RESPONSE PROBABILITY WEIGHT ADJUSTMENTS USING LOGISTIC REGRESSION

VG Iannacchione, JG Milne, and RE Folsom
Research Triangle Institute, PO Box 12194, RTP, NC 27709

KEY WORDS: Survey nonresponse; Sampling weights; Raking; Receiver operating characteristic curve

1. Introduction

Most adjustments for survey nonresponse are made by adjusting the sampling weights of respondents in a way that accounts for the nonrespondents. For example, weighting-class adjustments are based on the assumption that sample members can be partitioned into cells, or weighting classes, within which the responses of nonrespondents, had they been obtained, would be similar to those of respondents. Within each weighting class, the inverse of the weighted response rate is applied to the sampling weights of respondents so that the adjusted weight sums of respondents reproduce the unadjusted weight sums over respondents and nonrespondents.

The variables used to define weighting classes must be known for all sample members, i.e., nonrespondents as well as respondents. If the variables are correlated with the zero-one response indicator and with attributes of the planned survey estimates, the use of adjusted weights will substantially reduce the potential for nonresponse bias. However, this reduction in bias can be overshadowed by an increase in the variance if the number of respondents in a weighting class is small (Hansen, Horwitz, & Madow 1953). As a result, weighting classes may have to be combined to provide adequate numbers of respondents per cell.

If preserving marginal totals is of greater importance than cross-classifications, the "raking" procedure for nonresponse (Chapman 1976) allows the use of a large number of variables to define weighting classes simultaneously, without considering the number of respondents in each cross-classification. Instead, only the numbers of respondents in the marginal categories of each variable are of concern. In raking, the iterative-proportional-fitting algorithm is used to adjust the weights of survey respondents to the marginal totals of each of the variables.

Alternatively, nonresponse adjustments may be considered in the context of multiple regression where the zero-one response indicator is regressed on a set of independent variables which are available for both respondents and nonrespondents. The predicted value obtained from the regression equation is the estimated response probability for population members with the same values of the independent variables. As in the weighting-class adjustment, the inverse of this estimated response rate is used to adjust the sampling weight of respondents.

When the predicted response probabilities of a regression model are used for nonresponse adjustment, logistic regression models are preferred to linear regression models. Logistic regression has been shown to provide more accurate probability estimates than linear discriminant analysis when the assumptions of the latter (i.e., multivariate normality of predictor variables with common covariance matrix) are violated (Press & Wilson 1975). In addition, with the logistic model, the predicted probabilities will necessarily range between zero and one.

In this paper, the design-weighted logistic regression algorithm is modified so that, like the raking procedure, the adjusted weight sums for specified analysis domains restricted to respondents reproduce the corresponding unadjusted weight sums over respondents and nonrespondents. Unlike raking however, the logistic adjustment algorithm extends this property to achieve equalization of respondent and full sample weighted means for continuous predictors. An application of this procedure to adjust the weights of spouse participants in the Army Family Research Program is described.

2. Response Probability Weight Adjustments

Logistic regression analysis consists of fitting a linear logistic model to an observed proportion or rate to measure the relationship between the outcome variable and one or more predictor variables. Procedures have been developed at RTI (Shah 1989) for the specific problem of fitting logistic regression models to survey data such that the model parameters estimates and their variance-covariance matrix take the survey design into account. LaVange (1986) presents an application of these methods to predict high cost users of medical care based on data from the National Medical Care Utilization and Expenditure Survey.

Folsom (1991) has modified the design-weighted logistic regression algorithm to derive a response probability adjustment procedure for the sampling weights of survey respondents. The procedure is described as follows. Let

$$Y_i \equiv \begin{cases} 1 & \text{if sample member i responds, and} \\ 0 & \text{otherwise.} \end{cases}$$

Assume the following logistic model for the probability that sample member i responds:

$$\Pr[Y_i=1 \mid \underset{\sim}{X}_i, \underset{\sim}{\beta}] = [1 + \exp(-\underset{\sim}{X}_i\underset{\sim}{\beta})]^{-1} \qquad (1)$$

where

$\underset{\sim}{\beta} \equiv$ A vector of logistic regression coefficients.

and

$\underset{\sim}{X}_i \equiv (1, X_{1i}, \ldots, X_{pi})$, a p+1 element vector with a one followed by p predictor variables.

The logistic regression coefficients are estimated iteratively by solving the following estimation equations using the Newton-Raphson Method:

$$\sum_i (W_i \div \hat{\gamma}_i)\underset{\sim}{X}_i^T Y_i = \sum_i (W_i \div \hat{\gamma}_i)\underset{\sim}{X}_i^T \hat{\gamma}_i, \qquad (2)$$

where

W_i = The sampling weight assigned to sample member i,

and

$$\hat{\gamma}_i = [1 + \exp(-\underset{\sim}{X}_i\hat{\underset{\sim}{\beta}})]^{-1} .$$

Then, the response probability adjusted weight is

$$W_i^* = W_i Y_i \div \hat{\gamma}_i .$$

That is, each unadjusted weight is divided by the estimated probability of response. The adjusted weight of sample members who did not participate is zero. Notice that for any zero-one predictor X_{ik}, the estimation equations in (2) require that

$$\sum_i w_i^* \cdot x_{ik} = \sum_i w_i \cdot x_{ik} \quad . \qquad (3)$$

Because the first element of $\underset{\sim}{X}_i$ is uniformly one, the constraint equations in (3) force the adjusted weight sums for responding sample members (i.e. $Y_i = 1$) to equal the corresponding unadjusted weight sum across all sample members. In addition, the equality of the weight sums holds for any sample subset identified by an zero-one indicator in $\underset{\sim}{X}_i$. Finally, it is clear from equation (3) that the adjusted and unadjusted weighted total and mean for any continuous X_{ik} are equalized.

3. Implementation of the Weight Adjustment Algorithm

Although the logistic weight adjustment algorithm can be applied to any set of categorical or continuous variables that are known for both respondents and nonrespondents, its appeal is its connection to response probability modeling. That is, if the response probabilities of sample members can be assumed to be nonzero over conceptual repetitions of the survey and the predicted probabilities are asymptotically unbiased, then the logistic weight adjustments eliminate nonresponse bias. In addition, response probability modeling provides a formal statistical setting for the evaluation of variables believed to be related to response. This evaluation is particularly useful when there is a large number of potential response predictors.

Once the final logistic regression model is determined, there are two ways to implement the weight adjustment algorithm. The first method uses conventional design-weighted logistic regression software, while the second requires specially written software. Each are described below.

The solution equations for the weight adjustment algorithm in (2) differ from the standard design-weighted logistic regression in that the adjusted weights are substituted for the unadjusted weights. Thus, existing logistic regression software can be run iteratively by using the following updated set of weights at the (k+1)th run:

$$w_i^{k+1} = w_i \div \hat{\gamma}_i^k \qquad (4)$$

where $\hat{\gamma}_i^0$ = the predicted response probability for sample member i using the unadjusted weight.

Convergence is obtained when the logistic regression coefficients do not change measurably between runs.

Alternatively, a logistic mean matching algorithm (Folsom 1991) may be used to adjust the weights for nonresponse. The algorithm is specified as follows. Let,

$$\hat{N}_n \equiv \sum_i (1-Y_i) w_i$$

(i.e., the sum of the nonrespondent weights)

and

$$\bar{\underset{\sim}{X}}_n \equiv \sum_i (1-Y_i) w_i \underset{\sim}{X}_i \div \hat{N}_n$$

(i.e., the nonrespondent means)

where $\underset{\sim}{X}_i$ is now a p-element vector without a one as its first element. Now initialize the vector of logistic regression coefficients and the set of updated weights as follows:

$\hat{\underset{\sim}{\beta}}^{(0)} \equiv$ a p-element vector of zeros, and

$w_i^{(0)} \equiv w_i$

At the (k+1)th iteration, these values are:

$$\hat{\underset{\sim}{\beta}}^{(k+1)} = \hat{\underset{\sim}{\beta}}^{(k)} + \left[\underset{\sim}{S}^{(k)}\right]^{-1} \underset{\sim}{T}^{(k)} \qquad (5)$$

and

$$w_i^{(k+1)} = \exp\left[-\underset{\sim}{X}_i \hat{\underset{\sim}{\beta}}^{(k+1)}\right] w_i Y_i \qquad (6)$$

where

$$\underset{\sim}{S}^{(k)} = \sum_i w_i^{(k)} (\underset{\sim}{X}_i - \bar{\underset{\sim}{X}}_n)^T (\underset{\sim}{X}_i - \bar{\underset{\sim}{X}}_n)$$

and

$$\underset{\sim}{T}^{(k)} = \sum_i w_i^{(k)} (\underset{\sim}{X}_i - \bar{\underset{\sim}{X}}_n) .$$

If convergence is obtained after k^* iterations the final adjusted weight is

$$w_i^* = w_i Y_i + \frac{\hat{N}_n}{\sum_i w_i^{(k^*)}} w_i^{(k^*)} \qquad (7)$$

Notice that the sum of the adjusted weights for respondents equals the sum of the unadjusted weights for all sample members.

While conventional logistic regression software is readily available, it must be used repeatedly to obtain the updated sets of weights in (4). Thus, it may require a substantial amount of computer resources. The logistic mean matching algorithm, however, updates the weights at each iteration (equation (6)) instead of between runs. As a result, it requires significantly fewer iterations (and matrix inversions) than conventional logistic software.

Both methods were used to adjust the sampling weights of the 3,277 participants to the AFRP Spouse Survey described in the next section. The SAS Logistic procedure (SAS 1990) was run seven times with each run taking either five or six iterations to converge. As a result, almost 40 matrix inversions were required to adjust the weights. In contrast, the logistic mean matching algorithm required only five iterations to achieve the same level of convergence.

4. Adjusting for Nonresponse to the AFRP Spouse Survey

The Army Family Research Program (AFRP) was designed to examine the role of family factors in retention, readiness, and sense of community among Army personnel. Toward this end, data were collected from a large, cross-sectional, probability sample of active-duty soldiers and their spouses. In the Soldier Survey, soldiers were asked about their work, their community, and the preparedness of their unit to perform its mission. In addition, soldiers were asked about their family characteristics and their perceptions of Army and civilian life alternatives.

Each married soldier was asked to provide his/her spouse's mailing address for use in a mail survey. In the Spouse Survey, spouses of participating soldiers were mailed a questionnaire asking about their opinions of Army life, their opportunities for work, and their relocations. Spouses were also asked about their finances, their family and friends, and the chances of their spouse staying in the Army.

4.1 Soldier Survey

The Soldier Survey was administered during 1989 in group sessions at Army installations by survey teams working with designated Army liason officers. Whenever possible, all sampled soldiers from a unit were scheduled for the same time slot. However, some soldiers, and in some cases, whole units were unable to attend their scheduled group session because of special assignment, field training exercises, or because of the nature of their mission (e.g., military police or health services personnel). The questionnaire packets of units or individual soldiers who were unable to attend the group administration sessions were routed to the appropriate unit liason officer who was briefed on the distribution, administration, and confidentiality procedures of the questionnaires.

Soldiers were informed that their participation was voluntary and that the data they provided would be kept confidential and used for research purposes only. The Soldier Questionnaire took an average of two hours to complete. (Married soldiers were asked to provide at least 400 distinct responses.) Spouse addresses were requested of married soldiers on the last page of the Soldier Questionnaire. However, the completion of the address form could not be verified in the field because the confidentiality procedures developed for the survey required that the soldier instrument be sealed by the soldier before being returned.

4.2 Spouse Survey

Unlike soldiers in the sample, spouses could not be tasked to attend survey administrative sessions. Instead, a self-administered, mail-out/mail-back questionnaire was developed for spouses of sample soldiers. In addition, a Korean version of the Spouse Questionnaire was developed to encourage the participation of Korean-speaking spouses of soldiers stationed in Korea. An introductory letter from the commanding general of the U.S. Army Community and Family Support Center accompanied each questionnaire. A postage-paid return envelope was also included in the packet.

Each participating soldier was asked to provide his/her spouse's mailing address. If the married soldier did not provide an address, no attempt was made to obtain it from another source (e.g., the soldier's unit). Spouse addresses were only obtained through soldiers to be consistent with the voluntary participation and informed consent policy conveyed to participating soldiers.

Originally, plans were made to use postcard reminders and an intensive telephone followup to increase the response rate to the spouse survey. However, budget constraints precluded both activities. Instead, three additional mailings of the letters and the questionnaire were made to nonrespondents. Questionnaires that were returned by the Postal Service as undeliverable were re-mailed if a forwarding address was provided.

4.3 Development of the Spouse Response Model

The requirement that a spouse address be obtained directly from the soldier made soldier participation a prerequisite for spouse participation. That is, if the soldier was selected for the Soldier Survey but either did not participate or did not provide his/her spouse's address, the spouse was not sent a questionnaire and was considered nonresponding. This pre-condition for spouse response motivated the development of two logistic regression models: the first model was used to predict the probability of obtaining a spouse address from the soldier; the second model was used to predict the probability of spouse response conditional on obtaining a mailing address. Taken together, these models identified the set of known factors most related to spouse participation.

The models were specified by assigning the following zero-one indicators to each soldier of the 7,792 married soldiers who participated in the soldier survey:

$$A_i \equiv \begin{cases} 1 & \text{if soldier i provided a spouse address,} \\ 0 & \text{otherwise,} \end{cases}$$

and

$$R_i \equiv \begin{cases} 1 & \text{if the spouse of soldier i responded,} \\ 0 & \text{otherwise.} \end{cases}$$

Then, the overall probability of spouse response was written as:

$$\begin{aligned} P[R_i=1] &= P[A_i=1] \cdot P[R_i=1 \mid A_i=1] \\ &= \alpha_i \cdot \rho_i \, . \end{aligned} \tag{8}$$

The underlying assumption of the response probability models was that the probability of a spouse's participation could be predicted by knowing the demographics, attitudes, and data collection environment of each participating soldier. To test this assumption, a set of potential predictor variables was constructed from the Army personnel files, the Soldier Questionnaire, and the survey's control system. Because the significance of most of the predictors was expected to vary by rank, all predictor variables were interacted with a six-level categorical variable that identified three enlisted rank groups: Enlisted, NCO, and Senior NCO, and three officer rank groups: Warrant, Company-Grade, and Field-Grade.

Some of the predictor variables derived from the Soldier Questionnaire were relatively straightforward and could be used as zero-one indicators (e.g., Does your spouse work?). However, most of the information required multiple indicators to assure maximum sensitively of measurement (e.g., family adjustment to Army life or Army-civilian job comparisons). When two or more questionnaire items were designated as components of a scale, the relationship among the items were analysed using factor analysis procedures. If the questionnaire items loaded within the same factor, a scaled variable was created by adding the items together. Those scales believed to be related to response were used extensively as predictor variables in both the spouse address and spouse response models.

Each response model was parsed by eliminating any predictor variables that were not significant at the 0.05 level for at least one paygrade group. The intercept was retained so that the sum of the respondents adjusted weights would equal the unadjusted weight sum across respondents and nonrespondents.

Item nonresponse among the predictor variables in the final models caused about 17 percent of the observations to be deleted from the models. (Most of these observations had just one missing predictor variable.) Because the response probability adjustment factors require non-missing predictor values, a weighted sequential hot deck imputation procedure (Cox, 1981) was used to impute missing values. The significance levels of the final models were basically unaffected by the addition of observations with imputed predictors.

4.4 Model Evaluation

Generalized Wald statistics, adjusted for design effects (Rao and Scott 1981), were used to test the goodness of fit of each model and were found to be highly significant (i.e., at least one regression parameter not zero) at the 0.001 level of significance. However, the overall predicted probability of a spouse's participation (i.e., the predicted value produced by the spouse address model multiplied by the predicted value produced by the spouse response model) was not amenable to conventional regression analysis because of the lack of independence between the models. Instead, a Receiver Operating Characteristic (ROC) curve was constructed to assess the overall predictive ability of the combined model.

ROC curves are used to judge the discrimination ability of statistical methods that combine various clues, test results, etc. into a prediction. For example, in a signal detection experiment using the two-alternative forced choice technique, subjects are asked to use available evidence to decide which of two stimuli is "noise" and which is "signal plus noise." For the spouse participation models, the predicted response probability provides the evidence for detecting response which acts as the signal.

A point on a ROC curve is constructed by considering a given predicted probability as a cutoff point for deciding whether a spouse is a respondent or a nonrespondent. For a given cutoff, such as the one shown in Exhibit 1, a point on the ROC curve is obtained by plotting the proportion of respondents with a predicted probability greater than the cutoff (i.e., the proportion of true positives) versus the proportion of nonrespondents with a predicted probability greater than the cutoff (i.e., the proportion of false positives). The points on a ROC curve are obtained by computing the proportion of true and false positives for the entire range of possible cutoff points: from always predicting response (i.e., cutoff less than lowest predicted response probability), to never predicting response (i.e., cutoff greater than highest predicted response probability).

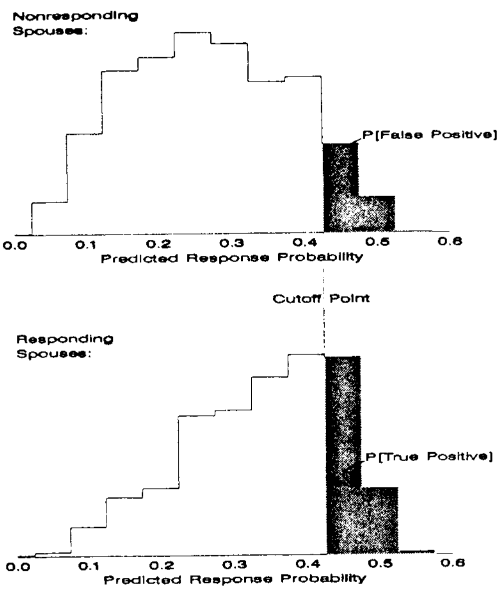

Exhibit 1. Illustration of a cutoff point for deciding the response status of spouses of junior enlisted personnel

The area under a ROC curve measures the probability that a randomly chosen pair of observations, one respondent and one nonrespondent, will be correctly ranked (Hanley and McNeil 1982). This probability of a correct pairwise ranking is the same quantity that is estimated by the nonparametric Wilcoxon statistic which is usually computed to test whether the levels of a quantitative variable in one population tend to be greater than in a second population. No assumptions about how the variable is distributed in the populations are required for the test.

The null hypothesis associated with the Wilcoxon test is that the variable is not a useful discriminator between the populations. For the spouse response model, this corresponds to a null hypothesis that the predicted response probability of a respondent is just as likely to be smaller than the predicted response probability of a nonrespondent as it is to be greater. Thus, if the null hypothesis is true, the ROC curve will be a diagonal line that reflects the equally likely chance of making a correct or incorrect decision and the area under the curve will be 0.5. If the null hypothesis is not true, the ROC curve will rise above the diagonal and the area under the curve will be significantly greater than 0.5.

Six ROC curves were developed for the predicted probability of spouse participation, one for each rank group. The ROC curve for spouses of junior enlisted personnel is shown in Exhibit 2. The areas under each curve were approximated using Simpson's rule. The levels of significance associated with the Wilcoxon tests for each paygrade group were found to be highly significant ($p < 0.0001$). However, the curves indicate that the predicted probabilities discriminate most effectively for spouses of junior enlisted persons and junior NCOs and least effectively for spouses of field-grade officers.

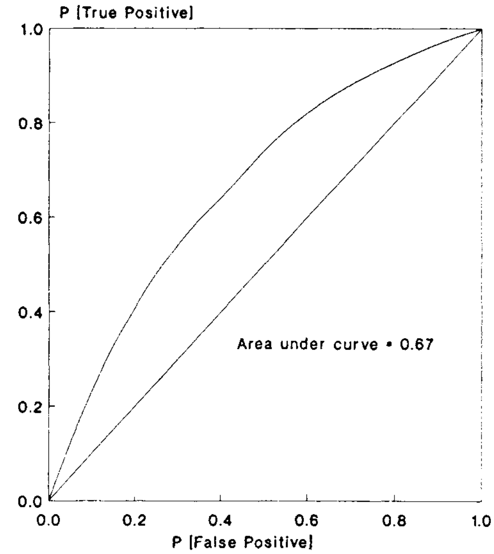

Exhibit 2. ROC curve of the predicted response probabilities for the spouses of junior enlisted personnel

4.5 Factors Affecting a Soldier's Propensity to Provide a Spouse Address

The single most significant predictor of a soldier's propensity to provide his/her spouse's address was the completeness of the last section of the Soldier Questionnaire. The positive regression coefficients associated with this predictor within each paygrade group indicate that, the more complete the section, the more likely the soldier was to provide an address. Because the spouse address form followed the last section, a plausible explanation for not obtaining a spouse address was that the soldier developed "respondent fatigue" and stopped answering questions before he/she reached the address form.

Among the demographic predictors, junior and mid-grade personnel were less likely to provide an address if they were stationed in Europe than elsewhere, and, with field-grade officers, were more likely to provide an address if they were living with their spouse. All other paygrades were unaffected by region of the world and living arrangement.

Among the attitudinal predictors, young enlisted persons and young officers were less likely to provide an address if they felt they could get a better job as a civilian than in the military. Except for senior NCOs, soldiers were less likely to provide an address if they indicated that their family had frequent disagreements and/or difficulty adjusting to Army life.

4.6 Factors Affecting a Spouse's Propensity to Respond

Although the time between soldier interview and the first mailing of a Spouse Questionnaire ranged from less than a month to almost six months, it was not a significant predictor of spouse response. Instead, a combination of demographic and attitudinal variables comprised the final spouse response model.

Among the demographic predictors, living arrangement was found to be an important predictor among the spouses of older personnel in that spouses of senior NCOs and field-grade officers were less likely to respond if they lived apart from the soldier. Male spouses of enlisted personnel and warrant officers were less likely to respond than female spouses of personnel in these ranks. Finally, the spouses of black soldiers other than warrant officers were less likely to respond than the spouses of non-black soldiers.

As with soldier response, family adjustment to Army life was a key attitudinal predictor of spouse response. The spouses of junior officers and enlisted persons who indicated that their families had difficulty adjusting and/or that there was a risk of marital separation were less likely to respond than other spouses of soldiers in these paygrades. In addition, the spouses of senior NCOs who indicated that they were not in control of their situation were less likely to respond than other spouses of senior NCOs.

4.7 Spouse Weight Adjustments

Sampling weights were computed to reflect the three-stage, hierarchical sample design used to select the AFRP sample (Iannacchione & Milne 1991). Initial sampling weights were assigned to each sampling unit as the inverse of its selection probability. Adjustment factors were applied to the sampling weights of participating soldiers compensate for survey ineligibility and nonresponse.

The logistic coefficients of the final models were used to compute a predicted response probability for each participating spouse. These predicted probabilities then were used to adjust the weights of the 3,277 spouses who participated in the Spouse Survey. Continuing the notation in Section 4.3, let:

$W_i \equiv$ The unadjusted weight assigned to spouse i in S,

$\hat{\alpha}_i \equiv$ The predicted probability of obtaining an address for spouse i

$\hat{\rho}_i \equiv$ The predicted probability that spouse i responds given that an address is obtained.

Then, the response probability adjusted weight is

$$W_i^* = W_i R_i \div (\hat{\alpha}_i \hat{\rho}_i) \tag{9}$$

That is, each unadjusted weight is divided by the overall predicted probability of response. The adjusted weight of spouses who did not participate is zero. The mean adjustment factors applied to the spouse weights are shown by rank group in Exhibit 3.

5. Discussion

A noticeable distinction between response probability weight adjustments and weighting class adjustments is that response probability adjustment factors are applied to individual sample members rather than a group or class of sample members. In the extreme case, when each respondent possesses a unique vector of response predictors, it is possible for every respondent to have a different adjusted weight. The effect of this unequal weighting is beneficial from the standpoint of nonresponse error reduction. However, the gain in precision may be offset if the variation among the weights is excessive. Because of this, response probability adjusted weights, like all weights, should be examined for excessive variation (Potter 1990).

The variance inflation attributable to unequal weighting was examined for the AFRP Spouse Survey by calculating unequal weighting effects for the unadjusted and adjusted spouse weights. The unequal weighting effects for each rank group are shown in Exhibit 3 along with the mean of the adjustment factors applied to the unadjusted spouse weights. In general, the effect of unequal weighting increases as the magnitude of the adjustment factor increases. However, the increase in unequal weighting attributable to the response probability adjustments is not excessive and probably does not warrant further investigation.

6. References

Chapman, D.W. (1976), "A Survey of Nonresponse Imputation Procedures," Proceedings of the American Statistical Association, Section on Survey Research Methods.

Cox, B.G. (1980), "Weighted Sequential Hot Deck Imputation," Proceedings of the American Statistical Association, Section on Survey Research Methods.

Folsom, R.E. (1991), "Exponential and Logistic Weight Adjustments for Sampling and Nonresponse Error Reduction," Proceedings of the American Statistical Association, Social Statistics Section.

Hansen, M.H., Hurwitz, W.N., Madow, W.G. (1953), Sample Survey Methods and Theory, Vol. II. John Wiley & Sons, New York.

Hanley, J.A. and McNeil, B.J. (1982), "The Meaning and Use of the Area under a Receiver Operating Characteristic (ROC) Curve," Diagnostic Radiology, 143, 29-36.

Iannacchione, V.G. and Milne J.G. (1991), "Sampling Weights for the AFRP Core Research Effort," Research Triangle Institute, Research Triangle Park, NC.

LaVange, L.M., Iannacchione, V.G., Garfinkel, S.A. (1986), "An Application of Logistic Regression Methods to Survey Data: Predicting High Cost Users of Medical Care," Proceedings of the American Statistical Association, Section on Survey Research Methods.

Potter, F. J. (1990), "A Study of Procedures to Identify and Trim Extreme Sampling Weights," Proceedings of the American Statistical Association, Section on Survey Research Methods.

Press, S.J., and Wilson, S. (1978), "Choosing Between Logistic Regression and Discriminant Analysis," Journal of the American Statistical Association, 73, 699-705.

Rao, J.N.K, and Scott, A.J. (1984), "The Analysis of Categorical Data from Complex Sample Surveys: Chi-Squared Tests for Goodness of Fit and Independence in Two-Way Tables," Journal of the American Statistical Association, 76, 221-230.

SAS Institute Inc. (1990), SAS/STAT User's Guide, Version 6, Fourth Ed., Vol. 2, SAS Institute Inc., Cary, NC, 1071-1126.

Shah, B.V. et al. (1989), "SUDAAN: Procedures for Descriptive Statistics User's Guide," Research Triangle Institute, Research Triangle Park, NC.

Exhibit 3. AFRP Spouse Survey Weighting Summary

Soldier's Rank	Eligible Spouses in Population[1]	Spouse Sample			Mean Adj Factor[2]	Unequal Wgt Effect[3]	
		Sample Size	Address Provided	Ques Provided		Unadjusted Weights	Adjusted Weights
Spouses of Enlisted Persons							
Enlisted	83,113	2,690	2,113	826	3.44	1.19	1.51
NCO	109,998	1,750	1,315	613	2.79	1.18	1.34
Senior NCO	38,970	524	397	226	2.37	1.17	1.26
	232,081	4,964	3,825	1,665	3.06	1.37	1.47
Spouses of Officers							
Warrant	8,202	170	124	87	1.96	1.12	1.43
Company grade	19,799	1,345	1,111	752	1.82	1.25	1.38
Field grade	17,959	1,313	1,063	773	1.69	1.16	1.19
	45,960	2,828	2,298	1,612	1.76	1.43	1.77
All Spouses	278,041	7,792	6,123	3,277	2.42	1.72	2.11

1 The sum of the adjusted weights of responding spouses equals the number of eligible spouses in the population.

2 Mean adjustment factor applied to the unadjusted weights of responding spouses.

3 The unequal weighting effect measures the increase in sampling variance attributable to unequal weighting.