EXPERIENCE WITH THE GENERALIZED EXPONENTIAL MODEL FOR WEIGHT CALIBRATION FOR THE NATIONAL HOUSEHOLD SURVEY ON DRUG ABUSE

P. Chen, M.A. Penne, and A.C. Singh, Research Triangle Institute
P. Chen, Statistics Research Division, RTI, RTP, NC 27709 pchen@rti.org

Key Words: Adjustment Bounds, Predictor Variables, Newton-Raphson, Extreme Values

1. The 1999 NHSDA: a brief description

The National Household Survey on Drug Abuse (NHSDA) is designed to estimate prevalence of both licit and illicit drug use in the U.S. for various demographic and geographic domains. Since 1999, it has become a statewide survey that includes 50 States and the District of Columbia. The target population includes civilian, noninstitutionalized persons aged 12 or older. Eight States (California, Florida, Illinois, Michigan, New York, Ohio, Pennsylvania, and Texas), referred to as the "big" States, have a sample designed to yield 3600 respondents per State, while the remaining 43 "small" States have a sample designed to yield 900 respondents per State. The total sample size is 66706 persons (corresponding to 51821 dwelling units (DUs) selected at the second phase out of 169166 DUs screened at the first phase) with a low of 756 for Nevada to a high of 1280 for Utah among "small" States, and a low of 2669 for New York and a high of 4681 for California among "big" States.

In the NHSDA design since 1999, States serve as the primary strata, and field interviewer (FI) regions within each State serve as secondary strata. In the "small" States, 12 FI regions are created, while in the "big" States, 48 FI regions are formed. Segments within FI regions form first-stage sample units drawn with probabilities proportional to composite size measures using Chromy's algorithm (Chromy, 1981). DUs within segments form the second- stage units that are drawn according to a random systematic scheme with an EPSEM (equal probability selection method) goal. Within each FI region, segments are formed to contain a minimum of 150 DUs. From each FI region, two segments are drawn per quarter with a total of eight per year. On average, about 30 DUs are selected per segment with an objective of 10 completed person-level interviews. This average of three selected DUs per completed person interviews reflects various levels of attrition such as DU eligibility to the target population, DU-level nonresponse, and person-level nonresponse. The NHSDA design is a multistage design with deep stratification, which can be viewed as a two-phase design with the second-phase units of persons nested within the first-phase DUs. After the DU selection, first-phase information (e.g., eligibility, age, race/ethnicity, and gender) for all members of the DU is collected, and then age is used to define deep stratification variables for the second-phase sample of persons within eligible DUs. At this phase, 0, 1, or 2 persons are selected within each DU using an adaptation of Brewer's sampling scheme.

2. Sampling Weight Calibration

The sample weighting of the 1999 NHSDA posed new challenges because of the sheer magnitude of the number of State-specific predictors for use in nonresponse (nr) and poststratification (ps) adjustments. With the 51-State survey, it was not practical to use a single model for each of the adjustments. Also, treating each State separately was not desirable because individual State sample sizes are not large enough to support reliable estimation of a fair number of parameters. It was decided to group the 51 States into nine model groups corresponding to the nine Census divisions. This helped to keep a substantial number of predictor variables in each model while reducing computational time that would be associated with fitting a large number of models.

The nine weight components at the phase I DU level and six at the phase II person level are shown in Exhibit 1. The generalized exponential model (GEM) of Folsom and Singh (2000) (see Section 4) for sampling weight calibration was used to adjust for extreme values (ev), nr, and ps. Note that in the weight component 8, the screener DU weights in the first phase are poststratified to population counts by adjusting DU's weighted contribution of person counts to various demographic domains. This will reduce coverage bias resulting from the first phase. The weight component 12 is a bit unusual in that it is ps of the selected persons (this includes respondents and nonrespondents) in the second phase to estimated controls from the large first phase sample of persons for various predictor variables at the segment, DU, and person levels. This will give stable controls for the step of nr adjustment of respondent weights. It may be noted that this would not have been possible in the absence of screener data information on demographics of members of the selected households.

3. Covariates for Modeling Adjustment Factors

Typical predictors used for the screener DU nr adjustment are State/region, quarter, group quarters indicator, population density, percentage hispanic in segment, percentage black in segment, percentage owner-occupied DUs in segment, and socioeconomic status (SES) indicator. Similarly, the predictors for the person-level nr adjustments include, in addition to those stated above, age group, gender, race, hispanicity, and relation to head of household. For ps, predictors typically used are State/region, age, race, gender, Hispanicity, and quarter, and the model consists of main effects and some interactions of these predictors. For the separate adjustment for ev, we could use the same predictors as used in ps.

Note that it is desirable (whenever possible) to include ps predictors (correlated with the outcome variable) as part of the nr predictors (correlated with the response variable) because of the potential variance reduction as an offset to the varaince inflation due to random controls used in the nr adjustment. In general, this is not possible due to the unavailability of information about demographic variables (often used for ps) for nonrespondents. However, with two phase designs such as NHSDA, this is not a problem because the screener data from the first phase have the necessary information.

We used some general guidelines on choosing an initial set of State-specific covariates and modify the initial set iteratively as we faced problems in meeting them. We would begin with the baseline model of one-factor effects and then attempt to add higher order effects (two and three) although some collapsing might be needed along the way depending on the individual State sample sizes. In collapsing effects, every effort was made to include as many important state-specific covariates (typically defined by socio-demographic domains by state geography) as possible in models for nr and ps weight adjustments in order to obtain more precise state-level estimates. However, keeping a lot of state-specific covariates was not possible because individual state sample sizes were not large enough to support stable estimation of a fair number of model parameters. We, therefore, used a hierarchical structure in grouping states for covariate inclusion in the model; the order being covariates at the national level, followed by covariates at the census division level within the nation, then covariates at the combined -state level within the census division, and finally whenever possible covariates at the state level within the combined states. In situations where model parameters were inestimable due to insufficient sample sizes when certain additional covariates were included in the model, the hierarchy strategy mentioned above was used to combine states within a division so that covariates at the combined level could be included. The levels of the covariates were collapsed (or coarsened) only when combining states was not a feasible alternative. The reason for this is that ultimate estimation domains of interest are defined by covariate levels, and if stable estimates can be obtained at the combined state level, then this will be beneficial in obtaining more reliable state level estimates using the small area estimation techniques. The eight big states were not combined with other smaller states to the extent possible in order to get direct state-level estimates without relying on the small area estimation technique.

The objective criterion of measuring the gain in efficiency (for a set of study variables) can be used to check for the suitability of the number of controls, assuming that they can be met by the sample. This should be done in addition to keeping a watch on the increase in the unequal weighting effect (UWE). For the NHSDA, it became apparent that the number of controls could be very high such as in excess of 1,000, which would be computationally prohibitive. The implementation of ps involves matrix inversion at each iterative step whose dimension corresponds to the number of controls. A solution is to use separate models within groups rather a single overall model. It can be shown under the assumption that whenever we need to collapse an effect (two factor or higher order), if it is always done within a group of States, then fitting an overall model is equivalent to fitting separate models for each group. In this way, we can reduce the computational problem when faced with too many controls. We, therefore, partitioned the sample into nine model groups corresponding to the nine census divisions.

4. Modeling of Weight Adjustment Factors by GEM

The methodology of generalized exponential model (GEM) has several features:

(i) It allows for different bounds on the adjusted weights for different cases (or sample units). Thus, adjustments on initially identified extreme weights can be kept under control by having a separate built-in control for extreme weights.

(ii) The above built-in control on the adjustments to ev is often adequate in that the frequency of ev after nr and ps is not high. However, if this were not the case, GEM can also be used for a separate ev after ps such that sample distribution of weights obtained after the initial ps is preserved.

(iii) GEM allows for a unified approach to extreme value treatment, nr adjustment, and ps of initial design weights. The differences are only in terms of bounds and control totals for the adjusted weights.

(iv) GEM is a generalization of the commonly used raking-ratio method in which a distance function is minimized such that the initial weights are perturbed only a little, lie within certain bounds, and control totals are met. It is also a generalization of Deville and Särndal's (1992) logit method in that bounds on weights are not required to be uniform. Moreover, the lower bound can be set to 1, often desirable for the nr adjustment.

(v) Like the raking-ratio method, GEM fitting requires iterations (such as in Newton-Raphson).

Let T_x denote the p-vector of control totals corresponding to predictor variables (x_1, ..., x_p, say). Then, the model parameters λ are estimated by solving the calibration constraints.

$$\sum_{k\in s} x_k d_k a_k(\lambda) = T_x$$

$$a_k(\lambda) = \frac{\ell_k(u_k - c_k) + u_k(c_k - \ell_k)\exp(A_k x_k' \lambda)}{(u_k - c_k) + (c_k - \ell_k)\exp(A_k x_k' \lambda)}.$$

For more details, see Folsom and Singh (2000).

5. GEM Implementation

5.1 Definition of Extreme Values of Sampling Weights

An important aspect of GEM is the built-in provision of extreme value 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 pre-specified domains defined usually by design strata corresponding to deep stratification. For example, the DU (DU) level weight adjustment for 1999 CAI NHSDA uses the FI region as the domain. Person-level weight adjustment uses the hierarchy of three domains: (a) State x age group, (b) FI Region, and (c) State. The hierarchy is used in the sense that a minimum of 30 observations is required for defining the extreme boundaries or critical values. If this is not met at the lower level, then the next level up in the hierarchy is used. In fact, the deep strata are field interviewer (FI) region x age group, which turn out to be unsuitable for outlier definition-domains because of insufficient sample sizes. So, collapsing FI regions within a State gives rise to such domains as State x age group. Even at this level, sample sizes may be insufficient, and then FI regions, and later States themselves can be used as outlier definition domains. The critical values for low and high ev are denoted by $b_{k\,(l)}$ and $b_{k\,(u)}$ in the sequel.

5.2 Definition of Lower and Upper Bounds

For implementing extreme value control via GEM, the variable m_k is defined as the minimum of ($b_{k(u)}$ /w_k) and 1 for high ev, and the maximum of ($b_{k(l)}$ /w_k) and 1 for low ev, where w_k is the sampling weight before adjustment, and ($b_{k(u)}$, $b_{k(l)}$) denote the threshold or the critical value for the ev. Note that for the high ev, the more extreme the weight is, the smaller the m_k will be. While for the low ev, the smaller the weight is, the bigger the m_k will be. Non-extremes have a value of 1 for m_k. Now, the upper and lower bounds for the adjustment factor are defined respectively as the product of m_k and the upper and lower boundary parameters of GEM. GEM allows inputs of three different upper and lower boundary parameters (L1 and U1, L2 and U2, L3 and U3, respectively) for high, non-, and low ev. By applying small upper boundary parameter for high ev, and high lower boundary parameter for low ev, the ev could be controlled in the modeling. GEM also requires specification of centers C1, C2, and C3 such that L < C < U. For nr, it is desirable to require all adjustments to be greater than 1 because they represent the inverse of response propensities. Thus, all the three L1, L2, and L3 are set to 1. The value of C in this case is chosen as the inverse of the overall response propensity. For ps, C's are set to 1 because we want to adjust the weights not too far away from the original design weight. Here L2 is chosen to be less than 1, and U2 greater than 1 because the control totals could be larger or smaller than the estimated totals based on the design weights. Subsection 5.7 gives guidelines for the choice of L, C, and U parameters. The case of ev treatment is analogous to ps.

5.3 Definition of Control Totals

GEM modeling for nr, ps, and ev involves estimation of parameters of the adjustment factor model such that certain control totals are satisfied. There are three types of control totals. For nr, the control totals are from the full sample (i.e., respondents and nonrespondents), while for ps, control totals are obtained from external sources, such as the Census Bureau or a large first-phase sample, and for ev, from the sample of respondents. For example, in the 1999 NHSDA, the control totals for various domains for the (selected) person-level ps were obtained from the first-phase sample containing roster information, and the control totals for the (respondent) person-level ps were obtained from the Census Bureau's postcensal population estimates for various demographic domains.

5.4 Efficient Computation by using Grouped Data

In view of the fact that adjustment factors remain the same for units (DUs or persons) that have common values for explanatory variables used in the model, one can reduce the sample data size by grouping units having common values of the explanatory variables. This can save computation time significantly, especially if the original sample size is large. The units with ev are grouped such that in addition to the common explanatory variables, they also have common values of m_k. Note, however, that for GEM with grouped data, the UWE and t-test statistics normally produced in the output would be misleading because the weights in grouped data are sums of the weights for the individual units within each group. Also the definition of variance estimation strata (VESTR) and replicates (VEREP) required for variance calculation would not be proper. To avoid these misleading results from using the grouped data, the final model should be rerun with the full (i.e., ungrouped) data.

5.5 Steps in GEM Fitting

After specifying the GEM parameters, such as the initial upper and lower bounds, the number of the Newton-Ralphson iterations and half-steps, and the type of the weight adjustment (nr, ps, or ev), a forward selection method for modeling can be used. The model with only the main effects is first fit to obtain the baseline upper and lower bounds for extreme and non-ev and to calculate a baseline UWE. Without unduly increasing the UWE and the ev proportion, as many higher order interactions as possible should be added to the model to help reduce coverage bias. Convergence problem can be addressed by loosening Ls and Us, and collapsing or dropping variables. In GEM, t-tests and p-values for significance of various effects can be computed for a previously converged model, which is helpful in deciding about collapsing of effects when convergence problems arise with tighter bounds.

Collapsing implies combining the values of one variable with another variable explicitly present in the model, while dropping implies combining with the reference levels of the variables not explicitly present in the model. Collapsing or dropping of lower order interactions have a direct impact on the number of higher order interactions. When adding higher order terms, all previously selected explanatory variables are retained in the model to the extent possible. Possible reasons for nonconvergence include explanatory variables corresponding to domains with small sample sizes, or the domains with large discrepancy between estimated totals based on the initial weights and the target control totals. The variables causing problems of convergence can be identified by high magnitude of the estimated model parameters. Once the explanatory variables are finalized, finer adjustments of Us and Ls can optimize the model in the sense of small UWE and the ev proportion.

5.6 Quality Control Checks

Comparisons of the weight distribution, maximum over mean factor, and UWE across various domains before and after the adjustment are conducted to uncover any unusual impact of weight adjustment on the initial weights. The ev proportion after adjustment is checked to see how effective the modeling was on controlling ev. Coverage bias analysis based on the slippage rates is also conducted to check the impact of ps on various non-controlled domains. In addition, after the final weight adjustment, point estimates for the main drug use variables as well as their standard errors computed using a sandwich variance formula (see Vaish, Gordek, and Singh, 2000), are compared with the corresponding estimates and standard errors for the baseline (or the main effects) model.

5.7 Practical Guidelines in using GEM

(i) Collapsing Checks for Small Domains: Examine the number of observations or the sample sizes in various domains defined by levels of the factor effects. If the domain sample size is zero, then the corresponding factor level is either dropped or collapsed with some other level. This would automatically collapse the corresponding factor level with reference level. However, if the corresponding control total is not zero, then this would effectively prevent the control total to be met by the reference level. This may not be desirable if the reference level involves big states because we would like to meet the big state level controls as much as possible.

If a domain sample size is small, then collapse the corresponding factor effect with another effect based on substantive considerations. If state is involved, then it would be better in general to collapse over states than over other factor levels as discussed in Section 3. Need for collapsing should be checked at each stage of model enlargement in the forward selection of factors. The reason for this is that if some collapsing was done at a previous stage, the corresponding factor levels should also be collapsed at succeeding stages involving higher order factor effects under the hierarchy principle.

(ii) Singularity Checks: As in the case of collapsing checks, singularity (linear dependence of columns of values of the predictors) checks should also be done at each stage of model enlargement because it depends on what other predictors are in the model. Note that although all variables are linearly independent of each other, it is possible that the columns of their realized values may be linearly dependent.

(iii) Finding the Initial Factor Set: After the collapsing and singularity checks, the remaining factor effects at a given stage of model enlargement form the initial factor set.

(iv) Baseline Model: Start with the model consisting of all one-factor effects from the initial factor set and find a convergent version (after some collapsing if necessary) and the corresponding p-values under no bound restrictions. Use p-values to decide about collapsing, and optimize it with respect to the model characteristics (see guideline vii below) (i.e., try to reduce the UWE, and tighten the bounds).

(v) Baseline plus Two-factor Effects: Start with the baseline model, and add all the two-factor effects from the initial factor set. Find a convergent version and the corresponding p-values under no bound restrictions. Proceed as in guideline (iv) above. In our application, first add the non-State two-factor effects, and then in a separate step add the State two-factor effects.

(vi) Baseline with Two and Higher Order Factor Effects: Start with the optimized model from guideline (v), and add the higher order factor effects. Now, proceed similar to guideline (v) to get an optimum version.

(vii) Optimizing a Model with respect to the Target Model Characteristics: These are summarized in the following points based on empirical and heuristic considerations:

1. For each step of model enlargement, compute the UWE for the initial weights, and as a guideline, allow tolerance in the resulting UWE increase up to 20% or the maximum allowable UWE (generally under 6), whichever is lower.

2. Use the following guidelines for setting bounds. In the case of ps, set L1 = L2, and U2 = U3, and C1 = C2 = C3 =1. Now for high ev, start with 1.5 for U1 and try to keep it close to 1 and ; for nonev, start with (.5,3) for (L2, U2) and vary them inside (.3, 4); and for low ev, start with .8 for L3 and try to keep it close to 1. These guidelines are only rough in that a lot depends on the realized values of maximum and minimum adjustment factors under loose bounds. The critical points for ev within GEM modeling are defined as median +/- 2.5 times the interquartile range (IQR) which is conservative in that weights lying on the boundary but inside are also controlled. However for QC purposes, the original definition with the factor of 3 for IQR should be used for checking the ev proportion.

3. In the case of nr, set L1 = L2= L3 =1, and U2 = U3. All the C's are set equal to the common value of the overall inverse response propensity. Now for ev, start with a value of U1 slightly bigger than C1, and try to stay close to C1, and for nonev, start with (1, 3) for (L2, U2) and vary them inside (1, 4).

4. For target percentage ev and outwinsors within GEM for nr and ps, the general guideline is to try to keep respectively the unweighted ev, weighted ev, and outwinsors under 3%, 15%, and 5%; actually, these percentages are rather liberal and serve as upper bounds only. (The term outwinsor is used to signify the proportion of weight-sum out of the total weight-sum that would be trimmed if weights were winsorized). In practice, it is preferable to be able to reduce them by half. If these guidelines cannot be met, a separate GEM for ev after ps is implemented. Note that the bound specification for ev is similar to that for ps.

(viii) Evaluation Measures: After each stage of model enlargement, it is important to examine various characteristics for large values such as UWE, maximum over mean weight factor, percentage ev and outwinsors, and other characteristics such as distance between the total sample weighted count and the target population count, i.e., slippage rates by different domains, weight summary statistics, and distributions of adjustment factors for highly asymmetric tails. Also with the bounds realized for the final model, run the baseline model, and then compare point estimates and SEs for selected outcome variables for the two models. Generally, the two estimates are likely to be close, but not the SEs. The SE for the final model is expected to be smaller but at times may be larger. Larger SEs should be identified and examined. In situations where the SEs of the estimates are much larger than the ones for the baseline model, there is an indication of the instability in model parameter estimates due to possible overfitting and insufficient sample sizes. In such situations, the final model should be revised to get a more parsimonious model.

6. Numerical Results

The 1999 NHSDA data for the East South Central Census Division is used to illustrate results obtained by fitting GEM at both the first phase for nr and ps of the screener DUs and the second phase for ps for selected persons followed by nr and ps for respondent persons. Table 1 gives the summary of characteristics of various models fitted, while Table 2 gives summary statistics in terms of UWE, ev and outwinsor proportions, as well as distributional characteristics of the weight distribution. It is seen that % ev is reasonable after ps, and therefore, there is no need for an extra ev step.) Table 3 compares point estimates and SEs for baseline and final models across a set of drug use variables. For confidentiality reasons, only ratios of point estimates (final over the baseline) are presented. However, individual RSEs (SE over the estimate) are presented. Two types of RSE are given; one is unadjusted signifying no adjustment for calibration, and the other is ps-adjusted RSE denoting a sandwich type formula to adjust for ps as given in Vaish, Gordek, and Singh (2000). It is seen that the two (baseline and final) RSE are quite comparable. Interestingly, the final RSE can be lower than the baseline RSE although it has more covariates. In cases where the final is higher, it is only marginally so showing no problems of overfitting.

References

Chromy, J.R. (1981). Variance estimation for a sequential selection procedure. In Current Topics in Survey Sampling (eds. Krewski, D., Platek, R, and Rao, J.N.K.), *Academic Press*, New York, pp 329-347.

Deville, J.-C., and Särndal, C.E. (1992). Calibration estimation in survey sampling. *JASA*, 87, 376-382.

Folsom, R.E., Jr., and Singh, A.C. (2000). A generalized exponential model for sampling weight calibration for extreme values, nonresponse, and poststratification. *ASA Proc. Surv. Res. Meth. Sec.*

Vaish, A.K., Gordek, H., and Singh, A.C. (2000). Variance estimation adjusted for weight calibration via the generalized exponential model with application to the National Household Survey on Drug Abuse. *ASA Proc. Surv. Res. Meth. Sec.* (to appear).

Table 1. Summary of Model Characteristics (Census Division: East South Central)					
	SDU Level		Person Level		
	sdu.nr	sdu.ps	sel.per.ps	res.per.nr	res.per.ps
Sample Size	8,903	8,230	4,664	3,688	3,688
Resp. Rate(unwtd,wtd)	(92.4% , 92.3%)		(79.1% , 72.1%)		
Bounds(L<C<U)					
High-extreme values	(1.00,1.08,1.35)	(0.30,1.00,1.01)	(0.30,1.00,1.70)	(1.00,1.39,1.40)	(0.30,1.00,1.20)
Non-extreme values	(1.00,1.08,1.45)	(0.30,1.00,2.20)	(0.30,1.00,3.50)	(1.00,1.39,3.50)	(0.30,1.00,3.00)
Low-extreme values	(1.00,1.08,1.45)		(0.40,1.00,3.50)		
# Variables in the Model					
Total	85(of 204)	111(of 144)	164(of 238)	138(of 238)	98(of 158)
1 factor	20(of 22)	16(of 16)	33(of 34)	33(of 34)	16(of 16)
2 factor	52(of 86)	53(of 57)	92(of 119)	74(of 119)	48(of 57)
3 factor	13(of 96)	42(of 71)	39(of 85)	32(of 85)	34(of 85)

Table 2. Summary Statistics for 1999 NHSDA CAI Weight Adjustments - (Census Division: East South Central)										
	Screener DU Level				Person Level					
	sdu.nr		sdu.ps		sel.per.ps		res.per.nr		res.per.ps	
	Before	After	Before	After	Before	After	Before	After	Before	After
UWE	1.11	1.12	1.12	1.15	3.02	3.27	3.28	3.87	3.87	3.87
Extreme Values										
Unwtd	1.49%	0.27%	0.27%	0.49%	0.00%	1.09%	0.89%	1.14%	1.44%	0.38%
Wtd	1.94%	0.61%	0.61%	1.01%	0.00%	2.29%	1.96%	2.66%	3.09%	0.97%
Outwinsor	0.29%	0.06%	0.06%	0.08%	0.00%	0.35%	0.40%	0.40%	0.52%	0.09%
Weight Distribution										
Wt. Adj. Factor										
Min		0.59		0.3		0.3		0.73		0.3
25%		1.03		0.98		0.83		1.07		0.97
Median		1.07		1.08		0.96		1.16		1.01
75%		1.11		1.22		1.12		1.34		1.05
Max		1.45		2.2		3.49		3.5		2.96
Weight Product										
Min	353.1	365.7	365.7	112.5	131.6	39.9	39.9	40	40	13.5
25%	502.2	552.9	552.9	570.6	713.4	675.8	667.1	767	767	775.4
Median	583.7	693.6	693.6	746	1,135.00	1,117.40	1,079.40	1,327.50	1,327.50	1,347.00
75%	819.8	875.8	875.8	975.1	2,714.20	2,762.70	2,487.20	3,161.70	3,161.70	3,093.10
Max	2,708.20	2,523.60	2,523.60	3,914.60	32,261.60	66,015.70	66,015.70	70,614.60	70,614.70	62,606.80

Table 3. Ratio of Estimates under Baseline and Final models and comparison of Unadjusted and ps-adjusted RSE(%)					
	Ratio of Estimates (F/B)	Unadjusted RSE		ps-adjusted RSE	
		Baseline	Final	Baseline	Final
Marijuana Past Year					
Overall	1	6.75	6.81	6.4	6.88
17-Dec	1.02	8.61	8.81	8.3	8.63
18-25	1	6.5	6.59	6.52	6.55
26-34	0.98	16.56	16.57	16.65	16.89
35+	1	16.53	17.02	17.65	18.31
Alcohol Past Year					
Overall	1	3.15	3.18	3.08	3.55
17-Dec	1.02	4.91	4.94	4.55	4.51
18-25	0.99	3.45	3.54	3.41	3.28
26-34	1	3.84	3.83	3.71	3.77
35+	1	5.28	5.3	5.12	6.43
Cocaine Past Year					
Overall	1	21.46	21.72	20.96	19.9
17-Dec	1.01	28.76	29.26	31.49	42.94
18-25	1	15.44	15.6	15.63	15.34
26-34	0.99	31.55	31.12	32.26	30.31
35+	1.01	48	48.08	46.67	45.81
Cigarette Past Year					
Overall	1	4.43	4.38	4.22	4.48
17-Dec	1.02	5.49	5.55	5.41	5.49
18-25	1	3.76	3.84	3.74	3.75
26-34	0.99	7.17	7.14	7.32	6.16
35+	1	7.42	7.36	7.06	7.16

Exhibit 1. The Weight Components at the dwelling unit and person levels

Phase I 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 Dwelling Unit
5 Inverse Probability of Added Dwelling Unit
6 Dwelling Unit Subsampling / Release Adjustment

7 Dwelling Unit Nonresponse Adjustment (*sdu.nr*)
8 Dwelling Unit Poststratification Adjustment (*sdu.ps*)
9 Dwelling Unit Extreme Value Check/ Adj (*sdu.ev*)

Phase II Person Level

Design Weight Components

10 Inverse Probability of Selecting a Person Within a Dwelling Unit
11 Person Subsampling Adjustment

12 Person(selected) Level Poststratification (*sel.per.ps*)
13 Person(respondent) Level Nonresponse Adjustment (*res.per.nr*)
14 Person(respondent) Level Poststratification Adjustment (*res.per.ps*)
15 Person(respondent) Level Extreme Value Check/ Adj (*res.per.ev*)