6. SAMPLE WEIGHTING

After the data collection and editing phases of the 2001 National Survey of Veterans (NSV 2001) were completed, we constructed the sampling weights for the data collected from the sampled veterans so that the responses could be properly expanded to represent the entire veteran population. The weights were the result of calculations involving several factors, including original selection probabilities, adjustment for nonresponse, households with multiple residential telephones, and benchmarking to veteran population counts from external sources. We produced a separate set of weights for the List and the RDD Samples and then combined them to produce the composite weights for use with the combined List and RDD Samples.

Our objectives in carrying out the sample weighting in the NSV 2001 were to:

- Enable the production of tabulations that provide estimates of the number of veterans in the population for the various categories selected;
- Compensate for disproportionate sampling of various subgroups in the List Sample;
- Compensate for the higher chance of selection of households with multiple residential telephone lines;
- Reduce biases arising from the fact that nonrespondents may be different from those who participated;
- Compensate, to the extent possible, for noncoverage in the sample due to inadequacies in the sampling frame or other reasons for noncoverage, such as veterans in households without telephones; and
- Reduce variances of the estimates by using auxiliary variables that are highly correlated with the study variables.

We also constructed a set of replicate weights for each respondent veteran and appended them to each record for use in estimating variances. This chapter describes the calculation of the full sample composite weights and replicate composite weights. We start with a description of the List and RDD Sample weights because the two sets of weights were constructed independently.

6.1 List Sample Weights

The List Sample weights are used to produce estimates from the List Sample that represent the population of veterans who are on the list frame. As described in Chapter 3, the list frame was constructed from the VHA Healthcare enrollment file and the VBA Compensation and Pension (C&P) file. The steps involved in constructing the List Sample weights are the calculation of a base weight, poststratification adjustment to known list frame population counts, and adjustments to compensate for veterans with unknown eligibility, and for nonresponse. These steps are described in detail below.

Calculation of List Sample Base Weights

The base weight for each veteran is equal to the reciprocal of his/her probability of selection. The probability of selection of a veteran is the sampling rate for the corresponding sampling stratum. If n_h out of N_h veterans are selected from a stratum denoted by h, then the base weight (or design weight) assigned to the veterans sampled from the stratum was obtained as

$$w_{hi} = \frac{N_h}{n_h}; \quad i \in h. \tag{6-1}$$

Properly weighted estimates using the base weights (as given above) would be unbiased if the eligibility status of every sampled veteran could be determined and every eligible sampled veteran agreed to participate in the survey. However, the eligibility status of each and every sampled veteran could not be determined (for example, some sampled veterans could not be located). Moreover, nonresponse is always present in any survey operation, even when participation is not voluntary. Thus, weight adjustment was necessary to minimize the potential biases due to unknown eligibility and nonresponse. In order to improve the reliability of the estimates we also applied a poststratification adjustment. Normally, the poststratification adjustment is applied after applying the nonresponse adjustment, but we carried this out before the nonresponse adjustment because determining the eligibility status of every veteran on the list frame would not have been feasible.

Poststratification Adjustment

Poststratification is a popular estimation procedure in which the base weights are adjusted so that the sums of the adjusted weights are equal to known population totals for certain subgroups of the population. We defined the poststrata to be the cross classification of three age categories (under 50, 50-64, over 64), gender (male, female), and census regions (Northeast, Midwest, South, and West), which resulted in 24 poststrata.

Let N_g denote the number of veterans on the list frame that belong to the poststratum denoted by g $(g = 1,2,\ldots,24)$ as obtained from the list frame, and let $\hat{N}_g$ be the corresponding estimate obtained by using the List Sample base weights. Then the ratio $N_g / \hat{N}_g$ is used as an adjustment to define the poststratified weight $w_{hi}^{(pst)}$ as

$$w_{hi}^{(pst)} = \left(\frac{N_g}{\hat{N}_g} \right) w_{hi}; \ (hi) \in g. \tag{6-2}$$

The superscript (*pst*) denotes that it is a poststratified weight. Because a veteran denoted by (hi) can belong to one and only one of the poststrata, the poststratified weights are uniquely defined. The advantage of poststratified weighting is that the reliability of the survey estimates is improved. The minimum sample size for poststratification cells was set at 30 veterans. For 2 out of the 24 poststrata, the sample sizes were fewer than 30 veterans. The two deficient cells were female veterans in the age group 50-64 in census regions “Northeast” and “Midwest.” Their sample sizes were equal to 16 and 29, respectively. We collapsed these two cells in order to achieve the sample size of more than 30 in the collapsed poststratum. Thus, the poststratified weights were computed using the auxiliary veteran population counts from the list frame for 23 poststrata.

For the sake of simplicity we will denote by w_i the poststratified weight of the i^{th} List Sample veteran. These weights are the input weights for adjustments for unknown eligibility and nonresponse.

Adjustments for Unknown Eligibility and Nonresponse

The List Sample cases can be divided into respondents and nonrespondents. Further, the respondents can be either eligible or ineligible (out of scope) for the survey. The eligibility of the nonrespondent veterans could not always be determined. For example, a sampled veteran who could not be located could have been deceased and hence ineligible for the survey. Or, an eligible veteran might have moved and new contact information (address and telephone number) might not be obtainable. Therefore, the nonrespondents were classified into two categories: (1) eligible nonrespondents and (2) nonrespondents with unknown eligibility. In order to apply the adjustments for unknown eligibility and nonresponse, the List Sample cases were grouped into four response status categories (Figure 6-1):

> Category 1: **Eligible Respondents.** This group consists of all eligible sampled veterans who participated in the survey, namely those who provided usable survey data. The category includes the final result codes CE and CX.
>
> Category 2: **Ineligible or Out of Scope.** This group consists of all sampled veterans who were ineligible or out of scope for the survey, such as veterans who had moved abroad and were therefore ineligible for the survey. The information that was obtained was sufficient to determine that these veterans were indeed ineligible for the survey.
>
> Category 3: **Eligible Nonrespondents.** This group consists of all eligible sampled veterans who did not provide usable survey data. The information that could be obtained was sufficient to ascertain that the veteran was eligible for the survey.
>
> Category 4: **Eligibility Unknown.** This group consists of all sampled veterans whose eligibility could not be determined. For example, sampled veterans who could not be located were placed in this category.

We used the final List Sample extended interview result codes *(MAINRSLT)* and the variable *"MCURSECT"* to assign the sampled veterans to one of the four response categories defined above. The groupings of the extended interview result codes and the *"MCURSECT"* values that define the above response categories are given in Appendix E. A list of final extended interview result codes are included in Appendix H. For incomplete cases, the variable *"MCURSECT"* indicates the point at which the interview broke off and that the interview could not be completed after that. We should note that the eligibility could not be determined when the Military Background module of the extended interview was not completed. Also, we interpreted cases with result code *"IA"* (not a veteran) as "hidden" refusals and assigned them to the category "eligible nonrespondents" irrespective of the *"MCURSECT"* value.

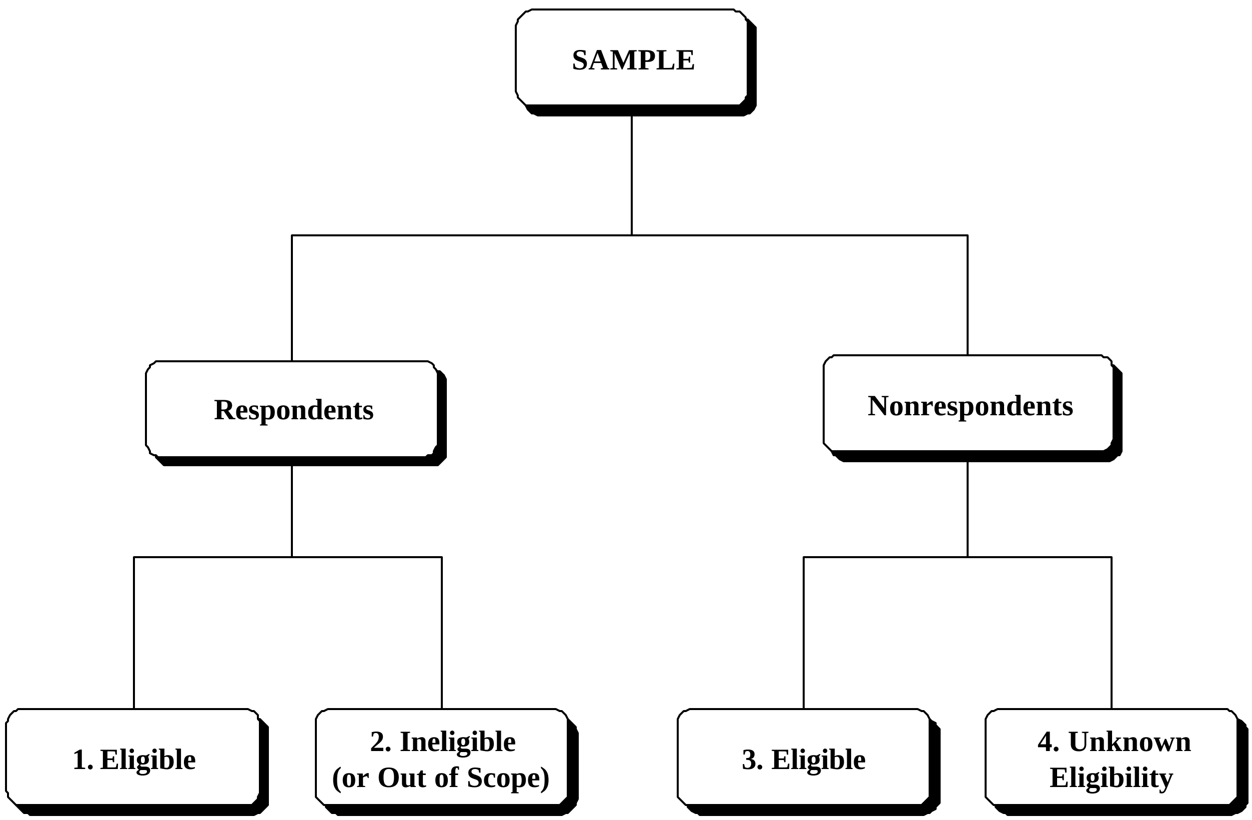

Figure 6-1. Categories of response status

The nonresponse adjustment was applied in two steps. In the first step the poststratified weights of the veterans with unknown eligibility (Category 4) were distributed proportionally over those with known eligibility (Categories 1, 2, and 3). As stated earlier, let w_i denote the poststratified weight of the i^{th} veteran sampled from the list frame. Then the adjustment for unknown eligibility was obtained as

$$A^{(ue)} = \frac{\sum_{i \in R} w_i + \sum_{i \in N} w_i + \sum_{i \in O} w_i + \sum_{i \in U} w_i}{\sum_{i \in R} w_i + \sum_{i \in N} w_i + \sum_{i \in O} w_i}, \tag{6-3}$$

where R represents veterans who were survey respondents (Category 1), O represents out-of-scope or ineligible veterans (Category 2), N represents eligible nonrespondents (Category 3), and U represents the sampled veterans whose eligibility could not be determined (Category 4). The adjustment factor $A^{(ue)}$ to account for unknown eligibility was applied to the poststratified weights of the eligible respondents (Category 1), out-of-scope or ineligible veterans (Category 2), and eligible nonrespondents (Category 3). Thus, the List Sample weight w_i^* adjusted for unknown eligibility was computed as

$$w_i^* = A^{(ue)} w_i \text{ if the } i^{th} \text{ veteran belongs to response Category 1, 2 or 3.} \tag{6-4}$$

The weights of the veterans with unknown eligibility (Category 4) were set to zero.

The adjustment for unknown eligibility was applied within homogeneous adjustment classes. These adjustment classes were determined with *CHAID* (Chi-square Hierarchical Automatic Interaction Detector) software described in Appendix F.

In the second step, we calculated an adjustment factor to account for the eligible nonrespondent veterans. The extended list interview nonresponse adjustment factor was calculated as the ratio of the sum of the weights (adjusted for unknown eligibility) for eligible respondents and eligible nonrespondents to the sum of the weights for only the eligible respondents. Thus, we calculated the nonresponse adjustment factor $A^{(nr)}$ to be the ratio of the sums as

$$A^{(nr)} = \frac{\sum_{i \in R} w_i^* + \sum_{i \in N} w_i^*}{\sum_{i \in R} w_i^*}, \tag{6-5}$$

where w_i^* is the weight obtained after applying the adjustment for unknown eligibility, R represents eligible respondents (Category 1), and N represents eligible nonrespondents (Category 3). The adjustment factor $A^{(nr)}$ is applied only to the weights of the eligible respondents (Category 1) in the sample. That is, the nonresponse-adjusted weight w_i^{**} is computed as

$$w_i^{**} = A^{(nr)} w_i^* \text{ if the } i^{th} \text{ sampled veteran is a respondent (Category 1).} \tag{6-6}$$

We applied the nonresponse adjustment, $A^{(nr)}$, within homogeneous nonresponse adjustment classes, which were also defined using *CHAID* software. The final List Sample weight for each eligible respondent was computed by multiplying the weight w_i^* by the appropriate nonresponse adjustment factor as defined above. The final List Sample weight for the eligible nonrespondent veterans was set to zero. The final List Sample weight of the out-of-scope/ineligible veterans is the weight obtained after applying the adjustment factor for unknown eligibility. The weights for the out-of-scope/ineligible veterans could be used to estimate the ineligibility rate of the list frame that we used to select the List Sample.

6.2 RDD Sample Weights

The calculation of the RDD Sample weights consisted of five main steps. The steps included computing the base weight and various adjustments at the screener interview level and the extended interview level. In summary, we:

- Computed base weight as the inverse of the probability of selection of the telephone number associated with the household;
- Applied an adjustment to account for household level nonresponse during screening;
- Applied an adjustment for multiple telephone lines as the reciprocal of the number of "regular residential" telephone numbers used by the household (excluding telephone numbers used only for business purposes, fax machines, cellular phones, pagers, or mobile phones);
- Applied an adjustment to correct for the nonresponse to the extended interview; and
- Benchmarked to known veteran population counts from the Census 2000 Supplementary Survey (C2SS) that the U.S. Bureau of the Census conducted.

The final RDD Sample weights were obtained as the product of the base weight and the various adjustments applied to the base weights. The steps involved in computing these weights are described in detail below.

RDD Sample Base Weights

As described in Chapter 3, the RDD Sample was selected with the list-assisted RDD sampling methodology except for the Puerto Rico RDD Sample, for which an RDD sample of telephone numbers was selected from all possible telephone numbers for Puerto Rico. The base weights for the two RDD Samples were defined accordingly.

List-assisted RDD Sample Base Weights

The base weight is defined as the reciprocal of the probability of selection. With the list-assisted RDD methodology, the telephone numbers were selected with equal probabilities of selection. We used a systematic sampling scheme to select telephone numbers, and the probability of selecting a telephone number when n telephone numbers from a pool of N numbers is selected is given by $f = n/N$. Because the national RDD Sample was selected from two RDD frames constructed at two different times (see Chapter 3) the selection probabilities were computed according to whether a telephone number was eligible for selection from both frames or from only one of the frames. Let F_1 and F_2 denote the RDD frames constructed at the two time periods, and N_1 and N_2 be their corresponding sizes. A random sample of n_1 (=240,000) telephone numbers was selected from the frame F_1 and a random sample of n_2 (=60,000) telephone numbers was selected from the frame F_2. The selection probabilities of the sampled telephone numbers were computed as follows.

$$\Pr ob(t) = \begin{cases} \dfrac{n_1}{N_1} + \left(1 - \dfrac{n_1}{N_1}\right)\dfrac{n_2}{N_2} & \text{if t was in both } F_1 \text{ and } F_2 \\ \dfrac{n_1}{N_1} & \text{if t was in } F_1 \text{ only} \\ \dfrac{n_2}{N_2} & \text{if t was in } F_2 \text{ only} \end{cases} \qquad (6\text{-}7)$$

where t denotes a sampled telephone number. The base weight of a telephone number selected from the RDD frames is given by the reciprocal of the corresponding probability of selection.

Puerto Rico Sample Base Weights

The Puerto Rico RDD Sample was a pure RDD sample due to the fact that information was not available on the telephones to construct the sampling frame for list-assisted RDD methodology. The base weight was defined to be the inverse of the selection probability.

RDD Sample Weight Adjustments

RDD Sample weight adjustments include weight adjustments for the national (list-assisted) RDD Sample and the Puerto Rico RDD Sample.

List-assisted RDD Sample Weight Adjustments

List-assisted RDD Sample weight adjustments were applied as screener interview nonresponse adjustment, adjustment for multiple telephone lines, and an adjustment for nonresponse at the extended interview.

Screener Nonresponse Adjustment. The base weights were adjusted to account for the households (telephones) with unknown eligibility during the screening interview. We defined the four categories listed below and assigned sampled telephone numbers to each based on the final screener result (*SCRNRSLT*) codes as given in Appendix E.

Category 1: **Eligible Respondents.** This category consists of all sample households that completed the screening questionnaire and contained at least one veteran eligible for the extended interview. The category includes the RDD screener final result codes CO or household selection flag (HSF) equal to 1 (YES).

Category 2: **Ineligible or Out of Scope.** This category consists of all sample households (telephones) that were ineligible or out of scope for the survey. For example, these included households with no veterans or telephone numbers that were business numbers.

Category 3: **Eligible Nonrespondents.** Although we defined an "eligible nonrespondents" category, no cases were assigned to it because once someone in the household responds to the screening questionnaire, the persons enumerated in that household can be categorized as either eligible respondents or ineligible/out of scope. Otherwise, the household is assigned to the category "Eligibility Unknown."

Category 4: **Eligibility Unknown.** This category consists of all sample telephones for which sufficient information could not be collected to determine whether or not there was a veteran in the household.

The assignment of sampled households (telephones) to the three response categories (categories 1, 2 and 4) was based on the final screener result codes (*SCRNRSLT*) and household selection flag (HSF) as given in Appendix E. A list of final screener result codes is included in Appendix H.

The base weights corresponding to the households (telephones) with unknown eligibility (Category 4) were distributed proportionally over those with known eligibility (Categories 1 and 2). To carry out the adjustment for unknown eligibility, the telephones with unknown eligibility were divided into two sub-categories: (1) those that we could determine were residential and (2) those for which we could not make a residential determination.

The adjustment for unknown eligibility was then applied in two separate steps. In the first step, we adjusted for those telephones whose type – residential, business, or nonworking – could not be determined. The weight adjustment was applied within homogeneous adjustment classes that were determined through the *CHAID* analysis.

In the second step, nonworking and business telephone numbers were removed and the weights were adjusted to account for the residential telephone numbers for which the eligibility for the NSV 2001 could not be determined. The adjustment for unknown eligibility in the second step was computed as the ratio of the sum of the weights adjusted in the first step of all residential sample cases (both with known and unknown eligibility) to those with known eligibility. It should be noted that the nonworking and business telephone numbers had been eliminated at this stage. The weights of those with known eligibility were adjusted by multiplying with the adjustment factor for the second step of unknown eligibility, and the weights of those with unknown eligibility were set to zero. The adjustment for the second step of unknown eligibility was also applied within homogeneous adjustment classes defined using the *CHAID* software.

Adjustment for Multiple Residential Lines. If every household had exactly one residential telephone number, then the weight for a household would be the same as the base weight of the corresponding telephone number. The adjustment for multiple residential telephone households prevents households with two or more residential telephone numbers from receiving a weight that is too large by reflecting their increased probability of selection. In theory, the household weight would be obtained by dividing the base weight by the number of residential telephone lines in the household. We assigned an adjustment factor of ½ to the households with more than one residential telephone number because the number of households with more than two residential telephone numbers would be small. A weighting factor of unity was assigned to households reporting only one telephone number in the household, and an adjustment factor of ½ was assigned to households with more than one residential telephone number.

RDD Extended Interview Nonresponse Adjustment. The RDD Sample required administration of both a household screening questionnaire and the extended NSV 2001 questionnaire, and included the possibility of identifying multiple veterans in a single household. Because the screener survey interview screened for the households with potential veterans, a small fraction of persons who were screened in were not actually eligible for the NSV 2001. Once the extended interview began, it was still necessary to establish with certainty that the selected person was indeed a veteran, so further screening took place at the beginning of the extended interview in the Military Background module. If the responses to the set of eligibility questions during the extended interview indicated that the person was not an eligible veteran, the interview was terminated. Moreover, for some cases that were screened in, no information could be collected from the extended interview to ascertain their eligibility (e.g., the potential veteran could not be contacted for the extended interview). Thus, the screened-in sample contained cases with unknown eligibility as well as eligible and ineligible cases. Further, the eligible cases contained respondents and nonrespondents. Therefore, the screened-in RDD Sample cases were grouped into the same four categories as the List Sample cases.

Category 1: **Eligible Respondents.** This group consists of all eligible sample veterans who participated in the survey, namely those who provided usable survey data. The category includes the final result codes CE and CX.

Category 2: **Ineligible or out of scope.** This group consists of all sample cases that were determined to be ineligible or out of scope for the survey, such as a screened-in person who was not a veteran and hence was ineligible for the survey.

Category 3: **Eligible Nonrespondents.** This group consists of all eligible sample veterans who did not provide usable survey data. The information that could be obtained was sufficient to ascertain that the veteran was eligible for the survey.

Category 4: **Eligibility Unknown.** This group consists of all sample cases whose eligibility could not be determined. For example, sample persons who could not be contacted were placed in this category.

The screened-in sample cases were assigned to the four response categories on the basis of final extended interview result codes (*MAINRSLT*) and the variable *"MCURSECT."* The groupings of the extended result codes, along with the *"MCURSECT"* values corresponding to the four response categories, are given in Appendix E. These categories are very similar to those for the List Sample extended interviews. We should note that the extended result code *"IA"* (not a veteran) for the List Sample cases was interpreted as a "hidden" refusal and hence was assigned to the category "eligible nonrespondents." The RDD Sample cases with the result code *"IA"* were assigned to the "ineligible" category because the eligibility status for RDD Sample cases was determined during the extended interview.

The weights of the cases with unknown eligibility (Category 4) were proportionally distributed over the other 3 categories (Categories 1, 2, and 3). These adjustment factors were calculated separately for homogeneous classes defined with *CHAID* analysis.

The next step in the RDD Sample weighting was the extended interview nonresponse adjustment. The RDD extended interview nonresponse adjustment factor was calculated as the ratio of the sum of weights for eligible RDD extended interview respondents and eligible RDD extended interview nonrespondents to the sum of the weights for only the eligible RDD extended interview respondents. Separate nonresponse adjustment factors were computed within homogeneous nonresponse adjustment cells. The nonresponse adjustment cells were determined with the *CHAID* software.

Puerto Rico Sample Weight Adjustments

We screened 96 households with potentially 102 veterans for which extended interviews were attempted. We completed only 51 extended interviews from the Puerto Rico RDD Sample. The nonresponse adjustment factors for the screener interview and extended interview were computed similarly to those for the national RDD Sample except that the screener nonresponse adjustment was computed separately for two age groups (under 60, over 59) and a single nonresponse adjustment was computed for the extended interviews. This was due to the small sample size for the Puerto Rico RDD Sample.

After applying the screener interview and extended interview nonresponse adjustments, the national (list-assisted) RDD and the Puerto Rico RDD Samples were combined into one RDD Sample. The base weights adjusted for nonresponse were further adjusted in a raking procedure, discussed in a later section. The raked weights were the final RDD Sample weights that were used to compute the composite weights for the combined List and RDD Samples.

Comparison of RDD Estimates with VA Population Model Estimates

As a check, we compared the RDD Sample estimate of number of veterans based on the weights before raking with the estimate from the Vetpop 2000 model[1], VA population projection model. The NSV 2001 target population includes only noninstitutionalized veterans living in the U.S. The reference period for the NSV 2001 is the year 2000[2]. The VA population model estimates are also for the year 2000 and these are based on the 1990 Census. These estimates are derived by incorporating survival rates and information on veterans leaving military service. The VA population model estimate for the entire veteran population is 25,372,000 veterans, whereas the estimate from the RDD Sample is 23,924,947 veterans, which is 5.7 percent lower than the VA population model estimate. The difference of 5.7 percent can be attributed to the combination of the differences from exclusion of the institutionalized veterans and RDD undercoverage of nontelephone households and households with unlisted telephone numbers belonging to "zero-listed telephone banks."

The portion of undercoverage due to nontelephone households and households with unlisted numbers belonging to "zero-listed telephone banks" was addressed with the raking procedure, described in the next section. The control total of veteran population for the raking procedure was 25,196,036 veterans. Thus, the estimated undercoverage due to nontelephone households and households with

---

[1] The Vetpop 2000 is a veteran population projection model developed by the office of the Actuary, Department of Veterans Affairs. It is the official VA estimate and projection of the number and characteristics of veterans as of September 30, 2000. Details of all aspects of the development and content of the model are available from the office of the Actuary, Department of Veterans Affairs, 810 Vermont Avenue NW, Washington DC 20420.

[2] The data collection field period for the survey was February through November 2001. Nearly all of the survey items that address use or nonuse of VA Health Care Services use a reference period of " during the past 12 months," and individual and household income questions are for the year 2000.

unlisted telephone numbers belonging to "zero-listed telephone banks" would be only about 5.0 percent. After correcting for the undercoverage from these two sources, the difference between the NSV 2001 and the Vetpop 2000 estimates is less than one percent, which is from institutionalized veterans and veterans living abroad.

Raking Ratio Estimation/Undercoverage Adjustment

The raking ratio estimation procedure is based on an iterative proportional fitting procedure developed by Deming and Stephan (1940), and involves simultaneous ratio adjustments to two or more marginal distributions of the population counts. Raking was proposed by Deming and Stephan (1940) as a way to ensure consistency between complete counts and sample data from the 1940 U.S. Census of population. The methodology is referred to as raking ratio estimation because weights are raked using ratio adjustments based on the known marginal population totals. Typically, raking is used in situations where the interior cell counts of cross-tabulation are either unknown or sample sizes in some cells are too small for efficient estimation. The purpose of the raking procedure in this survey is to improve the reliability of the survey estimates, and to correct for the bias due to missed households, namely, households without telephones and households with unlisted telephone numbers belonging to "zero-listed telephone banks." As described in Chapter 3, households without telephones and households with unlisted telephone numbers belonging to the "zero-listed telephone banks" are not included in the list-assisted RDD sampling frame.

The raking procedure is carried out in a sequence of adjustments. First, the base weights are adjusted to one marginal distribution and then to the second marginal distribution, and so on. One sequence of adjustments to the marginal distributions is known as a cycle or iteration. The procedure is repeated until convergence is achieved. The criteria for convergence can be specified either as maximum number of iterations or absolute difference (or relative absolute difference) from the known marginal population totals.

We used a two-dimensional raking procedure for the RDD Sample. The computational details of the two-dimensional raking procedure are given in Appendix G. We formed the two raking dimensions from the cross classification of veterans according to the demographic/education/region characteristics of the veterans. These characteristics were also obtained during the screening interview. The first dimension was formed from the cross classification of three age categories (under 50, 50-64, over 64) with four education levels (no high school diploma, high school diploma, some college, bachelor's degree or higher) and four race categories (Hispanic, Black, Other, and White), resulting in 48 cells. The second dimension was formed from the cross classification of gender (male, female) and the four census regions (Northeast, Midwest, South, and West), resulting in 8 cells. By using a set of cross classified variables for each raking dimension, the internal correlation structure of the data could be better preserved. The sample sizes for the race categories "Hispanics," "African American," and "Other" in the age group under 50, and education "no high school diploma" were 21, 15, and 17, respectively. These three cells in the first raking dimension were collapsed to achieve sufficient cell sample size. Thus, the number of cells for the first raking dimension was reduced to 46 after collapsing the three cells with deficient sample sizes. The sample sizes were more than 25 for all cells used for the raking.

We used the Census 2000 Supplementary Sample (C2SS) data from the U.S. Bureau of the Census to define the control totals for the raking procedure. We also included the Puerto Rico RDD Sample in the raking procedure. Because the C2SS did not include Puerto Rico in the survey target population, we estimated the Puerto Rico veteran population counts for the year 2000 from the Census 1990 population counts based on a model. The methodology for the veteran population counts to be used as control totals for the raking procedure is discussed briefly in the next section.

We applied the convergence criteria in terms of percent absolute relative difference, which was specified to be no more than 0.01 percent for all marginal population counts. The raking procedure converged in 8 iterations.

The above variables were chosen as the raking variables due to significant differences in the telephone coverage by categories of these variables, and hence maximum bias reduction would be achieved. The sample sizes at the adjustment cell level would become very small if we had used too many variables in the cross classification to define marginal distributions for raking.

Veteran Population Counts for the Raking Procedure

The independent estimates of veteran population counts for the raking procedure were obtained from the Census 2000 Supplementary Survey (C2SS). The C2SS sample does not cover Puerto Rico, and we used the 1990 Census data to obtain model-based estimates of the Puerto Rico veteran population counts for the year 2000. The methodology of the model-based estimates for Puerto Rico is discussed later in this section. For the purpose of the raking procedure Puerto Rico data were combined with the census region "South."

Estimates from the Census 2000 Supplementary Survey (C2SS)

The U.S. Bureau of the Census conducted the Census 2000 Supplementary Survey (C2SS). The survey covers the 50 states and the District of Columbia. The sample for the C2SS used a two stage stratified design with a sample of approximately 890,000 housing units designed to measure socioeconomic and demographic characteristics of housing units and their occupants. The C2SS sample of housing units was selected from the Master Address File (MAF). The MAF was created by combining the 1990 Census Control file, the Delivery Sequence File of the United States Postal Service (USPS), and addresses listed for the Census 2000. The first stage sampling involved dividing the United States into primary sampling units (PSUs) and grouping these PSUs into homogeneous strata. The C2SS design employed 1,925 PSUs. The strata were constructed so that they are as homogeneous as possible with respect to social and economic characteristics that are considered important by C2SS data users. A pair of PSUs was selected from each stratum with probability proportional to size (PPS) sampling. In the second stage of sampling, a sample of housing units within the sampled PSUs was drawn using a systematic sampling procedure.

The data were collected from more than 700,000 housing units. Our assumption was that 1 in 4 households contains a veteran and hence, the estimates of veteran population counts from the C2SS data will be based on approximately 175,000 interviewed veterans. The Census 2000 Supplementary Survey universe is limited to the household population and excludes the population living in institutions, college dormitories, and other group quarters. Because the NSV 2001 also excludes the institutionalized veteran population and veterans living abroad, the estimated veteran population counts from the C2SS could be used to benchmark the NSV 2001 estimates.

Model-based Estimates for Puerto Rico

The C2SS sample does not cover Puerto Rico, and external data for Puerto Rico for the raking variables is not available from any other source. We used the 1990 Census data to obtain the distribution of the Puerto Rico veteran population by the variables used for raking. We made the assumption that the distribution of the Puerto Rico veteran population by the raking variables has not changed between 1990 and 2000. Thus, we could use the total Puerto Rico veteran population in 2000 and the 1990 Census distribution to obtain the veteran population counts for the cells defined for the raking procedure. We used the Puerto Rico total veteran population for 2000 as derived from the veteran population model developed by VA (Vetpop 2000). According to the Veteran Population model, the Puerto Rico veteran population for 2000 was 142,680 veterans. We used these model-based estimates of Puerto Rico veteran population counts for 2000 to adjust the veteran population control totals obtained from C2SS so that the raking procedure could be used with the RDD sample, including the Puerto Rico RDD Sample. The Puerto Rico data were assigned to the census region "South" for raking.

6.3 Composite Weights

Integration of samples from multiple frames into a single micro-data file with a single weight requires, at a minimum, the ability to tell which of the veterans had more than one chance of selection. This is enough to create unbiased weights. The Social Security numbers (SSNs) of all the veterans on the list frame were known. To identify the RDD Sample veterans on the list frame, we needed to obtain their SSNs during data collection so that the overlap RDD Sample would be identified by matching the SSNs of the veterans in the RDD Sample with the list frame. However, out of 12,956 completed extended RDD interviews (including Puerto Rico), we were able to obtain an SSN from only 6,237 veterans, which is 48.1 percent of the RDD completed extended interviews. The veterans sampled as part of the RDD Sample could thus only be categorized as belonging to the overlap RDD Sample or nonoverlap RDD Sample if the SSN was reported. For others (those who did not report their Social Security numbers), we used a prediction model to impute the overlap status. The imputation of the overlap status and the construction of composite weights are discussed in the following sections.

Imputation of Overlap Status of Veterans Not Reporting SSN

We used the following model to predict the probability that a veteran in the RDD Sample for whom an SSN could not be obtained would actually belong to the overlap domain.

$$prob(Overlap) = prob(SSN) \times prob(Overlap \mid SSN) + prob\left(\overline{SSN}\right) \times prob\left(Overlap \mid \overline{SSN}\right), \qquad (6\text{-}8)$$

where $prob(Overlap)$ is equal to the probability that a veteran in the completed RDD Sample belongs to the overlap domain; $prob(SSN)$ is the probability that a veteran in the completed RDD Sample reported the SSN; $prob(Overlap \mid SSN)$ is the conditional probability that a veteran in the completed RDD Sample with a reported SSN belongs to the overlap domain; $prob(\overline{SSN})$ is equal to the probability that a veteran in the completed RDD Sample did not report a SSN, which is given by $\{1 - prob(SSN)\}$; and $prob(Overlap \mid \overline{SSN})$ is the conditional probability that a veteran in the completed RDD Sample with unreported SSN belongs to the overlap domain.

We needed to determine the probability of overlap that was conditional on not reporting an SSN (i.e., $prob(Overlap \mid \overline{SSN})$). This can be computed from the above expression because all other probabilities are known. We used *CHAID* analysis to determine homogeneous classes of overlap for those reporting SSNs in order to impute the overlap status within each class for those not reporting an SSN. We used demographic and socioeconomic variables, such as age, gender, race, education, income, and priority group as predictor variables in the *CHAID* model. The probability of overlap conditional on not reporting an SSN (i.e. $prob(Overlap \mid \overline{SSN})$) was determined independently for each cell, and the overlap status was imputed by taking a random sample of the veterans out of those who did not report an SSN. In other words, the overlap status of the veterans with an unreported SSN within a class was imputed as belonging to the overlap domain such that the proportion belonging to the overlap was as close to the desired probability as possible. The proportion belonging to the overlap domain was based on the weighted counts. Thus, the above approach is an imputation approach that effectively uses auxiliary variables, such as demographic variables and enrollment priority groups, to impute the overlap status of the RDD Sample veterans who did not provide Social Security numbers.

The veterans in the overlap RDD Sample (including the imputed cases) also had a chance of being selected in the List Sample, and hence, had an increased chance of selection. These RDD cases are referred to as the overlap sample because they represent the portion of the RDD frame that overlaps with the list frame. A composite weight was created for the identified overlap RDD Sample (both observed and imputed) and List Sample cases using the principles of composite estimation so that the combined RDD and List Sample file could be used for analysis.

Calculation of Composite Weights

Composite weights were calculated using an approach developed by Hartley (1962). We selected this approach because it could be adapted to take into account the design effects of the RDD and List Sample designs when combining the two samples. The List and RDD Samples were combined into one file, consisting of 12,956 completed extended interviews from the RDD Sample, and 7,092 completed extended interviews from the List Sample, resulting in a combined sample of 20,048 completed extended interviews.

In composite estimation, the estimates being combined are assumed to be independent, and are unbiased estimates of the same population parameter. In other words, the List Sample and the overlap RDD Sample cases theoretically represent the same population (i.e., veterans on the list frame). Therefore, a linear combination of the two independent estimates would also produce an unbiased estimate. The parameter for constructing the composite weights is chosen so that the variance is minimized. The composite weight for each veteran in the RDD Sample and List Sample was calculated as

$$w_{comp} = \begin{cases} l \times w_1 & \text{if veteran is in the List Sample} \\ (1-l) \times w_2 & \text{if veteran is in the overlap RDD Sample} \\ w_2 & \text{if veteran is in the nonoverlap RDD Sample} \end{cases} \qquad (6\text{-}9)$$

where

w_1 = original List Sample weight; and

w_2 = original RDD Sample weight.

The parameter l ($0 < l < 1$) defines the composite weight that is used to produce the composite estimate as a linear combination of the List Sample estimate and the overlap domain RDD Sample estimate. The optimum value of the parameter l for estimating a proportion is given by

$$l = \frac{s_2^2}{s_1^2 + s_2^2}, \qquad (6\text{-}10)$$

where

s_1^2 = variance of a proportion from the List Sample; and

s_2^2 = variance of a proportion from the overlap RDD Sample.

The composite weight gives increased weight to the estimates with smaller variance, namely a smaller value of s^2. Thus, the weight assigned to each of the estimates is inversely proportional to the corresponding variance. In practice, the survey estimates of proportions are produced for several characteristics and each would have its own optimum value of the parameter l. It would not be practical to have a separate set of weights for these characteristics and a common l value is highly desirable for the sake of internal consistency of the estimates. Therefore, the l values corresponding to these estimates were averaged according to the formula

$$l = \frac{\sum_i l_i \left(n^{(RDD)} \big/ deff_i^{(RDD)} + n^{(List)} \big/ deff_i^{(List)} \right)}{\sum_i \left(n^{(RDD)} \big/ deff_i^{(RDD)} + n^{(List)} \big/ deff_i^{(List)} \right)}, \tag{6-11}$$

where

l_i	=	l for the i^{th} estimated proportion;
$deff_i$	=	design effect for the i^{th} estimated proportion;
n	=	number of responding veterans;
RDD	=	overlap RDD Sample; and
$List$	=	List Sample.

In the above formula, the sample size when divided by the design effect represents the effective sample size as compared with simple random sampling because of such design features as clustering and unequal probabilities of selection. Thus, the value of l is obtained by taking the weighted average of the individual l values where the weights are proportional to the corresponding effective sample sizes. The rationale for the above averaging formula was that it gave more weight to the l values that are based on larger effective sample sizes.

The composite weight gives increased weight to the estimate with the smaller variance (or larger effective sample size). There would be some loss of variance efficiency from using a common l value for all of the characteristics instead of optimum l for each of the characteristics. The increase in the variance for a characteristic would depend on the absolute difference between the common (average) l value and the optimum l value for the particular characteristic.

We computed the estimates of proportions and their variances for 16 statistics identified as key variables by the VA for the List Sample and the overlap portion of the RDD Sample. These variables are listed in Table 6-1.

Table 6-1. VA key variables

MB24:	Combat or War Zone Exposure (Yes/No)
DIS1:	Ever Applied for VA Disability Benefits (Yes/No)
HB21:	Currently Covered by Medicare (Yes/No)
HC1:	Emergency Room Care During Last 12 Months (Yes/No)
HC4a:	VA Paid for Emergency Room Care (Yes/No)
HC5:	Outpatient Care During Last 12 Months (Yes/No)
HC6:	VA Facility for Outpatient Care (Yes/No)
HC9:	Hospitalized Overnight in a VA Hospital (Yes/No)
SD14d:	VA Service Connected Disability Compensation in 2000 (Yes/No)
SD14e:	VA Non-Service Connected Pension in 2000 (Yes/No)
SD14j:	Income Source: Public Assistance in 2000 (Yes/No)
ET1:	Ever Received Any VA Education or Training Benefits (Yes/No)
ML3a:	Ever Used VA Loan Program to Purchase Home (Yes/No)
ML3b:	Ever Used VA Loan Program for Home Improvement (Yes/No)
ML3c:	Ever Used VA Loan to Refinance Home (Yes/No)
PRIORITY:	Priority Group (Mandatory/Discretionary)

The weighted average of the individual l's based on the variables in the above table was computed according to the formula given in equation 6-11. The average l value turned out to be 0.7272 and was used to construct the composite weights for the combined sample. The individual l values ranged from 0.56 to 0.88.

Raked Composite Weights

The composite weights obtained by combining the List and RDD Samples were also raked using the same two dimensional raking procedure that was used for the RDD sample raking. The only difference was that we did not need to collapse the cells in the first raking dimension, which was defined by cross classification of age, education, and race/ethnicity. The RDD Sample sizes for three cells in the first raking dimension were not sufficient and these cells had to be collapsed for the raking procedure.

The combined RDD and List Sample sizes were more than 30 for all 48 cells used for the first raking dimension and hence we did not need to collapse cells.

The RDD Sample was raked mainly to correct for undercoverage because of nontelephone households and households with unlisted numbers in the "zero-listed telephone banks" that were missed in the list-assisted RDD sampling methodology. The composite weights were raked to achieve consistency with the C2SS estimates, and to improve the precision of the survey estimates. The improvement in the precision of the survey estimates would depend on the strength of correlation between the study variable and the variables employed in the raking procedure. The raking procedure is most beneficial if the estimation domains are defined on the basis of the raking variables, or if these variables are highly correlated with the study variables. We used the first raking dimension by cross classification of the variables age, education, and race/ethnicity to preserve the correlation structure among these variables. Similarly, the second dimension was defined by cross classification of the variables gender and census region. The variances of the national level estimates of totals of the variables used in the raking procedure would be identically equal to zero, which is an additional benefit of the raking procedure.

6.4 Replicate Weights

A separate set of replicate weights was created for the RDD Sample and the List Sample. These were then combined to construct the preliminary composite replicate weights. The final composite replicate weights were obtained by using the same two dimensional raking procedure with the preliminary composite replicate weights as the input weights that were used for the composite full sample weights.

List Sample Replicate Weights

A set of 51 Jackknife 1 (JK1) replicate weights was created for the List Sample for use in variance estimation. The replicate weights were designed for the JK1 replication method. To create the replicate weights, the entire List Sample, including ineligible and nonresponding veterans, was sorted by the twelve sampling strata, and by the order of selection within strata. The strata were not explicitly used in the assignment of replicates but the gains due to stratification were reflected in sorting the sample cases by strata. Records 1, 1+51, 1+2*51, 1+3*51, and so on were assigned to the first replicate group. Records 2, 2+51, 2+2*51, 2+3*51, and so on were assigned to the second replicate group. The same approach was used with each succeeding replicate group without regard for strata boundaries, until all records were assigned to one of the 51 replicate groups. The replicate base weights for the r^{th} replicate were created by setting to zero the base weights for the records in the r^{th} replicate group and reweighting the base weights in the remaining replicate groups by the factor 51/50.

The same adjustments applied to the full List Sample base weights to obtain the full List Sample final weights were applied to the replicate base weights to obtain the List Sample replicate final weights. This included poststratification and the extended interview nonresponse adjustments that were recalculated for each replicate, so that the sampling variability in the response rates would be captured in the replicate weights. The randomness in the number of sampled ineligible cases is also reflected in the varying number of sampled eligible veterans in each replicate.

RDD Sample Replicate Weights

A set of 51 JK1 replicate weights was also created for the veterans identified from the RDD Sample. JK1 replicates were assigned by first sorting the entire RDD Sample of telephone numbers, both eligible and ineligible, in the order of selection of the 10-digit numbers that determined each original RDD Sample. Records 1, 1+51, 1+2*51, 1+3*51, and so on were assigned to the first replicate group. Records 2, 2+51, 2+2*51, 2+3*51, and so on were assigned to the second replicate group. The same approach was used with each succeeding group, until all records were assigned to one of the 51 replicate groups. The replicate base weights for the r^{th} replicate were created by setting to zero the base weights for the records in the r^{th} replicate group and reweighting the base weights in the remaining replicate groups by the factor 51/50. The replicate base weights for the Puerto Rico RDD Sample were computed in the same way as those for the national (list-assisted) RDD Sample.

The replicate base weights were adjusted following the same steps as those applied to the full sample base weights. These included the screener level nonresponse adjustment, adjustment for multiple residential telephone lines, extended interview level nonresponse adjustment, and raking to the external veteran population counts obtained from the Census 2000 Supplementary Survey. By raking the replicate weights in the same manner as the full sample weights, the sampling variability in the raking adjustment factors would be reflected in the replicate weights, and hence included in the overall variance estimate. The raking procedure was carried out on the combined national and Puerto Rico RDD Samples.

If there were two or more veterans in a household, each respondent in the household received the same set of replicate base weights but the adjusted weights could differ because they could belong to different adjustment cells.

Composite Replicate Weights

To create the composite replicate weights, each replicate weight from the List Sample was multiplied by the same value of parameter l (=0.7272) that was used for creating the composite full sample weight. For the overlap RDD Sample cases, each replicate weight was multiplied by a factor of $(1-l)$. The remaining RDD Sample cases were assigned composite replicate weights equal to their original RDD Sample replicate weights. Finally, the composite replicate weights were raked to the veteran population counts estimated from the C2SS in a two dimensional raking procedure as was done for the composite full sample weights. The convergence criteria for the composite replicate weights was modified so that the percent absolute relative difference was no more than 0.1 percent for all marginal population counts. We recall that the convergence criteria for the composite full sample weights was that the percent absolute relative difference was no more than 0.01 percent for all marginal population counts.

6.5 Reliability of the Survey Estimates

Because estimates are based on sample data, they differ from figures that would have been obtained from complete enumeration of the veteran population using the same instrument. Results are subject to both sampling and nonsampling errors. Nonsampling errors include biases from inaccurate reporting, processing, and measurement, as well as errors from nonresponse and incomplete reporting. These types of errors cannot be measured readily. However, to the extent possible, each error has been minimized through the procedures used for data collection, editing, quality control, and nonresponse adjustment. The variances of the survey estimates are used to measure sampling errors. The variance estimation methodology is discussed in the next section.

Estimation of Variances of the Survey Estimates

The variance of an estimate is inversely proportional to the number of observations in the sample. Thus, as the sample size increases, the variance decreases. For the NSV 2001 the variance estimation methodology for estimates of totals, ratios (or means) and difference of ratios is based on the JK1 replication method, and the corresponding variance is given as:

$$v(\hat{q}) = \frac{R-1}{R} \sum_{r=1}^{R} \left(\hat{q}_{(r)} - \hat{q} \right)^2 , \tag{6-12}$$

where

q	is an arbitrary parameter of interest;
$\hat{q}$	is the estimate of q based on the full sample;
$\hat{q}_{(r)}$	is the estimate of q based on the observations included in the r^{th} replicate;
R	is the total number of replicates formed; and
$v(\hat{q})$	is the estimated variance of $\hat{q}$.

We have constructed the composite full sample and composite replicate weights for the combined List and RDD Samples corresponding to the JK1 replication methodology. The WesVar[3] variance estimation system can be used to produce the survey estimates based on the composite full sample weights and the corresponding variances of these estimates using the variance formula given in equation 6-12.

Construction of Confidence Intervals

Each of the survey estimates has an associated standard error, which is defined as the square root of the variance of the estimate. Consider the example of estimating the proportion of veterans with a certain characteristic, such as a service-connected disability. We denote by $\hat{p}$ the estimated proportion of

[3] WesVar is software for analyzing data from complex surveys. The software was developed by Westat and can be downloaded from Westat's website (www.westat.com/wesvar) for a 30-day free trial.

veterans with the particular characteristic of interest and let $v(\hat{p})$ be the corresponding variance estimate. Then the standard error of the estimated proportion $\hat{p}$ is given by

$$se(\hat{p}) = \sqrt{v(\hat{p}).} \tag{6-13}$$

The 95 percent confidence interval is the interval such that the unknown proportion p would have a 95 percent probability of being within the interval. The 95 percent confidence interval is given by

$$\hat{p} \pm t_{(0.025,50)} \times se(\hat{p}) \,. \tag{6-14}$$

The lower limit of the interval is $\hat{p} - t_{(0.025,50)} \times se(\hat{p})$, and the upper limit of the interval is $\hat{p} + t_{(0.025,50)} \times se(\hat{p})$. The width $t_{(0.025,50)} \times se(\hat{p})$ is known as half-width of the 95 percent confidence interval. The factor $t_{(0.025,50)}$ is the t-value at $a = 0.025$ with 50 degrees of freedom, which is approximately equal to 2.0. The smaller the half-width of the confidence interval, the more precise is the survey estimate.

Alternatively, the precision of the survey estimate can also be expressed in terms of the coefficient of variation (cv) of the estimate. The cv of an estimate is defined as the ratio of the standard error of the estimate and the magnitude of the estimate expressed in percent. Thus, the cv of the estimated proportion $\hat{p}$ is given by

$$cv(\hat{p}) = 100.0 \times {}^{se(\hat{p})}\!/_{\hat{p}} \,, \tag{6-15}$$

where $se(\hat{p})$ is the standard error of the estimated proportion $\hat{p}$. The smaller the cv of the estimate, the more precise is the estimate. The percent margin of error at the 95 percent confidence level can also be obtained by multiplying the cv of the estimate by the factor $t_{(0.025,50)}$.

6.6 Bias and Precision in the Combined Sample

We investigated two main issues associated with the use of the combined sample versus the separate RDD and List Samples. These were: (1) potential biases incurred in the estimates as a result of the matching involved in creating the composite weights, and (2) the gains in precision from the increased sample sizes of the combined sample. The reason that both of these issues are important is that the total mean square error (*MSE*) of a survey estimate is equal to the sum of its variance and the square of the bias, ($MSE = Variance + (Bias)^2$). In surveys with large sample sizes, the *MSE* may be dominated by the bias term. When sample sizes are small, the variance may be a greater cause for concern.

To address the first issue of bias, the potential risk of bias would be due mainly to imputing the overlap status of those RDD sample respondents who did not provide their Social Security numbers. We obtained an SSN from only 48 percent of the RDD Sample respondent veterans. Thus, the overlap status had to be imputed for those who did not report their SSNs. The question arises as to whether the cases that reported an SSN are different from those that did not. To answer this question, statistical comparisons were made for the two groups to see whether their distributions differed with respect to age and to other key statistics. Pairwise t-tests showed that those not reporting an SSN are:

More likely to be in the over 50 age group;

More likely to be in the higher income group;

More likely to have a higher education; and

More likely to belong to a discretionary priority group.

All comparisons are significant at the alpha = 0.05 level. For those who reported an SSN, we compared the characteristics of those who were on the list frame with those who were not on the list frame. The significant variables for this comparison were priority group, income, outpatient care, VA loan, and VA service-connected disability compensation. We used these variables as predictor variables in the *CHAID* analysis to determine homogeneous cells for imputing the overlap status for those who did not report their SSN. Therefore, the risk of potential bias was minimized due to imputing the overlap status within homogenous imputation cells.

The precision of the estimates can be evaluated by comparing the standard errors (SEs) of the estimates from the combined sample with those from the RDD Sample alone. In this situation, the population of analytical interest is the population of all noninstitutionalized veterans living in the U.S. The statistics of interest for the purpose of this analysis are proportions for various key statistics identified by the VA. As can be seen from the comparison of SEs in Table 6-2, the increased sample sizes of the combined sample always result in a significant reduction in sampling variability. The standard errors of the combined estimates are always lower than the standard errors of the corresponding estimates from the

RDD Sample alone. The design effects for the combined sample would generally be higher than the corresponding RDD Sample design effects due to increased variation in the sampling weights. The standard error of a survey estimate is inversely proportional to the square root of the effective sample size, where effective sample size is defined as the number of cases sampled divided by the design effect. Thus, the standard errors of the combined estimates would be lower than the RDD estimates as long as the increase in the design effect is less than the increase in the sample size. The ratio of the sample sizes for the combined sample and the RDD Sample alone is 1.54 (combined sample size divided by RDD Sample size). The standard error of the combined estimates therefore would be less than the standard error of the estimate from the RDD Sample alone as long as the design effect ratio is less than 1.54. We note from Table 6-2 that the design effect ratios for all the variables are less than 1.54. In fact, the design effect ratios are less than 1 for priority groups 1 through 4 and the service-connected disability (SD14d).

We recall that the List Sample design is a stratified design, where stratification is based on the health care priority groups (groups 1 through 6) and gender (male, female). The List Sample covered only the mandatory priority groups (groups 1 through 6). The gains from stratification for priority groups 1 through 4 more than offset the losses due to increased variation in the combined sample weights. Hence, the combined sample design effects are less than the RDD Sample design effect. The gains from stratification for priority groups 5 and 6 were not very large because of "strata jumpers." Many veterans who were thought to belong to priority groups 5 and 6 were actually observed as belonging to priority group 7. Therefore, the combined sample design effects for priority groups 5 and 6 are higher than the RDD Sample design effects. The combined sample design effect for the variable "SD14d" (service-connected disability) is lower than the RDD Sample design effect because of a high correlation between "SD14d" and the priority groups.

The efficiency of the combined sample as compared with the RDD Sample can also be defined as the ratio of the corresponding variances expressed as percentage. We denote by $Eff\left(Combined \;\; vs. \;\; RDD\right)$ the efficiency of the combined sample as compared with the RDD Sample alone, then

$$Eff\left(Combined \;\; vs. \;\; RDD\right) = 100 \times \frac{\mathrm{var}\left(RDD\right)}{\mathrm{var}\left(Combined\right)}, \qquad (6\text{-}16)$$

Table 6-2. Comparison between RDD and composite estimates

Question Number	Variable Description	Value	RDD only				Composite				RDD vs. Composite	
			Est(%)	SE	n	Deff	Est(%)	SE	n	deff	Deff Ratio	Var Ratio
	Priority Group	1	3.3	0.15	425	0.95	3.2	0.11	1974	0.83	0.87	1.86
		2	2.6	0.16	337	1.30	2.7	0.09	1666	0.60	0.46	3.16
		3	9.0	0.25	1162	0.96	9.0	0.18	3048	0.82	0.86	1.93
		4	0.1	0.03	10	1.17	0.1	0.02	96	0.55	0.47	2.25
		5	18.5	0.31	2296	0.82	18.4	0.26	3599	0.89	1.08	1.42
		6	10.9	0.31	1432	1.30	10.9	0.29	1887	1.77	1.36	1.14
		1 – 6	44.4	0.43	5662	0.99	44.3	0.38	12270	1.20	1.22	1.28
		7	55.6	0.43	7294	0.99	55.7	0.38	7778	1.20	1.22	1.28
MB24	COMBAT1	Yes	39.2	0.47	5145	1.18	39.2	0.46	9253	1.74	1.48	1.04
		No	60.8	0.47	7811	1.18	60.8	0.46	10795	1.74	1.48	1.04
ET1	EDUCTRG1	Yes	40.2	0.48	5369	1.25	40.2	0.42	8266	1.46	1.17	1.31
		No	59.8	0.48	7587	1.25	59.8	0.42	11782	1.46	1.17	1.31
HC1	ERYOU1	Yes	24.1	0.33	3107	0.75	24.1	0.32	5628	1.10	1.46	1.06
		No	75.9	0.33	9849	0.75	75.9	0.32	14420	1.10	1.46	1.06
HB21	MEDICARE1	Yes	39.3	0.23	5356	0.29	39.4	0.20	8789	0.35	1.22	1.32
		No	60.7	0.23	7600	0.29	60.6	0.20	11259	0.35	1.22	1.32
SD14d	VADISCMP	Yes	11.1	0.29	1436	1.10	11.2	0.18	5991	0.68	0.61	2.60
		No	88.9	0.29	11520	1.10	88.8	0.18	14057	0.68	0.61	2.60
SD14j	WELFARE	Yes	2.1	0.13	266	1.11	2.0	0.12	451	1.54	1.39	1.17
		No	97.9	0.13	12690	1.11	98.0	0.12	19597	1.54	1.39	1.17

where $\text{var}(RDD)$ and $\text{var}(Combined)$ are, respectively, the variances of the RDD Sample alone and the combined sample. The efficiency values of more than 100 percent imply that the combined sample estimates are more efficient than the estimates based on the RDD Sample alone. We notice that efficiencies are greater than 100 percent for all variables in Table 6-2 and the efficiency values range from 104 percent to 316 percent. Thus, the combined sample with the corresponding composite weights should be used for all VA analyses.