUCTION

Any statistician with experience in the field of surveys knows that essentially every survey suffers from some nonresponse. That is, in practical surveys, some items in the survey instrument are not answered by all units included in the survey. Commonly, the items likely to be unanswered are the more sensitive ones, such as those concerning personal income. Because nonresponse creates missing values, the complete-data statistics that would have been used in the absence of nonresponse can no longer be calculated. An obvious desire of both the data collector and the data analyst is to get rid of the missing values and thereby restore the ability to use standard complete-data methods to draw inferences.

1.1 Imputation

It is not surprising, therefore, that a very common method of handling the missing values created by nonresponse is to fill them in, or impute them. That is, when using imputation to handle nonresponse each missing value is replaced with a real value. Many different procedures have been proposed for imputation, for instance, filling in the respondents' mean for that variable or a value predicted from the modelling of the missing variable given observed variables using respondent data; as a specific example, when the missing value is personal income, a linear regression model predicting log(income) from demographic characteristics such as age, sex, education and occupation might be regarded as reasonable.

1.2 Advantages and Disadvantages of Single Imputation

In addition to the obvious advantage of allowing complete-data methods of analysis, imputation by the data collector (e.g. the Census Bureau) also has the important advantage of being able to utilize information available to the data collector but not available to an external data analyst such as a university social scientist analyzing a public-use file. This information may involve detailed knowledge of interviewing procedures and reasons for nonresponse that are too cumbersome to place in public-use files, or may be facts, such as street addresses of dwelling units, that cannot be placed on public-use files because of confidentiality constraints. This kind of information, even though inaccessible to the user of a public-use file, can often narrow the possible range of imputed values.

[1] Donald B. Rubin, Department of Statistics, Harvard University, Science Center, One Oxford Street, Cambridge, Massachusetts, 02138, U.S.A.

Just as there are obvious advantages to imputing one value for each missing value, there are obvious disadvantages of this procedure arising from the fact that the one imputed value cannot itself represent any uncertainty about which value to impute: If one value were really adequate, then that value was never missing. Hence, analyses that treat imputed values just like observed values generally systematically underestimate uncertainty, even assuming the precise reasons for nonresponse are known. Equally serious, single imputation cannot represent any additional uncertainty that arises when the reasons for nonresponse are not known.

1.3 Multiple Imputation to the Rescue

Multiple imputation, first proposed in Rubin (1977, 1978), retains the two major advantages of single imputation and rectifies its major disadvantages. As its name suggests, multiple imputation replaces each missing value by a vector composed of $M \geq 2$ possible values. The M values are ordered in the sense that the first components of the vectors for the missing values are used to create one completed data set, the second components of the vectors are used to create the second completed data set and so on. The first major advantage of single imputation is retained with multiple imputation, since standard complete-data methods are used to analyze each completed data set. The second major advantage of imputation, that is, the ability to utilize data collectors' knowledge in handling the missing values, is not only retained but actually enhanced. In addition to allowing data collectors to use their knowledge to make point estimates for imputed values, multiple imputations allow data collectors to reflect their uncertainty as to which values to impute. This uncertainty is of two types: sampling variability assuming the reasons for nonresponse are known, and variability due to uncertainty about the reasons for nonresponse. Under each posited model for nonresponse, two or more imputations are created to reflect sampling variability under that model; imputations under more than one model for nonresponse reflect uncertainty about the reasons for nonresponse. The multiple imputations within one model are called repetitions and can be combined to form a valid inference under that model; the inferences under different models can be contrasted to reveal sensitivity of answers to posited reasons for nonresponse.

Before reviewing some more general results in Section 3, Section 2 illustrates essential ideas in a highly artificial example used in Rubin (1986a), which is a comprehensive treatment of multiple imputation. Other references on multiple imputation include Rubin (1979, 1980, 1986b), Herzog and Rubin (1983), Li (1985), Schenker (1985), Rubin and Schenker (1986), and Heitjan and Rubin (1986).

2. AN ARTIFICIAL EXAMPLE ILLUSTRATING MULTIPLE IMPUTATION

Suppose we have taken a simple random sample of $n = 10$ units from a large population. The objective of the survey is to estimate $\bar{Y}$ the mean of Y in the population. We know the mean value of a covariate X in the population, and the survey attempts to record both X and Y for each of the n units included in the sample.

Table 1 presents the observed values of (Y, X) for the ten units in the sample where the question marks indicate missing Y data due to nonresponse.

2.1 Multiply Imputing for the Missing Values

Suppose the missing values in Table 1 are to be multiply imputed using two values drawn under each of two models (i.e. two repetitions per model). In general, any number of models can be used with any number of repetitions within each model. Model 1 is an "ignorable" model for nonresponse; ignorable is defined precisely in Rubin (1976), but essentially it means that a nonrespondent is only randomly different from a respondent with the same value of X. Model 2 is a nonignorable model and posits a systematic difference between respondents and nonrespondents with the same value of X. The repeated imputations under each model are based on a simple procedure closely related to the hot-deck, which can be improved upon but is useful to illustrate ideas.

For each nonrespondent, the two closest matches among the respondents are found, where the distance for matching is defined by the values of X. For the first nonrespondent, unit 2, the two closest matches are units 1 and 3, and for the second nonrespondent, unit 4, the closest matches are 3 and 5. The repeated imputations are created by drawing at random from the two closest matches. For the ignorable model, we simply impute the value Y provided by the matching respondent: the first two columns of Table 2 give the result. For the nonignorable model, we suppose that the nonresponse bias is such that a nonrespondent will tend to have a value of Y 20% higher than the matching respondent's value of Y: the last two columns of Table 2 give the result where the Y values have been rounded to the nearest integer. The repeated imputations within each model allow the user to draw a valid inference under that model. The use of two models, an ignorable one and a nonignorable one, allows the display of sensitivity of inference to assumptions about nonresponse. Generally such assumptions are untestable using the data at hand.

Table 1

Observed Data

Unit	Y	X
1	10	8
2	?	9
3	14	11
4	?	13
5	16	16
6	15	18
7	20	6
8	4	4
9	18	20
10	22	25

Table 2

Multiple Imputations for Data of Table 1

	Model 1 Repetition		Model 2 Repetition	
	1	2	1	2
Unit 2	10	14	12	17
Unit 4	16	14	19	17

2.2 Analyzing the Resultant Multiply-Imputed Data Set

Each set of imputations, that is each column of Table 2, can be used with the incomplete data in Table 1 to create a completed data set. Since there are four sets of imputations, four completed data sets can be created; these are displayed in Tables 3 to 6. Each completed data set is analyzed just as if there had been no nonresponse.

Assume that with complete data, the ratio estimator $\bar{X}\bar{y}/\bar{x}$ would be used with associated variance SE^2, where $\bar{X}$ is the known mean of X in the population, say 12, $\bar{y}$ and $\bar{x}$ are the means of Y and X in the random sample of n units, and

$$SE^2 = \sum (Y_i - X_i\bar{y}/\bar{x})^2/[n(n - 1)]$$

Table 3

Complete Data Set 1 (Model 1, Rep. 1)
For Multiply Imputed Data Set of Tables 1 and 2

Unit	Y	X
1	10	8
2	10	9
3	14	11
4	16	13
5	16	16
6	15	18
7	20	6
8	4	4
9	18	20
10	22	25
means	14.5	13

Table 4

Complete Data Set 2 (Model 1, Rep. 2)
For Multiply Imputed Data Set of Tables 1 and 2

Unit	Y	X
1	10	8
2	14	9
3	14	11
4	14	13
5	16	16
6	15	18
7	20	6
8	4	4
9	18	20
10	22	25
means	14.7	13

Table 5

Complete Data Set 3 (Model 2, Rep. 1)
For Multiply Imputed Data Set of Tables 1 and 2

Unit	*Y*	*X*
1	10	8
2	12	9
3	14	11
4	19	13
5	16	16
6	15	18
7	20	6
8	4	4
9	18	20
10	22	25
means	15	13

Table 6

Complete Data Set 4 (Model 2, Rep. 2)
For Multiply Imputed Data Set of Tables 1 and 2

Unit	*Y*	*X*
1	10	8
2	17	9
3	14	11
4	17	13
5	16	16
6	15	18
7	20	6
8	4	4
9	18	20
10	22	25
means	15.3	13

Table 7

Ratio Estimates and Associated Variances of Estimates
for the Complete Data Sets of Tables 3-6

	Model 1 Repetition		Model 2 Repetition	
	1	2	1	2
Estimate	13.38	13.57	13.85	14.12
Variance	2.96	3.19	3.38	3.84

Table 8

Combined Estimates and Variances for the Multiply Imputed Data Sets of Tables 1 and 2

	Model 1	Model 2
Estimate	13.48	13.98
Variance	3.10	3.66

where the sum is over the units in the sample. Table 7 presents the estimates and variances associated with each of the four completed data sets given in Tables 3-6.

The two answers obtained under the same model can be combined to obtain one inference for $\bar{Y}$ under each model. The results are displayed in Table 8: the estimate is the average of the estimates and the variance associated with this estimate has two components: (i) the average within-imputation variance associated with the estimate and (ii) the between-imputation variance of the estimate. Thus, under Model 1, the estimate is $(13.38 + 13.57)/2 = 13.48$; the associated estimated average within variance is $(2.96 + 3.19)/2$, and the associated estimated between variance is $[(13.38 - 13.48)^2 + (13.57 - 13.48)^2)]$. The estimated variances are combined as: (estimated total variance) = (estimated average within variance) + $(1 + M^{-1})$ × (estimated between variance), where the factor $(1 + M^{-1})$ multiplying the usual unbiased estimate of between variance is an adjustment for using a finite number of imputations. The associated 95% interval estimate for $\bar{Y}$ is (10.0, 16.9) under Model 1 and (10.2, 17.7) under Model 2. In practice, better intervals can be formed by calculating degrees of freedom as a simple function of the variance components and using the 95% points appropriate to the corresponding t-distribution; when either M is large or the between variance component is small relative to the total variance (as in this artificial example), the degrees of freedom will be large and thus the normal 95% points will be used. Details are given in Section 3.

The essential feature to notice in this illustrative example is that only complete-data methods of analysis are needed. We merely have to perform the complete-data analysis that would have been used in the absence of nonresponse on each of the completed data sets created by the multiple imputations. The resultant answers under each model are then easily combined to give one inference under each model. Although not illustrated here, diagnostic analyses using complete-data techniques can be applied to each completed data set; Heitjan and Rubin (1986) provides several examples.

3. GENERAL PROCEDURES

The example in Section 2 illustrated methods for creating multiple imputations and analyzing the resultant multiply-imputed data set in a special case. We now outline the methods needed for general practice.

3.1 Proper Imputation Methods

Multiple imputations ideally should be drawn according to the following general scheme. For each model being considered, the M imputations of the missing values, Y_{mis}, are M repetitions from the posterior predictive distribution of Y_{mis}, each repetition being an independent drawing of the parameters and missing values under an appropriate Bayesian model for the posited response mechanism. In practice, implicit models such as illustrated in Section 2 can often be used in place of explicit models. Both types of models are illustrated in Herzog and Rubin (1983), where repeated imputations are created using an explicit regression model and an implicit matching model, which is a modification of the Census Bureau's hot-deck.

Procedures that incorporate appropriate variability among the repetitions within a model are called *proper*, which is defined precisely in Rubin (1986a). The essential idea of proper imputation methods is to properly reflect sampling variability when creating repeated imputations under a model. For example, assume ignorable nonresponse so that respondents and nonrespondents with a common value of X have Y values only randomly different from each other. Even then, simply randomly drawing imputations for nonrespondents' from matching respondents' Y values ignores some sampling variability. This variability arises from the fact that the sampled respondents' Y values at X randomly differ from the population of Y values at X. Properly reflecting this variability leads to repeated imputation inferences that are valid under the posited response mechanism.

In the context of simple random samples and ignorable nonresponse, Rubin and Schenker (1986) study hot-deck imputation (i.e. simply randomly drawing imputed values from respondents), which is *not* proper, and a variety of proper imputation methods based on both explicit and implicit models, including a fully normal model, the Bayesian Bootstrap (Rubin, 1981), and an approximate Bayesian Bootstrap. The Approximate Bayesian Bootstrap (ABB) can be used to illustrate how an intuitive imputation method, such as the simple random hot-deck, can be modified to be proper.

3.2 Example of a Proper Imputation Method with Ignorable Nonresponse - The ABB

Consider a simple random sample of size n with n_R respondents and $n_{NR} = n - n_R$ nonrespondents. The ABB creates M ignorable repeated imputations as follows. For $\ell = 1, \ldots, M$, create n possible values of Y by first drawing n values at random with replacement from the n_R observed values of Y, and second drawing the n_{NR} missing values of Y at random with replacement from those n values. The drawing of the n_{NR} missing values from a possible sample of n values rather than the observed sample of n_R values generates appropriate between imputation variability, at least in large samples, as shown by Rubin and Schenker (1986). The ABB approximates the Bayesian Bootstrap by using a scaled multinomial distribution to approximate a Dirichlet distribution.

3.3 Analysis - The Repeated Imputation Inference

The general methods for analyzing a multiply imputed data set implicitly assume proper imputation methods have been used to create the multiple imputations. As illustrated in Section 2, the repeated imputations within each model are analyzed as a collection to create one *repeated-imputation* inference as follows. Each data set completed by imputation is analyzed using the same complete-data method that would be used in the absence of nonresponse. More precisely, let $\hat{\Theta}_\ell, U_\ell, \ell = 1, \ldots, M$ be M complete-data estimates and their associated variances for a parameter Θ, calculated from the M data sets completed by repeated imputations under one model for nonresponse. The final estimate of Θ is

$$\bar{\Theta}_M = \sum_{\ell=1}^{M} \hat{\Theta}_\ell / M.$$

The variability associated with this estimate has two components: the average within-imputation variance,

$$\bar{U}_M = \sum_{\ell=1}^{M} U_\ell / M,$$

and the between-imputation component,

$$B_M = \sum (\hat{\Theta}_\ell - \bar{\Theta}_M)^2 / (M-1)$$

where with vector Θ, $(\bullet)^2$ is replaced by $(\bullet)^{\mathrm{T}}(\bullet)$. The total variability associated with $\bar{\Theta}_M$ is then

$$T_M = \bar{U}_M + (1 + M^{-1}) B_M.$$

With scalar Θ, the reference distribution for interval estimates and significance tests is a t-distribution.

$$(\Theta - \bar{\Theta}_M) \, T_M^{-1/2} \sim t_v,$$

where the degrees of freedom,

$$v = (M - 1) \{ 1 + [(1 + M^{-1}) B_M / \bar{U}_M]^{-1} \}^2$$

is based on a Satterthwaite approximation (Rubin and Schenker 1986 and Rubin 1986a). The within to between ratio $\bar{U}_M / B_M$ estimates the population quantity $(1 - \gamma)/\gamma$, where γ is the fraction of information about Θ missing due to nonresponse. In the case of ignorable nonresponse with no covariates, γ equals the fraction of data values that are missing.

3.4 Significance Levels for Multicomponent Θ

For Θ with k components, significance levels for null values of Θ can be obtained from M repeated complete-data estimates, $\hat{\Theta}_\ell$, and variance-covariance matrices, U_ℓ, using multivariate analogues of the previous expressions.

A simple procedure described in Li (1985) and Rubin (1986a) that works well for M large relative to k is to let the p-value for the null value Θ_0 of Θ be Prob $\{F_{k,v} > D_M\}$ where $F_{k,v}$ is an F random variable and $D_M = (\Theta_0 - \bar{\Theta}_M) T_M^{-1} (\Theta_0 - \bar{\Theta}_M)^{\mathrm{T}}$ with v defined by generalizing $B_M / \bar{U}_M$ to be the average diagonal element of $B_M \bar{U}_M^{-1}$, $\mathrm{trace}(B_M \bar{U}_M^{-1})/k$. Better procedures are described in Rubin (1986a). Less precise p-values can be obtained directly from M repeated complete-data significance levels; also see Rubin (1986a).

4. DISCUSSION

4.1 Frequency Evaluations

Although repeated imputation inferences are most directly motivated from the Bayesian perspective, they can be shown to possess good frequency properties. In fact, the definition of proper imputation methods means that in large samples infinite-M repeated imputation inferences will be valid. Since the finite-M adjustments are derived using approximations to Bayesian posterior distributions, however, deficiencies can arise with finite M. For example, the large sample relative efficiency of $\bar{\Theta}_M$ to $\bar{\Theta}_\infty$ that is, the efficiency of the finite-M repeated imputation estimator using proper imputation methods relative to the infinite-M estimator in units of standard errors is $(1 + \gamma/M)^{-1/2}$. Even for relatively large γ, modest values of M result in estimates $\bar{\Theta}_M$ that are nearly fully efficient.

4. 2 Confidence Coverage

In large samples the confidence coverage of proper imputation methods using the t-reference distribution can be tabulated as a function of M, γ and the nominal level, $1 - \alpha$. Table 9 is from Rubin (1986a) and is also partially reported in Rubin and Schenker (1986) and Schenker (1985). Also included are results for single imputation, where the between component of variance is set to zero, since it cannot be estimated, and the reference distribution is the normal, since υ cannot be estimated without B_M. Even in extreme cases, two or three repeated imputations yield nearly valid confidence coverages; this is in striking contrast to using only one imputation. Even worse coverages for single imputation would have been obtained using best prediction methods, such as "fill in the mean".

Table 9

Coverage probabilities in % of interval estimates based on the t-reference distribution as a function of the number of proper repeated imputations, $M \geq 2$, the fraction of missing information, γ, and the nominal level, $1 - \alpha$. Also included for contrast are results based on single imputation $M = 1$, using the normal reference distribution with the between component of variability set to zero.

		γ								
$1-\alpha$	M	.1	.2	.3	.4	.5	.6	.7	.8	.9
	1	46	42	38	34	30	26	22	18	12
	2	50	50	51	51	50	50	50	50	50
50%	3	50	50	50	50	50	50	50	50	50
	5	50	50	50	50	50	50	50	50	50
	∞	50	50	50	50	50	50	50	50	50
	1	75	70	65	60	54	48	41	33	23
	2	80	80	80	79	78	77	76	76	76
80%	3	80	80	80	80	79	79	79	79	79
	5	80	80	80	80	80	80	80	80	80
	∞	80	80	80	80	80	80	80	80	80
	1	86	82	77	72	66	59	51	42	29
	2	90	90	89	88	87	86	85	84	83
90%	3	90	90	90	89	89	88	88	88	88
	5	90	90	90	90	90	90	89	89	89
	∞	90	90	90	90	90	90	90	90	90
	1	92	89	85	80	74	67	59	49	35
	2	95	95	94	93	92	91	89	88	87
95%	3	95	95	95	94	94	93	93	92	92
	5	95	95	95	95	95	94	94	94	94
	∞	95	95	95	95	95	95	95	95	95
	1	98	96	94	91	86	80	72	61	45
	2	99	99	98	98	97	96	95	93	92
99%	3	99	99	99	98	98	98	97	97	96
	5	99	99	99	99	99	99	98	98	98
	∞	99	99	99	99	99	99	99	99	99

4.3 Significance Levels

Work on accurately obtaining significance levels is at an early stage of development. Table 10 is from Rubin (1986a) and is also partially reported in Li (1985). It indicates that if $M > k$ and γ is modest, accurate tests can be obtained using D_M. Better procedures are considered by Li (1985), Rubin (1986a) and in current thesis work by T.E. Raghunathan.

Table 10

Level in % of D_M with $F_{k,\,\upsilon}$ reference distribution as a function of: nominal level, α; number of components being tested, k; number of repeated proper imputations, M; and fraction of missing information, γ.

			$\alpha = 1\%$				$\alpha = 5\%$				$\alpha = 10\%$			
k	M	$\gamma =$	.1	.2	.3	.5	.1	.2	.3	.5	.1	.2	.3	.5
	2		1.0	1.2	1.6	2.5	4.9	5.3	5.9	7.5	9.9	10.3	11.0	12.9
	3		1.0	1.0	1.0	1.3	4.9	4.9	5.0	5.5	9.9	9.8	10.0	10.9
2	5		1.0	1.0	1.1	1.2	5.0	5.0	5.1	5.6	10.0	10.0	10.2	10.9
	10		1.0	1.0	1.1	1.2	5.0	5.1	5.3	5.7	10.1	10.2	10.4	11.0
	25		1.0	1.0	1.0	1.0	5.0	5.0	5.0	5.0	10.0	9.9	9.9	10.0
	50		1.0	1.0	1.0	1.0	5.0	5.0	5.0	5.0	10.0	9.9	9.9	10.0
	100		1.0	1.0	1.0	1.0	5.0	5.0	5.0	5.0	10.0	10.0	10.0	10.1
	2		1.0	1.1	1.3	1.7	5.1	5.3	5.6	6.3	10.3	10.6	11.1	12.0
	3		1.0	1.0	1.0	1.0	5.1	5.2	5.3	5.7	10.2	10.5	10.9	12.3
3	5		1.0	1.0	1.1	1.3	5.0	5.2	5.4	6.2	10.1	10.3	10.8	12.2
	10		1.0	1.0	1.1	1.2	5.0	5.2	5.3	5.9	10.1	10.3	10.6	11.6
	25		1.0	1.0	1.1	1.2	5.0	5.1	5.2	5.6	10.1	10.2	10.4	10.9
	50		1.0	1.0	1.0	1.0	5.0	5.0	5.0	5.1	10.0	10.0	10.0	10.2
	100		1.0	1.0	1.0	1.0	5.0	5.0	5.1	5.1	10.0	10.0	10.1	10.2
	2		0.9	0.8	0.8	0.9	5.1	4.8	4.5	4.0	10.5	10.4	10.1	9.2
	3		1.0	1.0	1.0	0.9	5.2	5.5	5.7	6.1	10.5	11.3	12.1	14.4
5	5		1.1	1.1	1.2	1.4	5.2	5.6	6.1	7.7	10.4	11.1	12.2	15.4
	10		1.0	1.1	1.2	1.5	5.1	5.3	5.6	6.9	10.1	10.4	11.1	13.1
	25		1.0	1.0	1.1	1.3	5.0	5.2	5.3	6.0	10.1	10.3	10.6	11.5
	50		1.0	1.0	1.0	1.1	5.0	5.1	5.1	5.4	10.0	10.1	10.2	10.7
	100		1.0	1.0	1.0	1.1	5.0	5.0	5.1	5.2	10.0	10.1	10.1	10.4
	2		0.8	0.5	0.3	0.1	5.1	4.0	2.9	1.5	10.8	10.1	8.5	5.4
	3		1.1	0.9	0.6	0.3	5.6	5.9	5.7	4.9	11.3	12.7	13.8	16.2
10	5		1.1	1.2	1.3	1.4	5.4	6.3	7.4	11.0	10.7	12.4	14.8	22.7
	10		1.1	1.2	1.4	2.2	5.2	5.8	6.8	10.3	10.4	11.4	13.1	19.0
	25		1.0	1.1	1.2	1.6	5.0	5.2	5.6	7.1	10.0	10.4	11.0	13.4
	50		1.0	1.0	1.1	1.3	5.0	5.1	5.4	6.1	10.0	10.2	10.6	11.8
	100		1.0	1.0	1.1	1.2	5.0	5.2	5.3	5.8	10.1	10.2	10.5	11.3

5. CONCLUSION

In conclusion, multiple imputation is a very promising new tool for helping to handle nonresponse in surveys. Although much work remains to be done before it will become a commonplace method, many interesting theoretical and practical results suggest effort expended in its development will be well rewarded by important contributions to applied work.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation Grant Number SES-8311428. I wish to thank M.P. Singh and a referee for their helpful editorial comments on the earlier draft of this article.

REFERENCES

HEITJAN, D.F., and RUBIN, D.B. (1986). Inference for coarse data using multiple imputation. *Proceedings of the 18th Symposium on the Interface of Computer Science and Statistics.*

HERZOG, T.N., and RUBIN, D.B. (1983). Using multiple imputations to handle nonresponse in sample surveys. In *Incomplete Data in Sample Surveys , Volume 2: Theory and Bibliography*, New York: Academic Press, 209-245.

LI, K.H. (1985). *Hypothesis Testing in Multiple Imputation - with Emphasis on Mixed-up Frequencies in Contingency Tables.* Ph.D. Thesis, Department of Statistics, University of Chicago.

RUBIN, D.B. (1976). Inference and missing data. *Biometrika*, 63, 581-592.

RUBIN, D.B. (1977). The design of a general and flexible system for handling nonresponse in sample surveys. Unpublished paper prepared for the U.S. Social Security Administration.

RUBIN, D.B. (1978). Multiple imputations in sample surveys - a phenomenological Bayesian approach to nonresponse. *Proceedings of the Survey Research Methods Section of the American Statistical Association*, 20-34. Also in *Imputation and Editing of Faulty or Missing Survey Data*, U.S. Dept. of Commerce, 1-23.

RUBIN, D.B. (1979). Illustrating the use of multiple imputations to handle nonresponse in sample surveys. *Proceedings of the 1979 Meetings of the ISI-IASS, Manila.*

RUBIN, D.B. (1980). *Handling Nonresponse in Sample Surveys by Multiple Imputations.* U.S. Dept. of Commerce, Bureau of the Census Monograph.

RUBIN, D.B. (1981). The Bayesian bootstrap. *Annals of Statistics*, 9, 130-134.

RUBIN, D.B. (1986a). *Multiple Imputation for Nonresponse in Surveys.* New York: John Wiley & Sons.

RUBIN, D.B. (1986b). Statistical matching using file concatenation with adjusted weights and multiple imputations. *Journal of Business and Economic Statistics*, 87-94.

RUBIN, D.B., and SCHENKER, N. (1986). Multiple imputation for interval estimation from simple random samples with ignorable nonresponse. *Journal of the American Statistical Association*, 81, 366-374.

SCHENKER, N. (1985). *Multiple Imputation for Interval Estimation from Surveys with Ignorable Nonresponse.* Ph.D Thesis, Department of Statistics, University of Chicago.