2002 NATIONAL SURVEY ON DRUG USE AND HEALTH

QUESTIONNAIRE DWELLING UNIT–LEVEL AND PERSON PAIR–LEVEL SAMPLING WEIGHT CALIBRATION

Contract No. 283-98-9008
RTI Project No. 7190.474.300
Deliverable No. 28

Project Director: Thomas G. Virag

Authors:

Patrick Chen
Lanting Dai
Harper Gordek
Eric Grau
Avinash C. Singh
Matthew Westlake

Prepared for:

Substance Abuse and Mental Health Services Administration
Rockville, Maryland 20857

Prepared by:

RTI International
Research Triangle Park, North Carolina 27709

October 8, 2004

2002 NATIONAL SURVEY ON DRUG USE AND HEALTH

QUESTIONNAIRE DWELLING UNIT-LEVEL AND PERSON PAIR-LEVEL SAMPLING WEIGHT CALIBRATION

Contract No. 283-98-9008
RTI Project No. 7190.474.300
Deliverable No. 28

Authors:

Patrick Chen
Lanting Dai
Harper Gordek
Eric Grau
Avinash C. Singh
Matthew Westlake

Project Director: Thomas G. Virag

Prepared for:

Substance Abuse and Mental Health Services Administration
Rockville, Maryland 20857

Prepared by:

RTI International
Research Triangle Park, North Carolina 27709

October 8, 2004

Preface

This report documents the method of weight calibration used for producing the final set of questionnaire dwelling unit (QDU) and pair weights for the National Survey of Drug Use and Health data from 2002. The weighting team faced several challenges in this subproject and was able to address them by resorting to innovative modifications of certain basic statistical ideas. These are listed below.

- Under Brewer's method, high weights may occur due to small pair selection probabilities. In any calibration exercise, some treatment of extreme value (ev) in weights is needed, but there is a danger of introducing too much bias by over-treatment. In GEM, extreme value control is built-in, but one needs to define suitable ev domains so that not too many evs are defined. If too many design variables are used to define ev domains, then each domain will be very sparse and not be of much use in defining thresholds for ev. As for 1999-2001 NHSDA pair data, a hierarchy of domains was defined using pair age (each pair member being in one of the three categories: 12 to 25, 26 to 49, and 50+), number of persons aged 12 to 25 in the household, State, and clusters of States (see Section 5.2 for details).

- Control of extreme values in weights helps to reduce instability of estimates to some extent, but there is a need for alternative methods which do not introduce as much bias. Following the famous suggestion of Hajek (1971) in his comments on Basu's fabled example of circus elephants, we performed ratio adjustment (a form of poststratification) to estimated totals obtained from the household data on number of persons belonging to the pair domain of interest. This was implemented in a multivariate manner to get one set of final sets.

- In the absence of a suitable source of poststratification controls for the person pair–level weights and the household level weights, the inherent two-phase nature of the NSDUH design was capitalized upon to estimate these controls from first phase of the large screener sample. The first phase sample weight was poststratified to person-level U.S. Census counts to get more efficient estimated counts for pair and household data.

- The problem of multiplicities complicated the issue of providing one set of final weights. When dealing with person-level parameters involving drug-related behaviors among members of the same household, it is possible for an individual to manifest himself or herself in the pair sample through different pairs. To avoid overcounting, the pair weights have to be divided by multiplicity factors which tend to be domain-specific. For this reason, multiplicity factors for a key set of pair analysis domains are also produced along with a set of final calibrated pair weights.

- Missing items in the respondent questionnaire led to imputation for deriving pair relationships, multiplicity factors, and household counts for Hajek adjustments. The general method of predictive mean neighborhood (PMN) was used for this purpose with suitable modifications.

The subproject required enduring efforts by a dedicated team consisting of Patrick Chen, Matt Westlake, Harper Gordek and Lanting Dai for weight calibration, and Shijie Chen and Eric

Grau for imputation. Some results from this calibration task were also presented at 2001 JSM (Chromy and Singh on estimation in the presence of multiplicities and extreme weights; Singh, Grau, and Folsom on the use of PMN; and Penne, Chen, and Singh on GEM calibration of pair weights). The authors would like to take this opportunity to thank a number of individuals for useful discussions and suggestions: Doug Wright, Joe Gfroerer, and Art Hughes of SAMHSA, and Jim Chromy and Ralph Folsom of RTI.

Avi Singh, Task Leader
Research Triangle Park, NC
February 2004

Table of Contents

Chapter **Page**

Table of Contents (continued)

Chapter **Page**

Appendix

Table of Contents

Appendix **Page**

List of Tables

List of Tables (continued)

Table **Page**

List of Tables (continued)

List of Exhibits

List of Exhibits (continued)

List of Terms and Abbreviations

DU	Dwelling unit.
ev	Extreme value. See Section 5.1 for more detail.
GEM	Generalized exponential model. See Chapter 3 for more detail.
half-step	This refers to halving the increment in the Newton-Raphson iterative process for fitting GEM.
Household-level person count	The number of pairs associated with a given domain in a given household. These counts are used as control totals in the poststratification step. See Chapter 6 for details on how these counts are created, and Chapter 4 for details on their use in poststratification.
IQR	Interquartile range.
Multiplicity factor	The number of pairs associated with a given respondent in a given domain. See Chapter 6 for more detail.
nr	Nonresponse.
Outwinsor	The proportion of weights trimmed after extreme value adjustment via winsorization.
Pair domain	A pair relationship where the target population is defined by one of the pair members, conditional on the attributes of the other pair member.
Pair relationship	The relationship between selected pair members.
Parent-child	A pair relationship where either both pair members identify the other as part of a parent-child relationship, or both pair members otherwise are determined to form a parent-child pair (either through other evidence or through imputation).
ps	Poststratification.
QDU	Questionnaire dwelling unit: a household where at least one member responded to the questionnaire.
SDU	Screener dwelling unit: a household where screener information is available.
sel.pr.ps	Selected person pair poststratification adjustment step. See Section 7.3.2 for more detail.
sel.qdu.ps	Selected questionnaire dwelling unit poststratification adjustment step. See Section 7.2.2 for more detail.
Sibling-sibling	A pair relationship where the pair members are siblings (either reported to be so, or otherwise determined to be so).

List of Terms and Abbreviations (continued)

Spouse-spouse	A pair relationship where the pair members are either married or living together as though married (either reported to be so, or otherwise determined to be so).
res.pr.nr	Respondent pair nonresponse adjustment step. See Section 7.3.3 for more detail.
res.qdu.nr	Respondent questionnaire dwelling unit nonresponse adjustment step. See Section 7.2.3 for more detail.
res.pr.ev	Respondent pair extreme value adjustment step. See Section 7.3.5 for more detail.
res.qdu.ev	Respondent questionnaire dwelling unit extreme value adjustment step. See Section 7.2.5 for more detail.
res.pr.ps	Respondent pair poststratification adjustment step. See Section 7.3.4 for more detail.
res.qdu.ps	Respondent questionnaire dwelling unit poststratification adjustment step. See Section 7.2.4 for more detail.
sandwich SE	Sandwich standard error. See Section 8.5 for more detail.
SE	Standard error.
SES	Socioeconomic status indicator. See Exhibit 4.1 for more detail.
UWE	Unequal weighting effect. It refers to the contribution in the design effect due to unequal selection probability and is defined as 1+[(n-1)/n]*CV2, where CV = coefficient of variation of weights and n is the sample size.
Winsorization	A method of extreme value adjustment that replaces extreme values with the critical values used for defining low and high extreme values.

1. Introduction

Traditionally, most household surveys have been designed either to measure characteristics of the entire household or to focus on a randomly selected respondent from among those determined to be eligible for the survey. Selecting more than one person from the same household was considered ill-advised since persons from the same household tended to repeat the same general information characteristic of the entire household. Selecting only one person per household totally avoided the clustering effect on the variance. The "one person per household" sampling approach, however, precludes the opportunity to gather information about the relationships among household members. In the National Survey on Drug Use and Health[1] (NSDUH), we allow for a richer analytic capability of a survey designed to assure a positive pairwise probability of selection among all eligible household members in each sample household. Achieving positive probabilities for all pairs within sampled households permits unbiased estimation of the within dwelling unit component of variance. Besides providing efficient data collection, this sampling method also facilitates the study of the relationships of social behaviors among members of the same household. This report documents the methodology and development of calibrated weights for the second objective, the study of behavioral relationships among persons residing in the same household. The report also describes the development of questionnaire dwelling unit (QDU) weights which are of independent interest for studying household level characteristics and are also needed for producing household count estimates of the number of persons belonging to pair relationship domains for use as poststratification controls for pair weights.

The NSDUH allows for estimating characteristics at the person level, pair level, and household or QDU level. This report describes the weight calibration methods used for the pair-level and QDU-level respondents. As described in the person-level report, the NSDUH is an annual survey of about 70,000 persons selected from the civilian noninstitutional population aged 12 and older from all 50 States and the District of Columbia. Under a stratified design with States serving as the primary strata and FI regions as secondary, segments and dwelling units for the screener questionnaire are selected using probability proportional to size sampling in two stages. A large number of screener dwelling units (SDUs) (about 200,000) are selected to ensure that various age groups (five in all: 12 to 17, 18 to 25, 26 to 34, 35 to 49, and 50+) of eligible individuals are adequately represented in the final sample. Information collected from SDUs also provides estimates of population controls (as in two-phase sampling) for calibration at levels (such as pair and QDU) for which suitable Census-based controls are not available. From each selected SDU, 0, 1, or 2 persons are selected using a modification of Brewer's method such that prescribed sampling rates for the five age groups by each State are achieved with high selection rates for youth (12 to 17) and young adults (18 to 25). Table 1.1 shows the eligible number of selected and responding SDUs, QDUs, pairs, and persons for each of the four years 1999–2002. The distribution of pair data for different pairs of age groups may vary considerably; see Chapter 2 for details. It is seen that for certain age group domains, the realized sample size may not be sufficient to yield reliable estimates. Also, there may be problems of extreme weights due to small pair selection probabilities under Brewer's method that may cause instability of estimates.

---

[1] This report presents information from the 2002 National Survey on Drug Use and Health (NSDUH). Prior to 2002, the survey was called the National Household Survey on Drug Abuse (NHSDA).

These and some other estimation issues related to pair data are discussed below, along with some possible solutions that were adopted.

Table 1.1 1999, 2000, and 2001 NHSDAs, and 2002 NSDUH Sample Sizes

Sample unit		**1999**	**2000**	**2001**	**2002**
SDU	**Selected**	187,842	182,576	171,519	150,162
	Completed	169,166	169,769	157,471	136,349
QDU	**Selected**	67,385	69,522	66,697	55,686
	Completed	51,821	55,924	53,134	48,088
Pair	**Selected**	22,498	22,439	23,048	24,895
	Completed	14,885	15,840	15,795	20,038
Person	**Selected**	89,883	91,961	89,745	80,581
	Completed	66,706	71,764	68,929	68,126

First, we note that for studying drug-related behavioral relationships among members of the same household, pair data is required because the outcome variable is generally defined with respect to the specific other member selected from the household. However, the parameter of interest is generally at the person level and not at the pair level. For example, in the parent-child pairs, one may be interested in the proportion of children that have used drugs in the past year who have parents that report talking to their child about drugs. Here the target population only consists of children, and not all possible pairs. Note the pair-level (2 persons per QDU) sample forms a subsample of the larger person-level (1 or 2 persons per QDU) sample, with the QDUs themselves selected from the larger sample of SDUs. The NSDUH has features of a two-phase design which turns out to be useful for estimating calibration controls for poststratification of household-level weights and person pair–level weights. No other outside source is available for obtaining these controls. For this purpose, the screener-level household weights are poststratified to person-level U.S. Census counts to obtain more efficient estimated controls for pair and household data.

In estimation for pair domains, two major problems arise: one is that of multiplicities because, for a given domain defined by the pair relationship, when the parameter of interest is at the person level, several pairs in the household could be associated with the same person. The other problem is that of extreme weights that may arise due to small selection probabilities for certain pair-age groups, which may lead to unstable estimates. Each of these issues is discussed in turn.

If several pairs in the household are associated with the same person, it is necessary to use the average measure of behavior relationships for each member, which gives rise to multiplicities. Thus, the design weights need to be divided by the person-level multiplicity factors for each domain of interest. Therefore, multiplicity factors need to be produced along with the final set of calibrated weights. Because it is not straightforward to create these multiplicities, analyses would have to be necessarily limited to pair relationships where the multiplicities were produced a priori. It was anticipated that analyses of interest would be limited to 14 pair domains, listed in Table 1.2. Since no multiplicity was necessary for the spouse-spouse/partner-partner pair relationships (by definition, each pair member could only have one partner or one spouse), multiplicity factors were produced for only 12 of these domains. Note that a single pair relationship might have two domains associated with it, since the parameter of interest might be associated with only one member of the pair (the "focus" member), and the multiplicity would differ depending upon which pair member was the focus member.

Table 1.2 Pair Domains

Pair relationship	Focus
Parent-child, child aged 12-14	Parent
Parent-child, child aged 12-14	Child
Parent-child, child aged 12-17	Parent
Parent-child, child aged 12-17	Child
Parent-child, child aged 12-20	Parent
Parent-child, child aged 12-20	Child
Parent-child, child aged 15-17	Parent
Parent-child, child aged 15-17	Child
Sibling-sibling, older sibling 15-17, younger sibling 12-14	Older sibling
Sibling-sibling, older sibling 15-17, younger sibling 12-14	Younger sibling
Sibling-sibling, older sibling 18-25, younger sibling 12-17	Older sibling
Sibling-sibling, older sibling 18-25, younger sibling 12-17	Younger sibling
Spouse-spouse and partner-partner	No multiplicity necessary
Spouse-spouse and partner-partner, with children aged 0-17	No multiplicity necessary

Some of the multiplicities, including counts of all possible pairs in a household for a given domain, were used for poststratification. Details are given in Chapter 4.

A resolution to the extreme-weight problem is to use a Hajek-type modification. This modification essentially entails calibration (like poststratification) to controls for the number of persons in households belonging to each domain of interest; these controls can be obtained from the larger sample of singles and pairs (i.e., one or two persons selected from dwelling units [DUs]). Note, however, that the multiplicity factor, being domain-specific, renders the calibration adjustment factor domain-specific. This raises the question of finding one set of calibration weights for use with all domains or outcome variables. To get around this problem, we perform a multivariate calibration with respect to a key set of pair domains. This type of poststratification is then followed by a repeat poststratification to further control the extreme weights by imposing separate bound restrictions on the initially identified extreme weights.

The generalized exponential model (GEM) method (Folsom & Singh, 2000) was used for calibration of both QDU- and pair-level design weights through several steps of adjustment as shown in Exhibit 1.1. In GEM, treatment of extreme value (ev) weights is built-in via the definition of lower and upper bounds for the extreme weights. For pair data, there was a problem defining suitable domains for defining extreme weights, as explained in the following.

In dealing with extreme weights, it is assumed that they arise due to design (due to an imperfect frame, assignment of very small selection probabilities to some units, or a small weight adjustment factor after calibration), so they do make the sample a representative of population and, hence, do not introduce bias. The only problem is that they may lead to highly unstable estimates similar to the problem of Basu's circus elephants[2] (see Hajek, 1971). So, we need to perform some treatment (such as winsorization) within suitably defined extreme weight domains such that these domains contain units possibly from different strata but with similar sample selection probabilities to avoid the occurrence of extreme weights due to a mix of different designs. The domains must be large enough (say at least of size 30) to be able to define extreme values according to the domain-specific weight distribution. Any extreme value treatment to increase precision of estimates would introduce some bias. However, this bias can be reduced considerably if the ev treatment is performed under calibration controls. This is what the built-in ev control in GEM tries to accomplish.

It follows that the definition of extreme weight domains should depend on factors that affect the selection probabilities of units in the sample, such as State- and age-specific sampling rates, segment selection probabilities, pair age–specific selection probabilities, and household composition. If one tries to define extreme weight domains by taking account of all these factors via cross-classification, it will lead to too many domains without having a sufficient number of observations. That is why it is difficult to define suitable extreme weight domains for pair data. In the case of person-level weights it was less difficult, since State by age group suitably captured the extreme weight domain requirements. The definition of extreme weight domains used in the 2002 NSDUH was the same as in 1999–2001. The domains were defined as the cross-classification of State, new pair age, and number of persons aged 12 to 25. In particular, the new pair age was defined by the age groups of each pair member according to the age categories 12 to 25, 26 to 49, and 50+ (resulting in six pair age categories), and the number of persons aged 12 to 25 were categorically defined as 0, 1, and 2 or more. For more details, see Chapter 5.

[2] A circus owner had 50 elephants, and wanted to estimate the total weight to help him make arrangements for shipping. To save time, he only wanted to weigh Sambo (an average sized elephant), and use 50 times its weight as an estimate. However, the circus statistician, being highly conscious of the optimality and unbiasedness of the Horvitz-Thompson (HT) estimator, objected about the potential bias of his estimate because of the purposive selection. Instead, he suggested random selection of an elephant with a very high probability of 99/100 for Sambo, and the rest including Jumbo (the biggest in the herd) with probability 1/4900 each. The circus owner was very unhappy with the statistician's response of 100/99 times the Sambo's weight as the estimate if Sambo got selected in this random draw, and was outraged with the response of 4900 times the Jumbo's weight if Jumbo happened to get selected. It was obvious to the owner that this new estimator was extremely poor, although he didn't know anything about its unbiasedness. The story had an unhappy ending with the circus statistician losing his job. To alleviate the instability of the HT-estimator, Hajek suggested to multiply it by 50 divided by inverse of the selection probability, which reduces simply to 50 times the weight of the selected elephant.

Exhibit 1.1 QDU and Pair Sampling Weight Calibration Steps

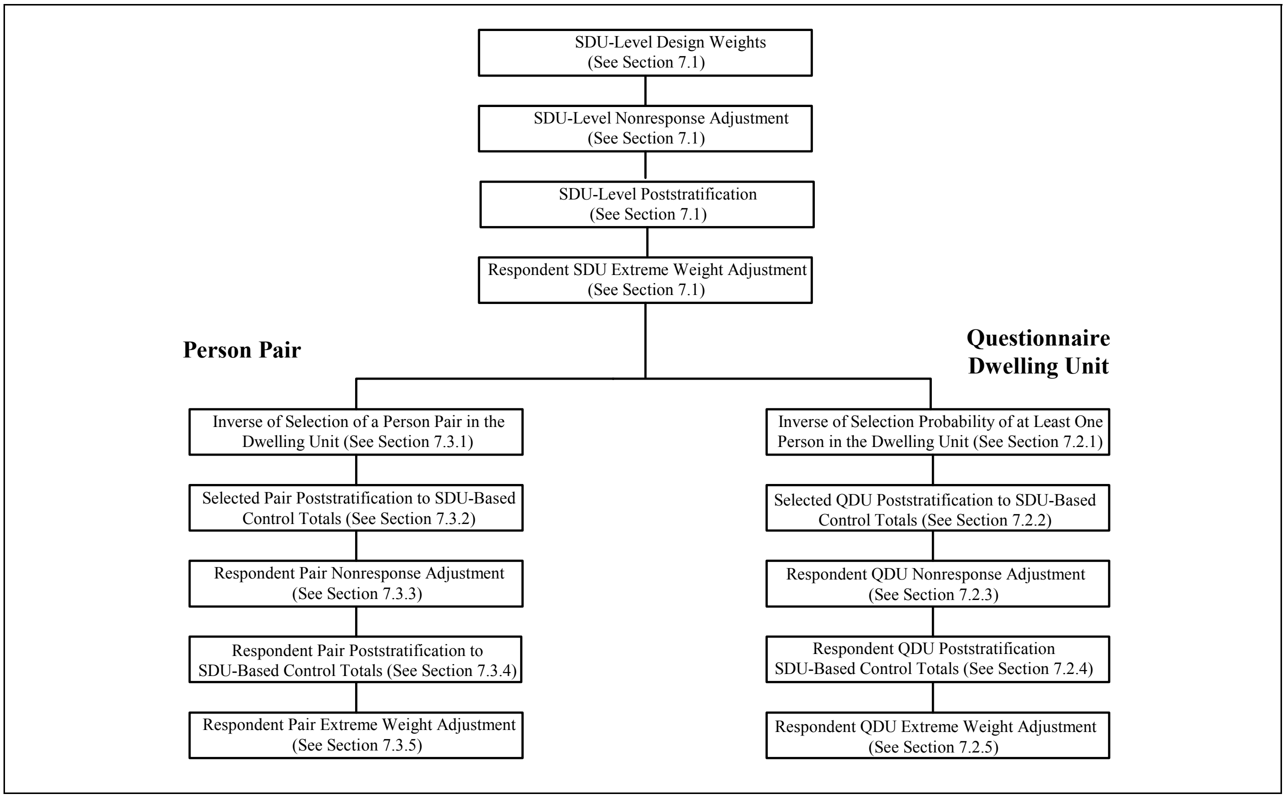

2. Questionnaire Dwelling Unit and Pair Selection Probabilities

Similar to the 1999, 2000, and 2001 National Household Survey on Drug Abuse (NHSDA) computer-assisted interview (CAI) surveys, the 2002 National Survey on Drug Use and Health (NSDUH) has a two-phase design and uses a CAI method. There were three stages of selection: segments, dwelling units within segments, and persons within dwelling units. Any two survey eligible persons had some nonzero chance of being selected and, when both are selected, they form a within household pair. This design feature is of interest to NSDUH researchers because, for example, it allows analysts to examine how the drug use propensity of an individual (in a family) relates to the drug use propensity of other members residing in the same dwelling unit (Odom et al., 2004).

For previous years (1999–2001), the method used for selecting pairs was as follows. For a given dwelling unit (DU), if the sum of the age-specific selection probabilities was larger than 2, then the individual person-selection probabilities were ratio adjusted downward to make the sum equal to 2. If the sum was less than 2, the difference between 2 and the sum of the probabilities was evenly distributed over 3 dummy persons, so that the sum of the person probabilities was made exactly 2. Brewer's method was then applied to select a person pair using the pair selection formula (2.1). If the selected pair consisted of two real persons, then both persons were selected. If the selected pair consisted of one real person and one dummy person, then the real person was selected. If the selected pair consisted of two dummy persons, no one was selected from that dwelling unit.

However, in order to increase the number of pairs selected in the sample, the pair-sampling algorithm for the 2002 NSDUH was modified from the one used for the previous NHSDAs (1999–2001) as explained below. For dwelling units with sum of person-level selection probabilities (denoted by S) less than 2, the earlier algorithm was modified (see Case II below) to increase the chance for selecting a pair. However, for dwelling units with $S \geq 2$, there was no need to change the algorithm denoted below by Case I.

2.1 Pair Selection Probability

Case I: DUs with $S \geq 2$

For a given DU, if the sum of the age-specific person selection probabilities (S) was larger than 2, then the selection probability was ratio adjusted by a multiplicative adjustment factor so that all probabilities were scaled down to sum to exactly 2. Now, Brewer's method sets the pairwise selection probabilities at

$$P_{h(ij)} = \left[\frac{P_{h(i)}P_{h(j)}}{K}\right]\left[\frac{1}{1-P_{h(i)}} + \frac{1}{1-P_{h(j)}}\right] \qquad (2.1)$$

by setting K at

$$K = 2 + \sum \frac{P_{h(i)}}{1 - P_{h(i)}} \tag{2.2}$$

where i = i^{th} person in the household (whose selection probability depends on his or her age category: 0, 1, 2, 3, 4, or 5)

j = j^{th} person in the household (whose selection probability depends on his or her age category: 0, 1, 2, 3, 4, or 5)

where age category 0 corresponds to dummy persons, 1 to persons aged 12 to 17, 2 to persons aged 18 to 25, 3 to persons aged 26 to 34, 4 to persons aged 35 to 49, and 5 to persons aged 50+.

The sum of the pairwise selection probabilities taken over all unique pairs will be guaranteed to be exactly one.

$$\sum_{i} \sum_{j>i} P_{h(ij)} = 1 \tag{2.3}$$

It also guarantees that

$$\sum_{j \neq i} P_{h(ij)} = P_{h(i)} \tag{2.4}$$

for all values of i.

As noted earlier, the above scheme will always select a pair, but it allows all combinations of eligible persons and dummy persons so that 0, 1, or 2 eligible persons are selected.

Case II: DUs with S<2

If the sum S of person-level selection probabilities is less than 2, the earlier method used in previous NHSDAs consists of dividing 2-S equally among the three dummy persons added to the household, and then using Brewer's method as in Case I. However, if the household has 2 or more persons, we would like to have more chances for a pair to be selected. To achieve this goal, the individual selection probabilities $P_{h(i)}$ are scaled upward by a factor F_s such that their sum comes close to but does not exceed 2, and that each person-selection probability does not exceed the maximum allowed probability of 0.99. Thus, denoting the revised person selection probabilities by $P'_{h(i)}$, the factor F_s is given by

$$F_S = Min\left\{\frac{T(\lambda)}{S}, \frac{0.99}{\max\{P_{h(i)}\}}\right\} \tag{2.5}$$

where $T(\lambda) = S + \lambda(2\text{-}S)$ and λ is set to 0.5. Note that if λ is chosen as 0, then F_s =1 as in Case I. The individual person probabilities are scaled upward by the factor F_s to sum to 2 or as close to 2 as possible. If, after scale adjustment, the sum, S', is exactly 2, then dummy persons are not needed. If S' is less than 2, three dummy persons are added as before.

Now, for Brewer's method, we set the pairwise selection probabilities similar to (2.1), as

$$P'_{h(ij)} = \left[\frac{P'_{h(i)}P'_{h(j)}}{K'}\right]\left[\frac{1}{1-P'_{h(i)}} + \frac{1}{1-P'_{h(j)}}\right] \tag{2.6}$$

by setting K' at

$$K' = 2 + \sum_i \frac{P'_{h(i)}}{1-P'_{h(i)}} \tag{2.7}$$

where $P'_{h(i)}$ and $P'_{h(j)}$ are the selection probabilities adjusted by the scaling factor (F_s)
i = i^{th} person in the household (whose selection probability depends on his or her age category 0, 1, 2, 3, 4, 5)
j = j^{th} person in the household (whose selection probability depends on his or her age category 0, 1, 2, 3, 4, 5)

Note that we now have $\sum_{j \neq i} P'_{h(ij)} = P'_{h(i)}$. To maintain the original person selection probabilities despite the scale adjustment by F_s, we modify the above Brewer's method as follows. Draw a random number R from uniform (0,1); if $R \leq 1/F_s$, then select a pair using Brewer's method based on formula (2.6). However, if $R > 1/F_s$, then no persons are selected from the household. In this way, the probability for selecting a pair (i,j) in household h becomes $P^*_{h(ij)} = P'_{h(ij)}/F_s$, which, in turn, gives the original person selection probabilities, $P_{h(i)}$.

It can be seen from the Table 2.1, the number of pairs selected for the 2002 NSDUH increased under the new pair selection algorithm, particularly for pairs where both members are aged 26 or older.

2.2 Questionnaire DU Selection Probability

A dwelling unit was considered a selected questionnaire dwelling unit (QDU) if it had completed the screening interview and had at least one person selected for the questionnaire interview. QDUs with at least one respondent were considered respondent QDUs.

The QDU selection probability was defined as

$$P^*_h = \left(1 - P^*_{h(00)}\right) \tag{2.8}$$

where $P^*_{h(00)}$ is the probability of not selecting any person. For the DUs with the sum of age-specific selection probabilities larger than or equal to 2 (Case I), $P^*_{h(00)}$ is 0. It follows from Section 2.1, under Case II, $P^*_{h(00)}$ can be calculated as

$$P^*_{h(00)} = \left(1 - \frac{1}{F_S}\right) + \frac{3}{F_S}\left[\frac{P'_{h(0)}P'_{h(0)}}{K'}\right]\left[\frac{1}{1-P'_{h(0)}} + \frac{1}{1-P'_{h(0)}}\right] \tag{2.9}$$

where $P'_{h(0)}$ is the selection probability of a dummy person when person-selection probabilities are adjusted by F_s.

Table 2.1 Distribution of Pair Age Groups for the 1999, 2000, and 2001 NHSDAs and the 2002 NSDUH

	1999 NHSDA			2000 NHSDA			2001 NHSDA			2002 NSDUH		
Domain	**Selected[1]**	**Respondents[2]**	**Rate[3]**	**Selected[1]**	**Respondents[2]**	**Rate[3]**	**Selected[1]**	**Respondents[2]**	**Rate[3]**	**Selected[1]**	**Respondents[2]**	**Rate[3]**
DUs												
Total DUs Screened	169,166	n/a	n/a	169,769	n/a	n/a	157,469	n/a	n/a	136,349	n/a	n/a
QDUs												
Total QDUs	67,385	51,821	76.90%	69,522	55,924	80.44%	66,697	53,134	79.66%	55,686	48,088	86.36%
Persons												
Total Persons	89,883	66,706	74.21%	91,961	71,764	78.04%	89,745	68,929	76.81%	80,581	68,126	84.54%
12 to 17	32,011	25,357	79.21%	31,242	25,717	82.32%	28,188	23,133	82.07%	26,230	23,645	90.14%
18 to 25	30,439	21,933	72.06%	29,424	22,613	76.85%	30,304	22,658	74.77%	27,216	23,066	84.75%
26 to 34	10,870	7,878	72.47%	12,403	9,552	77.01%	8,825	6,893	78.11%	7,672	6,374	83.08%
35 to 49	8,825	6,246	70.78%	9,583	7,158	74.69%	13,663	10,036	73.45%	12,076	9,620	79.66%
50+	7,738	5,292	68.39%	9,309	6,724	72.23%	8,765	6,209	70.84%	7,387	5,421	73.39%
Pairs												
Total Pairs	22,498	14,885	66.16%	22,439	15,840	70.59%	23,048	15,795	68.53%	24,895	20,038	80.49%
None, 12 to 17	11,742	10,955	93.30%	11,559	11,014	95.29%	9,696	9,417	97.12%	8,548	8,327	97.41%
None, 18 to 25	13,304	11,386	85.58%	12,859	11,488	89.34%	12,150	11,056	91.00%	9,903	9,220	93.10%
None, 26 to 34	8,175	6,164	75.40%	9,147	7,295	79.75%	6,608	5,326	80.60%	3,677	3,287	89.39%
None, 35 to 49	5,019	3,781	75.33%	5,551	4,382	78.94%	7,802	6,125	78.51%	4,418	3,815	86.35%
None, 50+	6,647	4,650	69.96%	7,967	5,905	74.12%	7,393	5,415	73.24%	4,245	3,401	80.12%
12 to 17, 12 to 17	5,906	4,391	74.35%	5,630	4,363	77.50%	4,772	3,724	78.04%	4,667	4,192	89.82%
12 to 17, 18 to 25	4,092	2,684	65.59%	3,740	2,600	69.52%	3,534	2,475	70.03%	3,245	2,742	84.50%
12 to 17, 26 to 34	1,103	730	66.18%	1,184	880	74.32%	836	604	72.25%	826	694	84.02%
12 to 17, 35 to 49	2,787	1,893	67.92%	2,941	2,108	71.68%	4,054	2,848	70.25%	3,795	3,121	82.24%
12 to 17, 50+	475	313	65.89%	558	389	69.71%	524	341	65.08%	482	377	78.22%
18 to 25, 18 to 25	5,450	3,299	60.53%	5,212	3,506	67.27%	5,921	3,716	62.76%	5,520	4,419	80.05%
18 to 25, 26 to 34	915	575	62.84%	1,085	729	67.19%	875	602	68.80%	975	806	82.67%
18 to 25, 35 to 49	808	460	56.93%	815	506	62.09%	1,329	792	59.59%	1,449	1,042	71.91%
18 to 25, 50+	420	230	54.76%	501	278	55.49%	574	301	52.44%	604	418	69.21%
26 to 34, 26 to 34	285	172	60.35%	415	279	67.23%	177	129	72.88%	774	559	72.22%
26 to 34, 35 to 49	65	42	64.62%	87	50	57.47%	111	72	64.86%	450	346	76.89%
26 to 34, 50+	42	23	54.76%	70	40	57.14%	41	31	75.61%	196	123	62.76%
35 to 49, 35 to 49	61	29	47.54%	75	44	58.67%	152	85	55.92%	807	543	67.29%
35 to 49, 50+	24	12	50.00%	39	24	61.54%	63	29	46.03%	350	210	60.00%
50+, 50+	65	32	49.23%	87	44	50.57%	85	46	54.12%	755	446	59.07%

[1] Selected pairs are based on the screener age.
[2] Respondent pairs are based on questionnaire age.
[3] These rates are unweighted and based only on the total selected and total responding counts of pairs.

3. Brief Description of Generalized Exponential Model (GEM)

In survey practice, design-based weights are typically adjusted in three steps: (1) for extreme values (ev), via winsorization, (2) for nonresponse adjustment (nr) via weighting class, and (3) for poststratification (ps) via raking-ratio adjustments. If weights are not treated for extreme values, the resulting estimates, although unbiased, will tend to have low precision. The bias introduced by winsorization is alleviated to some extent through ps. The nr adjustment is a correction for bias introduced in estimates based only on responding units; ps is an adjustment for coverage (typically undercoverage) bias, as well as for variance reduction due to correlation between the study and control (usually demographic) variables.

There are limitations in the existing methods of weight adjustment for ev, nr, and ps. It would be desirable to adjust for bias introduced in the ev step (when extreme weights are treated via winsorization) in that the sample distribution for various demographic characteristics is preserved. For the nr step, there are general raking-type methods, such as the scaled constrained exponential model developed by Folsom and Witt (1994), where the lower and upper bounds can be suitably chosen by use of a separate scaling factor. The factor is set as the inverse of the overall response propensity. It would be desirable to have a model for the nr adjustment factor so that the desired lower and upper bounds on the factor are part of the model. Note that the lower bound on the nr adjustment factor should be 1, as it is interpreted as the inverse of the probability of response for a particular unit. For the ps step, on the other hand, the general calibration methods of Deville and Särndal (1992), such as the logit method, allow for built-in lower (L) and upper (U) bounds (for ps, typically $L<1<U$). However, it would be desirable to have nonuniform bounds (L_k, U_k) depending on the unit k such that the final adjusted weights, w_k, could be controlled within certain limits. An important application of this feature would be weight adjustments in the presence of ev such that the user will have some control on the final adjustment of the initially identified extreme values.

A modification of the earlier method of scaled constrained exponential model of Folsom and Witt (1994), termed as the method of GEM and proposed by Folsom and Singh (2000), provides a unified approach to the three weight adjustments for ev, nr, and ps, and it has the desired features mentioned above. The functional form of the GEM adjustment factor is given in Appendix A. It generalizes the logit model of Deville and Särndal (1992), typically used for ps, such that the bounds (L, U) may depend on k. Thus, it provides a built-in control on ev during both ps and nr adjustments. In addition, the bounds are internal to the model and can be set to chosen values (e.g., $L_k = 1$ in the nr step). If there is a low frequency of ev in the final ps, then a separate ev step may not be necessary.

In fitting GEM to a particular problem, choice of a large number of predictor variables along with tight bounds will have an impact on the resulting unequal weighting effect (UWE) and the proportion of extreme values. In practice, this leads to somewhat subjective considerations of trade-off between the target set of bounds for a given set of factor effects, and the target UWE and the target proportion of extreme values. It may also be beneficial to look at the proportion of "outwinsors" (a term coined to signify the extent of residual weights after winsorization), which is probably more realistic in determining the robustness of estimates in the presence of extreme values.

A large increase in the number of predictor variables in GEM typically would result in a higher UWE, thus indicating a possible loss in precision. This was checked by comparing SUDAAN-based standard errors of a key set of estimates computed from two sets of calibration models, one baseline using only the main effects and other the final model. The results are presented in Chapter 8.

To implement GEM, several steps need to be followed: (1) define and create all the covariates; (2) define the extreme weights; (3) fit GEM model. The details of practical aspects of GEM implementation can be found in Chapter 4 of Chen et al. (2004).

4. Predictor Variables for the Questionnaire Dwelling Unit (QDU) and Pair Weight Calibration via the Generalized Exponential Model (GEM)

We note that unlike the person-level weight calibration, the control totals for the QDU-level and person pair–level poststratification are not available from the U.S. Census. A way around this potential problem is to take advantage of the two-phase nature of the design, in which the screener data provides a large sample containing demographic information that can be used to derive control totals for the QDU-level and person pair–level sampling weight calibrations, as well for the selected person poststratification adjustment. The stability of control totals from the screener dwelling unit (SDU)–level data can be improved by poststratification of the SDU sample using person-level counts from the U.S. Census. This was indeed done as documented in the person-level weight calibration reports (Chen et al., 2004).

4.1 QDU Weight Calibration

After the nonresponse and poststratification adjustments at the SDU level, which are common to the person-level weight calibration, the QDU sample weights were adjusted in three steps: poststratification of selected QDUs, nonresponse adjustment of respondent QDUs, and poststratification of respondent QDUs. The set of initially proposed predictor variables for these adjustments using GEM were set to be common and correspond to those used for the SDU nonresponse and poststratification adjustments. The variables are of two types: those used for SDU nonresponse adjustment are 0/1 indicators, while those used for SDU poststratification adjustment are counting variables. The variables of the first type (0/1 indicators) are population density, group quarters, race of householder, percentage of persons in segment who are black, percentage of persons in segment who are Hispanic, percentage of owner-occupied dwelling units (DUs) in segment, segment-combined median rent and housing value, and household type. Variables of the second type (counting variables) represent the number of eligible persons within each DU who fall into the various demographic categories of race, age group, Hispanicity, and gender. Note that the State and quarter variables are represented as both binary and counting variables. Thus, not only are DU counts within a specific State or quarter in the QDU sample controlled to the corresponding totals obtained from the SDU sample, but also counts of persons living in the DUs in the QDU sample are controlled to totals from the SDU sample. These person-level totals match the U.S. Census estimates because of the SDU-level poststratification to U.S. Census counts. It may be noted that in the poststratification of selected QDUs and the nonresponse adjustment of the respondent QDUs steps, demographic information from screener data was used in defining covariates, whereas in the poststratification of the selected QDUs step, questionnaire demographic information was used.

Exhibit 4.1 lists all predictor variables proposed for QDU-level calibration and identifies them as counting, binary, or both. Various main effects and higher level factor effects based on the predictor variables were included in the GEM modeling. As stated previously, all adjustment steps at the QDU level used a common set of proposed predictor variables.

4.2 Pair Weight Calibration

Like QDU, the initial set of weight components in pair weight calibration are the same as the set obtained from the SDU-level weight calibration. The SDU-calibrated weight is multiplied by the pair-level design weight, which in turn was adjusted in four steps: poststratification of selected pairs, nonresponse adjustment of respondent pairs, poststratification of respondent pairs, and the extreme weight adjustment of respondent pairs. All the adjustment steps for pair weights utilized the same set of initially proposed predictor variables, which included a subset of those used for the person-level nonresponse adjustment. This included segment characteristic variables, such as population density, percentage of persons in segment who are black, percentage of persons in segment who are Hispanic, percentage of owner-occupied DUs in segment, and segment-combined median rent and housing value. Also included were pair-specific covariates, such as the demographic characteristics pair age, pair race, and pair gender, as well as dwelling unit characteristics, such as race of householder, household type, household size, and group quarters indicators. State and quarter indicators were included as well. However, for two-factor effects, instead of individual State, State/region was used due to insufficient sample size. This resulted in a 12-level variable where the eight large sample States were kept separate, and the remainder of States were grouped according to the four Census regions. All variables were defined as (0/1) indicators. These proposed predictor variables and their levels are shown in Exhibit 4.2.

In the poststratification of selected pairs and the nonresponse adjustment of respondent pairs, screener data were used in the definition of the pair-specific variables pair age, pair race, and pair gender, whereas in the poststratification and extreme weight adjustment of respondent pairs, these variables were obtained from the questionnaire. For the latter case, in addition to the variables described above, indicator covariates corresponding to selected pair domains were included to perform Hajek-type ratio adjustments via weight calibration, as mentioned in Chapter 1. The selected pair domains were limited to 10 of the 14 pair domains listed in Chapter 1. (Parent-child pairs where the child was in the 15-to-17 age range and sibling-sibling-younger-sibling-focus pairs were not included in the poststratification.) The inclusion of these pair domain covariates led to the use of two sets of control totals in the modeling. Details of the construction of these control totals may be found in Appendix B.

Exhibit 4.1 Definitions of Levels for QDU-Level Calibration Modeling Variables

Age [c]
1: 12–17, 2: 18–25, 3: 26–34, 4: 35–49, 5: 50+ [1]
Gender [c]
1: Male, 2: Female [1]
Group Quarter Indicator [b]
1: College Dorm, 2: Other Group Quarter, 3: Non-Group Quarter [1]
Hispanicity [c]
1: Hispanic, 2: Non-Hispanic [1]
Household Size [c]
Continuous variable - count of individuals rostered with DU.
Household Type (ages of persons rostered within DU) [b]
1: 12–17, 18–25, 26+, 2: 12–17, 18–25, 3: 12–17, 26+, 4: 18–25, 26+, 5: 12–17, 6: 18–25, 7: 26+
Percentage of Owner-Occupied Dwelling Units in Segment (% Owner) [b]
1: 50%–100%, [1] 2: 10%–50%, 3: <10%
Percentage of Segments That Are Black (% Black) [b]
1: 50%–100%, 2: 10%–50%, 3: <10%[1]
Percentage of Segments That Are Hispanic (% Hispanic) [b]
1: 50%–100%, 2: 10%–50%, 3: <10%[1]
Population Density [b]
1: MSA 1,000,000 or more, 2: MSA less than 1,000,000, 3: Non-MSA urban, 4: Non-MSA rural [1]
Quarter [b,c]
1: Quarter 1, 2: Quarter 2, 3: Quarter 3, 4: Quarter 4 [1]
Race (3 levels) [c]
1: white, [1] 2: black, 3: Other
Race (4 levels) [c]
1: white, [1] 2: black, 3: Native American, 4: Asian
Race of Householder [b]
1: Hispanic white, [1] 2: Hispanic black, 3: Hispanic others, 4: Non-Hispanic white,
5: Non-Hispanic black, 6: Non-Hispanic others
Relation to Householder [b]
1: Householder or Spouse, 2: Child, 3: Other Relative, 4: Nonrelative [1]
Segment Combined Median Rent and Housing Value (Rent/Housing) [b,2]
1: First Quintile, 2: Second Quintile, 3: Third Quintile, 4: Fourth Quintile, 5: Fifth Quintile [1]
States [b,c,3]
Model Group 1: 1: Connecticut, 2: Maine, 3: Massachusetts, 4: New Hampshire, 5: New Jersey, 6: New York, 7: Pennsylvania, 8: Rhode Island, 9: Vermont
Model Group 2: 1: Illinois, 2: Indiana, 3: Iowa, 4: Kansas, 5: Michigan, 6: Minnesota, 7: Missouri, 8: Nebraska, 9: North Dakota, 10: Ohio, 11: South Dakota, 12: Wisconsin
Model Group 3: 1: Alabama, 2: Arkansas, 3: Delaware, 4: District of Columbia, 5: Florida, 6: Georgia, 7: Kentucky, 8: Louisiana, 9: Maryland, 10: Mississippi, 11: North Carolina, 12: Oklahoma, 13: South Carolina, 14: Tennessee, 15: Texas, 16: Virginia, 17: West Virginia
Model Group 4: 1: Alaska, 2: Arizona, 3: California, 4: Colorado, 5: Idaho, 6: Hawaii, 7: Montana, 8: Nevada, 9: New Mexico, 10: Oregon, 11: Utah, 12: Washington, 13: Wyoming

[1] The reference level for this variable. This is the level against which effects of other factor levels are measured.
[2] Segment-combined Median Rent and Housing Value is a composite measure based on rent, housing value, and percentage owner-occupied.
[3] The State or district assigned to a particular model is based on Census regions.
[b] Binary variable.
[c] Counting variable.

Exhibit 4.2 Definitions of Levels for Pair-Level Calibration Modeling Variables

Group Quarter Indicator
1: College Dorm, 2: Other Group Quarter, 3: Non-Group Quarter [1]
Household Size
1: DU with 2 persons, 2: DU with 3 persons, 3: DU with >=4 persons
Pair Age (15 levels)
1: 12–17 & 12–17, [1] 2: 12–17 & 18–25, 3: 12–17 & 26–34, 4: 12–17 & 35–49, 5: 12–17 & 50+,
6: 18–25 & 18–25, 7: 18– 25 & 26–34, 8: 18–25 & 35–49, 9: 18–25 & 50+, 10: 26–34 & 26–34,
11: 26–34 & 35–49, 12: 26–34 & 50+, 13: 35–49 & 35–49, 14: 35–49 & 50+, 15: 50+ & 50+
Pair Age (6 levels)
1: 12–17 & 12–17, [1] 2: 12–17 & 18–25, 3: 12–17 & 26+, 4: 18–25 & 18–25, 5: 18–25 & 26+, 6: 26+ & 26+.
Pair Age (3 levels)
1: 12–17 & 12–17, [1] 2: 12–17 & 18+, 3: 18+ & 18+
Pair Gender
1: Male & Female, [1] 2: Female & Female, 3: Male & Male
Pair Race (10 levels)
1: white & white, [1] 2: white & black, 3: white & Hispanic, 4: white & other, 5: black & black,
6: black & Hispanic, 7: black & other, 8: Hispanic & Hispanic, 9: Hispanic & other,
10: other & other.
Pair Race (5 levels)
1: Mixed race pair, 2: Hispanic pair, 3: black pair, 4: white pair, [1] 5: other pair.
Pair Race (4 levels)
1: Mixed race pair or other & other, 2: Hispanic pair, 3: black pair, 4 white pair [1]
Percentage of Owner-Occupied Dwelling Units in Segment (% Owner)
1: 50%–100%, [1] 2: 10%–>50%, 3: 0–>10%
Percentage of Segments That Are Black (% Black)
1: 50%–100%, 2: 10%–>50%, 3: 0–>10% [1]
Percentage of Segments That Are Hispanic (% Hispanic)
1: 50%–100%, 2: 10%–>50%, 3: 0–>10% [1]
Segment-Combined Median Rent and Housing Value (Rent/Housing) [2]
1: First Quintile, 2: Second Quintile, 3: Third Quintile, 4: Fourth Quintile, 5: Fifth Quintile [1]
Population Density
1: MSA 1,000,000 or more, 2: MSA less than 1,000,000, 3: Non-MSA urban, 4: Non-MSA rural [1]
Quarter
1: Quarter 1, 2: Quarter 2, 3: Quarter 3, 4: Quarter 4 [1]
Race of Householder
1: Hispanic white, [1] 2: Hispanic black, 3: Hispanic others, 4: Non-Hispanic white,
5: Non-Hispanic black, 6: Non-Hispanic others

Exhibit 4.2 Definitions of Levels for Pair-Level Calibration Modeling Variables (continued)

State / Region

1: Connecticut, Maine, Massachusetts, New Hampshire, New Jersey, Rhode Island, Vermont,
2: Indiana, Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota, South Dakota, Wisconsin,
3: Alabama, Arkansas, Delaware, District of Columbia, Georgia, Kentucky, Louisiana, Maryland, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Virginia, West Virginia,
4: Alaska, Arizona, Colorado, Idaho, Hawaii, Montana, Nevada, New Mexico, Oregon, Utah, Washington, Wyoming,
5: New York,
6: Pennsylvania,
7: Michigan,
8: Illinois,
9: Ohio,
10: Florida,
11: Texas,
12: California

States [3]

Model Group 1: 1: Alabama, 2: Arkansas, 3: Connecticut, 4: Delaware, 5: District of Columbia, 6: Florida, 7: Georgia, 8: Kentucky, 9: Louisiana, 10: Maine, 11: Maryland, [1] 12: Massachusetts, 13: Mississippi, 14: New Hampshire, 15: New Jersey, 16: New York, 17: North Carolina, 18: Oklahoma, 19: Pennsylvania, 20: Rhode Island, 21: South Carolina, 22: Tennessee, 23: Texas, 24: Vermont, 25: Virginia, 26: West Virginia

Model Group 2: 1: Alaska, 2: Arizona, [1] 3: California, 4: Colorado, 5: Idaho, 6: Illinois, 7: Indiana, 8: Iowa, 9: Hawaii, 10: Kansas, 11: Michigan, 12: Minnesota, 13: Missouri, 14: Montana, 15: Nebraska, 16: Nevada, 17: New Mexico, 18: North Dakota, 19: Ohio, 20: Oregon, 21: South Dakota, 22: Utah, 23: Washington, 24: Wisconsin, 25: Wyoming

Pair Relationship Associated with Multiplicity

1: Parent-child (12-14)*
2: Parent-child (12-17)*
3: Parent-child (12-10)*
4: Parent*-child (12-14)
5: Parent*-child (12-17)
6: Parent*-child (12-20)
7: Sibling (12-14) - sibling (15-17)
8: Sibling (12-17) - sibling (18-25)
9: Spouse-spouse/partner-partner
10: Spouse-spouse/partner-partner with children under 18

[1] The reference level for this variable. This is the level against which effects of other factor levels are measured.
[2] Segment-combined Median Rent and Housing Value is a composite measure based on rent, housing value, and percentage owner-occupied.
[3] The States or district assigned to a particular model is based on combined Census regions.
* The pair member focused on.

5. Definition of Extreme Weights

An important feature of the generalized exponential model (GEM) is the built-in provision of extreme value (ev) treatment. For this purpose, sampling weights are classified as extreme (high or low) if they fall outside the interval, median ± 3*interquartile range (IQR), for some prespecified domains defined usually by design variables corresponding to deep stratification.[3] The critical values for low and high extreme values will be denoted by $b_{k(l)}$ and $b_{k(u)}$, respectively. Within GEM modeling, these critical values were defined as median ± 2.5*IQR, which were conservative when compared with the commonly used standard of median ± 3*IQR. This is because in order to better prevent the adjusted weights from crossing the standard boundary, weights near but below it (i.e., those that have the most potential to become extreme) were treated as extreme by GEM as well.

For implementing extreme value control via GEM, the variable m_k was defined as the minimum of ($b_{k(u)} / w_k$) and 1 for high extreme weights, and the maximum of ($b_{k(l)} / w_k$) and 1 for low extreme weights, where w_k represents the sampling weight before adjustment, and ($b_{k(u)}$, $b_{k(l)}$) denote the critical values for the extreme weights. (Note that under this definition, for high extreme weights, the more extreme the weight is, the smaller m_k will be; conversely for low extreme weights, the more extreme the weight is, the bigger m_k will be.) Non-extreme weights had a value of 1 for m_k. The upper and lower bounds for the adjustment factors were defined, respectively, as the product of m_k and the upper and lower boundary parameters of GEM. GEM allows inputs of up to three different upper and lower boundary parameters (L_1 and U_1, L_2 and U_2, L_3 and U_3) for high, non-, and low extreme weights. By applying a small upper boundary parameter for high extreme weights and a large lower boundary parameter for low extreme weights, the extreme weights can be controlled in the modeling process.

5.1 Questionnaire Dwelling Unit (QDU) Extreme Weight Definition

For the QDU-level weight adjustment, extreme weights are defined using a nested hierarchy of six domains:

1. State,
2. Field interviewer region,
3. State by household type,

Levels of household type indicate whether household has members who are youth, young adult, or adult, where youth signifies 12 to 17, young adult 18 to 25, and adult 26+.

I. Youth, Young Adult, Adult;
II. Youth, Young Adult;
III. Youth, Adult;
IV. Young Adult, Adult;
V. Youth Only;
VI. Young Adult Only;

[3] Deep stratification refers to the stratification that was used in the sample design. In the case of the 2002 NSDUH, deep stratification refers to the cross-classification of FI region by age group.

VII. Adult Only.

4. Field interviewer region by household type,
5. State by household type by household size (1, 2, 3, 4+),
6. Field interviewer region by household type by household size.

The hierarchy is used to satisfy the minimum of 30 observations for defining the boundaries for extreme values. If this sample size requirement is not met at the lower level, then the next level up in the hierarchy is used.

5.2 Person Pair Extreme Weight Definition

The pair selection probability is a function of the selection probability of each person in the pair given by formula (2.1) or (2.6), depending on the sum of the person selection probabilities within the household as discussed in Section 2.1. This probability could be very small if the selection probabilities of individual members are small. For example, consider a selected DU (ID=FL22180127) from the 2002 NSDUH. This DU gave rise to a selected pair of respondents, both aged 50+. The selection probability for a respondent aged 50+ was 0.084345, and, using the formula (2.6) in Chapter 2, the pair-selection probability was computed to be 0.0001757. Therefore, the inverse of the probability, the pair-level design weight, was 5691.21. Thus, a small pair selection probability could create a high initial weight, which is the product of the screener dwelling unit (SDU) weight and the person-pair design-based weight.

As mentioned in the introduction, it turns out to be difficult to select suitable domains for defining extreme weights for pair-level data. However, as was done for 1999–2001, the extreme weight definition was based on the following hierarchy of domains.

1. Pair age group (with three age categories, 12 to 25, 26 to 49, and 50+) by number (0, 1, 2+) of persons aged 12 to 25 in the household,
2. State cluster (with 5 levels [explained below]) by Pair age group by number (0, 1, 2+) of persons aged 12 to 25 in the household,
3. State cluster (with 3 levels [explained below]) by Pair age group by number (0, 1, 2+) of persons aged 12 to 25 in the household,
4. State by pair age group by number of persons aged 12 to 25 (0, 1, 2+) in the household.

The hierarchy is used to satisfy the minimum of 30 observations for defining the boundaries for extreme values. If this sample size requirement is not met at the lower level, then the next level up in the hierarchy is used.

We now briefly introduce the considerations behind the above definition for extreme weight domains. The sample design prespecified the person-level selection probability within State by five age groups (12 to 17, 18 to 25, 26 to 34, 35 to 49, 50+). Age groups 12 to 17 and 18 to 25 have a relatively similar selection probability, and the same is true for age groups 26 to 34 and 35 to 49. The 50+ group, however, has a quite different selection probability from the other groups. Furthermore, since the 12 to 17 and 18 to 25 age groups have large selection probabilities, they have a very high chance of being selected if the household has persons in these age groups. Therefore, the number of persons aged 12 to 25 in the household has a significant impact on the type of pair selected and the pair selection probability. Taking into consideration these design-related features, a suitable domain to define the pair-level extreme weight seems to be given by State by pair age group by number of persons aged 12 to 25 in the household.

The hierarchy of domains mentioned above was used to satisfy the minimum of 30 observations; however, it was found that for many ev domains the minimum sample size requirement was not met. To get around this problem, States were grouped into a small number of clusters, say 3 or 5. The assignment of States to clusters was determined by the clustering algorithm in PROC CLUSTER in SAS, where the clustering variable was defined as the average person-level weight (ANALWT) for each of the 5 age groups within each State. The choice of the average person-level weight for each group for each State was motivated from the objective of finding a single variable that would reflect the design-based difference in pair-selection probabilities across States. Even with clustering of States, the ev domain sample size may be insufficient, so the most general level of the hierarchy, the national level, is required. Even at the national level, we had to collapse some pair age categories in forming domains of reasonable sample size to define extreme weights. More specifically, for the national level, we collapsed all levels of number of persons aged 12 to 25 for the pair age groups (50+, 50+), and (26 to 49, 50+). In addition, levels 1 and 2+ of number of persons aged 12 to 25 were combined for the pair age group (26 to 49, 26 to 49).

6. Editing and Imputation of Pair Relationships, Multiplicity Factors, and Household-Level Person Counts for Poststratification

6.1 Introduction

"Pair data" are used to study outcome variables among members of the same household. These outcome variables are measured using the "pair relationship," the relationship between selected pair members. For these analyses, the outcome variables may be at either the person level or the pair level. The most common type of analysis is the person-level analysis, where the inferential population is defined by one of the pair members. This pair member is the "focus" pair member. An example of an outcome at the person level is the proportion of youths who use drugs and whose parents report talking to them about drugs, where the focus is on the youth in a parent-child pair. An example at the pair level is child-parent drug behavior for all possible parent-child pairs (within the youth's age group). Knowledge of the pair relationship and the inferential population gives rise to the "pair domain."

For analyses at the pair level, the pair domain is completely defined by the pair relationship, whereas the pair domain for a person-level analysis depends upon which pair member is the focus. "Multiplicity" is an issue that arises in the analysis of pair data where the analysis is at the person level for a given pair domain. Several pairs in the household could be associated with the same person. Consider the previous example where we are interested in the proportion of children who use drugs and whose parent reports talking to them about drugs. In this case, if the household has two parents, the selected child has two inclusion possibilities (one with each parent) in the set of all such parent-child pairs. Since children form the target population for this example, it is desirable to assign one observation per child. A reasonable way to achieve this is to take an average of the two responses, which together correspond to the two pairs associated with the child (i.e., one for each parent in this example). In other words, the response for each child-parent pair from two-parent households is divided by the number of parents; this divisor is known as the "multiplicity factor." The multiplicity problem does not arise if there is only one inclusion possibility (e.g., a single-parent household, if the child is the focus), or if the analysis is a pair-level analysis, (e.g., child-parent pair drug behavior).

To illustrate how multiplicities appear in the definitions of parameters and estimates, consider estimation of the total number of children who used drugs in the past year, where a parent reported talking to them about drugs. Let $y_{hip}(d)$ be defined as the drug-related behavior outcome for pair p containing the individual i belonging to domain d in household h. Now, for the population of all individuals who belong to the domain d, the total parameter is defined as (Chromy & Singh, 2001)

$$\tau_y(d) = \sum_{h=1}^{H} \sum_{i=1}^{N_h(d)} \sum_{p=1}^{M_{hi}(d)} \frac{y_{hip}(d)}{M_{hi}(d)}$$

i.e., total of averages over pairs (p) associated with the individual i, over all i in domain d and in the household h. Here $M_{hi}(d)$ denotes the multiplicity (i.e., the number of pairs associated) for the person i in domain d, and $N_h(d)$ can be thought of as the multiplicity count for the household h, (i.e., the number of persons in the household that are in domain d). This latter multiplicity count is equivalent to the household-level person count described in the next paragraph. For the sake of simplicity, the weights are not shown in the above estimator.

In order to obtain more stable pair-level analysis weights, in the respondent pair poststratification (ps) step, in addition to the predictor variables proposed for all previous generalized exponential model (GEM)[4] adjustment steps, ten covariates derived from ten pair domains were included in the weight adjustment process. (The ten pair domains are identified in the following section.) Each covariate was defined by the appropriate pair relationship divided by its associated multiplicity. In this ps step, for these ten pair domains, the nonresponse (nr)-adjusted weights were poststratified to the final questionnaire dwelling unit (QDU) weights. The household-level person counts, which are counts of the number of pairs in the household belonging to a given pair domain, were used to form the control totals in the ps step for these domains. For other domains, the control totals were formed by the screener dwelling unit (SDU) weights from all the possible screener pairs associated with the number of possible pairs in the dwelling unit.

In the process of setting up variables for analyses at the pair level, three types of variables, which are not weights, required editing and imputation. The procedures associated with these three types of variables are referred to as stages. Stage one refers to the creation and imputation of the variables that identify the pair relationships. The multiplicity and household-level person counts that are described in the preceding paragraph were created and imputed in stages two and three respectively. Missing values in all three stages were imputed using the semiparametric predictive mean neighborhoods (PMN) imputation procedure, which uses predicted means from models to find donors in a nearest neighbor hot deck. The hot deck is described in Appendix M, and the PMN procedure is described in detail (in its application to drug use variables) in Appendix N.

6.2 Stage One: Creation and Imputation of Pair Relationships

6.2.1 Editing the Household Roster of Each Pair Member

Prior to the identification of the relationships between selected pair members, a key step was to edit the questionnaire household rosters for each pair member. This involved identifying situations where the relationship listed in the roster, for a particular roster member, was not possible given the roster member's age and relationship to the respondent. In many cases, this resulted in setting the relationship code to bad data, and sometimes the roster member's age was also set to bad data. In general, no effort was made to try to match values of roster-derived

[4] The GEM macro, which was written in SAS/IML® software, was developed at RTI for weighting procedures.

household composition variables between pair members, since interviews of the different members of the same household could have taken place at different times. However, information from other pair members was sometimes used to change a relationship code from one value to another, instead of setting the relationship code to bad data. The editing of the household roster is described in detail in Chapter 8 of the Imputation report (Grau et al., 2004).

6.2.2 Creation of Pair Relationship Variable, PAIRREL

Because the creation of the multiplicity factors was not automatic, multiplicities could not be created for all possible pair relationships. The following pair relationships were considered "of interest," requiring the creation of multiplicities in each case.

a. Parent-child, child aged 12 to 14
b. Parent-child, child aged 12 to 17
c. Parent-child, child aged 15 to 17
d. Parent-child, child aged 12 to 20
e. Sibling-sibling, younger aged 12 to 14, older aged 15 to 17
f. Sibling-sibling, younger aged 12 to 17, older aged 18 to 25
g. Spouse-spouse (includes partner-partner), with children under 18
h. Spouse-spouse (includes partner-partner), with no children under 18

Even though these pair relationships were of the most interest, all types of pairs were selected. The identification of the relationships was limited by the relationship codes that were available: parent, child, grandparent, grandchild, sibling, spouse, live-in partner, roommate, parent-in-law, child-in-law, roommate, boarder, other relative, and other non-relative. This precluded the possibility of identifying an uncle-nephew relationship, for example. The various pair relationships that could be identified are given in the variable PAIRREL, the levels of which are summarized in Table 6.1. The levels in PAIRREL do not correspond exactly with those given above, but the relevant pair relationships can be derived from the value of PAIRREL. For example, a value of PAIRREL = 3 indicates that, among the pair relationships given above, the pair relationship was a parent-child pair with a child between 12 and 20 years old.

The process of identifying the pair relationships was a two-step process: (1) match the household rosters of the pair members; and (2) determine the pair relationship using the relationship codes and ages of the matched rosters.

Table 6.1 Levels of the Variable PAIRREL

Value of PAIRREL	Interpretation	Domain of Interest?
1	The respondent is part of a parent-child (12–14) pair.	Yes
2	The respondent is part of a parent-child (15–17) pair.	Yes
3	The respondent is part of a parent-child (18–20) pair.	Yes, indirectly
4	The respondent is part of a parent-child (21+) pair.	No
5	The respondent is part of a sibling (12–14)-sibling (15–17) pair.	Yes
6	The respondent is part of a sibling (12–17)-sibling (18–25) pair.	Yes
7	The respondent is part of another sibling-sibling pair.	No
8	The respondent is part of a spouse-spouse[1] pair, with children in the household under the age of 18.	Yes
9	The respondent is part of a spouse-spouse[1] pair, with no children in the household under the age of 18.	Yes
10	The respondent is part of a spouse-spouse[1] pair, but it is unclear whether children under the age of 18 in the household belong to the pair.	Yes
11	The respondent is part of a grandparent-grandchild pair.	No
12	The respondent is part of another clearly identifiable pair.	No
13	The respondent is part of a pair that is not clearly identifiable, but it is clear from the relationship codes that it is not within codes 1 through 11.	No
14	The respondent is part of a pair that is not clearly identifiable, and it could be any pair relationship.	Maybe
15	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 1 or 12.	Maybe
16	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 2 or 12.	Maybe
17	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 3 or 12.	Maybe
18	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 4 or 12.	No
19	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 5 or 12.	Maybe
20	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 6 or 12.	Maybe

(continued)

Table 6.1 Levels of the Variable PAIRREL (continued)

Value of Pairrel	Interpretation	Domain of Interest?
21	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 7 or 12.	No
22	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 8 or 12.	Maybe
23	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 9 or 12.	Maybe
24	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 8, 9, or 12.	Maybe
25	The respondent is part of a pair that is not clearly identifiable, but it could be pair codes 11 or 12.	No
99	The respondent is not a member of a pair.	No

[1] The pair relationship labeled "spouse-spouse" includes partner-partner pair relationships.

6.2.2.1 Matching the Household Rosters

To match the household rosters of the pair members, let the pair members be identified as pair member "A" and pair member "B." For the household roster of pair member A, it was necessary to determine which listed household member in A's roster corresponded to the other selected pair member. The same had to be done for pair member B. This was accomplished using the age and gender of the pair members, in addition to a variable (hereafter referred to as MBRSEL) that was supposed to identify the roster member corresponding to the other selected pair member. In a perfect setting, the questionnaire age and gender of pair member B (AGE and IRSEX respectively) would have corresponded exactly to the age and gender entered for one of the members of pair member A's household roster (RAGE and RSEX). Moreover, the value of MBRSEL for this roster member would have been 1, and the value of MBRSEL for all other roster members would have been 0. In this perfect setting, exact matches with exactly one MBRSEL=1 correctly identifying the other pair member would have also been found with pair member B's roster. This did not always occur, of course, so some effort was required to determine the roster member most likely to correspond to the other selected pair member.

In fact, the quality of the match varied depending upon the quality of the roster entries, and the time between interviews. There are a number of if-then-else conditions, called priorities (due to the hierarchical nature of the conditions), each of which gave a pair match that was considered valid in the vast majority of cases. These conditions are given in Appendix O. In general, the conditions matched IRSEX and AGE for the one pair member against the age and sex of the roster members in the other pair member's roster, using MBRSEL to help identify the appropriate roster member. These conditions in general terms are given in Table 6.2. It was necessary that at least one of the two pair members have a match as good as that given below.

Table 6.2 Measures of Quality of Definitive Roster Matches

Measure Number	Description
0	Age and gender matched exactly, with exactly one MBRSEL correctly identifying the other pair member
1	Age and gender matched exactly, with MBRSEL correctly identifying the other pair member, but there was more than one MBRSEL[1]
2	Age within one, gender matched exactly, with exactly one MBRSEL correctly identifying the other pair member
3	Age within two, gender matched exactly, with exactly one MBRSEL correctly identifying the other pair member
4	Age and gender matched exactly, with MBRSEL missing for all roster members
5	Age matched exactly, gender off, with exactly one MBRSEL correctly identifying the other pair member
6	Age within one, gender matched exactly, with MBRSEL correctly identifying the other pair member, but there was more than one MBRSEL[1]
7	Age within two, gender matched exactly, with MBRSEL correctly identifying the other pair member, but there was more than one MBRSEL[1]
8	Age within one, gender matched exactly, with MBRSEL missing for all roster members
9	Age within two, gender matched exactly, with MBRSEL missing for all roster members[3]
10	Age within 10, gender matched exactly, with exactly one MBRSEL correctly identifying the other pair member[2]

[1] Since the 2001 survey, it was technically impossible to identify more than one roster member as the "other pair member selected," resulting in either 0 or 1 MBRSEL for each responding pair. As a result, measures #1, #6, and #7 did not occur in the 2002 survey.

[2] For pairs where one pair member had a match corresponding to measures #9 or #10, if the other pair member had a match no better than measure #9, an additional requirement was implemented where the reported household sizes for both pair members had to be equal to 2.

Given that at least one side had a match according to one of the measures given in Table 6.2, the other side could have a match that was weaker (i.e., not definitive), using the measures in Table 6.3. Additional columns are given in Table 6.3, giving the weakest match allowed (as denoted by the measure) for the other pair member. The new column titled "In Code" shows the weakest measure allowed in the code, and the new column titled "Observed" shows the weakest measure that was actually observed for the other pair member.

In the cases where a single roster member had to be selected among duplicates (measures #14, #15, #17, and #18), where the duplicates had the same relationship code, it was necessary that the relationship codes be limited to child or sibling.

Table 6.3 Measures of Quality of Roster Matches that are not Definitive, Given One Side had a Definitive Match (as Given by the Conditions Given in Table 6.2)

Measure Number	Description	Weakest Measure Allowed for Other Pair Member	
		In Code	Observed
11	Age within 10, gender matched exactly, with MBRSEL missing for all roster members, provided another roster member with a closer age could not have been chosen	8	0
12	Everything missing, but the other pair member had good data	9	2
13	Age missing, gender matched exactly, with exactly one MBRSEL correctly identifying the other pair member, household sizes equal	9	0
14	Age and gender matched exactly for two roster members, both with the same relationship code, with two MBRSELs identifying the two roster members (one was randomly selected) [1]	8	Not observed
15	Age and gender matched exactly for two roster members, both with the same relationship code, but MBRSEL was missing for all roster members (one was randomly selected).	8	Not observed
16	Age and gender matched exactly for two roster members, where one was a sibling or child but the other was not, and MBRSEL was missing for all roster members (the one selected depended upon the relationship code given by the other pair member).[2]	8	4
17	Age within one and gender matched exactly for two roster members, both with the same relationship code, with two MBRSELs identifying the two roster members (one was randomly selected)[1]	8	Not observed
18	Age within one and gender matched exactly for two roster members, both with the same relationship code, with two MBRSELs identifying the two roster members (one was randomly selected)[3]	8	Not observed
20	No matches possible, but relationship codes indicate the pair is not a part of a domain of interest	As with other pair member	As with other pair member

Table 6.3 Measures of Quality of Roster Matches that are not Definitive, Given One Side had a Definitive Match (as Given by the Conditions Given in Table 6.2) (continued)

Measure Number	Description	Weakest Measure Allowed for Other Pair Member	
21	Age matches exactly, gender off, with MBRSEL missing for all roster members	9	0
22	No matches possible	9	5

[1] Since the 2001 survey, it was technically impossible to identify more than one roster member as the "other pair member selected," resulting in either 0 or 1 MBRSEL for each responding pair. As a result, measures #14 and #17 did not occur in the 2002 survey.

[2] Measure #16 actually refers to a number of situations where a variety of conditions were used to determine which of the multiple matches should be considered the correct match.

[3] Since the 2001 survey, it was technically impossible to identify more than one roster member as the "other pair member selected," resulting in either 0 or 1 MBRSEL for each responding pair. As a result, measures #14 and #17 did not occur in the 2002 survey.

In some cases, due to the poor quality of the rosters of the pair members, it was not possible to locate the listed household member in A's roster that corresponded to pair member B, and vice versa. The determination of the pair relationships for these cases was left to imputation. Even when a pair of roster members was successfully identified, it was not always possible to successfully determine the pair relationship, as is pointed out in the next section.

6.2.2.2 Determining the Pair Relationship Using the Relationship Codes of the Matched Rosters

Once the pair was identified, two observations per household resulted, each with a relationship code corresponding to the other selected pair member. The relationship codes for these two observations had to be matched to determine the pair relationship. For example, suppose a 15 year old and a 38 year old were selected. If the 38 year old was identified as the parent on the 15 year old's roster, and the 15 year old was identified as child of the 38 year old on the 38 year old's roster, then the pair relationship would be identified according to the levels of PAIRREL given in Table 6.1 as PAIRREL = 2. Thus, these two individuals would belong to the following pair relationships of interest: child (15 to 17)-parent, child (12 to 17)-parent, and child (12 to 20)-parent. As noted earlier, the pair relationship of interest can be derived from the values of PAIRREL. In particular, the child (12-17)-parent and child (12 to 20)-parent domains can be derived from pair relationships created using 12 to 14 year olds, 15 to 17 year olds, and 18 to 20 year olds, the levels referenced in the levels of PAIRREL. Moreover, the overall spouse-spouse domain can be derived from the two spouse-spouse pair relationships with and without children.[5]

As with the procedure used to match the household rosters, a series of if-then-else conditions were used to identify the relationship between pair members. These conditions, also called priorities because of their hierarchical nature, used ages and relationship codes to identify

[5] The spouse-spouse pair relationship includes partner-partner pair relationships.

the pair relationships, and are summarized in Appendix O. In a perfect setting, like the example given in the first paragraph of this section, the relationship codes would be nonmissing and in agreement between the pair members. In some instances, however, either the relationship codes were missing or did not agree across the pair members. The detailed conditions given in Appendix O present a method for interpreting the relationship codes in such cases.

A few points summarize the strategies used to identify a pair relationship in an imperfect setting:

1. If a relationship code was missing on one side but not on the other, the pair relationship was assumed to be identified by the nonmissing relationship code. The exception to this rule occurred if the identified relationship was child-parent with a child under 18, the "parent" was less than 10 years older than the child, and the "parent" answered the parenting experiences question (FIPE3) by saying that the other respondent was not his/her child. In this case, the nonmissing relationship code was considered spurious, and the relationship was left missing.

2. If it was not possible to definitively determine the relationship between the pair members using the relationship codes, but the relationship codes on both sides indicated that the unknown pair relationship was not a relationship of interest, the pair relationship was identified as such and no imputation was required. For example, if pair member "A" identified the pair member "B" as a "boarder," but pair member "B" identified pair member "A" as "other relative," the relationship was not a relationship of interest and code '13' would have been applied in the variable PAIRREL.

3. If it was not possible to definitively determine the relationship between the pair members using the relationship codes, but a parent-child relationship was possible given the relationship code in one of the pair member's rosters, the FIPE3 variables were used to assist in the determination of a pair relationship. An example of a case where this would have been useful: a pair member who was a stepparent refers to his/her stepchild as "child," but the child refers to the stepparent as "other non-relative." Membership in a parent-child relationship where the child was under 18 was indicated if the stepparent answered FIPE3 affirmatively, thereby entering the parenting experiences module. On other hand, if the stepparent answered FIPE3 negatively, then the stepparent was not considered the parent. A third scenario arose if the FIPE3 answer was not given. In this case, a parent-child relationship was assumed if the stepparent was legally married, and the child identified the spouse of the other pair member as "parent."

The quality of the match for PAIRREL levels 1 through 25 is indicated by the variable RELMATCH, the levels of which are summarized in Table 6.4.

In general, imputation was required for values of RELMATCH of 0 or 4, or if PAIRREL = 10. PAIRREL = 10 was a special case, since it was clear that a relationship "of interest" would always have been involved. (The pair relationships "of interest" are outlined in Section 6.2.2.) For this value of PAIRREL, the value of RELMATCH was equal to 1 or 2. However, imputation was still required, since it was not clear whether children were in the household. The number of cases that were matched or not matched, as indicated by the RELMATCH variable (or PAIRREL = 10), is given in Table 6.5 for the 1999–2002 surveys. As is apparent from this table, the amount of imputation required decreased from earlier surveys to more recent surveys. This was due to the improvement, with each survey, in the quality of the rosters. The improvement from the 1999 survey to the 2000 survey could be traced to the introduction of a requirement of exactly one self, and the improvement from the 2000 survey to the 2001 and 2002 surveys could be traced to the introduction of a handful of consistency checks. The attributes of the roster, including the consistency checks, are described in Chapter 8 of the imputation report (Grau et al., 2004).

Table 6.4 Values of PAIRREL that Correspond to the Levels of the Variable RELMATCH

Value of RELMATCH	Values of PAIRREL	Interpretation
0	14	FAILURE: The relationship was not identifiable, and could have been a relationship of interest.
1	1–9, 11–13	SUCCESS: The relationship was clearly identifiable using information from both pair members, or was unmistakably not a relationship of interest.
1	10	FAILURE: A spouse-spouse[1] relationship was definitively established using information from both pair members, but it was unclear whether the pair had children in the household.
1.5	8	SUCCESS: A spouse-spouse[1] relationship was definitively established using information from both pair members, and children under 18 were in both rosters. Relationship codes on one side indicated children belonged to the pair; on the other side the relationship codes corresponding to the children were missing.
2	1–13	SUCCESS (unless PAIRREL=10): The relationship was clearly identifiable using information from one pair member, while the relationship code from the other pair member was missing.
3	1, 2, 3, 8, 13	SUCCESS: Relationship information was conflicting between the pair members, but conclusions were drawn anyway for some parent-child pairs and some spouse-spouse[1] pairs using either information outside the household roster, or logical reasoning.[2]

Table 6.4 Values of PAIRREL that Correspond to the Levels of the Variable RELMATCH (continued)

Value of RELMATCH	Values of PAIRREL	Interpretation
4	15-25	FAILURE: Relationship was not identifiable. Information was in conflict between the pair members, where one pair member indicated relationship of interest and the other did not. However, ages supported a relationship of interest (may be used to limit imputation).

[1] The pair relationship labeled "spouse-spouse" includes partner-partner pair relationships.

[2] In the case of potential parent-child pairs, further evidence that a parent-child relationship was involved or not involved was obtained by looking at the FIPE3 variable, by whether a stepparent had a spouse that corresponded to a child's parent, or by the ages of the respondents. For spouse-spouse relationships, two situations occurred: in the case where the respondents were not legally married, the children of one pair member were considered the children of the pair in the household, even though they were not identified as such by the other pair member. In the case where only one pair member referred to the other as "married" or "living together as though married," if both had the same children they were considered "spouse-spouse with children." The other pair member was usually referred to as a "roommate" or "other nonrelative."

Table 6.5 Frequencies of the Levels of the Variable RELMATCH, 1999–2002 Surveys

RELMATCH	Frequency, 1999 survey	Frequency, 2000 survey	Frequency, 2001 survey	Frequency, 2002 survey
0	80 (0.54%)	43 (0.27%)	14 (0.09%)	22 (0.11%)
1 (PAIRREL ≠ 10)	14,064 (94.48%)	15,241 (96.22%)	15,411 (97.57%)	19612 (97.87%)
1 (PAIRREL = 10)	28 (0.19%)	45 (0.28%)	43 (0.27%)	27 (0.13%)
2 (PAIRREL ≠ 10)	460 (3.09%)	274 (1.73%)	109 (0.69%)	86 (0.43%)
2 (PAIRREL = 10)	8 (0.05%)	5 (0.03%)	1 (0.01%)	2 (0.01%)
3	80 (0.54%)	100 (0.63%)	87 (0.55%)	127 (0.63%)
4	150 (1.01%)	128 (0.81%)	123 (0.78%)	157 (0.78%)

6.2.3 Creation of Covariates for Imputing Pair-Level Variables

For pairs where the relationship was not clear due to missing pieces of the household roster, or where pairs could not be determined because the relationship codes did not match, imputation was required. In stages two and three, imputation was also required for missing multiplicities and household-level person counts. In all three stages, the PMN method was used to impute missing values, which required the fitting of models. Since the imputation was performed at the household level rather than at the respondent level, it was necessary to have classing variables (i.e, variables forming imputation classes) and model covariates that were defined at the household level. Segment-level covariates were used for this purpose, since they were automatically defined at the household level, using external information that was constant regardless of when the interviews were conducted. However, it would also be useful to have information from the questionnaire. Logical choices for questionnaire-derived variables would be the household composition variables IRHHSIZE (household size), IRKID17 (number in household under the age of 18), IRHH65 (number in household aged 65 years or over), and IRFAMSKP (indicator whether other family members in household). However, because interviews between pair members could have been conducted at different times, these variables are not necessarily consistent across pair members. New count variables were needed that were consistent across the pair members within a household, which used the screener information to reconcile disagreements between pair members. These variables were created in two steps: (1) create the count variables for each pair member, and (2) attempt to reconcile disagreeing values between pair members. The following sections describe these two steps in the creation of household size, household composition age count variables, and household composition age count variables for males only, each of which were consistent across pair members. These variables also had to be created for respondents who were not part of a pair, for the purposes of creating and imputing the household-consistent person counts of various domains.

6.2.3.1 Household Size

The new variable created to represent a household size that was consistent across the pair members is called HHSIZE. The first step was to compare the edited household size, TOTPEOP, between pair members. If the values for TOTPEOP agreed across pair members, and were both nonmissing and greater than 1, then HHSIZE was simply set to that value. There were two ways that TOTPEOP would disagree across pair members. In the first case, if the count for one pair member was missing, and the count for the other was not and was greater than 1, a natural choice for HHSIZE would have been the nonmissing value. In the second case, the household size counts disagreed across pair members. The tools used to determine the final value of HHSIZE in these cases included the reported and edited household size variables previously mentioned, as well as other measures of household size and "quality of roster" measures. These "other measures" included the screener household size and two sums of total valid ages within a pair member's roster. The first sum was a simple total count of the number of roster members with valid ages, obtained by summing the counts within certain age groups. The second sum adjusted the first by accounting for the minimum number within each age category given the questionnaire ages of the two pair members. It differed from the first if the number of valid ages in a given age category was less than the minimum possible in that age category, given the ages of the two pair members selected. For example, suppose a household roster had one 12 to 17 year old, but two 12 to 17 year olds were selected. The second sum was determined by replacing the number of 12 to 17 year olds by the minimum number possible, two. An additional situation occurred where the household size counts could not be easily determined by looking at both pair members. If the counts for both pair members were missing, the screener household size was used to define HHSIZE. In some cases, disagreement between pair members with regard to the true household size could not be easily resolved. The screener household size did not support either household size in these cases, and the age counts mentioned above also did not resolve the disagreement. A decision had to be made as to which pair member's household size should be believed. This decision depended upon the "quality of the roster," where the household size was determined by the pair member with a better "roster quality." One obvious way to measure roster quality was by noting the number of cases where the ages, relationship codes, or genders were missing in the roster. Clearly, if a roster was missing one or more of these three variables for some of the roster members, the roster was of "poorer quality" than a roster with these variables nonmissing for all roster members.

If only one household member was selected as a respondent, known colloquially as a "nonpair household," the rules for creating HHSIZE were the same as those that were used if two household members were selected in a pair, but only one of the pair members had a nonmissing, acceptable value for a reported household size, with one important exception. If only one household member was selected as a respondent, it was obviously permissible to have a reported household size of 1, whereas in a selected pair a reported household size of 1 was considered "bad data," necessitating the use of the screener household size as the source variable for HHSIZE.

In summary, the variables used to determine HHSIZE included, for each pair member, the reported and edited household sizes, the number of cases with valid ages in the roster, the number of cases with valid ages with the count in some age categories replaced by the minimum possible in that age category, and a quality of roster count of the number of roster members with missing information. The screener household size, which was the same for each pair member, was also used. Using all of these tools, HHSIZE did not have any missing values in the 2002 survey, nor did it in surveys from previous years. General points about the creation of the household size variable are given in Appendix P.

6.2.3.2 Household Composition Age Count Variables

It would seem logical to assert that the ages of other household members would be good predictors for the domain to which a pair might belong. Such variables would also be important for imputing multiplicity and household-level domain counts. The household-consistent age counts were limited to the following age ranges: under 12 years old, 12 to 14 years old, 15 to 17 years old, 12 to 17 years old, 12 to 20 years old, 18 to 25 years old, 26 to 34 years old, 35 to 49 years old, and 50 years old or older. These variables were called AGE011, AGE1214, AGE1517, AGE1217, AGE1220, AGE1825, AGE2634, AGE3549, and AGE50P respectively.

The first step in this process was to count the nonmissing ages for roster members in the household for each pair member. In some cases, it was necessary to adjust the count, since the ages could not match exactly. For example, suppose a 38 year old and a 17 year old were interviewed, where the 17 year old was interviewed first. Suppose also that the 17 year old turned 18 (i.e., had his 18th birthday) before the 38 year old was interviewed. Hence, the 17 year old would have had an age of 18 in the 38 year old's roster. Because the ages for the pair domains were defined at the time of each pair member's interview, the ages of interest for pair domains would have been 17 and 38. Hence, it was necessary to account for this, by creating a new roster age variable which matched the age given in the other pair member's questionnaire. The age counts using this new roster age variable were equivalent to subtracting 1 from the previously obtained 18-to-25 count and adding 1 to the previously obtained 12-to-17 count in the 38 year old's roster. These adjustments were made for all cases where a match was made between one pair member's roster and another pair member's interview age and sex, and the ages did not match exactly.

If no roster ages were missing, the sum of these counts was equal to the edited household size, TOTPEOP. Note that the raw household size was not considered here, since the counts were obtained from an edited roster. As with household size, a series of if-then-else conditions were used to obtain the most likely count within each age group. These conditions are called priorities due to their hierarchical nature. If the appropriate count was ambiguous due to disagreement between the pair members, the quality of the roster and the age of the respondent (in that order) were used to determine the appropriate count. The roster quality was determined by the number of bad or missing roster entries (as indicated in the previous section) and the quality of the match between the pair member's roster and the other pair member's questionnaire age and sex. If only one household member was selected as a respondent, the rules were the same as when two household members were selected in a pair, but only one of the pair members had nonmissing data for the roster ages, with one important exception. When determining minimum possible counts for various age groups, it was obviously not necessary to incorporate information from another pair member to increment the minimum for that pair member. General points about the creation of the age variables are given in Appendix P.

6.2.3.3 Household Composition Age Count Variables for Males Only

For some pair variables, particularly spouse-spouse pairs, knowledge of the gender of the roster member was important in imputing missing values. In a similar manner to that used in the creation of the household composition age count variables, variables counting the number of males within the given age ranges were created. Disagreements between pair members were resolved in a similar manner to what was done with the household composition age count variables, as described in the previous section. For a given age range, the number of females could be obtained by subtracting the number of males from the total number within that age range. The names of the male age counts are MALE011, MALE1214, MALE1517, MALE1217, MALE1220, MALE1825, MALE2634, MALE3549, and MALE50P.

6.2.4 Creation of Imputation-Revised Pair Relationship Variable, IRPRREL

It was not always possible to definitively determine the pair relationship for the selected pair. In some cases, the relationship codes between the two pair members could not be reconciled. In other cases, no information was available about the type of pair relationship. This section describes how those missing pair relationships were imputed using the PMN method described in Appendix N. In this section, the application of the PMN method to the imputation of pair relationships is described. Since only the pair relationship was imputed, the imputation was univariate in the sense that no sequential models were necessary. However, in some cases the outcome variable was multinomial, which meant that matching was done on more than one predicted mean for each recipient pair.

6.2.4.1 Setup for Model Building

Pair relationships varied greatly according to the age of the respondent. Table 6.6 presents 11 age group pairs, followed by the pair relationships prevalent within each age group pair. The widely varying distributions of pair relationships within each age group pair are evident in this table. Because of the different prevalence of pair relationships within age group pairs, PMN was applied separately within each age group pair. Imputations were done one variable at a time, so no hierarchy of variables was required to set up a sequence of models, as is normally done with PMN. The first step, therefore, was to define respondents, nonrespondents, and the item response mechanism, within each age group pair. For a pair to be considered a complete data responding pair, the pair relationship must be definitively established. In terms of the variable PAIRREL, this meant that the pair had to have a value of PAIRREL within the range of 1 to 9, or equal to 11 or 12. Response propensity adjustments were then computed for each age group pair in order to make the respondent pair weights representative of the entire sample of pairs. (Because the modeling of the final pair weight adjustments was not completed at the time of the pair imputations, the pair-level sample design weights were adjusted to account for nonresponse at the household level using a simple ratio adjustment.)[6] These adjustments were calculated using an item response propensity model. This model is a special case of the GEM, which is described in greater detail in Appendix A.

Table 6.6 Age Group Pairs with Associated Possible Pair Relationships

Age group pair number	Age group pair	Pair relationships appearing in age group pair (in order of prevalence)	
		>=10% prevalence[1]	<10% prevalence[2]
0	12–14/12–14	sibling-sibling	other relationship[4]
1	12–14/15–17	sibling-sibling	other relationship
2	12–14/18–25	sibling-sibling	other relationship; parent-child; spouse-spouse**[5]
3	15–17/15–17	sibling-sibling	other relationship; spouse-spouse*
4	15–17/18–25	sibling-sibling	other relationship; spouse-spouse; parent-child*
5	18–20/18–25	other relationship; sibling-sibling; spouse-spouse	parent-child**
6	21–25/21–25	spouse-spouse; other relationship; sibling-sibling	parent-child**
7	12–14/26+	parent-child	other relationship; grandparent-grandchild; sibling-sibling*
8	15–17/26+	parent-child	other relationship; grandparent-grandchild; sibling-sibling; spouse-spouse**
9	18–20/26+	parent-child	other relationship; sibling-sibling; spouse-spouse; grandparent-grandchild

[6] In subsequent text, the use of the word "weights" will refer to the ratio-adjusted design weights.

Table 6.6 Age Group Pairs with Associated Possible Pair Relationships (continued)

Age group pair number	Age group pair	Pair relationships appearing in age group pair (in order of prevalence)	
		>=10% prevalence [1]	<10% prevalence[2]
10	21+/26+	spouse-spouse; parent-child; other relationship; sibling-sibling	grandparent-grandchild*

[1] The pair relationship labeled "spouse-spouse" includes partner-partner pair relationships.
[2] The pair relationships in this column each form at least 10% of the overall total number of pair relationships with the given age group pair, and the total is at least 85% of the overall total.
[3] Pair relationships followed by stars occur rarely, in less than 1% of the overall total number of pair relationships. Two stars indicate such rarity that the pair relationship did not appear in the age group pair in every survey year.
[4] "Other relationship" refers to a relationship other than sibling-sibling, parent-child, grandparent-grandchild, or spouse-spouse.
[5] The spouse-spouse domain as listed here actually consists of two domains that have been collapsed for the purposes of making the table easier to read.

6.2.4.2 Model Building and Determination of Predicted Means

The PMN method is a two-step process. The first step is the modeling step, followed by a hot-deck step where imputed values replace missing relationships. As stated earlier, each age group pair acted as an imputation class, within which the modeling and hot-deck steps were performed separately. The different attributes of the 11 models, corresponding to the 11 age group pairs, are described in this subsection.

Response categories. Ideally, we would like each type of pair relationship within an age group pair to constitute a response category in a multinomial response model. However, the numbers of cases corresponding to some pair relationships within each age group pair were very small, as is apparent in Table 6.6. Hence, it was not feasible to fit multinomial models that cover all the possible pair relationships for a given age group pair. Rather, in the modeling step, some of the response categories were combined, with separate assignments of imputed values within each of the 11 age group pairs. Priority was placed on placing the pair relationships "of interest" into separate categories. (Pair relationships "of interest" are defined in Section 6.2.2.) In some cases, pair relationships that were not of interest were combined with other categories, even if there were sufficient numbers to have a separate category in the multinomial model. Table 6.7 presents the response categories that were used for modeling. The delineation between categories that were combined for modeling was left to the hot deck step.

Table 6.7 Modeled Pair Relationships within Age Group Pairs

Age group pair number	Age group pair	Number of levels in response	Levels of modeled response
0	12–14/12–14	2	sibling-sibling; all others
1	12–14/15–17	2	sibling-sibling; all others
2	12–14/18–25	2	sibling-sibling; all others
3	15–17/15–17	2	sibling-sibling; all others
4	15–17/18–25	2	sibling-sibling; all others
5	18–20/18–25	3	both spouse-spouse pair relationships[1]; all others
6	21–25/21–25	3	both spouse-spouse pair relationships[1]; all others
7	12–14/26+	2	parent-child; all others
8	15–17/26+	2	parent-child; all others
9	18–20/26+	2	parent-child; all others
10	21+/26+	3	both spouse-spouse pair relationships[1]; all others

[1] The two spouse-spouse pair relationships are the spouse-spouse and the spouse-spouse with children under 18 pair relationships. The pair relationships labeled "spouse-spouse" include partner-partner pair relationships.

As an example, consider age group pair #5. In this age group pair, there are typically four types of pair relationships that have a sufficient number of respondent pairs to fit a satisfactory model, including spouse-spouse domains, sibling-sibling pairs, and all others. However, it is always easier to fit a good model with a smaller number of levels in the response. Since only two of those four were pair relationships of interest, these two (the two spouse-spouse domains) were used as levels in the response variable. The third level was obtained by combining the sibling-sibling and other relationship pairs. There are typically a small number of parent-child pairs, which were also combined in with the other relationship pairs.

Covariates in Models. After the weights were adjusted using the item response propensity model within each age group pair, binomial and multinomial logistic models were fitted using the adjusted weights, with the response variable defined as in Table 6.7. As noted in previous sections, the number of covariates at the household level was limited. The pool of covariates to be used in the item response propensity model included the following variables:

1. Household size (HHSIZE, as defined in Section 6.2.3.1)
2. Age category of older respondent (where applicable)
3. Race of older respondent
4. Sex of older respondent
5. Sex of younger respondent
6. Marital status of older respondent (where applicable)

7. Marital status of younger respondent (where applicable)
8. Education of older respondent (where applicable)
9. Education of younger respondent (where applicable)
10. Employment status of older respondent (where applicable)
11. Employment status of younger respondent (where applicable)
12. Region
13. Population density
14. Categorical percentage Hispanic in segment
15. Categorical percentage black in segment
16. Categorical percentage owner-occupied households in segment

In some cases, due to the ages of the pair members, the education, employment status, and marital status did not apply to one or both members of a pair. In order to increase the ability to obtain convergent models, some of the cells in the categorical covariates were collapsed. For all models, "employment status" was a binary response, full-time employed versus not full-time employed. The marital status cells also were collapsed. For models where the response variable involved spouse-spouse pairs, the covariate's levels were married versus not married. If the response variable involved child-parent pairs, the covariate's levels were married at least once versus never married.

Additional variables defined in Section 6.2.3.2 were used to adjust the weights in the final response models for each of the 11 age group pairs, in those cases where the variables were nonmissing. The variables follow:

a. Number in household aged 0 to 11
b. Number in household aged 12 to 17
c. Number in household aged 18 to 25
d. Number in household aged 26 to 34
e. Number in household aged 35 to 49
f. Number in household aged 50+

In the cases where these variables were all nonmissing, they were put into the pool of covariates for the final response model in place of HHSIZE. However, there were a handful of cases for which these variables could not be determined. In those cases, 11 additional final response models were fitted without the household composition age count variables listed above, using the same pool of covariates that were used for the item response propensity models.

Building of the Models. For age group pairs 0 through 4 and 7 through 9, binary logistic regression models were built. Since there were three outcomes with age group pairs 5, 6, and 10, multinomial polytomous logistic models were fitted for these age group pairs. All the models incorporated the design pair weights that were ratio adjusted for unit nonresponse (where a pair was selected but did not respond to the survey), and calibrated to account for item nonresponse (where a pair responded to the survey, but the pair relationship was unknown) using the item response propensity models, as described in Section 6.2.4.1. Naturally, not all the covariates in the original pool could be included in each model, due to convergence problems. The final set of covariates corresponding to each model is given in Appendix Q.

Determination of Predicted Means. Although models were built using respondent pairs where the pair relationship was known definitively, predicted probabilities were required for all pairs. Once the models were fitted, predicted means were determined for both respondent pairs and nonrespondent pairs, using the parameter estimates from the models.

6.2.4.3. Constraints on Hot Deck Neighborhoods and Assignment of Imputed Values

If possible, donor pairs in the hot deck step of PMN were chosen with predictive means within delta[7] of the recipient pair's predicted mean(s), where the value(s) of delta varied depending upon the value of the predictive means. In this case, delta was defined as 5 percent of the predicted probability if the probability was less than 0.5, and 5 percent of 1 minus the predicted probability if the probability was greater than 0.5. This allowed a looser delta for predicted probabilities close to 0.5 and a tighter delta for predicted probabilities close to 0 or 1. The range of values for delta across various predicted probabilities is given in Table 6.8. If no donor pairs were available with predictive means within delta of the recipient pair's predicted mean, the neighborhood was abandoned and the donor pair with the closest predictive mean was chosen.

Table 6.8 Values of Delta for Various Predicted Probabilities

Predicted probability (p)	Delta
$p \leq 0.5$	$0.05 * p$
$p > 0.5$	$0.05 * (1 - p)$

In general, the members of the neighborhoods were restricted to satisfy two types of constraints: "logical constraints" and "likeness constraints." Constraints that made the imputed values consistent with preexisting values of other variables were called logical constraints and were required for the candidate donor pair to be a member of the neighborhood. Likeness constraints were implemented to make donor pairs and recipient pairs as much alike as possible. Although logical constraints could not have been loosened, likeness constraints could have been loosened if they had forced the donor pool to be too sparse. Details of these imputation procedures are given in Appendix N.

In addition to the likeness constraint defined by delta, other likeness constraints were also included in the neighborhoods. These constraints follow:

Older pair member age constraint, 26+-year-old pair members. The 26+ age group associated with age group pairs 7 through 10 was split up into three groups: 26 to 34, 35 to 49, and 50+. This was most useful to delineate child-parent pairs.

Marital status likeness constraints. Each respondent's marital status, as entered in the core section of the questionnaire, was closely related to the relationship between the pair

---

[7] "Delta" refers to the value that defined the neighborhood of donor pairs that were "close" to the recipient pair. The difference between the predictive mean of the recipient pair and the predictive means of the donor pairs must have been less than delta. See Appendix N for more details.

members.[8] This marital status variable had four levels among respondents ages 15 or over: married, widowed, separated or divorced, and never married. Marital status likeness constraints combined the information from this variable for both pair members, where the levels were collapsed in different ways depending upon the age group pair. For age group pairs where only child-parent pair relationships were involved, two classes were required: both respondents never married, and one respondent never married. Three classes were required for age group pairs where only spouse-spouse pair relationships were possible: both respondents not currently married, one respondent not currently married, and both respondents currently married. Finally, six classes were attempted if both the spouse-spouse and child-parent pair relationships were possible: a) both respondents never married; (b) one respondent never married, the other formerly married (widowed or divorced); (c) one respondent never married, the other currently married; (d) both respondents formerly married (widowed or divorced); (e) one respondent formerly married, the other currently married; and (f) both respondents currently married. It should be noted that not all of these classes would need donor pairs if no recipient pairs were within the class. It should be noted that marital status could not have been considered a logical constraint where spouse-spouse pairs were involved, since many live-in partners, who were considered spouse-spouse pairs, answered the marital status question as "never married."

Gender makeup of pair likeness constraints. For donors who formed a spouse-spouse pair, the vast majority were male-female. Hence, in those cases where a spouse-spouse pair was possible, the gender likeness constraint required that the donor pair and recipient pairs both be either of the same gender or both of a different gender. This meant that the likelihood of same-sex spouse-spouse pair relationships were equally likely (more or less, depending upon the model) among donors and recipients.

Age constraints on 15-to-17-year-old pair members. For the 15 to 17 age group, the likelihood of being in a spouse-spouse relationship was very small. Nevertheless, the likelihood that a 17 year old was married was considerably greater than the likelihood for a 15 year old. Hence, for the age group pairs where at least one pair member was between 15 and 17, the younger pair member of both the donor pair and recipient pair had to be of the exact same age.

Constraints on number of children. In Section 6.2.3.2, a covariate was defined for the number of children in the household under 12, AGE011, and one was defined for children in the household between 12 and 17, AGE1217. If there was disagreement between pair members on the values of these covariates, the pair member with information agreeing with the screener was used if possible. For the imputation of spouse-spouse relationships with and without children, these covariates were used to restrict donor pairs, where AGE011 was used for potential parents under 18, and AGE011+AGE1217 for potential parents 18 or over. If the recipient pair had no children according to the relevant covariate or covariates, donor pairs also did not have children. If the recipient pair had children, the same was true for the donor pair. In almost all cases, when

---

[8] Pairs that include a pair member with an imputed marital status were not eligible to be donor pairs. If a recipient pair had a pair member with an imputed marital status, donor pairs had any marital status, unless one of the pair members in the recipient pair had a non-imputed marital status indicating married, widowed, or divorced.

there was disagreement between pair members regarding whether the pair had children in the household or not, the imputation used information that was closer to the screener.[9]

The likeness constraints were loosened in the following order (where applicable) (1) for the age group pairs where 6 marital status classes were used, collapse to two classes (as with child-parent pairs) or three classes (as with spouse-spouse pairs) depending on the response that was most common; (2) abandon the neighborhood, and choose the donor pair with the closest predicted mean or means; (3) loosen age constraint (26+ groups); (4) loosen the marital status restrictions; and finally (5) simultaneously loosen the age constraints on 15-to-17-year-old pair members and the gender makeup likeness constraints. The constraint on the number of children in the household was never loosened. For the multinomial logistic models, a Mahalanobis distance was used to define the distance across the multiple predicted probabilities.

Logical constraints were limited to the information that was already known about the pair, as denoted by the level of the variable PAIRREL. If, for example, PAIRREL = 14, then no information was available about the identity of the pair relationship, and no logical constraint was needed. On the other hand, if PAIRREL = 15, this meant that the pair relationship was either a child-parent pair where the child was aged 12 to 14, or it was some relationship other than spouse-spouse, parent-child, grandparent-grandchild, or sibling-sibling. One could argue that the household composition age counts be considered logical constraints. However, these variables did not exist for all respondent pairs, and in some cases the values were set for these variables in a somewhat arbitrary manner. Moreover, due to the timing of the interviews, it was conceivable that an unexpected pair relationship could occur even though the household composition age counts would seem to preclude it.

6.2.4.4. Additions to Analytic File

The imputation-revised pair relationship variable is called IRPRREL, with an accompanying imputation indicator IIPRREL. In addition to these variables, the edited pair relationship variable PAIRREL, the quality-of-match indicator RELMATCH, and the pair indicator PAIRMEM, which simply indicates whether a respondent in the analytic file was part of a responding pair, were released to the analytic file. Four additional variables were released to the analytic file to aid in pair analyses. These included the variables PRNTIND, AGEOTHER, SEXOTHER, and PAIRID. PRNTIND identifies whether the respondent was a parent in a parent-child relationship; AGEOTHER gives the age of the other respondent in the pair, SEXOTHER gives the gender of the other respondent in the pair, and PAIRID gives the questionnaire ID (QUESTID) of the other pair member.

6.3 Stage Two: Creation and Imputation of Multiplicities

As stated earlier, multiplicities were required to account for analyses that were made at the person level, even though the pair weights were calculated at the pair level. The multiplicities were relevant only at the person level, so naturally the definition of a multiplicity required the

[9] This will not always be true, because it is not always possible that the screener can be used to determine the value for AGE011 and AGE1217 when the pair members' information disagrees.

identification of the focus member of the pair. Using the pair relationships determined in Section 6.2, the following domains were considered:

1. parent-child (child 12 to 14), parent focus
2. parent-child (child 12 to 14), child focus
3. parent-child (child 15 to 17), parent focus
4. parent-child (child 15 to 17), child focus
5. parent-child (child 12 to 17), parent focus
6. parent-child (child 12 to 17), child focus
7. parent-child (child 12 to 20), parent focus
8. parent-child (child 12 to 20), child focus
9. sibling (12 to 14)-sibling (15 to 17), sibling (15 to 17) focus
10. sibling (12 to 14)-sibling (15 to 17), sibling (12 to 14) focus
11. sibling (12 to 17)-sibling (18 to 25), sibling (18 to 25) focus
12. sibling (12 to 17)-sibling (18 to 25), sibling (12 to 17) focus
13. spouse-spouse (includes partner-partner) with children under 18
14. spouse-spouse (includes partner-partner)

Determining the multiplicity entailed finding the number of roster pairs in the domain of interest that contained the focus member in the pair. In broad terms, the process of determining the multiplicity count was a three-step process: (1) determine the multiplicity count for each pair member; (2) use the screener, quality of roster, and other means to figure out the appropriate count if each pair member's counts did not match; and (3) impute multiplicities that could not be otherwise determined.

Since the pair weights reflected selection done at the time of screening, the multiplicity count should have reflected the household makeup at that time. However, this was not entirely possible, since the screener roster was not as complete as the questionnaire roster, and recorded relationships in the screener roster were relative to the head of household, rather than to each pair member. Hence, no account was made for cases where a change in the household makeup occurred between the screening time and the time of both interviews. The change in household makeup could have occurred because of an intervening birthday, or because a roster member left or entered the household after screening. Technically, adjustments should have been made to account for this. However, the number of cases where this occurred was small, and to implement such an adjustment would have been extremely complicated, especially for the household counts discussed in Section 6.4. Nevertheless, in cases where there were disagreements between pair members on the value of the multiplicity count, the screener was used to resolve those disagreements.

6.3.1 Determining Multiplicity Count for Each Pair Member

The multiplicity counts for each pair member consisted of a direct count and an indirect count. The direct count was obtained by looking at the pair member who was the focus. It was simply a count of the roster members that could have been selected, where the same pair domain would have resulted. The indirect count was obtained by looking at the pair member who was not the focus. It was a count of the pair member himself or herself, plus other roster members who, by virtue of their relationship code, would have had the same pair relationship had they been selected. A summary of the ways of determining the direct count and indirect count for each pair domain are given in Table 6.9. For the domains given in Table 6.9, neither the direct nor the indirect count could be 0, since the pair member who was not the focus had to be part of the count. For spouse-spouse counts, no work was necessary to determine multiplicity counts. If a respondent was in a spouse-spouse pair, the multiplicity count was necessarily 1, since only one spouse-spouse pair could have been selected that included that pair member. Note that other spouse-spouse pairs in the household (one spouse's parents, for example) would have been of interest in the household counts discussed in subsequent sections.

Table 6.9 Multiplicity Counts for Each Pair Member

Pair relationship	Focus member	Direct count	Indirect count
Parent-child	Child	from child: number of parents	from parent: self + spouse/partner
Parent-child	Parent	from parent: number of children in appropriate age range	from child: self + number of siblings in the appropriate age range
Sibling-sibling	Older sibling	from older sibling: number of siblings in younger age range	from younger sibling: self + number of siblings in younger age range
Sibling-sibling	Younger sibling	from younger sibling: number of siblings in older age range	from older sibling: self + number of siblings in older age range

6.3.2 Matching Multiplicity Counts across Pair Members

Once the counts were determined for each pair member, it was necessary to resolve differences between these counts across pair members. In most cases, the direct and indirect counts agreed, with no bad relationship codes for either pair member, resulting in an easy determination of the final multiplicity count. An easy determination was usually possible if one pair member had bad relationship codes or had a count of 0, which meant that the final multiplicity count came from the pair member with good data.[10] Exceptions to this rule are discussed in Appendix R. For some cases, both pair members had bad relationship codes, which meant that the final multiplicity was left to imputation. Some of the remainder of cases could be reconciled, and some could not. In the cases where reconciliation was possible, many of the disagreements between the pair members were resolved by going to the screener. The method used to reconcile differing counts depended upon the domain. In addition to the screener, for the parent-child domains, the FIPE3 variable was used to help reconcile differences. Detailed rules for reconciling differences between pair members are given in Appendix R.

If reconciliation between the counts from the two pair members in the household and the screener was not possible, upper and lower bounds within which the imputed value had to reside were determined from the counts for each pair member, and the counts for the screener. The

[10] There were some provisions to this rule. If the bad relationship codes were only within the relevant age ranges, then the count from the good side was only used if the age ranges in the good side matched the screener.

amount of imputation required for the multiplicity counts is given in Table 6.10 for the 2002 survey year as well as previous survey years. From this table, it is apparent that the greatest degree of uncertainty came with the determination of the number of parents in the child-focus parent-child domains. This occurred because, even though the parent-child pair relationship had been established, it was often unclear whether there was a second "parent" in the household. Other domains had very little uncertainty; the counts of the number of children in the parent-focus parent-child domain were almost always definitively determined.

6.3.3 Creation of Imputation-Revised Multiplicity Variables

In many cases where the pair relationships were not defined, multiplicity counts were also not defined. In addition, there were a handful of cases where multiplicity counts were not determined even when the pair relationship was known. In all of these cases, imputation was required to determine the multiplicity count. As with the pair relationship imputation, missing multiplicities were imputed using the PMN method described in Appendix N. In this section, the application of PMN to the imputation of multiplicities is described. Since only the multiplicity in the second stage was imputed for each pair, the imputation was univariate in the sense that no sequential models were necessary. However, in some cases, several variables were associated with a single model, as is described below.

6.3.3.1 Setup for Model Building

Multiplicity counts were only defined within the relevant domain, which, in turn, depended upon the pair relationship. For the sibling-sibling pairs, four separate imputations were conducted for the multiplicities associated with the four sibling-sibling pair domains. The parent-child domains are hierarchical, however, where the imputations could not have been conducted independently if consistency was to be maintained. Hence, only two models were fitted to the child-parent pairs, using just the domains with children 12 to 20 years old. One set of models was for the number of the parent's children and the other set was for the number of parents of the child. Using the predicted means from these models, a single donor pair for each focus was selected from which the multiplicity counts were determined for 12 to 14, 12 to 17, 15 to 17, and 12 to 20 child-parent pairs. No imputation was required for the spouse-spouse multiplicity counts, since a selected respondent in a spouse-spouse pair naturally had only one spouse.

The first step for these six models was to define respondents, nonrespondents, and the item response mechanism for each model, separately. For a pair to be considered a complete data responding pair with regard to multiplicities, the multiplicity had to be nonmissing for all the variables being imputed. For the parent-child pairs, this meant that the multiplicity had to be nonmissing for the domains with 12-to-20-year-old children. A nonmissing multiplicity for this domain would automatically guarantee nonmissing multiplicities for the subset parent-child domains. Response propensity adjustments were then computed for each of the six models, in order to make the respondent pair weights representative of the entire sample of pairs. (Because the modeling of the final pair weight adjustments was not completed at the time of the pair imputations, the pair-level design weights were adjusted to account for nonresponse at the household level using a simple ratio adjustment.) These adjustments were calculated using an item response propensity model. This model is a special case of the GEM, which is described in greater detail in Appendix A.

Table 6.10 Amount of Imputation Required for Multiplicities in Various Pair Domains

Pair Domain	Multiplicity	Missing cases			
		1999	2000	2001	2002
Parent-child (12-14) child focus	number of parents	38	40	50	74
Parent-child (12-14) parent focus	number of children	0	0	0	0
Parent-child (15-17) child focus	number of parents	26	50	42	66
Parent-child (15-17) parent focus	number of children	0	0	0	2
Parent-child (12-17) child focus	number of parents	64	90	92	140
Parent-child (12-17) parent focus	number of children	2	0	0	4
Parent-child (12-20) child focus	number of parents	76	92	110	170
Parent-child (12-20) parent focus	number of children	4	0	0	4
Sibling (12-14)-sibling (15-17), older sibling focus	number of younger siblings	16	2	0	4
Sibling (12-14)-sibling (15-17), younger sibling focus	number of older siblings	14	2	2	2
Sibling (12-17)-sibling (18-25), older sibling focus	number of younger siblings	18	20	4	8
Sibling (12-17)-sibling (18-25), younger sibling focus	number of older siblings	22	14	10	6

6.3.3.2 Model Building and Determination of Predicted Means

The PMN method is a two-step process. The first step is the modeling step, followed by a hot-deck step where imputed values replace missing multiplicities. The different attributes of the six multiplicity models, corresponding to the six pair domains, are described in this subsection.

Response Categories. The response categories for the six multiplicity final response models were simply the multiplicity counts for each domain among the complete data cases.

Covariates in Models. The pool of covariates for the response propensity models was the same pool that was used for the pair relationship response propensity models. By the same token, this pool was also used for the final response multiplicity models when the household composition age count variables were missing. When these variables were not missing, the same pool was again used as with the pair relationship models. Naturally, the final set of covariates differed from the initial pool; the final set of covariates that were used in the models is given in Appendix Q.

Building of the Models. For the child-focus parent-child domains, the count being modeled was the number of parents. Since we have already established the pair relationship, only two responses are possible within the parent-child pair relationship: 1 parent or 2 parents. For these multiplicity counts, the fitted models were binomial logistic regression models. Only respondents who had a non-imputed pair relationship with non-missing multiplicity counts was eligible for the model-building data set.

The other responses (parent-focus parent-child and sibling-sibling multiplicity counts) were counts, where Poisson regression models were used. However, the data were underdispersed for a Poisson distribution, so that the data had to be scaled using the observed variance.

Determination of the Predicted Means. Although models were built using respondent pairs where the multiplicity was known definitively, predicted means were required for all pair domains where imputation was required. Once the models were fitted, predicted means were determined for both respondent pairs and nonrespondent pairs, using the parameter estimates from the models.

6.3.3.3 Constraints on Hot Deck Neighborhoods and Assignment of Imputed Values

Like the pair relationship response variables, the child-focus parent-child domains had a binary response (1 parent or 2 parents). Hence, in the same manner as with the pair relationship imputations, donor pairs in the hot deck step of PMN for these domains were chosen with predictive means, if possible, within delta of the recipient pair's predicted mean. The value of delta varied depending on the value of the predicted mean. The values of delta for predicted probabilities are given in Section 6.2.4.3. For the other domains (parent-focus parent-child and the sibling-sibling domains), an error in the software meant that the formula given in Section 6.2.4.3 was applied to the predicted values, even though they were not probabilities. This prevented delta from being applied for these domains except in the strictest of circumstances. However, the impact of this error was very small, since the amount of imputation for these domains was considerably less than it was for the child-focus parent-child domains (see Table 6.10). Moreover, the other constraints that were placed on the neighborhoods ensured that donor pairs and recipient pairs were very much alike. Finally, for variables where the delta constraint was used, it was often the first constraint to be loosened.

Wherever necessary and feasible, logical and likeness constraints (as defined in Section 6.2.4.3) were placed on the membership in the hot-deck neighborhoods. The hot deck step and the accompanying constraints are described separately for each of the variables in turn.

Parent-child pairs, child focus. The donor pairs and recipient pairs had to have the same pair relationship, excluding the restrictions on ages. This acted as a logical constraint. (Donor pairs had to have non-imputed pair relationship data.) In addition, the number of parents was restricted by the number in the household of the appropriate age. An additional constraint, therefore, was that donor pairs and recipient pairs have the same number of individuals in the household aged 26 and over, provided this information was available for the recipient pair. (Donor pairs had to have complete data on all the household composition age count variables.) If the recipient pair had only one person in the household in this age range, then the number of parents in the household could still have been two, if the other parent was under 26 years old. However, this constraint ensured that donor pairs and recipient pairs have the same household age pattern. This was a likeness constraint that was never loosened. Besides delta, additional likeness constraints all involved the household composition.

In addition to the 26-or-over constraint, the neighborhoods were further restricted by requiring donor pairs and recipient pairs to have the same number of household members within the age ranges 26 to 34, 35 to 49, and 50 and over. Other likeness constraints included requirements that donor pairs and recipient pairs to have the same (1) number of household members under 12 years of age; (2) number of household members between the ages of 12 and 17 (inclusive); and (3) household sizes; and (4) values for IRPRREL. This latter constraint strengthened the requirement of matching pair relationships to include the restrictions on the ages. It meant that, for example, that donor pairs and recipient pairs within the domain involving 12 to 17 year olds both involved 12 to 14 year olds or both involved 15 to 17 year olds.

The likeness constraints were loosened in the following order (where applicable) (1) abandon the neighborhood, and choose the donor pair with the closest predicted mean or means; (2) abandon the requirement that donor pairs and recipient pairs have the same number of household members under 12 and between 12 and 17 (inclusive); (3) abandon the requirement that donor pairs and recipient pairs must have the same number of household members within the age ranges 26 to 34, 35 to 49, and 50 and over, and drop the household size constraint. The IRPRREL constraint was never loosened.

Parent-child pairs, parent focus. As with the child-focus pairs, donor pairs and recipient pairs had to have the same pair relationship, and donor pairs were required to have non-imputed pair relationship data. For the parent-focus pairs, the counts could have taken on more than two values. If the counts from the two pair members did not get reconciled, but both pair members had valid rosters, then the two counts acted as upper and lower bounds for the imputation, acting as additional logical constraints. The counts were limited anyway, however, since the age ranges of the children were, by definition, constrained. Specifically, donor pairs and recipient pairs had to have the same number of household members within the relevant age ranges (12 to14, 12 to 17, 15 to 17, or 12 to 20, depending upon the recipient pair's value for IRPRREL). (As before, donor pairs had to have complete data on the roster age variables.) Additional likeness constraints included a requirement that donor pairs and recipient pairs have the same number of household members under 12, and a requirement that household sizes be the same. The constraint on IRPRREL was also included.

The likeness constraints were loosened in the following order (where applicable) (1) abandon the neighborhood, and choose the donor pair with the closest predicted mean or means; (2) abandon the requirement that donor pairs and recipient pairs have the same number of household members under 12 years old; (3) abandon the requirement that donor pairs and recipient pairs must have the same household size; and (4) loosen the IRPRREL constraint.

Sibling-sibling pairs. As with the parent-child pairs, donor pairs and recipient pairs had to have the same value for IRPRREL, and donor pairs were required to have non-imputed pair relationship data. As with the parent-child parent-focus pairs, the counts from the two pair members acted as upper and lower bounds for the imputation, as additional logical constraints, provided both pair members had valid rosters. Donor pairs and recipient pairs were also required, as a logical constraint, to have the same number of household members within relevant age ranges. For example, for a sibling-sibling pair with ages 12 to 14 and 15 to 17, with a focus on the younger member, the donor pair and recipient pair were required to have the same number of 15 to 17 year olds. (As before, donor pairs had to have complete data on the roster age variables.) Additional likeness constraints included a requirement that donor pairs and recipient pairs have the same (1) number of household members under 12; (2) household sizes; (3) number of household members in the age group corresponding the pair member of focus; and (4) number in the household between ages 12 and 17, for the sibling-sibling pairs where one member was between 12 and 14 (inclusive) and the other was between 15 and 17 (inclusive).

The likeness constraints were loosened in the following order (where applicable) (1) abandon the neighborhood, and choose the donor pair with the closest predicted mean or means; (2) abandon all likeness age count constraints; and (3) abandon the requirement that donor pairs and recipient pairs must have the same household size.

6.3.3.4 Additions to Analytic File

The imputation-revised versions of the parent-child multiplicity variables are called IRMPCCxx and IRMPCPxx, where the final C and P refer to the focus in the domain. The "xx" refers to the age range of the children, which is the upper bound if the lower bound is 12, or "57" if the range is 15 to 17. The edited version of these variables, MCPCCxx and MCPCPxx, were also released to the analytic file. The sibling-sibling imputation-revised variables are called IRMSxxxx, where the second two x's in the "xxxx" refer to the upper bound of the age range corresponding to the focus pair member, and the first two x's refer to the upper bound of the age range corresponding to the remaining pair member. The edited version of these variables is given by MCSxxxx. The imputation indicators were also released to the analytic file, with II prefixes instead of IR prefixes. Finally, the spouse-spouse counts are called MCSPSP and MCSPSPWC. These are simply indicators of whether the pair was a spouse-spouse pair, or whether the pair was a spouse-spouse pair with children under 18. No imputation was required for these variables.

6.4 Stage Three: Creation and Imputation of Household-Level Person Counts in Each Domain for the Purposes of Pair Weight Calibration

In order to improve the quality of the estimates from the pair data through post-stratification, it was necessary to identify the household-level person counts for each domain. This entailed finding the number of individuals in the household that belonged to a particular domain, given one member of a domain was selected as the focus. These counts were more difficult to derive than the multiplicity counts since all households were considered. Within each household, counts for any of the domains of interest were derived, regardless of whether that household belonged to that domain, or even whether a pair was selected at all. The counts were derived for 10 of the 14 pair domains given in Section 6.3. For two of the remaining domains, the parent-child counts where the child was between 15 and 17, calculating the household counts was unnecessary.[11] For the other two remaining sibling-sibling domains, the reasons are historical: they were added after the procedures were first developed, and there was insufficient time to develop the household counts for those domains. The domains where these counts were created are listed below:

1. parent-child (child 12 to 14), parent focus
2. parent-child (child 12 to 14), child focus
3. parent-child (child 12 to 17), parent focus
4. parent-child (child 12 to 17), child focus
5. parent-child (child 12 to 20), parent focus
6. parent-child (child 12 to 20), child focus
7. sibling (12 to 14)-sibling (15 to 17), sibling (15 to 17) focus
8. sibling (12 to 17)-sibling (18 to 25), sibling (18 to 25) focus
9. spouse-spouse (includes partner-partner) with children under 18
10. spouse-spouse (includes partner-partner)

Determining the household-level person counts was a three-step process: (1) determine the household count for each respondent, whether a member of a pair or a single respondent; (2) use the screener, quality of roster, and other means to figure out the appropriate final count, either by attempting to reconcile differing counts between pair members or by attempting to determine the appropriate count when information from only one roster was available; and (3) impute missing counts. For households where only one respondent was selected, the matching step (step 2) was unnecessary.

Since the pair weights reflected selection done at the time of screening, the household-level person counts should have reflected the household makeup at that time. As with the multiplicity counts, however, this was not entirely possible, so no account was made for cases where a change in the household makeup occurred between the screening time and the time of both interviews. An explanation for why this was not possible is given for the multiplicity counts in the introduction to Section 6.3. Moreover, as stated in that section, to implement such an adjustment would have been extremely complicated for the household-level person counts. Nevertheless, in cases where there were disagreements between pair members on the value of the household-level person count, the screener was used to resolve those disagreements.

6.4.1 Determining the Household-Level Person Count for Each Respondent

The multiplicity count was a count of the number of pairs in the household that could be associated with the person of focus. The household-level person counts asked a different question: how many persons of focus are there for a given pair domain, provided such a pair domain existed in the household, regardless of what pair (or whether a pair) was actually selected? For a parent-child pair, for example, if two parents are in the household with three children aged 12 to 14, then the household person count for the parent focus is the same as the

[11] Since household counts were defined for everybody, it was possible to derive these counts using the counts for the parent-child domains where the child was between 12 and 14, and where the child was between 12 and 17. The multiplicity counts for the parent-child (15 to 17) domain had to be calculated, however, and could not have been derived in this easy way. This was due to the fact that multiplicity counts were only defined if the pair relationship corresponded to the pair domain of interest.

multiplicity count for the child focus: 2. Similarly, the household person count for the child focus is the same as the multiplicity count for the parent focus. Household person counts would also have been obtained for the various sibling-sibling and spouse-spouse domains in this example, even though the relationship was parent-child.

6.4.1.1 Parent-Child Domains

When obtaining household-level person counts for parent-child domains, the six parent-child domains given in the introduction to this section are what were under consideration. In any household, the household-level person counts for parent-child domains were nonzero if at least one parent was present in the household with children within the relevant age range. In this instance, the child-focus count would have been simply the number of children in the household within that age range that belonged to the parent in the household, and the parent-focus counts would have been the number of parents. If more than one "family unit" (mother and/or father with children) lived within the household, the child-focus counts should have counted children from more than one set of parents, and the parent-focus counts should have counted more than two parents. One situation where this occurred was where three generations lived within the same household, with children in both the youngest and the second generations within the relevant age range. Using the youngest generation as the reference point, some of the parent's siblings (the grandparents' other children) were within the relevant age range. In this instance, the parent-child domains of the number of children would have included both children of the parents and the children of the grandparents who were in that age range. The count of the number of parents included both the parents and grandparents (and exceeded two). Identifying more than one family unit in a household with children within the relevant age range under other scenarios (e.g., two sisters both with children within the relevant age range, both living within the same household) could not be determined from the data, and had to be disregarded. Regardless of how many family units were in the household, counts had to be determined in different ways depending upon whether a parent-child pair "of interest" was selected or not. Descriptions of how to obtain the household-level person counts are given below for the parent-child domains outlined above, first for parent-child pairs of interest, with parent-focus and child-focus domains considered together. In this instance, the pair actually belonged to a pair relationship where analysis using one or more of the domains listed was possible. This was followed by descriptions for other pairs and single respondents, with parent-focus and child-focus domains considered separately.

6.4.1.1.1 Obtaining Counts for Parent-Child Domains (Parent-Focus and Child-Focus): Parent-Child Pairs, Child Under 21

If the pair was identified as parent-child and the three-generation situation described above was not apparent, the household-level child-focus person count was given by the parent-focus multiplicity count. Similarly, the household-level parent-focus person count was given by the child-focus multiplicity count. If a three-generation situation was identified and the grandparent also had children within the relevant age range, the number of children and number of parents were adjusted appropriately. The final household count in this instance was greater than the imputation-revised multiplicity count, which did not include all the children in the household within the relevant age range.

6.4.1.1.2 Obtaining Counts for Child-Focus Parent-Child Domains: Other Pairs and Single Respondents[12]

For other pairs and single respondents, the following conditions were required to determine the household count of the number of children of parents in the household:

1. If the age of the respondent was within the relevant age range, and that child had at least one parent, then the child-focus counts were determined in the same way as the parent-focus multiplicity counts: the count was of the self plus the child's siblings within the relevant age range. If the child's parents were not identified as living with him/her in the household, the count was set to 0.

2. If the respondent had children within the relevant age range, then the count was of the respondent's children within that range. If the respondent also had older children who had children of their own within the relevant age range, then the count was of the respondent's children and grandchildren within the relevant age range.

3. If the age of the respondent was outside the relevant age range, but the respondent had parents living with them in the household and had siblings within the relevant age range, then the count was of the number of the respondent's siblings.

4. If the respondent had grandchildren within the relevant age range, and the respondent also had children over 25 or children-in-law living with them, then the count was the number of the respondent's grandchildren. (The assumption was that the respondent's children or children-in-law are the parents of the respondent's grandchildren. The likelihood of this not being the case was small. In the case where a pair was selected, this can be resolved by looking at the count of the other pair member.)

6.4.1.1.3 Obtaining Counts for Parent-Focus Parent-Child Domains: Other Pairs and Single Respondents[12]

For other pairs and single respondents, the following conditions were required to determine the household count of the number of parents of children in the household:

1. If the age of the respondent was within the relevant age range, then the count was of the number of the respondent's parents (which could be 0).

2. If the age of the respondent was outside the relevant age range, but the respondent had siblings within the relevant age range, then the count was of the number of the respondent's parents (again, this could be 0).

3. If the respondent had children within the relevant age range, then the parent-focus counts were determined in the same way as the parent-focus multiplicity counts: the

[12] "Other pairs" included pairs that were not within a domain of interest because the age of at least one of the pair members was outside the relevant age range. For parent-child pairs, this applies to a pair with a child that was 21 or over. For sibling-sibling pairs, this applies to siblings where both were within the same age range (both were 12 to 14, 15 to 17, or 18 to 25), or where at least one of the siblings was older than 25 years of age.

count was of the self plus the spouse or live-in partner. If the respondent also had older children who had children of their own within the relevant age range, and had a child over 25 and a child-in-law living with him/her, the count was two plus the self and spouse or live-in partner. If the respondent had a child over 25 but no child-in-law, the count was one plus the self and spouse or live-in partner. (Note that, under these scenarios, the number of parents could exceed two.)

4. If the respondent had grandchildren within the relevant age range but no children in that range, but the respondent had a child over 25 or a child-in-law living with them, the count was two if both the child over 25 and child-in-law were living in the household, one if not.

6.4.1.2 Sibling-Sibling Domains

When obtaining household-level person counts for parent-child domains, the two sibling-sibling domains given in the introduction to this section are what were under consideration. As with the parent-child counts, the household-level person counts for sibling-sibling domains were nonzero if at least one sibling-sibling pair was present in the household within the relevant age ranges, where the count was simply the number of appropriately-aged siblings. If sets of siblings from more than one "family unit" (sets of siblings from different parents) resided within the same household, the sibling-sibling counts should have counted possible pairs from within each set. However, sets of siblings that did not involve the respondent's family unit could not have been identified from the data. Regardless of how many sets of siblings were in the household, counts had to be determined in different ways depending upon whether a sibling-sibling pair "of interest" was selected or not. Descriptions of how to obtain the household-level person counts are given below for the sibling-sibling domains outlined above, first for sibling-sibling pairs of interest. In this instance, the pair actually belonged to a pair relationship where analysis using one or more of the domains listed was possible. This was followed by descriptions for other pairs and single respondents. In each case, the descriptions apply regardless of which sibling-sibling domain is under consideration.

6.4.1.2.1 Obtaining Counts for Sibling-Sibling Domains: Sibling-Sibling Pairs of Interest

If the pair was identified as sibling-sibling within a relevant domain, the multiplicity count was simply given by the number of younger siblings, since the older sibling was the focus. For the household-level sibling-sibling person count, we are interested in the number of older siblings. The counts are determined in a similar manner to the multiplicity count, except that the older siblings are now of interest. If the pair member is the older sibling, then the household count is the self plus the number of siblings in the older age range. The count for the younger sibling pair member is simply the number of siblings within the same older age range. Unlike the case with the parent-child household-level counts, inconsistencies in the sibling-sibling counts when the pair selected was sibling-sibling still need to be resolved. However, the rules for resolving inconsistencies can follow directly from those used for the multiplicity counts when counting the number of younger siblings, given in Appendix R. Note that a pair that is within one sibling-sibling pair domain had to be outside the other sibling-sibling pair domain.

6.4.1.2.2 Obtaining Counts for Sibling-Sibling Domains: Other Pairs and Single Respondents[12]

For other pairs and single respondents, the following conditions were required to determine the household count of the number of siblings within the older age ranges of the domains of interest in the household:

1. If the age of the respondent was within the age range of the older sibling, and that child had at least one sibling in the younger age range, then the counts were given as the self plus the child's siblings within the older age range. If the child did not have any siblings within the younger age range, the count was set to 0.

2. If the age of the respondent was within the age range of the younger sibling, and that child had at least one sibling in the older age range, then the counts were given by the number of child's siblings in the older age range.

3. If the age of the respondent was outside the age range of the older or younger sibling, but had at least one sibling in each of the older and younger age ranges, the counts were given by the number of siblings in the older age range.

4. If the age of the respondent was outside the age range of the older or younger sibling, but the respondent had children both within the older and the younger age ranges, the count was of the number of respondent's children in the younger age range.

5. If the age of the respondent was outside the age range of the older or younger sibling, but the respondent had grandchildren within the older and younger age ranges, the count was of the number of grandchildren in the younger age range. (If the respondent's grandchildren were cousins rather than siblings, we had no way of deciphering this from the data. This had to be resolved by looking at the information from the other pair member, if another pair member was selected.)

6.4.1.3 Spouse-Spouse Domains

What is referred to as a "spouse-spouse domain" was actually derived from spouse-spouse and partner-partner pair relationships. The following conditions were required for the number of spouse-spouse (including partner-partner) pairs to be incremented by one. Some of these conditions were applied to the same household:

1. The respondent was part of a spouse-spouse (or partner-partner) pair.
2. The respondent was not part of a spouse-spouse pair, but had a spouse (or live-in partner).
3. The respondent had two parents living in the house.
4. The respondent had two parents-in-law living in the house.
5. The respondent had two grandparents living in the house.
6. The respondent had a child and a child-in-law living in the house.

The following conditions were required for the number of spouse-spouse pairs with children under 18 to be incremented by one. (These also include partner-partner pairs with children under 18.) Some of these conditions were applied to the same household:

1. The respondent was part of a spouse-spouse (or partner-partner) pair with children under 18.
2. The respondent was not part of a spouse-spouse pair,[13] but had a spouse (or live-in partner), and children under 18.
3. The respondent had two parents living in the house, and was either under 18 or had siblings under 18.
4. The respondent had a child and a child-in-law living in the house, and had grandchildren under 18.

6.4.2 Determining the Final Household-Level Person Count

For a particular type of household-level person count, there are three types of households from a sample selection perspective. For the first type, a pair was selected and both pair members responded, where the pair relationship corresponded directly to the pair domain being counted. In this case, the household-level person count was usually easy to obtain using the multiplicity counts, although an adjustment was sometimes required if more than one family unit was in the household. An example of this: a parent-child pair was selected where the child was 12 years old, and the household-level person count for the parent-focus parent-child (12 to 14) domain was required. In the second type of household, a pair was also selected and both pair members responded, but in this type the pair relationship did not correspond directly to the pair domain being counted. In this case, determining the final count was sometimes more difficult, particularly if one or more of the counts was a zero count. A zero count from a roster with good data did not necessarily mean that the final count should be zero. An example of this: suppose a household consisted of a man, his wife, brother, and two sons, and suppose one of the sons and his uncle (the man's brother) were selected. If the uncle's roster would have a zero count for all domains of interest, since all of the household members were "other relatives" from his perspective, so that no nonzero parent-child count could be obtained. The final count would have to be determined from imputation. In the third type of household, only one respondent was selected. In this case, it was not necessary to match counts from different pair members, but determining the final count could still be difficult if the count was a zero count for a household where the value was not truly zero.

For situations where a pair was selected and both pair members had good roster data, if the counts agreed between the pair members and were not zero, an easy determination of the final household-level count was possible. Surprisingly, this occurred in a majority of cases. If one pair member had a bad roster with no information in it and the other had a good roster, this was treated in the same way as if a single respondent was selected with a good roster. In either of these cases, the final count could be determined, provided a considerable number of conditions were satisfied. The conditions used to accept a good roster's count, when either the other pair

[13] We excluded all spouse-spouse pairs here since spouse-spouse pairs with children were already accounted for, and spouse-spouse pairs without children had already been defined, possibly by imputation, not to have children under 18.

member's roster was bad, or no pair was selected, are given in Appendix S. If these conditions were not meant, the final household-level person count was left to imputation. Imputation was also the way out if two pair members were selected, both with bad rosters.

For the remainder of cases, some could be reconciled, and some could not. In the cases where reconciliation was possible, some of the disagreements were caused by the pair members' rosters have different age-and-gender compositions. In these cases, many of the disagreements between the pair members were resolved by going to the screener. However, the screener did not provide much help if the age and gender composition of the pair members' rosters were identical, yet the counts still disagreed, as was the case with the nephew-uncle pair described above. In that example, one count was zero and the other nonzero. Under conditions set out in Appendix S, it was possible to determine that the disagreement in this case was due to the uncle not being able to identify the parent-child domains, and the nonzero count could be used. More detailed rules for reconciling differences between pair members are given in Appendix S.

If the attempt to reconcile differences in the household-level person counts between pair members was unsuccessful, upper and lower bounds within which the imputed value must reside were determined from the counts for each pair member, and the counts for the screener.

6.4.3 Creation of Imputation-Revised Household-Level Person Count Variables

Because of the difficulty in definitively determining household-level counts in many cases, imputation was not an uncommon proposition. As with the imputation of pair relationships and multiplicities, the imputation was conducted using the PMN method described in Appendix N. In this section, the application of PMN to the imputation of household-level person counts is described. Since only the household-level person count in the third stage was imputed for each household, the imputation was univariate in the sense that no sequential models were necessary. However, in some cases several variables were associated with a single model, as is described below.

6.4.3.1 Setup for Model Building

Household-level person counts of the domains listed in the introduction to Section 6.4 are defined for all respondents, regardless of what pair they belonged to, or even whether they were within a pair at all. Moreover, since a nonzero count did not depend upon the respondent being within the relevant age range, no logical constraints on age were necessary. However, the age of the respondent did have an impact on the final count. The biggest difference in the presence or absence of particular domains in a household was the presence of youth under 18. This was especially true if there were two or more youth in a household, in which case the household-level person counts would be considerably different from situations where this was not the case. As a result, both the pair and single-respondent samples were split by age. For the pairs, both pair members in one sample were under 18, and the remainder of pairs were in the other sample. For the single respondents, one sample consisted of respondents under 18, and the other consisted of the remainder. Separate imputations were conducted in the two samples.

Four separate imputations were conducted for the sibling-sibling domains, arising from four separate models. Unlike the multiplicity counts, no imputations were conducted for the younger focus sibling-sibling domains. Hence, only two of the sibling-sibling domains had household-level person counts imputed. However, four separate imputations were required since the sample was split into two subsamples for both pairs and single respondents.

The parent-child domains are hierarchical, so as with the multiplicities, the imputations could not have been conducted independently if consistency was to be maintained. Hence, like the multiplicities, only two models were fitted to the child-parent pairs, using just the counts for children (12 to 20 year olds). One set of models was for the number of the children who had at least one parent and the other set was for the number of parents who had a child aged 12 to 20. Using the predicted means from these models, a single donor pair was selected from which the household-level person counts were determined for 12-to-14, 12-to-17, and 12-to-20 child-parent pair domains. (The household-level person counts for the 15-to-17 child-parent domains were not determined, but could be easily derived.) Since the household-level person counts for specific domains were not dependent upon the pair relationship, it was not necessary to impute the parent-focus and child-focus counts separately, as we did with the multiplicities. Hence, although separate models were fit to the parent-focus and child-focus counts, the predicted values from these models were brought together in a single multivariate imputation.

The spouse-spouse household-level person counts were also hierarchical, in that knowledge of whether a spouse-spouse pair was in the household was required before one could say that the pair had children. It was somewhat more complicated than the parent-child hierarchical setup, however, as one model could not represent whether there was a spouse-spouse pair in the household, and whether that pair had children. As a result, the imputations were conducted in two stages, with the spouse-spouse pair imputations processed first, followed by the imputations of whether the pairs had children.

The first step for these models was to define respondents, nonrespondents, and the item response mechanism. For a pair or single respondent to be considered complete, the household-level person counts had to be nonmissing for all the variables being imputed. For the parent-child pair domains, this meant that the household-level person count had to nonmissing for the parent-focus and child-focus 12-to-20 domains. Nonmissing household-level person counts for these domains automatically guaranteed nonmissing counts for the subset parent-child domains. A single response propensity adjustments was calculated for all the parent-child domains within each subsample; separate response propensity adjustments were calculated for the remainder of domains. Separate response propensity adjustments were calculated for pairs and single respondents. For pairs, these adjustments were calculated in order to make the respondent pair weights representative of the entire sample of pairs. For single respondents, household weights were used. The adjustments were calculated in order to make the respondent household weights representative of the entire sample of households that were not part of a pair. Because the spouse-spouse imputations were conducted in two stages, the response propensity adjustment for the spouse-spouse with children domain adjusted weights to be representative of all spouse-spouse pairs. Missing counts for the spouse-spouse with children domain were not imputed until it was known definitively, after the hot deck step of the PMN imputation, whether a household had spouse-spouse pairs.

6.4.3.2 Model-Building and Determination of Predicted Means

The PMN method is a two-step process. The first step is the modeling step, followed by a hot-deck step where imputed values replace the missing household-level person counts. The different attributes of the models are described in this subsection.

Response categories. The response categories for the household-level person count final response models were simply the household-level person counts, corresponding to each domain, among the complete data cases. In some cases, two family units were in a household. If these resulted in unusual household-level person counts, they were excluded from the modeling step, and were considered nonrespondents for the purposes of weight adjustment. No predicted mean was calculated in these cases. This occurred with the parent-child parent focus counts and the spouse-spouse-with-children counts. For the parent-child parent-focus counts, two family units sometimes resulted in counts of 3 or 4 parents, which were extremely rare levels. The response categories for the models in the case of the parent-child parent focus counts were, therefore, limited to 0, 1, or 2. With the spouse-spouse with children counts, having two spouse-spouse pairs with children under 18 was also an extremely rare category. The response categories that resulted for the spouse-spouse with children models were, therefore, 0 or 1. Households with two family units did not need to be excluded from the spouse-spouse models, since having two spouse-spouse pairs in a household, though not common, was not rare.

Covariates in Models. The same pool of covariates that was used for the multiplicity models was also used for the household-level person counts. The same dual set of models were fitted, according to whether the household composition age count variables existed or not. Naturally, the final set of covariates differed from the initial pool; the final set of covariates that were used in the models is given in Appendix Q.

Building of the Models. The household-level person counts could have a value of 0, which distinguished them from the multiplicities from a modeling point of view. For the child-focus parent-child domains, the count modeled was the number of parents, which had three values for reasons explained earlier: 0, 1, or 2. The model for spouse-spouse pairs also had three levels: 0, 1, or 2. Both of these models (within each subsample) were fitted as multinomial logistic models. Also for reasons stated earlier, the spouse-spouse with children models had only two levels (0 or 1), so binomial logistic models were fitted to those data. Poisson regression was used to fit the models for the household-level person counts corresponding to the sibling-sibling domains, as well as the child-focus parent-child domains. The data were underdispersed for a Poisson distribution, so the data had to be scaled using the observed variance.

Determination of the Predicted Mean. Although models were built using respondent pairs and single respondents where the household-level person counts were known definitively, predicted means were required for all pairs and for all respondents who were not part of a pair. Once the models were fitted, predicted means were determined for respondent pairs and single respondents, as well as item nonrespondents among pairs and singles, using the parameter estimates from the models.

6.4.3.3 Constraints on Hot-Deck Neighborhoods and Assignment of Imputed Values

Like the child-focus parent-child multiplicities, the spouse-spouse with children models had a binary response (no pairs or one pair). Hence, in the same manner as the child-focus parent-child multiplicity and the pair relationship variables, donors (among pairs and single respondents) in the hot-deck step of PMN for the counts associated with this domain were chosen with predictive means, if possible, within delta of the recipient's (whether a pair or single respondent) predicted mean. The value of delta varied depending on the value of the predicted means. The values of delta for predicted probabilities are given in Section 6.2.4.3. For the counts associated with other domains (parent-focus parent-child, sibling-sibling, and spouse-spouse domains), an error in the software meant that the formula given in Section 6.2.4.3 was applied to the predicted values, even though they were not probabilities. This prevented delta from being applied for these domains except in the strictest of circumstances. However, the other constraints that were placed on the neighborhoods ensured that donor pairs and recipient pairs were very much alike. Moreover, for variables where the delta constraint was used, it was often the first constraint to be loosened.

Wherever necessary and feasible, logical and likeness constraints (as defined in Section 6.2.4.3) were placed on the membership in the hot-deck neighborhoods. The hot-deck step and the accompanying constraints are described separately for each of the variables in turn.

In those instances where an imputed value could not be found after loosening all the likeness constraints, the imputed value was determined by doing a random imputation within bounds derived from the household composition. One of the situations where this occurred was when the household had two or more family units in the household. Even though the counts were not included in the models, no predicted means were calculated. (This occurred with the parent-focus parent-child counts, as well as the spouse-spouse with children counts.) Hence, instead of matching donors and recipients using predicted means, the imputed value was determined using the random imputation described earlier. Even though two-family households were included in the model for the child-focus parent-child counts, the resulting predicted means were not used. This was due to the fact that the parent-focus parent-child counts were in the same multivariate set as the child-focus parent-child counts, and the predicted means could not be used in the imputation of the parent-focus parent-child counts when two families were in the household.

6.4.3.3.1 Parent-child counts

Since parent-focus and child-focus counts were so closely related, a logical constraint was placed on donors such that if the parent-focus count was nonmissing and nonzero, then the child-focus count had to exceed 0. Similarly, a nonzero, nonmissing child-focus count required that the donor's parent-focus count exceed 0. If the child focus counts were missing, donors and recipients had to have the same number of household members in the age range corresponding to the domain of interest. (Donors had to have complete data on all the roster age variables.) The same constraint was applied if the parent-focus counts were missing but the child-focus counts were nonmissing, with an additional requirement: it had to be possible that no parent-child pairs existed in the household. (If it was known that there were parents in the household for the appropriate domain, it was not necessary to limit donors to have the same child-age composition as the recipient.) These were likeness constraints that were never loosened. In addition, if a recipient had two family units in the household, a regular hot deck imputation could not be done, as stated earlier. For all missing counts, the counts from the two pair members (in the case of pair recipients) and the household composition were used to create upper and lower bounds, provided valid roster information was available. These bounds acted as additional logical constraints. Besides delta, additional likeness constraints all involved the household size and additional constraints on the household composition, which are described in the following paragraph.

An attempt was made to match donors and recipients in each of three age ranges that are commonly associated with children aged 12 to 20: 26 to 34, 35 to 49, and 50 or over. This likeness constraint was applied whether the child-focus or the parent-focus count was missing; however, its application in the case of a missing child-focus count and nonmissing parent-focus count required an additional condition: it had to be possible that no parent-child pairs existed in the household. (If it was known that there were children in the household who belonged to parents, it was not necessary to limit donors according to the parent age ranges.) A looser form of this constraint was to collapse the 26-to-34 and 35-to-49 age ranges into a single age range, and drop the 50-or-over constraint. Other household composition constraints required donors and recipients to have the same number of household members under 12 years old, and between the ages of 18 and 25 (inclusive).

The likeness constraints were loosened in the following order (where applicable) (1) abandon the neighborhood, and choose the donor with the closest predicted mean or means; (2) abandon the requirement that donors and recipients have the same household size; (3) abandon the requirement that donors and recipients have the same number of household members under 12, between 18 and 25 (inclusive), and 50-or-older, and collapse the 26-to-34 and 35-to-49 age constraints; (4) remove the 26-to-49 age constraint.

6.4.3.3.2 Sibling-sibling counts

For all missing counts, the counts from the two pair members (in the case of pair recipients) and the household composition were used to create upper and lower bounds, provided valid roster information was available. These bounds acted as additional logical constraints. If the sibling-sibling counts were missing, donors and recipients had to have the same number of household members in the age range corresponding to the domain of interest. (Donors had to have complete data on all the roster age variables.) Since imputations for the household-level person counts were only done on the sibling-sibling domains with the older sibling as focus, this meant that donors and recipients had to have the same number in the household aged 15 to 17 (for the 12-to-14/15-to-17 domain) or aged 18 to 25 (for the 12-to-17/18-to-25 domain). This was a likeness constraint that could be loosened to a logical constraint: the imputed count could not exceed the recipient's number of household members in the relevant age range. An additional likeness constraint recognized the correlation between parent-child domains and sibling-sibling domains; that is, the presence of parent-child domains in the household meant that a sibling-sibling domain would be more likely. Hence, donors and recipients both either had to have parent-child domains in the household, or not have such domains. Other likeness constraints were related to the household composition.

In addition to matching donors and recipients on household size, they also had to match on the number of household members under 12, and the number of household members within the younger sibling's age range. The age constraints corresponding to the age ranges of the siblings could be loosened so that the counts for the donor and recipient for the older sibling's age range had to be both zero or both nonzero.

The likeness constraints were loosened in the following order (where applicable) (1) abandon the neighborhood, and choose the donor with the closest predicted mean or means; (2) abandon the requirement that donors and recipients have the same household size, and the same number of household members under 12; (3) abandon the requirement that donors and recipients have the same number of household members under 12 years old; (4) remove age constraints corresponding to the age ranges of the siblings, so that the only age constraint is logical: donors' counts cannot exceed the total number in the older sibling's age range. At the same time, the constraint that required donors and recipients to have the same status with regard to parent-child domains was removed.

6.4.3.3.3 Spouse-spouse counts

For all missing counts, the counts from the two pair members (in the case of pair recipients) and the household composition were used to create upper and lower bounds, provided valid roster information was available. These bounds acted as logical constraints. In addition, if a recipient had two family units in the household, a regular hot-deck imputation could not be done, as stated earlier. The rest of the likeness constraints all used information from the household composition, with recognition of the fact that the vast majority of spouse-spouse pairs were male-female pairs.

For the spouse-spouse pairs, which also included partner-partner pairs, the constraints attempted to match donors and recipients as much as possible in their household age and gender pattern. This included some likeness constraints that were never loosened: both donors and recipients were required to have the same number of household members under 18 if that number was 0, 1, or 2. If the recipient had 2 or more members in his or her household under 18, the donor also had to have 2 or more. In addition, donors and recipients had to have the same number of household members of ages 15 or more, and the same number of males of ages 15 or more.

Likeness constraints that were loosened also were related to the age and gender composition of the household. In particular, donors and recipients had to match their household size, and the number of household members, as well as males, within the age ranges of 18 to 25, 26 to 34, 35 to 49, and 50 or over. Looser forms of these constraints required the same number in the household, and the same number of males in the household, within the age ranges of 18 to 34 and 26 to 49.

The likeness constraints were loosened in the following order (where applicable) (1) abandon the neighborhood, and choose the donor with the closest predicted mean or means; (2) abandon the requirement that donors and recipients have the same household size, and abandon the requirement that the donors and recipients had to have the same number of household members (male or female) within the age ranges of 18 to 25, 26 to 34, 35 to 49, and 50 or over;

(3) abandon the requirement that donors and recipients have the similar number of household members under 18 as described above, and loosen the requirement that donors and recipients had to have the same number of males within the age ranges of 18 to 25, 26 to 34, and 35 to 49, so that donors and recipients were required to have the same number of males in the age range of 18 to 34, as well as 26 to 49; (4) abandon the requirement that donors and recipients have the same number of males in the age range from 18 to 34.

6.4.3.3.4 Spouse-spouse-with-children counts

The constraints for the spouse-spouse-with-children counts were exactly the same as the spouse-spouse constraints, with one exception. When the requirement that donors and recipients have a similar number of household members under 18 was abandoned, another constraint replaced it and was never loosened. This constraint required that if the recipient did not have anyone in the household under 18, then the same should be true of the donor; however, if the recipient did have someone in the household under 18, the donor should too. Presumably this constraint was not necessary, since no imputation would be required if it was known that no children were in the household.

6.4.3.4 Additions to Analytic File

The imputation-revised versions of the parent-child household-level person count variables are called IRHPCCxx and IRHPCPxx, where the final C and P refer to the focus in the domain. The "xx" refers to the age range of the children, which is the upper bound, since the lower bound is always 12. The edited version of these variables, HCPCCxx and HCPCPxx, were also released to the analytic file. The sibling-sibling imputation-revised variables are called IRHSxxxx, where the second two x's in the "xxxx" refer to the upper bound of the age range corresponding to the focus pair member, and the first two x's refer to the upper bound of the age range corresponding to the remaining pair member. The edited version of these variables is given by HCSxxxx. Finally, the imputation-revised versions of the spouse-spouse counts are given by IRHCSPSP and IRHCSPWC for the spouse-spouse and spouse-spouse with children counts, respectively. The edited versions of these spouse-spouse counts are called HCSPSP and HCSPSPWC. The imputation indicators were also released to the analytic file, with II prefixes instead of IR prefixes.

7. Weight Calibration at Questionnaire Dwelling Unit (QDU) and Pair Levels

The 2002 National Survey of Drug Use and Health[14] (NSDUH) was based on probability sampling so that valid inferences can be made from survey findings about the target population. Probability sampling refers to sampling in which every unit on the frame is given a known, nonzero probability for inclusion in the survey. This is required for unbiased estimation of the population total. The assumption of nonzero inclusion probability for every pair of units in the frame is also required for unbiased variance estimation. The basic sampling plan involved three stages of selection across two phases of design: within Phase I, (1) the selection of subareas or segments (comprised of U.S. Census blocks) within State field interviewer (FI) regions; (2) the selection of dwelling units (DUs) within these subareas; and, finally, within Phase II, (3) the selection of eligible individuals within DUs. Specific details of the sample design and selection procedures for the sample can be found in the 2002 NSDUH Sample Design report (Odom et al., 2004).

As part of the post-survey data-processing activities, analysis weights that reflected the selection probabilities from various stages of the sample design were calculated for respondents. These sample weights were adjusted at the DU (screening sample), questionnaire dwelling unit, person, and paired respondent level (the latter three all based around the drug questionnaire sample) to account for bias due to extreme values (ev), nonresponse (nr), and undercoverage (i.e., poststratification [ps]).

The final sample weights for Phase I screener dwelling units (SDU) and Phase II QDU, person, and pair levels for the 2002 samples consist of products of several factors, each representing either a probability of selection at some particular stage or some form of ev, nr, or ps calibration adjustment. In the following sections, we describe the QDU and pair weight components in greater detail. In summary, the first nine factors are defined for all SDUs and reflect the fully adjusted SDU sample weight. The remaining components branch to reflect QDU and pair selection probabilities as well as additional adjustments for ev, nr, and ps. Note that the final QDU and pair weights for the 2002 NSDUH sample are the product of all weight components for each type of sample, illustrated in Exhibits 7.1 and 7.2.

For QDU data, generalized exponential modeling (GEM) calibration modeling was applied by partitioning the data into four groups of States: Northeast, South, Midwest, and West, based on Census regions, in the interest of computational feasibility. Previous experience showed that with current computing power, the large number of variables and records prevented any further reduction of modeling groups.

For pair data, GEM modeling was initially applied by partitioning the pair data into four groups based on U.S. Census regions. However, there were not enough observations in each group to fit a comprehensive model to reduce bias. Alternatively, a single model was attempted

---

[14] This report presents information from the 2002 National Survey on Drug Use and Health (NSDUH). Prior to 2002, the survey was called the National Household Survey on Drug Abuse (NHSDA).

for the whole pair data, but was not practical due to computational limitations. A compromise approach was adopted by combining U.S. Census regions into two groups: Northeast with South, and Midwest with West. This grouping proved both manageable and desirable as it assisted in bias reduction, ease of modeling, and workload reduction. Exhibit 7.3 provides more details of the data partition for GEM modeling.

Unlike the 1999, 2000, and 2001 NHSDAs, for pair data, it may be noted that the built-in feature of ev control in GEM was used for each adjustment step. In 1999, 2000, and 2001, the built-in ev control feature was not used until the final respondent pair ev adjustment step. The reason for this is that the definition for ev domain was not finalized before the pair data calibration process was begun.

Exhibit 7.1 Summary of 2002 NSDUH QDU Sample Weight Components

Phase I Screener Dwelling Unit Level

	Design Weight Components
#1	Inverse Probability of Selecting Segment
#2	Quarter Segment Weight Adjustment
#3	Subsegmentation Inflation Adjustment
#4	Inverse Probability of Selecting SDU
#5	Subsampling of Added SDU Adjustment
#6	SDU Release Adjustment

#7	SDU Nonresponse Adjustment ***(res.sdu.nr)****
#8	SDU Poststratification Adjustment ***(res.sdu.ps)****
#9	SDU Extreme Value Adjustment ***(res.sdu.ev)****

Phase II Questionnaire Dwelling Unit Level

	Design Weight Component
#10	Inverse of Selection Probability of at Least One Person in the Dwelling Unit

#11	Selected QDU Poststratification to SDU-based Control Totals ***(sel.qdu.ps)****
#12	Respondent QDU Nonresponse Adjustment ***(res.qdu.nr)****
#13	Respondent QDU Poststratification to SDU-based Control Totals ***(res.qdu.ps)****
#14	Respondent QDU Extreme Value Adjustment ***(res.qdu.ev)****

* These adjustments use the generalized exponential model (GEM), which also involves pre- and post-processing in addition to running the GEM macro. See Exhibit 4.1 (Chen et al., 2004). For computational feasibility, all weight adjustments were done using the four model groups based on U.S. Census regions defined in Exhibit 7.3.

Exhibit 7.2 Summary of 2002 NSDUH Person Pair Sample Weight Components

Phase I Screener Dwelling Unit Level

	Design Weight Components
#1	Inverse Probability of Selecting Segment
#2	Quarter Segment Weight Adjustment
#3	Subsegmentation Inflation Adjustment
#4	Inverse Probability of Selecting SDU
#5	Subsampling of Added SDU Adjustment
#6	SDU Release Adjustment

#7	SDU Nonresponse Adjustment ***(res.sdu.nr)****
#8	SDU Poststratification Adjustment ***(res.sdu.ps)****
#9	SDU Extreme Value Adjustment ***(res.sdu.ev)****

Phase II Person Pair Level

	Design Weight Component
#10	Inverse of Selection Probability of a Person Pair in the SDU

#11	Selected Pair Poststratification to SDU-based Control Totals ***(sel.pr.ps)****
#12	Respondent Pair Nonresponse Adjustment ***(res.pr.nr)****
#13	Respondent Pair Poststratification Adjustment to SDU-based Control Totals ***(res.per.ps)****
#14	Respondent Pair Extreme Value Adjustment ***(res.per.ev)****

* These adjustments use the generalized exponential model (GEM), which also involves pre- and post-processing in addition to running the GEM macro. See Exhibit 4.1 (Chen et al., 2004). For computational feasibility, all weight adjustments were done using the two model groups (Northeast + South, Midwest + West) based on combined U.S. Census regions defined in Exhibit 7.3.

Exhibit 7.3 U.S. Census Regions/Model Groups

Model Group	Census Region
QDU	
1	**Northeast (9 States)** Connecticut, Maine, Massachusetts, New Hampshire, New Jersey, New York, Pennsylvania, Rhode Island, Vermont
2	**Midwest (12 States)** Illinois, Indiana, Iowa, Kansas, Michigan, Minnesota, Missouri, Nebraska, North Dakota, Ohio, South Dakota, Wisconsin
3	**South (16 States and the District of Columbia)** Alabama, Arkansas, Delaware, District of Columbia, Florida, Georgia, Kentucky, Louisiana, Maryland, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia, West Virginia
4	**West (13 States)** Alaska, Arizona, California, Colorado, Idaho, Hawaii, Montana, Nevada, New Mexico, Oregon, Utah, Washington, Wyoming
Pair	
1	**Northeast + South (25 States and the District of Columbia)** Alabama, Arkansas, Connecticut, Delaware, District of Columbia, Florida, Georgia, Kentucky, Louisiana, Maine, Massachusetts, Maryland, Mississippi, New Hampshire, New Jersey, New York, North Carolina, Oklahoma, Pennsylvania, Rhode Island, South Carolina, Tennessee, Texas, Vermont, Virginia, West Virginia
2	**Midwest + West (25 States)** Alaska, Arizona, California, Colorado, Hawaii, Idaho, Illinois, Indiana, Iowa, Kansas, Michigan, Minnesota, Missouri, Montana, Nebraska, Nevada, New Mexico, North Dakota, Ohio, Oregon, South Dakota, Utah, Washington, Wisconsin, Wyoming

Table 7.1 Sample Size by Model Group at QDU and Pair Level

Model Group	1999		2000		2001		2002	
	Selected QDUs	Completed QDUs	Selected QDUs	Completed QDUs	Selected QDUs	Completed QDUs	Selected QDUs	Completed QDUs
QDU								
Northeast	12,667	9,289	14,415	11,341	14,208	11,155	11,436	9,724
South	20,790	16,459	20,765	17,215	19,814	16,029	15,582	13,489
Midwest	18,348	13,926	19,055	15,027	18,903	14,804	17,121	14,877
West	15,580	12,147	15,287	12,341	13,772	11,146	11,547	9,998
Total	67,385	51,821	69,522	55,924	66,697	53,134	55,686	48,088
Model Group	**1999**		**2000**		**2001**		**2002**	
	Selected Pairs	Completed Pairs	Selected Pairs	Completed Pairs	Selected Pairs	Completed Pairs	Selected Pairs	Completed Pairs
Pair								
Northeast+South	10,727	7,100	10,996	7,879	11,436	7,869	12,463	10,005
Midwest+West	11,771	7,785	11,443	7,961	11,612	7,926	12,432	10,033
Total	22,498	14,885	22,439	15,840	23,048	15,795	24,895	20,038

7.1 Phase I SDU-Level Weight Components

A total of nine weight components for the SDU level correspond to selection probabilities and nonresponse, poststratification, and extreme value adjustment factors. The first six components in the Phase I sample weights reflect the probability of selecting the DUs. These components were derived from (1) the probability of selecting the geographic segment within each State FI region, (2) a quarter segment weight adjustment, (3) a subsegmentation inflation factor, (4) the probability of selecting a DU from within each counted and listed sampled segment, (5) the probability of inclusion of added DUs, and (6) DU percent release adjustment. The three remaining weight components, #7–#9, are GEM calibration adjustments accounting for (7) DU nonresponse at the screening level, (8) DU poststratification to Census controls and (9) DU-level extreme value adjustment, although in 2002 extreme value adjustment at this stage was deemed unnecessary, and so Weight Component #9 was set to 1 for all respondent DUs. The person-level, QDU-level, and person pair–level weights use the product of the above nine weight components as the common initial weight before further adjustments. For more detailed information on Weight Components #1 and #3 through #6, refer to the 2002 NSDUH Sample Design Report (Odom et al., 2004), and for more detail on Weight Components #2 and #7-#9, see the 2002 Person-Level Sampling Weight Calibration report (Chen et al., 2004).

7.2 QDU Weight Components

7.2.1 QDU Weight Component #10: Inverse of Selection Probability of at Least One Person in the Dwelling Unit

Selection of a questionnaire dwelling unit from all completed SDUs is based on the outcome of a variant of Brewer's method, which may select zero, one, or two persons. Any pair of survey-eligible residents within the dwelling unit had some known, nonzero chance of being selected for the survey. The value for weight component #10 is equal to the inverse of the probability that at least one person in the dwelling unit is selected (see Section 2.2 for details).

7.2.2 QDU Weight Component #11: Selected QDU Poststratification to SDU-based Control Totals

This factor adjusts the weights for selected QDUs to the SDU-based control totals. The SDU-based control totals are obtained by using the calibrated SDU weights. This adjustment step provides more stable controls for the subsequent nonresponse adjustment (component #12). Exhibit 4.1 lists the initially proposed variables for GEM modeling. The predictor variables are either 0/1 indicators or counting variables representing the number of persons who fall into a given demographic domain. The counting variables are derived from the screener demographic information. It may be noted that during screening, the only required demographic information was the age of each person rostered. Thus, other demographic information necessary for weight calibration, such as race/ethnicity and gender may be missing for certain rostered eligible persons, and so imputation was done to replace this missing data. For more details on the imputation of screener demographic information, see Chen et al. (2004).

The details on the predictor variables retained in the model and model summary statistics can be found in Appendix C.

7.2.3 QDU Weight Component #12: Respondent QDU Nonresponse Adjustment

This nonresponse adjustment step accounts for the failure to obtain respondent person(s) from each and every selected QDU. The same set of initially proposed predictor variables were used as for the previous adjustment (#11).

See Appendix C for more details on the predictor variables retained in the model and model summary statistics.

7.2.4 QDU Weight Component #13: Respondent QDU Poststratification to SDU-based Control Totals

This final poststratification for all respondent QDUs utilized the same set of initially proposed predictor variables as previous adjustments. The corresponding control totals were obtained from SDU level sample, as for the weight component #11.

See Appendix C for more details on the predictor variables retained in the model and model summary statistics.

7.2.5 QDU Weight Component #14: Respondent QDU Extreme Value Adjustment

The extreme weight proportions for the final poststratified weights were acceptably low, and so it was decided that the extreme value adjustment was not needed. Weight component #14 was set to one for each responding QDU.

7.3 Pair-Level Weight Components

Exhibit 4.2 lists the initially proposed predictor variables for the following adjustment steps via GEM.

7.3.1 Pair Weight Component #10: Inverse of Selection Probability of a Person Pair in the DU

Selection of pairs of individuals from all eligible persons residing within the dwelling unit is based on the outcome of a variant of Brewer's method, which may select zero, one, or two persons. Any pair of survey-eligible residents within the dwelling unit has some known, nonzero chance of being selected for the survey. When two persons are selected a pair is formed. The pair selection probability is determined by the formula in Chapter 2. This weight component is the inverse of the selection probability discussed above.

7.3.2 Pair Weight Component #11: Selected Pair Poststratification to SDU-based Control Totals

Similar to QDU weight component #11, this step was motivated by the consideration that the larger sample of all possible pairs provides more stable control totals for the respondent pair nonresponse adjustment. The weights of selected pairs were poststratified to the control totals that derived from calibrated SDU weights of all possible pairs. The pair-level demographic variables for all selected pairs, such as pair age group, pair race, etc., were derived from screener demographic information.

The details on the predictor variables retained in the model and model summary statistics can be found in Appendix H.

7.3.3 Pair Weight Component #12: Respondent Pair Nonresponse Adjustment

If both persons in the selected pair completed interviews successfully, the pair then was considered a respondent pair. This adjustment step accounts for failure to obtain respondent pairs from all selected pairs. In this step, respondent pair weights were adjusted to the control totals based on the full sample of selected pairs. Due to the low response rate of person pairs, this step had a relatively large adjustment on the weights. The same set of proposed predictor variables was used as for the pair weight component #11. Similar to weight component #11, the pair level demographic variables for all selected pairs, such as pair age group, pair race etc., were derived from screener demographic information.

See Appendix H for more details on the predictor variables retained in the model and model summary statistics.

7.3.4 Pair Weight Component #13: Respondent Pair Poststratification to SDU-based Control Totals

This final poststratification utilized the same set of initially proposed predictor variables as previous adjustment steps. In addition, ten pair relationship domain-level indicator variables were added to the set of covariates. The control totals for GEM calibration were derived from the

SDU sample of all possible pairs of eligible persons, as for weight component #11. The calibration control totals for these ten domains used household-level person counts and the final QDU weights. As mentioned in the introduction, use of these household-level count totals for pair relationship domains in GEM calibration provided Hajek-type weight adjustment in the interest of obtaining more stable estimates. In setting up calibration covariates, multiplicity factors were needed. These factors, as discussed in the introduction, are used in constructing estimates for person-level parameters based on pair-related drug behavior. The factors depend on the pair domains of interest. For a selected set of pair domains, multiplicity factors are provided along with the pair-level analysis weights. See Chapter 6 for more detail on creation of and imputation of missing values in the pair relationship, multiplicity, and household-level person counts. See Chapter 4 for more detail on the use of multiplicities and household-level person counts in poststratification.

Unlike weight components #11 and #12, demographic covariates were based on data from the questionnaire instead of information pulled from the dwelling unit screener.

For more details on the predictor variables retained in the GEM model and model summary statistics, see Appendix H.

7.3.5 Pair Weight Component #14: Respondent Pair Extreme Weight Adjustment

We checked the extreme weight proportions for the weights up to weight component #13 using the extreme weight domains (see Section 5.2). Even though the previous adjustment steps utilized the built-in extreme weight control feature of GEM, the extreme weight proportions were still high enough to cause concern that they might produce unreliable estimates. Therefore, the extreme weight adjustment via GEM was implemented, using the same final set of predictor variables kept in the model for the previous pair weight component #14. This step was successful in reducing the extreme weight proportion in all model groups, across all three years. For details, see Appendix J, Tables J3, J6, and J9.

8. Evaluation of Calibration Weights

During the weight calibration process, several criteria for quality control were implemented to assess model adequacy. In this chapter, we describe the individual procedures and a summary of their results. All tables referred to in this chapter can be found in Appendices D–G and I–L.

8.1 Response Rates

Table D in Appendix D displays the final selected and responding questionnaire dwelling unit (QDU) sample sizes from the 2002 National Survey of Drug Use and Health[15] (NSDUH), for various national domains. This table also shows the weighted response rates. Comparing response patterns across domains shows that the domain with the greatest variation in response rates was household type, which ranged from 76.52 percent for households composed of persons aged 26+ only, to 92.06 percent for households composed of persons aged 12 to 17 only. Most domains reflect the overall 80.60 percent response rate, with only a few rates dropping below 80 percent, although the highest response rate is significantly different at 97.58 percent, from the group-level of the group quarters variable. The lowest response rate came from Census Region level Northeast, with 77.28 percent.

Table I in Appendix I displays the final selected and responding pair-level sample sizes from the 2002 NSDUH, for various national domains. Due to the nature of the pair data, the response rates were lower in all domains examined than at the QDU level, with an overall response rate of 71.50 percent. The response rates range from a low of 48.75 percent in the pair race other category to a high of 92.00 percent from what was the high response domain from the QDU level, the group-level of group quarters. This extreme range of response rates is probably due to a combination of small sample sizes and response burden due to selection of pairs within households among various domains. Like at the QDU level, the top response rates are among the younger respondents (as measured by household type for the QDU data and pair age for the pair data). This pattern may be related to the relatively high response rates in the group level of the variable group quarters, as it includes college dormitories.

8.2 Proportions of Extreme Value and Outwinsor Weights

During the stages of modeling adjustments (i.e., nonresponse [nr] and poststratification [ps]), one major issue of concern when deciding the adequacy of a particular model was the extent of the resulting proportion of extreme value (ev) and outwinsor weights; see Sections 5.1 and 5.2 for these definitions. For each weight adjustment step, these proportions are computed before and after the step for various domains. Prior to adjustment, the product of all weight components is used to compute proportions of evs and outwinsors, while after the adjustment the product includes the new adjustment factor. If the proportion of evs and outwinsors are deemed high, a separate ev treatment step after ps could be performed. This was done for the pair-level

[15] This report presents information from the 2002 National Survey on Drug Use and Health (NSDUH). Prior to 2002, the survey was called the National Household Survey on Drug Abuse (NHSDA).

weights. Details of this step are explained in Section 7.3.5. A separate ev treatment step was deemed unnecessary for the QDU-level weights.

Tables E1 to E3 and Tables J1 to J3 present percentages of evs at the QDU level and the pair level, respectively, for various domains. Unweighted percentages are the percentage of actual counts of units defined as evs relative to the total sample size. Weighted percentages reflect the percentage of total ev weights relative to the total sample weight, while outwinsor percentages represent the total amount of residual weight when the weights are trimmed to the critical values (used for ev definition), relative to the total sample weight. For evaluation purposes, the outwinsor percentage is considered the most important of the three percentages, as this gave a measure of the impact of winsorization (or trimming) of ev weights, if we performed this treatment. See Sections 5.1 and 5.2 for the domains that were used to define extreme values.

8.3 Slippage Rates

The slippage rate for a given domain is defined as the relative percentage difference between the sampling weights and the external control totals, both before and after ps. The control totals for QDU and person-pair ps are derived from the screener dwelling unit (SDU) weights which were poststratified to U.S. Census population estimates (Chen et al., 2004). Table F displays QDU national domain-specific weight sums for both before and after ps, as well as the desired totals to be met through ps. Table K shows the same for the pair sample. These tables also show the relative percentage difference, or the amount of adjustment necessary (positive or negative) to meet the desired totals. The first relative difference is used explicitly during the ps modeling procedure to identify potential problems for convergence. Large differences in domains with relatively small sample sizes are indicative of potential large adjustment factors, which may cause problems in convergence while satisfying bound constraints. The reason is that adjustments required for one domain may have an adverse effect on another domain when a unit belongs to both.

As an example, consider that Table F, for the 2002 QDU domain household size of one, indicates a sample size of 5,884 with a total design-based weight of 29,809,582 and a Census total of 29,297,697 with an initial slippage rate of 1.75 percent, which would imply a common weight adjustment of ≈1.02, if this were the only calibration control. Similarly, looking at pair data, in Table K, the pair race category white and other has a sample size of 548, a design-based weight of 6,192,585, and a Census total of 3,816,560. The resultant required adjustment would be ≈0.62 if this were the only control. However, in the generalized exponential model (GEM), all controls are simultaneously satisfied under a complex algorithm that allows for different adjustment factors for different units.

8.4 Weight Adjustment Summary Statistics

Tables G1 to G2 and L1 to L3 display summary statistics on the product of weight components before and after all stages of adjustment for the QDU and person pair respectively. The summary statistics include sample size (n), minimum (min), maximum (max), median (med), 25^{th} percentile (Q1), 75^{th} percentile (Q3), and the unequal weighting effect (UWE). Note in Tables L2 and L3 the sample size for pair age group, pair race, and pair gender are slightly different. This is because those variables were defined using screening demographic information in the nonresponse adjustment of respondent pairs, while in the poststratification of respondent pairs they were defined from questionnaire demographic information. Because UWE is directly affected by weight adjustment factors and extreme weights, these values, along with the percentage of extreme weights, as noted in Section 8.2, were used as guidelines for determining model adequacy.

8.5 Sensitivity Analysis of Drug Use Estimates

It is known that, in general, there is a trade-off between bias reduction and variance reduction. For instance, with GEM (for nr or ps), enlarging a simple model (such as the one with only main effects) has the potential of further reducing the bias. At the same time, this enlargement may also be associated with a corresponding increase in the variance of the estimate, due to additional variability caused by estimating the model parameters. To check for possible overfitting of the GEM model, we conducted a sensitivity analysis for respondent QDU poststratification for the QDU weights, respondent pair poststratification, and extreme weight adjustment for the person pair weights. A simple baseline model was fitted with the same bounds and maximum number of iterations as was used for the chosen (more complex) final model. We then looked for substantial changes in point estimates and standard errors (SEs). For the QDU weights, some household-level characteristics were selected, such as family income, number of kids, whether the household had health insurance coverage, and number of elders living in the household. The estimates and SEs are displayed in Table 8.1. For the person pair weights, selected licit and illicit drug use prevalence rates of 12 to 17 year olds were calculated from parent-child pairs, and estimates and SEs of the estimates based on pair weights are shown in Tables 8.2a to 8.7b.

As seen in Table 8.1, the estimates and their SEs for the two models (baseline and the final) are generally similar to each other for the QDU weights. However, among the person pair estimates and SEs there are some differences, but they do not seem significant in general.

Since the sensitivity analyses for both QDU- and pair-level calibrated weights seem to indicate that adding more covariates does not introduce an undesirable degree of instability in the estimates or their standard errors, the final, more complex GEM models were deemed reasonable.

Table 8.1 Estimates of Totals and SEs for Domains of Interest for Based on QDU Sample for 2002

	2002		
Domain	***n***	**Baseline**[1]	**Final**[2]
Households with Family Income			
$0 - $10,000	4,875	9,841,344 (328,749)	9,841,037 (329,462)
$10,000 - $20,000	6,527	15,687,345 (390,695)	15,662,818 (389,924)
$20,000 - $30,000	6,278	15,026,764 (374,418)	15,019,243 (374,532)
$30,000 - $40,000	6,172	14,385,709 (340,736)	14,381,866 (340,488)
$40,000 - $50,000	6,110	13,631,682 (319,675)	13,626,091 (319,892)
$50,000 - $75,000	8,498	18,643,995 (364,724)	18,659,638 (366,037)
$75,000+	9,628	22,276,486 (507,185)	22,302,632 (508,214)
Households with # of Kids (<18)			
0	18,887	68,977,335 (937,851)	68,977,868 (939,066)
1	11,771	16,822,892 (266,969)	16,815,037 (267,499)
2	10,521	15,068,999 (267,292)	15,059,713 (268,301)
3	4,611	6,085,091 (144,203)	6,097,080 (145,146)
4+	2,298	2,539,008 (83,025)	2,543,628 (83,670)
Households with Insurance Coverage			
Yes	39,653	94,250,960 (1,007,446)	94,258,548 (1,008,253)
No	8,435	15,242,365 (314,377)	15,234,778 (314,924)
Households with # of Elders (65+)			
0	44,011	85,323,328 (905,226)	85,342,474 (906,683)
1	2,815	16,084,959 (528,278)	16,069,317 (528,204)
2	1,240	7,973,627 (328,474)	7,968,651 (329,159)
3+	22	111,411 (29,846)	112,884 (30,148)

NOTE: Standard errors of estimated totals are provided in parentheses.

[1] Baseline refers to the weight obtained from using a main effects only model for the last step of calibration, res.qdu.ps, and a full model for those preceding.

[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.2a Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Tobacco and Alcohol among Mother-Child (12 to 17) Pairs by Mother Use for 2002

		2002		
Drug	**Mother User**	***n***	**Baseline**[1]	**Final**[2]
Alcohol				
Lifetime	Yes	2,047	42.32 (1.70)	42.31 (1.72)
	No	254	26.61 (4.87)	27.13 (4.71)
	Overall	2,301	39.90 (1.57)	39.97 (1.57)
Past Year	Yes	1,611	36.83 (1.83)	36.92 (1.87)
	No	690	20.96 (2.58)	20.88 (2.49)
	Overall	2,301	31.49 (1.52)	31.50 (1.51)
Past Month	Yes	1,186	21.11 (1.89)	21.45 (1.94)
	No	1,115	10.02 (1.17)	10.09 (1.22)
	Overall	2,301	15.53 (1.10)	15.73 (1.13)
Cigarettes				
Lifetime	Yes	1,707	35.31 (1.83)	35.36 (1.86)
	No	594	21.44 (2.86)	21.63 (2.78)
	Overall	2,301	31.28 (1.50)	31.44 (1.52)
Past Year	Yes	760	28.45 (2.61)	28.39 (2.65)
	No	1,541	14.99 (1.42)	15.04 (1.46)
	Overall	2,301	18.82 (1.27)	18.85 (1.30)
Past Month	Yes	687	18.66 (2.23)	18.46 (2.22)
	No	1,614	8.26 (1.04)	8.13 (1.05)
	Overall	2,301	10.88 (0.96)	10.75 (0.96)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.2b Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Tobacco and Alcohol among Father-Child (12 to 17) Pairs by Father Use for 2002

		2002		
Drug	**Father User**	***n***	**Baseline[1]**	**Final[2]**
Alcohol				
Lifetime	Yes	1,466	42.09 (2.16)	42.05 (2.18)
	No	97	23.00 (8.42)	23.05 (7.98)
	Overall	1,563	40.45 (2.26)	40.47 (2.23)
Past Year	Yes	1,190	37.02 (2.34)	37.06 (2.36)
	No	373	19.58 (4.13)	19.38 (3.97)
	Overall	1,563	32.63 (2.16)	32.66 (2.14)
Past Month	Yes	987	21.20 (2.41)	21.06 (2.40)
	No	576	11.86 (2.45)	11.81 (2.35)
	Overall	1,563	17.76 (1.97)	17.67 (1.92)
Cigarettes				
Lifetime	Yes	1,253	33.53 (2.33)	32.88 (2.30)
	No	310	19.37 (3.38)	19.71 (3.41)
	Overall	1,563	30.78 (2.07)	30.31 (2.05)
Past Year	Yes	525	23.43 (3.10)	22.75 (3.04)
	No	1,038	13.80 (1.94)	13.72 (1.88)
	Overall	1,563	16.79 (1.57)	16.52 (1.54)
Past Month	Yes	465	14.31 (2.24)	13.83 (2.15)
	No	1,098	7.75 (1.56)	7.63 (1.48)
	Overall	1,563	9.59 (1.18)	9.37 (1.14)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.3a Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Marijuana or Any Illicit Drug among Mother-Child (12 to 17) Pairs by Mother Use for 2002

		2002		
Drug	**Mother User**	***n***	**Baseline**[1]	**Final**[2]
Any Illicit				
Lifetime	Yes	1,389	34.03 (2.12)	34.34 (2.15)
	No	912	20.28 (2.42)	20.04 (2.33)
	Overall	2,301	28.13 (1.59)	28.24 (1.59)
Past Year	Yes	243	32.03 (4.07)	32.20 (4.13)
	No	2,058	17.21 (1.24)	17.23 (1.25)
	Overall	2,301	18.52 (1.21)	18.56 (1.22)
Past Month	Yes	116	25.63 (5.18)	25.15 (5.16)
	No	2,185	8.89 (0.91)	8.89 (0.91)
	Overall	2,301	9.51 (0.89)	9.49 (0.89)
Marijuana				
Lifetime	Yes	1,269	23.01 (1.68)	23.13 (1.71)
	No	1,032	10.01 (1.77)	10.06 (1.69)
	Overall	2,301	16.80 (1.20)	16.90 (1.19)
Past Year	Yes	136	30.96 (5.58)	31.42 (5.66)
	No	2,165	11.07 (0.92)	11.16 (0.94)
	Overall	2,301	12.03 (0.92)	12.13 (0.94)
Past Month	Yes	69	25.04 (6.86)	24.34 (6.84)
	No	2,232	5.53 (0.65)	5.61 (0.68)
	Overall	2,301	5.94 (0.66)	6.00 (0.68)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.3b Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Marijuana or Any Illicit Drug among Father-Child (12 to 17) Pairs by Father Use for 2002

		2002		
Drug	**Father User**	***n***	**Baseline**[1]	**Final**[2]
Any Illicit				
Lifetime	Yes	1,035	30.59 (2.50)	30.82 (2.48)
	No	528	18.80 (2.69)	18.81 (2.71)
	Overall	1,563	26.50 (1.94)	26.64 (1.93)
Past Year	Yes	205	33.53 (5.60)	32.53 (5.38)
	No	1,358	17.81 (1.73)	18.03 (1.74)
	Overall	1,563	19.39 (1.65)	19.49 (1.65)
Past Month	Yes	128	19.85 (7.28)	18.96 (6.77)
	No	1,435	8.67 (1.13)	8.70 (1.13)
	Overall	1,563	9.41 (1.12)	9.37 (1.11)
Marijuana				
Lifetime	Yes	965	22.38 (2.47)	22.47 (2.46)
	No	598	10.74 (1.86)	10.28 (1.77)
	Overall	1,563	17.84 (1.71)	17.71 (1.70)
Past Year	Yes	151	21.42 (6.45)	20.85 (6.05)
	No	1,412	12.57 (1.55)	12.53 (1.54)
	Overall	1,563	13.23 (1.47)	13.15 (1.46)
Past Month	Yes	101	16.01 (8.20)	15.26 (7.57)
	No	1,462	5.54 (0.89)	5.48 (0.89)
	Overall	1,563	6.12 (0.89)	6.02 (0.88)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.4 Percentages of Children (12 to 17) Living with a Parent Reporting Lifetime, Past Year, and Past Month Use of Tobacco and Alcohol among Parent-Child (12 to 17) Pairs Asked Whether Their Parents Had Spoken to Them About the Dangers of Tobacco, Alcohol, or Drug Use Within the Past Twelve Months for 2002

Drug	Parent talked about dangers with child	2002		
		n	Baseline[1]	Final[2]
Alcohol				
Lifetime	Yes	2,253	41.49 (1.92)	41.51 (1.93)
	No	1,578	40.23 (2.13)	40.24 (2.11)
	Overall	3,831	40.97 (1.44)	40.99 (1.43)
Past Year	Yes	2,253	33.19 (1.80)	33.06 (1.78)
	No	1,578	31.50 (2.05)	31.62 (2.04)
	Overall	3,831	32.49 (1.39)	32.47 (1.37)
Past Month	Yes	2,253	15.99 (1.29)	16.02 (1.29)
	No	1,578	18.12 (1.64)	18.26 (1.68)
	Overall	3,831	16.87 (1.09)	16.94 (1.09)
Cigarettes				
Lifetime	Yes	2,253	31.81 (1.88)	31.41 (1.87)
	No	1,578	32.62 (1.95)	32.79 (1.97)
	Overall	3,831	32.14 (1.36)	31.98 (1.35)
Past Year	Yes	2,253	18.75 (1.43)	18.51 (1.44)
	No	1,578	18.59 (1.71)	18.68 (1.73)
	Overall	3,831	18.68 (1.07)	18.58 (1.09)
Past Month	Yes	2,253	10.74 (0.97)	10.49 (0.95)
	No	1,578	10.20 (1.23)	10.05 (1.20)
	Overall	3,831	10.52 (0.76)	10.31 (0.75)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.5 Percentages of Children (12 to 17) Living with a Parent Reporting Lifetime, Past Year, and Past Month Use of Marijuana and Any Illicit Drug among Parent-Child (12 to 17) Pairs Asked Whether Their Parents Had Spoken to Them About the Dangers of Tobacco, Alcohol, or Drug Use Within the Past Twelve Months for 2002

Drug	Parent talked about dangers with child	2002		
		n	Baseline[1]	Final[2]
Any Illicit				
Lifetime	Yes	2,253	28.39 (1.90)	28.11 (1.85)
	No	1,578	29.12 (1.86)	29.63 (1.90)
	Overall	3,831	28.69 (1.37)	28.74 (1.36)
Past Year	Yes	2,253	18.68 (1.33)	18.54 (1.31)
	No	1,578	20.91 (1.76)	21.20 (1.79)
	Overall	3,831	19.60 (1.07)	19.63 (1.08)
Past Month	Yes	2,253	9.21 (0.92)	9.15 (0.92)
	No	1,578	11.14 (1.30)	11.02 (1.28)
	Overall	3,831	10.01 (0.76)	9.92 (0.75)
Marijuana				
Lifetime	Yes	2,253	18.38 (1.53)	18.16 (1.46)
	No	1,578	16.80 (1.42)	16.97 (1.45)
	Overall	3,831	17.73 (1.06)	17.67 (1.04)
Past Year	Yes	2,253	13.12 (1.14)	13.08 (1.14)
	No	1,578	11.91 (1.19)	11.98 (1.21)
	Overall	3,831	12.62 (0.82)	12.63 (0.83)
Past Month	Yes	2,253	6.24 (0.70)	6.28 (0.72)
	No	1,578	6.07 (0.85)	5.90 (0.84)
	Overall	3,831	6.17 (0.54)	6.13 (0.54)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.6a Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Tobacco and Alcohol among Mother-Child (12 to 17) Pairs For Mother in the Pair Asked Whether They had Spoken to Their Children About the Dangers of Tobacco, Alcohol, or Drug Use Within the Past Twelve Months for 2002

Drug	Mother talked about dangers with child	2002		
		n	Baseline[1]	Final[2]
Alcohol				
Lifetime	0 times	113	21.90 (5.80)	21.60 (5.96)
	1-2 times	324	39.33 (4.02)	39.13 (4.05)
	A few times	601	34.09 (3.22)	34.30 (3.22)
	Many times	1,152	44.12 (2.36)	44.13 (2.34)
	Overall	2,190	39.29 (1.62)	39.34 (1.62)
Past Year	0 times	113	17.59 (5.39)	17.82 (5.64)
	1-2 times	324	30.63 (3.83)	30.42 (3.87)
	A few times	601	25.87 (2.87)	25.61 (2.81)
	Many times	1,152	35.11 (2.27)	35.22 (2.26)
	Overall	2,190	30.84 (1.55)	30.83 (1.54)
Past Month	0 times	113	10.18 (4.55)	10.54 (4.96)
	1-2 times	324	11.71 (2.17)	11.66 (2.18)
	A few times	601	12.31 (2.06)	12.44 (2.11)
	Many times	1,152	18.45 (1.75)	18.66 (1.80)
	Overall	2,190	15.25 (1.13)	15.43 (1.16)
Cigarettes				
Lifetime	0 times	113	25.74 (6.44)	26.27 (6.60)
	1-2 times	324	25.54 (3.80)	25.16 (3.83)
	A few times	601	24.02 (2.80)	24.24 (2.84)
	Many times	1,152	37.16 (2.31)	37.38 (2.31)
	Overall	2,190	31.14 (1.53)	31.32 (1.55)
Past Year	0 times	113	19.17 (5.96)	19.73 (6.19)
	1-2 times	324	17.20 (3.48)	16.87 (3.53)
	A few times	601	11.20 (2.01)	11.10 (2.02)
	Many times	1,152	23.16 (2.03)	23.31 (2.07)
	Overall	2,190	18.79 (1.31)	18.84 (1.34)
Past Month	0 times	113	2.70 (1.20)	2.71 (1.17)
	1-2 times	324	6.85 (1.94)	6.54 (1.91)
	A few times	601	6.05 (1.44)	5.82 (1.40)
	Many times	1,152	15.35 (1.65)	15.27 (1.67)
	Overall	2,190	10.76 (0.98)	10.63 (0.99)

NOTE: Standard errors of prevalence estimates are provided in parentheses.

[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.

[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.6b Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Tobacco and Alcohol among Father-Child (12 to 17) Pairs For Father in the Pair Asked Whether He had Spoken to His Child About the Dangers of Tobacco, Alcohol, or Drug Use Within the Past Twelve Months for 2002

Drug	Father talked about dangers with child	2002		
		n	Baseline[1]	Final[2]
Alcohol				
Lifetime	0 times	157	43.43 (7.67)	42.68 (7.67)
	1-2 times	318	36.33 (4.98)	36.71 (4.85)
	A few times	490	38.16 (3.65)	38.54 (3.65)
	Many times	477	43.41 (3.89)	43.17 (3.90)
	Overall	1,442	40.02 (2.38)	40.12 (2.35)
Past Year	0 times	157	36.97 (7.99)	36.44 (7.96)
	1-2 times	318	24.81 (4.32)	25.05 (4.22)
	A few times	490	32.60 (3.51)	32.87 (3.52)
	Many times	477	34.57 (3.78)	34.32 (3.77)
	Overall	1,442	31.94 (2.29)	32.02 (2.26)
Past Month	0 times	157	20.98 (7.84)	21.03 (7.75)
	1-2 times	318	13.09 (3.94)	13.10 (3.83)
	A few times	490	15.54 (2.93)	15.37 (2.86)
	Many times	477	22.18 (3.10)	21.97 (3.06)
	Overall	1,442	17.75 (2.08)	17.68 (2.03)
Cigarettes				
Lifetime	0 times	157	18.30 (5.75)	17.81 (5.67)
	1-2 times	318	26.40 (4.48)	25.60 (4.26)
	A few times	490	32.18 (3.62)	31.92 (3.60)
	Many times	477	34.22 (3.92)	33.78 (3.90)
	Overall	1,442	30.36 (2.19)	29.91 (2.17)
Past Year	0 times	157	12.90 (5.50)	12.28 (5.39)
	1-2 times	318	12.75 (3.67)	12.37 (3.36)
	A few times	490	13.21 (2.18)	13.14 (2.20)
	Many times	477	21.93 (3.53)	21.58 (3.48)
	Overall	1,442	16.06 (1.62)	15.81 (1.59)
Past Month	0 times	157	2.06 (0.94)	2.05 (0.93)
	1-2 times	318	8.75 (3.55)	8.24 (3.20)
	A few times	490	8.15 (1.90)	8.11 (1.93)
	Many times	477	12.01 (2.34)	11.68 (2.24)
	Overall	1,442	9.06 (1.21)	8.81 (1.16)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.7a Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Marijuana and Any Illicit Drug among Mother-Child (12 to 17) Pairs For Mothers Asked Whether She had Spoken to Her Child About the Dangers of Tobacco, Alcohol, or Drug Use Within the Past Twelve Months for 2002

Drug	Mother talked about dangers with child	2002		
		n	Baseline[1]	Final[2]
Any Illicit				
Lifetime	0 times	113	19.99 (6.08)	20.76 (6.28)
	1-2 times	324	19.82 (3.04)	20.00 (3.05)
	A few times	601	25.66 (3.62)	25.82 (3.57)
	Many times	1,152	31.98 (2.25)	31.96 (2.25)
	Overall	2,190	27.68 (1.62)	27.80 (1.62)
Past Year	0 times	113	12.71 (4.87)	14.10 (5.43)
	1-2 times	324	15.53 (2.71)	15.67 (2.71)
	A few times	601	16.58 (2.66)	16.51 (2.60)
	Many times	1,152	20.80 (1.75)	20.73 (1.76)
	Overall	2,190	18.36 (1.23)	18.42 (1.24)
Past Month	0 times	113	2.86 (1.61)	2.90 (1.63)
	1-2 times	324	6.61 (1.51)	6.82 (1.55)
	A few times	601	6.34 (1.63)	6.27 (1.61)
	Many times	1,152	12.51 (1.44)	12.44 (1.45)
	Overall	2,190	9.35 (0.91)	9.34 (0.91)
Marijuana				
Lifetime	0 times	113	13.66 (5.18)	14.72 (5.58)
	1-2 times	324	12.81 (2.53)	12.94 (2.55)
	A few times	601	12.49 (2.25)	12.85 (2.34)
	Many times	1,152	19.98 (1.84)	19.82 (1.79)
	Overall	2,190	16.48 (1.22)	16.59 (1.21)
Past Year	0 times	113	9.75 (4.61)	10.97 (5.19)
	1-2 times	324	9.74 (2.17)	9.83 (2.17)
	A few times	601	9.11 (1.87)	9.25 (1.93)
	Many times	1,152	14.07 (1.38)	14.02 (1.39)
	Overall	2,190	11.81 (0.93)	11.92 (0.95)
Past Month	0 times	113	0.85 (0.58)	0.90 (0.62)
	1-2 times	324	3.55 (1.08)	3.61 (1.10)
	A few times	601	2.99 (0.97)	3.08 (1.04)
	Many times	1,152	8.62 (1.17)	8.64 (1.20)
	Overall	2,190	5.85 (0.68)	5.91 (0.70)

NOTE: Standard errors of prevalence estimates are provided in parentheses.

[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.

[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

Table 8.7b Percentages of Children (12 to 17) Reporting Lifetime, Past Year, and Past Month Use of Marijuana and Any Illicit Drug among Father-Child (12 to 17) Pairs For Father Asked Whether He had Spoken to His Child About the Dangers of Tobacco, Alcohol, or Drug Use Within the Past Twelve Months for 2002

Drug	Father talked about dangers with child	2002		
		n	Baseline[1]	Final[2]
Any Illicit				
Lifetime	0 times	157	20.13 (5.86)	19.69 (5.78)
	1-2 times	318	19.03 (3.08)	19.24 (3.04)
	A few times	490	30.22 (3.64)	30.48 (3.65)
	Many times	477	29.19 (3.50)	29.20 (3.50)
	Overall	1,442	26.49 (2.03)	26.63 (2.02)
Past Year	0 times	157	11.14 (5.02)	11.20 (5.14)
	1-2 times	318	12.91 (2.53)	13.09 (2.51)
	A few times	490	21.75 (3.28)	21.96 (3.28)
	Many times	477	23.57 (3.35)	23.39 (3.33)
	Overall	1,442	19.47 (1.74)	19.56 (1.73)
Past Month	0 times	157	2.47 (1.24)	2.59 (1.32)
	1-2 times	318	7.36 (1.94)	7.42 (1.90)
	A few times	490	7.88 (1.92)	7.77 (1.91)
	Many times	477	13.33 (2.66)	13.16 (2.61)
	Overall	1,442	9.14 (1.17)	9.08 (1.16)
Marijuana				
Lifetime	0 times	157	9.52 (3.40)	8.86 (2.97)
	1-2 times	318	11.14 (2.24)	11.27 (2.23)
	A few times	490	21.35 (3.44)	21.18 (3.41)
	Many times	477	21.19 (3.30)	20.92 (3.27)
	Overall	1,442	17.98 (1.80)	17.83 (1.79)
Past Year	0 times	157	4.71 (1.69)	4.74 (1.72)
	1-2 times	318	5.88 (1.48)	5.99 (1.50)
	A few times	490	13.98 (2.87)	13.72 (2.79)
	Many times	477	19.10 (3.25)	18.80 (3.21)
	Overall	1,442	13.11 (1.55)	12.99 (1.54)
Past Month	0 times	157	2.40 (1.24)	2.54 (1.32)
	1-2 times	318	3.60 (1.20)	3.67 (1.22)
	A few times	490	4.07 (1.10)	3.89 (1.03)
	Many times	477	9.67 (2.48)	9.32 (2.40)
	Overall	1,442	5.73 (0.91)	5.59 (0.90)

NOTE: Standard errors of prevalence estimates are provided in parentheses.
[1] Baseline refers to the weight obtained from using a main effects only model for the last two steps of calibration, res.pr.ps and res.pr.ev, and a full model for those preceding.
[2] Final refers to the weight obtained using a full model throughout all steps of calibration.

References

Chen, P., Dai, L., Gordek, H., Shi, W., Singh, A., & Westlake, M. (2004, January 23). Person-level sampling weight calibration for the 2002 NSDUH. In *2002 National Survey on Drug Use and Health: Methodological resource book* (Section 3, prepared for the Substance Abuse and Mental Health Services Administration, Office of Applied Studies, under Contract No. 283-98-9008, Deliverable No. 28, RTI/07190.474.100). Research Triangle Park, NC: RTI International. [Available as a PDF at http://oas.samhsa.gov/nhsda/methods.cfm]

Chen, P., Penne, M. A., & Singh, A. C. (2000). Experience with GEM for weight calibration for NHSDA. In *Proceedings of the American Statistical Association, Survey Research Methods Section* (pp. 604-607). Alexandria, VA: American Statistical Association.

Chromy, J. R. (1979). Sequential sample selection methods. In *Proceedings of the American Statistical Association, Survey Research Methods Section* (pp. 401-406). Alexandria, VA: American Statistical Association.

Chromy, J. R., & Singh, A. C. (2001, August). Estimation for person-pair drug-related characteristics in the presence of pair multiplicities and extreme sampling weights. In *Proceedings of the Survey Research Methods Section, American Statistical Association [CD-ROM]*. Alexandria, VA: American Statistical Association.

Cox, B. G. (1980). The weighted sequential hot deck imputation procedure. In *Proceedings of the American Statistical Association, Survey Research Methods Section* (pp. 721-726). Washington, DC: American Statistical Association.

Deville, J. C., & Särndal, C. E. (1992). Calibration estimating in survey sampling. *Journal of the American Statistical Association, 87*, 376-382.

Folsom, R. E., & Singh, A. C. (2000). The generalized exponential model for sampling weight calibration for extreme values, nonresponse, and poststratification. In *Proceedings of the American Statistical Association, Survey Research Methods Section* (pp. 598-603). Alexandria, VA: American Statistical Association.

Folsom, R. E., & Witt, M. B. (1994). Testing a new attrition nonresponse adjustment method for SIPP. In *Proceedings of the American Statistical Association, Social Statistics Section* (pp. 428-433). Washington, DC: American Statistical Association.

Grau, E. A., Barnett-Walker, K., Copello, E., Frechtel, P., Licata, A., Liu, B., & Odom, D. M. (2004, February). Imputation report. In *2002 National Survey on Drug Use and Health: Methodological resource book* (Section 4, prepared for the Substance Abuse and Mental Health Services Administration, Office of Applied Studies, under Contract No. 283-98-9008, RTI/7190.474.200). Research Triangle Park, NC: RTI International. [Available as a PDF at http://oas.samhsa.gov/nhsda/methods.cfm]

Hajek, J. (1971). Comment on D. Basu's paper, "An essay on the logical foundations of survey sampling, part one" [pp. 203-234]. In V. P. Godambe, & D. A. Sprott (Eds.), *Foundations of statistical inference: A symposium* (p. 236). Toronto, Ontario, & Montreal, Quebec, Canada: Holt, Rinehart and Winston of Canada.

Iannacchione, V. (1982). Weighted sequential hot deck imputation macros. In *Proceedings of the Seventh Annual SAS Users Group International Conference*. Cary, NC: SAS Corporation.

Little, R. J. A., & Rubin, D. B. (1987). *Statistical analysis with missing data*. New York: John Wiley & Sons.

Odom, D. M., Bowman, K. R., Chromy, J. R., & Martin, P. C. (2004, January). 2002 National Survey on Drug Use and Health: Sample design report. In *2002 National Survey on Drug Use and Health: Methodological resource book* (Vol. Volume 1, Section 2, prepared for the Substance Abuse and Mental Health Services Administration, Office of Applied Studies, under Contract No. 283-98-9008, Deliverable No. 10, RTI/7190.430.400). Research Triangle Park, NC: RTI.

Rubin, D. B. (1986). Statistical matching using file concatenation with adjusted weights and multiple imputations. *Journal of Business and Economic Statistics, 4* (1), 87-94.

Schafer, J. L. (1997). *Analysis of incomplete multivariate data* (No. 72, Monographs on Statistics and Applied Probability). Boca Raton, FL: Chapman and Hall/CRC.

Singh, A. C., & Mohl, C. A. (1996). Understanding calibration estimators in survey sampling. *Survey Methodology, 22*, 107-115.

Williams, R. L., & Chromy, J. R. (1980). SAS sample selection MACROS. In *Proceedings of the Fifth International SAS Users Group International Conference* (pp. 382-396). Cary, NC: SAS Corporation.