

ICF INTERNATIONAL

Sampling Plan for the 2015 National Youth Tobacco Survey (NYTS)

Submitted by ICF International

5/27/2014

Table of Contents

Chapter 1—Introduction 2

Chapter 2—NYTS Sampling Design 3

2.1	Overview.....	3
	Exhibit 2-1: Key Sampling Design Features.....	3
2.2	Design Updates and Modifications.....	4
	Exhibit 2-2 Historical Trends for Non-Hispanic black Students.....	4
	Exhibit 2-3 Historical Trends for Hispanic Students.....	4
	Exhibit 2-4 Impact of a School Threshold of 25.....	6
	Exhibit 2-5 Impact of a School Threshold of 40.....	6
2.3	Frame construction.....	7
	Exhibit 2-6 Increases in Coverage for Public Schools and Non-Public Schools by Using a Combined Frame.....	7
	Exhibit 2-7 Increases in Coverage for Students in Public Schools and Non-Public Schools by Using a Combined Frame.....	8
2.4	Design Features for Minority Student Oversampling.....	8
2.5	Sampling Stages.....	9
2.6	Stratification and Linking Schools.....	10
	Exhibit 2-8: Preliminary Stratum Definition and PSU Allocation to Strata for NYTS 2014.....	12
2.7	Sample Sizes.....	12
	Exhibit 2-9: Planned Sample Sizes for the 2015 NYTS.....	14

Chapter 3— Sampling Methods 16

3.1	PSU Selection.....	16
3.2	School Selection.....	17
3.3	Grades.....	18
3.4	Classes.....	18

Chapter 4—Weighting 20

4.1	Overview.....	20
4.2	Sampling Weights.....	20
4.3	Non-Response Adjustments.....	22
4.4	Weight Trimming.....	23
4.5	Post-Stratification to National Student Population Estimates.....	24

Chapter 1—Introduction

The objective of the NYTS sampling design is to support estimation of tobacco-related knowledge, attitudes, and behaviors in a national population of public and non-public school students enrolled in grades 6 through 12 in the United States. More specifically, the study is designed to produce national estimates at a 95% confidence level by school level (middle school and high school), by grade (6, 7, 8, 9, 10, 11, and 12), by sex, and by race/ethnicity for non-Hispanic white, non-Hispanic non-Hispanic black, and Hispanic students. Additional estimates, such as cross-tabulations of grade by sex and of race/ethnicity by school level, are also supported; however, precision levels will vary considerably according to differences in sub-population sizes.

The universe for the study consists of all public and non-public school students enrolled in regular middle schools and high schools in grades 6 through 12 in the 50 states and the District of Columbia. Alternative schools, special education schools, Department of Defense operated schools, vocational schools that serve only pull-out populations, and students enrolled in regular schools unable to complete the questionnaire without special assistance are excluded.

The NYTS study is a continuation of the NYTS survey cycles that took place in 1999, 2000, 2002, 2004, 2006, 2009, 2011, 2012, 2013 and 2014. The NYTS survey system employs a repeat cross-sectional design to develop national estimates of tobacco use behaviors and exposure to pro- and anti-tobacco influences among students enrolled in grades 6–12. The 2011 and 2013 survey cycles of the NYTS were coordinated with the national Youth Risk Behavior Survey (YRBS); the 2015 NYTS is also intended to be coordinated with the 2015 YRBS.

Chapter 2 provides a description of the sampling design planned for the 2015 NYTS. Chapter 3 provides a description of the sampling methods. Chapter 4 describes the planned weighting plan.

Chapter 2—NYTS Sampling Design

2.1 Overview

The sample will be a stratified, three-stage cluster sample with primary stratification by racial/ethnic concentrations and urban versus rural status. Primary sampling units (PSUs) are classified as “urban” if they are in one of the 54 largest metropolitan statistical areas (MSAs), in the U.S.; otherwise, they are classified as “rural.” Within each stratum, PSUs, defined as a county, a portion of a county, or a group of counties, will be chosen without replacement at the first stage.

At the next sampling stage, Secondary School Units (SSUs), formed from schools so as to support class selection at each grade, will be classified into strata based on enrollment size (Small, Medium and Large) and school level (middle- versus high-school). In subsequent sampling stages, a probabilistic selection of SSUs is followed by a selection of grades and students in subsequent sampling stages.

Exhibit 2-1 presents a summary of the sampling design features.

Exhibit 2-1: Key Sampling Design Features

Sampling Stage	Sampling Units	Sample Size (Approximate)	Stratification	Measure of Size
1	PSUs: Counties or groups of counties	85	Urban vs. non-urban (2 strata); Minority concentration (8 strata)	Aggregate school size in target grades
2	Schools	220 school selections: 170 large schools (2 per PSU), 20 medium schools and 30 small schools	Small, medium and large; High-school vs. middle-school	Aggregate eligible enrollment
3	Classes/ students	1 or 2 classes per grade (2 per grade in large high-minority schools)		

The next section, Section 2.2, describes the planned modifications to the sampling design as well as updates that are typically made in each cycle. Section 2.3 describes the frame construction methods that capitalize from different data sources. This section summarizes the coverage advantages of the improved, combined method of building the sampling frame of public and non-public schools. Section 2.4 describes the approaches for oversampling minority students (i.e., non-Hispanic black and Hispanic students). Section 2.5 describes the sampling stages, and Section 2.6 describes the stratification planned for primary sampling units and schools. Section 2.7 discusses the sample sizes planned for the study.

2.2 Design Updates and Modifications

We plan to replicate the main features of the 2013 and 2014 NYTS sample designs. As in the past few cycles, we will continue to adjust sampling parameters to reflect changing demographics of the in-school population.

2.2.1 Decreasing Need to Oversample non-Hispanic Black and Hispanic Students

In general, as the proportion of non-Hispanic black and Hispanic students in the study population increases and the minority population becomes more evenly distributed, the parameters that drive minority oversampling can be relaxed, allowing us to maintain yields while moving towards a statistically more efficient design.

Specifically, growing percentages of non-Hispanic black and Hispanic students have allowed the design to adjust two design dimensions towards greater design efficiency (i.e., closer to a self-weighting design):

- 1) The measure of size (MOS) will be eligible enrollment rather than a weighted MOS designed to oversample minority students;
- 2) The allocation to strata will be proportional, or nearly so, rather than oversampling strata with higher concentrations of minority students.

The historical data on the concentrations of non-Hispanic black and Hispanic students reinforce the finding that oversampling via the measure of size is no longer necessary to achieve sufficient numbers of non-Hispanic black and Hispanic students. Exhibits 2-2 and 2-3 present the percentages of public middle-school high-school students who are non-Hispanic black and Hispanic, respectively, for the years 2008-09, 2009-10, 2010-11 and 2011-12. The tables show that while the percentage of non-Hispanic black students has remained stable, the percentage of Hispanic students has been steadily increasing over the last few years. The percentage of Hispanic high-school students has increased from 19% in 2008-09 to nearly 22% in 2011-12.

Exhibit 2-2 Historical Trends for Non-Hispanic black Students

	2008-2009	2009-10	2010-11	2011-12
Middle School	16.55%	16.54%	16.10%	15.98%
High School	16.90%	16.79%	16.23%	15.94%

Exhibit 2-3 Historical Trends for Hispanic Students

	2008-2009	2009-10	2010-11	2011-12
Middle School	20.90%	21.52%	22.61%	23.19%
High School	19.04%	19.88%	20.99%	21.72%

2.2.2 Design Updates

Other design features are also routinely updated in each cycle such as:

- The stratum boundaries based on the percentage of minority students will be re-computed to minimize variances according to the cumulative square root rule (Dalenius-Hodges rule).¹
- We will adjust PSU definitions to account for school openings and closings, and may also adjust PSU sample sizes by one or two (in either direction) if the simulated yields indicate the need for adjusting sample sizes.

In addition, we continue the practice instituted in the 2014 NYTS of constructing a more comprehensive sampling frame from different data sources. This is discussed in more detail in Section 2.3.

2.2.3. School Size Threshold

Another modification will be imposing a threshold for school size so that the frame will not include very small schools, those with an aggregate school enrollment of 25 or less in the eligible grades. The school size threshold was established primarily for cost efficiency, but also due to concerns about confidentiality, in consultation with CDC. As discussed below, we considered that the cost of recruiting and collecting data from very small schools outweighed the benefit of adding a relatively small number of students that attend this subset of schools (less than 0.6% of all eligible students).

The eligible school frame will exclude schools with student enrollment below a minimum value, c , in order to enhance efficiencies in data collection. The gains in efficiency may come at the price of under-coverage of small schools, with the potential for associated biases. This section summarizes the results of our investigation of the under-coverage impact of requiring a minimum school size.²

This analysis looks at the percentage of students that would be left out of the frame for varying values of the threshold. To assess the potential bias that might be associated with these exclusions, we also examine the percentage of non-Hispanic black and Hispanic students who are left out of the frame when very small schools are not included in the school frame.³ The analysis shows that the bias potential is very small for the size threshold that we plan to use, $c=25$. For comparison purposes, we also consider a size threshold of $c=40$ that would lead to larger under-coverage losses.

1 Dalenius, T. and Hodges, K. (1959) "Minimum variance stratification." *Jour. Amer. Statist. Assoc.*, 54, 88-101.

2 The new method for frame construction improves coverage by using a frame that combines MDR and NCES data files rather than relying on a single source. This method, shown to improve coverage for the 2014 NYTS, adds a disproportionately large number of very small schools that used to be left out of the frames based solely on the MDR files.

3 In theory, bias due to loss of coverage of these very small schools might also be assessed by comparing selected estimates of risk behavior outcomes for students in these schools with estimates from the balance of the schools or with overall estimates). This comparison is not statistically possible, however, as the number of tiny schools is relatively small in recent cycles of the surveys, and so is the student yield in these schools.

We investigated the percentage of students that would be left out of the frame for different values of the school size threshold. The two tables below (Exhibits 2-4 and 2-5) show the percent of students omitted from the frame when schools below a given size threshold are dropped.⁴ The relative loss is addressed for thresholds of 25 (Exhibit 2-4) and 40 (Exhibit 2-5). The tables consider the single frame design used in the 2013 YRBS-NYTS cycle as well as the combined frame design used in the 2014 NYTS cycle.

Exhibit 2-4 Impact of a School Threshold of 25

<i>Cycle and Frame</i>	<i>Survey and Grade Range</i>	<i>Percent of Students Lost</i>	<i>Percent of Non-Hispanic black Students Lost</i>	<i>Percent of Hispanic Students Lost</i>
2013, Single Source Frame	NYTS (Grades 6 - 12)	0.37%	0.08%	0.10%
2013, Single Source Frame	YRBS (Grades 9 - 12)	0.30%	0.06%	0.08%
2014, Dual Source Frame	NYTS (Grades 6 - 12)	0.59%	0.43%	0.29%
2014, Dual Source Frame	YRBS (Grades 9 - 12)	0.51%	0.44%	0.30%

Exhibit 2-5 Impact of a School Threshold of 40

<i>Cycle and Frame</i>	<i>Survey and Grade Range</i>	<i>Percent of Students Lost</i>	<i>Percent of Non-Hispanic black Students Lost</i>	<i>Percent of Hispanic Students Lost</i>
2013, Single Source Frame	NYTS (Grades 6 - 12)	0.90%	0.30%	0.29%
2013, Single Source Frame	YRBS (Grades 9 - 12)	0.64%	0.13%	0.17%
2014, Dual Source Frame	NYTS (Grades 6 - 12)	1.27%	0.96%	0.66%
2014, Dual Source Frame	YRBS (Grades 9 - 12)	0.97%	0.83%	0.56%

Exhibit 2-4 shows that 0.37% of the students would have been excluded from the 2013 NYTS frame using a truncation threshold of 25 students, and 0.30% excluded from the 2013 YRBS frame at the same threshold. Exhibit 2-5 shows that for a threshold of 40, these percent exclusions go up to 0.90% and 0.64%.

⁴ For the NYTS, this is the aggregate over middle school grades and high school grades separately, and for the YRBS we consider aggregation over high school grades only.

In summary, the size threshold $c=25$ will lead to small levels of student-level under-coverage, and therefore, minimum impact on student-level estimates. At the same time, excluding these very small schools will lead to substantial efficiencies in recruitment efforts and in increased student yields per visited school.

2.3 Frame construction

For the 2015 NYTS, the sampling frame construction will follow the improved procedures adopted in the 2014 NYTS. In past cycles of NYTS, a single source of national schools has been used as the sampling frame. The source, a list produced by MDR, contains school information including enrollments, grades, race distributions within the school, district and county information, and other contact information for public and non-public schools across the nation.

For the 2015 NYTS survey we will use a combination of sources to create the school frame in order to increase school coverage. Along with the MDR dataset, we will use two files from the NCES; the CCD which is a national file of public schools and the PSS, a file of national non-public schools. The principle behind combining multiple data sources is to increase the coverage of schools nationally.

Exhibit 2-6 and Exhibit 2-7 present the coverage gains for public and non-public schools. Exhibit 2-6 shows the relative increases in the number of schools and Exhibit 2-7 shows the relative increases in the number of students. The baseline for our comparisons is the use of the MDR data files alone, the approach used in all cycles prior to 2014.

Exhibit 2-6 Increases in Coverage for Public Schools and Non-Public Schools by Using a Combined Frame




Exhibit 2-7 Increases in Coverage for Students in Public Schools and Non-Public Schools by Using a Combined Frame




Exhibit 2-6 and Exhibit 2-7 show that the coverage gains are larger for non-public schools than for public schools. Exhibit 2-6 shows a school coverage increase of about 38% for non-public schools and 12% for public schools. The figures are also broken down by school level. The gains are also larger for high schools than for middle schools. Exhibit 2-7 shows a coverage increase of about 15% in terms of students for non-public schools and about 2% of students for public schools. In summary, these two exhibits reflect the addition to the frame of a large number of schools, a majority of which are non-public small schools.

2.4 Design Features for Minority Student Oversampling

To facilitate accurate prevalence estimates among minority groups, prior cycles of the NYTS have employed multiple strategies to increase the number of non-Hispanic black and Hispanic students included in the sample. These have included over-sampling PSUs in high-minority strata, the use of a weighted measure of size, and double class selection in high-minority schools.

A weighted measure of size (MOS) was used in all NYTS cycles prior to 2013 to increase the probability of selection of high-minority PSUs and schools within a probability proportional to size (PPS) sampling design. As the use of an unweighted MOS increases the statistical efficiency of the design, the 2013 NYTS moved to the use of an unweighted MOS. That is, student enrollment was the measure of size used for the 2013 NYTS and the 2014 NYTS and will be used in the 2015 NYTS.

Prior to sample selection, we will conduct simulation studies to set design parameters, such as the class doubling thresholds, and to ensure that the resulting sampling process will yield adequate precision with a minimum variance.

In previous NYTS cycles, high-minority schools were subject to double class selection, such that two classes per grade were selected in these schools (compared to one class per grade in other schools) to increase the number of minority students sampled. The 2015 NYTS will continue this practice, with double class selection occurring in the subset of schools with largest concentrations of non-Hispanic black students.

Starting with the 2011 NYTS, the design moved to a proportional allocation of PSUs to first-stage sampling strata. This step represents the culmination of a trend in the design over the past several cycles, driven by changes in the underlying student population – that of a reduced need for oversampling PSUs in high non-Hispanic black and high Hispanic strata.

2.5 Sampling Stages

The three-stage cluster sample will be stratified by racial/ethnic composition and urban versus non-urban status at the first stage. PSUs are defined as a county, a group of smaller counties, or a portion of a very large county. PSUs are classified as “urban” if they are in one of the 54 largest MSAs in the U.S.; otherwise, they are classified as “non-urban.”

Additional, implicit stratification will be imposed by geography by sorting the PSU frame by state and by 5-digit ZIP Code (within state). Within each stratum, a PSU will be randomly sampled without replacement at the first stage.

In subsequent sampling stages, a probabilistic selection of schools and students will be made from the sample PSUs.

The sampling stages may be summarized as follows:

- Selection of PSUs—Eighty-five PSUs were selected from sixteen strata with probability proportional to the total number of eligible students enrolled in all eligible schools located within a PSU.
- Selection of Schools—At the second sampling stage, two large schools will be selected from each PSU. Among medium schools, 10 high schools and 10 middle schools will be selected from a sub-sample of 10 PSUs. Similarly, among small schools a separate random sample of 15 middle schools and 15 high-schools will be taken from 15 sub-sample PSUs. The PSU subsamples for Medium and Small schools will be selected with simple random sampling, and the schools will be drawn with probability proportional to the total number of eligible students enrolled in a school.
- Selection of Students—Classes are selected based on two very specific scientific parameters to ensure a nationally representative sample. First, classes have to be selected in such a way that all students in the school have a chance to participate. Second, all classes must be mutually exclusive so that no student is selected more than once. All students in a selected classroom will be selected for the study.

The selection of grades from SSUs is discussed separately (Sections 2.6 and 3.3); this selection is not considered a sampling stage but rather a mechanism for selecting classes and students within sampled schools.

Schools will be stratified into Large, Medium and Small schools based on their ability to support two, one, or less than one class selection per grade. We will select two classes per grade in selected large schools, and one class per grade in the remaining schools. The threshold for large schools, and for double class sampling, will be set based on the simulation study. This will ensure that the required numbers of minority students are achieved per school level.

The sampling approach utilizes PPS sampling methods with the MOS defined as the count of final-stage sampling units, students. Coupled with the selection of a fixed number of units, the design results in an equal probability of selection for all members of the universe; i.e., a self-weighting sample. For the NYTS, we approximate these conditions, and thus obtain a roughly self-weighting sample.

The measure of size is used also to compute stratum sizes and PSU sizes. Assigning an aggregate measure of size to PSU, the sample allocates the PSU sample in proportion to the student population.

2.6 Stratification and Linking Schools

This section describes frame preparation steps for the selection of the first- and second-stage samples of PSUs and schools. These steps include combining counties into PSUs, linking schools into SSUs, and the stratification and allocation methods at these stages.

Defining a PSU

In defining PSUs, several issues are considered:

- a. Each PSU should contain at least 4 middle and 5 high schools.
- b. Each PSU should be large enough to contain the requisite numbers of schools and students by grade, yet not so large as to be selected with near-certainty.
- c. Each PSU should be compact geographically so that field staff can go from school to school easily.
- d. There should be recent data available to characterize the PSUs.
- e. PSU definitions should be consistent with secondary sampling unit (school) definitions

Generally, PSUs are counties with two exceptions: (1) low population counties are combined to provide sufficient numbers of schools and students, and (2) counties that are very large may be split to avoid becoming certainty or near-certainty PSUs. The 2015 NYTS PSU definitions will be based on the 2014 NYTS definitions and updated to ensure that all PSU meet the criteria above. County population figures will be aggregated from school enrollment data for the grades of interest.

Stratification of PSUs

The PSUs are organized into 16 strata, based on urban/non-urban location (as defined above) and minority enrollment. The approach involves the computation of optimum stratum boundaries using the cumulative square root of “f” method developed by Dalenius and Hodges. The boundaries or cutoffs change as the frequency distribution (“f”) for the racial groupings change from one survey cycle to the next. These rules are summarized below.

- If the PSU is within one of the 54 largest MSA in the U.S. it is classified as “urban,” otherwise it is classified as “non-urban.”
- If the percentage of Hispanic students in the PSU exceeds the percentage of non-Hispanic black students, then the PSU is classified as Hispanic. Otherwise it is classified as non-Hispanic black.
- Hispanic urban and Hispanic non-urban PSUs are classified into four density groupings depending upon the percentages of Hispanics in the PSU. The stratum boundaries, or cutoffs, will be computed using new frame data.
- Non-Hispanic black urban and non-Hispanic black non-urban PSUs are also classified into four groupings depending upon the percentages of non-Hispanic blacks in the PSU. The stratum boundaries, or cutoffs, will be computed using new frame data.

Allocation of the PSU Sample

We will design and select a sample of 85 PSUs, allocated in proportion to student enrollment. Using simulations as in previous studies, we will then make adjustments to the initial allocation to meet minority targets. Specifically, the adjustments round fractional allocations and ensure that each stratum have at least two sampled PSUs.

As an example, Exhibit 2-8 presents the allocation of the PSU sample to strata used in the 2014 cycle. The stratum boundaries and allocation to strata will be updated using the new frame.

Compared to previous cycles, this allocation is closer to proportional and therefore more efficient statistically; i.e., it leads to smaller variances and tighter confidence intervals.

Exhibit 2-8: Preliminary Stratum Definition and PSU Allocation to Strata for NYTS 2014

Predominant Minority	Urban / Non-Urban	Density Group Number	stratum	Student Population	Percent of Student Population	Number of Sample PSU's
Non-Hispanic black	Urban	1	BU1	2,720,181	9%	9
		2	BU2	975,490	3%	3
		3	BU3	908,299	3%	3
		4	BU4	516,712	2%	2
	Non-urban	1	BR1	3,937,157	14%	12
		2	BR2	1,503,403	5%	5
		3	BR3	1,026,612	4%	4
		4	BR4	313,063	1%	2
Hispanic	Urban	1	HU1	3,530,556	12%	11
		2	HU2	2,429,442	8%	7
		3	HU3	1,865,988	6%	5
		4	HU4	2,106,242	7%	7
	Non-urban	1	HR1	4,427,215	15%	14
		2	HR2	1,284,402	4%	4
		3	HR3	988,655	3%	3
		4	HR4	523,491	2%	2

2.7 Sample Sizes

The original specifications for NYTS sample sizes were not given in terms of student yields; rather, they were specified in terms of the precision of the resulting estimates. Thus the NYTS was designed to produce the key estimates accurate to within $\pm 5\%$ at a 95% precision level. Estimates by grade, sex, and grade by sex should meet this standard. The same standard is used for the estimates for racial/ethnic groups by school level.

Specifically, the NYTS is designed to produce accurate estimation to within $\pm 5\%$ at a 95% precision level for the following key subgroup estimates:

- Middle and high school estimates (school level)—middle school students in total (grades 6–8 combined) and high school students in total (grades 9–12 combined)
- Grade estimates—Individual grades 6, 7, 8, 9, 10, 11, and 12, separately
- Sex group estimates—males and females in total, by school level (male middle school students, female high school students, etc.), and by individual grade (6th grade males, 6th grade females, etc.)
- Racial group estimates (race/ethnicity)—in total and by school level (e.g., Hispanic middle school students)

Over the past several cycles of the NYTS, we have confirmed that sample sizes, and resulting student yields, were sufficient to achieve design goals in terms of precision. For the 2013, 2014

and 2015 NYTS survey design, anticipated precision levels were developed in order to ensure that the design meets the original precision targets.

The 2015 NYTS sample design is consistent with the sample design used in past cycles, which includes adjusting sampling parameters to reflect changing demographics of the in-school population of middle and high school students. The balance of this section presents this development. The number of schools selected and students per school were calculated so that the study was projected to result in completed surveys from approximately 20,077 students.

As detailed earlier, linked schools are constructed so as to contain a full complement of grades—6 to 8 for middle schools and 9 to 12 for high schools. Schools are further classified by size based on grade-level enrollments. This allows us to ensure that a sampled school of a given size classification is able to support the student sample sizes provided earlier.

The NYTS sample size calculations are premised on the following assumptions:

- The main structure of the sampling design will be consistent with the design used to draw the sample for prior cycles of the NYTS.
- The selection of a minimum of one SSU at the high school level and one SSU at the middle school level within each PSU. Some PSUs are selected to provide up to four extra schools. A PSU is a county or a group of contiguous counties.
- SSUs with at least 56 students per grade are considered Large, and those among the others with at least 28 students per grade are considered Medium; otherwise they are considered Small.
- On average, each selected class will include 28 students (based on historical averages).
- For SSUs classified as large schools, we will take two sections of students in 46% of these schools.
- A 72% overall response rate (based on historical averages) calculated as the product of the school and student response rate.

Based on these assumptions, we will select a sample of 85 PSUs. Within each PSU, we will draw two large schools, one at the middle school level to supply students in grades 6 through 8, and one at the high school level to supply students in grades 9 through 12. In addition, 15 PSUs will be sub-sampled to supply small SSUs, and another 10 PSUs will be independently sub-sampled to supply medium SSUs.

The number of students selected from all sample schools, will be about 27,884 students (before non-response).

Exhibit 2-9 summarizes the designed sample sizes for each school type. This table details the number of schools that were specified to be drawn by the sample design along with the number of participating schools and students anticipated when we developed the sample design based on the given assumptions.

In this exhibit, SSUs are “virtual schools” created by combining actual, physical schools so that each virtual school unit has a complete set of grades for the level. The virtual schools will be expanded to physical schools.

Across the seven previous cycles of the NYTS, the school participation has averaged 90%, with a low of 83%. Student participation has averaged 91% with a low of 88%. In calculating the sample sizes for the 2015 NYTS, we make our approach more robust by assuming a conservative overall rate of 72%, lower than the historical response rate, 82% overall.

Exhibit 2-9: Planned Sample Sizes for the 2015 NYTS

	Middle School Student Selections	High School Student Selections	Total Selections	Participants
Small Schools	512	1,089	1,601	1,153
Medium Schools	840	1,120	1,960	1,411
Large Schools				17,513
With two classes	6,569	8,758	15,327	11,036
With one class	3,856	5,141	8,996	6,477
TOTAL	11,777	16,108	27,884	20,077

In summary, it is anticipated that between 235 and 245 physical schools will be selected for participation in the 2015 NYTS, a number inflated due to the linking of physical schools into SSUs. These schools are expected to yield approximately 20,077 students.

The large projected sample size permits analysis by individual grade and by sex without any special considerations in the sampling plan. Additionally, grade and sex subgroups both typically cut across schools. Design effects are typically smaller for subgroups such as groups that cut across schools. Sex group estimates will have better precision than other groups that are less evenly dispersed across schools (e.g., racial/ethnic groups).

Because the design yields a greater number of completed surveys from high school students than from middle school students, the estimation precision is expected to be better at the high school level than at the middle school level. Moreover, because within grade estimates by sex have slightly larger standard errors than those for estimates by grade alone, estimates of sex were expected within $\pm 5\%$.

Middle School and High School Estimates

Estimates by school level are required to support separate analysis of students across middle school grades (6, 7, and 8) and high school grades (9, 10, 11, and 12). However, schools tend to vary in their grade structures, an inconsistency that compromises the ability to easily and efficiently link schools for sampling purposes in a manner that also uniformly divides students by grade. For example, 9th grade students are served by both grades 7–9 junior high schools and by grades 9–12 high schools. As a result, we have developed the school linking approach described earlier, and with this approach being applied independently for high schools and middle schools.

Estimates by Grade

NYTS estimates are typically not reported by grade level but rather by school level (middle school and high school). Still, the design balances the sample sizes for grade level by targeting at least 3,000 students per grade. This ensures that estimates at the grade level achieve the required precision levels. It is worth noting that this design feature resulted in a larger student allocation to the high school stratum than to the middle school stratum as high schools have four grades versus three grades for middle schools.

Estimates by Sex

The large designed sample size permits analysis by sex without any special considerations in the sampling plan. During the class selection process, frames of eligible classes from co-educational schools in which classrooms were segregated by sex (i.e., an all-male or all-female class) are avoided if at all possible.

Estimates by Racial Group

In order to support separate analysis of the data for white, non-Hispanic black and Hispanic students, in total and by school level, adequate sample sizes are required by the designed for subgroups defined by school level by racial grouping or by sex by racial grouping. Sample sizes are not designed, however, to support detailed analyses by sex and school level within racial/ethnic subgroups (e.g., middle school Hispanic males).

Chapter 3— Sampling Methods

This chapter describes the methods traditionally used by the NYTS in the selection of PSUs, schools, grades, and classes of students. In this process, we define the probabilities of selection associated with the various sampling stages as follows:

- Probability of selecting PSUs
- Probability of selecting schools
- Probability of selection of grades
- Probability of selecting classes and students

These probabilities provide the basis for the sampling weights discussed in Chapter 4.

The overall probability of selection for a student is the product of the probability of selection of the PSU, which is a group of schools, multiplied by the conditional probability of selecting the student's school, multiplied by the conditional probability of selecting the student's class. These steps are detailed in the selection below.

3.1 PSU Selection

Selection

Within each first-stage stratum, the PSUs will be sorted by five-digit ZIP Code to attain a form of implicit geographic stratification. Implicit stratification, coupled with the probability proportional to size (PPS) sampling method described below, ensures geographic sample representation. With PPS sampling, the selection probability for each PSU is proportional to the PSU's measure of size.

The following systematic sampling procedures are applied to the stratified frame to select a PPS sample of PSUs.

- Select 85 PSUs with a systematic random sampling method within each stratum. The method applies within each stratum a sampling interval computed as the sum of the measures of size for the PSUs in the stratum divided by the number of PSUs to be selected in the stratum.
- Subsample at random 10 of the sample PSUs for the medium school sample for each school level (middle school and high school)
- Subsample at random 15 of the sample PSUs for the small school sample for each school level (middle school and high school)

Probability

If MOS_{klm} is the measure of size for school k in PSU l in stratum m and if K_m is the number of PSUs to be selected in stratum m , then P_{lm}^p is the probability of selection of PSU l in stratum m :

$$P_{lm}^P = K_m \left(\frac{MOS_{.lm}}{MOS_{..m}} \right)$$

For the PSUs subsampled for the selection of medium and small schools, the sub-sample PSUs have an additional factor in their selection probability for these schools. This factor is incorporated into the school sampling probability below, as it is more closely associated with school selection.

3.2 School Selection

Selection

For large schools, one high school and one middle school are selected with PPS systematic sampling within a PSU. The schools are selected into the sample with probability proportional to the measure of size.

Small and medium schools will be sampled independently from large schools; they will be set in two separate strata sampled at lower rates. This approach will be implemented by drawing a sub-sample of 15 PSUs for the sampling of small schools sampling and a subsample of 10 PSUs for medium school sampling at each grade level. Then one small school or medium school will be selected in each sub-sampled PSU with probability proportional to the measure of size.

Replacement of Schools/School Systems

We will not replace refusing school districts, schools, classes, or students; we will, however, replace schools found to be ineligible during the recruitment process. We allow for school and student non-response by inflating the sample sizes to account for non-response. With this approach, all schools can be contacted in a coordinated recruitment effort, which is not possible for methods that allow for replacing schools.

Probability

The probability of selecting large school k in PSU l and stratum m , P_{klm}^{LS} , at each level is computed as follows:

$$P_{klm}^{LS} = \left(\frac{MOS_{klm}}{MOS_{.lm}} \right)$$

For small schools, one school is drawn from sub-sampled PSU at each level, so the probability of selection of a small school, P_{klm}^{SS} , then becomes:

$$P_{klm}^{SS} = (15/85) \left(\frac{MOS_{klm}}{MOS_{.lm}} \right)$$

For medium schools, one school is drawn from sub-sampled PSU, so the probability of selection of a medium school at each level, P_{klm}^{MS} , then becomes:

$$P^{MS}_{klm} = (10/85) \left(\frac{MOS_{klm}}{MOS_{.lm}} \right)$$

Note the additional sampling factor in the probability of selection for small schools and medium schools is due to the PSU sub-sampling for these schools as noted above.

3.3 Grades

Selection

Except for linked schools, all eligible grades are included in the class selection for each school.

In linked schools, grades are selected independently. One component school is selected to provide classes at each grade level, and grades within component schools are drawn with probability proportional to grade enrollment.

Probability

Most SSUs in the sample contained one component school. In these cases, all eligible grades are selected so that the probability of selecting a grade is 1.0.

In SSUs that are made up of more than one component schools, the selection of the component school at each grade is made with PPS sampling. The school selections from the component school at each grade level are made independently.

We denote this P^G_{jklm} the probability of selecting grade j in SSU k , in PSU l , stratum m . For the j^{th} grade within SSU k , this probability is equal to the ratio of the number of students at grade j in the component school to the total enrollment in grade j across all component schools within the SSU

3.4 Classes

Selection

In large schools, we select an average of 1.46 classes per grade by selecting 2 classes per grade in a subset of these schools and one class per grade in the remaining schools. The selection of schools for double class sampling—a subset of large schools—is described earlier. One class per grade is selected in medium schools. In small schools, that is, those that could not support a full class selection at each grade, all students in all eligible grades are taken into the sample.

All students in a selected class who can complete the survey without special assistance are considered eligible and offered the opportunity to participate in the survey. Refusing students are not replaced. Non-response at the student level is accounted for in the sample size using an average per class yield that assumes student response rates derived from historical experience with the NYTS.

A set of classes is identified for each school at each grade level such that every student in a given grade level is enrolled in exactly one of the classes in the set. For example, a required English course might be used. Selections are made at all eligible grade levels in the school.

Probability

The probability of selection of a class when there are C_{jklm} classes at grade j in school k , PSU i , stratum m is just $1/C_{jklm}$ or $2/C_{jklm}$ depending on whether 1 or 2 classes are taken in the school. All students in a selected class were chosen so the probability of selection of a student is the same as the class (i.e., $1/C_{jklm}$ or $2/C_{jklm}$).

Note that the probability of student selection within a class does not vary by race, ethnicity or sex. We denote this probability as P_{ijklm}^C as the probability of selecting class i in grade j , school k , PSU l , stratum m . Since every student in a selected class is also selected, the probability of selecting any student in class i , grade j , school k , PSU l , stratum k , is also equal to P_{ijklm}^C .

Chapter 4—Weighting

4.1 Overview

This chapter describes the procedures planned for weighting the NYTS 2015 data. The process will involve the steps outlined below:

- Sampling weights
- Non-response adjustments
- Weight trimming
- Post-stratification to national estimates of racial totals by grade, sex and school type

The final student level response data are weighted to reflect the initial probabilities of selection and non-response patterns, to mitigate large variations in sampling weights, and to post-stratify the data to known sampling frame characteristics.

4.2 Sampling Weights

The sampling weight attached to each student response is the inverse of the probability of selection for that student. This basic weight can be adjusted to compensate for non-response, to alleviate excess weight variation, and to match the weighted data to known control totals. A convenient way of computing the basic weight is by inverting the probabilities of selection at each stage, to derive a partial weight or stage weight. The stage weights are then multiplied together to form the overall weight.

4.2.1 Adjusted Conditional Student Weights

The adjusted conditional student weight is the student weight given the selection of the PSU, school and grade. This weight is the product of the inverse of the probability of selection, a non-response adjustment and a ratio adjustment to control to known school enrollment totals. This three step process is simplified to the ratio of the number of enrolled students to the number of responding students in a given weighting class within a school.

We denote the student selection weight W_{cklm}^R , where the subscripts k , l , and m refer to the school, PSU and stratum as before. The subscript c refers to the weight computation class, described below. This weight is computed as below, where N is the number of enrolled students⁵ and R is the number of responding students in weighting class c within a given school:

$$W_{cklm}^R = \frac{N_{cklm}}{R_{cklm}}$$

The weighting class definition is set dynamically, as described next, so as to avoid extreme weights.

⁵ The student enrollment for each school used in this calculation is obtained from the school during data collection. These counts are obtained by grade and sex.

Weighting class c is defined by a sequence of rules that depends on the number of responding students. This is done to avoid large weights for classes with low numbers of respondents. This process operates entirely within school.

Initially the weighting class is defined by grade and sex within each school. We then combine weighting classes if the weight for the class exceeds a maximum value. This cap C is computed using the equation following.

$$C_{cklm} = 2 \frac{N_{cklm}}{\min(10, N_{cklm})}$$

The combination sequence first combines males and females within grade. Then both the cap and the weight are re-computed. If the weight still exceeds the cap, grades are combined. The process is repeated, and if the student weight still exceeds the cap, the school is taken as the weight class.

This has the effect, within school, of setting an upper limit on the weight in class C of 2 in weight classes with an enrollment of less than 10, and 20% of the enrollment in weight classes with an enrollment of more than 10⁶.

4.2.2 School Sampling Weights

For large schools the partial school weight is the inverse of the probability of selection of the school given that the PSU was selected:

$$W_{klm}^{LS} = \left(\frac{MOS_{.lm}}{MOS_{klm}} \right) \frac{1}{P_{klm}^{LS}}$$

For small schools the partial school weight is:

$$W_{klm}^{SS} = (85/15) \left(\frac{MOS_{.lm}}{MOS_{klm}} \right) \frac{1}{P_{klm}^{SS}}$$

For medium schools the partial school weight is:

$$W_{klm}^{MS} = (85/10) \left(\frac{MOS_{.lm}}{MOS_{klm}} \right) \frac{1}{P_{klm}^{MS}}$$

4.2.3 Grade Sampling Weights

The partial weight for a grade, given the selection of the linked school containing it, is simply the inverse of the probability of selection described in Chapter 2. In a non-linked school the weight is 1.0. We denote the grade weight as W_{jklm}^G .

⁶ The cap could be exceeded in cases where the weight class is collapsed to the school level.

4.2.4 PSU Sampling Weights

The weight of the PSU is the inverse of the probability of selection of that PSU:

$$W_{lm}^P = \frac{1}{K_m} \left(\frac{MOS_{.m}}{MOS_{lm}} \right) \frac{1}{P_{lm}^P}$$

For small schools and medium school selections, the enclosing PSU were drawn as a subsample. This PSU subsampling component of the PSU weight is accounted for in the school selection probability and corresponding weight.

4.2.5 Overall Sampling Weight


The overall sampling weight is formed as the product of the stage selection weights. This weight, W^{T1} , is then adjusted for non-response, trimmed, and post-stratified to control totals as described in the following sections. This weight is computed as:

$$\begin{cases} W_{hijklm}^{T1} = W_{lm}^P W_{klm}^{LS} W_{jklm}^G W_{hijklm}^R \\ W_{hijklm}^{T1} = W_{lm}^P W_{klm}^{MS} W_{jklm}^G W_{hijklm}^R \\ W_{hijklm}^{T1} = W_{lm}^P W_{klm}^{SS} W_{jklm}^G W_{hijklm}^R \end{cases}$$

for large schools, medium schools, and small schools respectively, where the weights in the right hand side of the equations are defined in the preceding sections.

4.3 Non-Response Adjustments

This section describes how weights are adjusted for nonparticipation by entire schools, using strata as weighting classes. The adjustment process is different in small schools than in medium and large schools, as represented by the following equations for the adjustment factor.



The first equation applies to large and medium schools combined, and the second applies to

small schools. Note that this adjustment is made within stratum for large and medium schools and across the whole sample for small schools. The student weight, adjusted for non-response, is $A_{lm}^{SS} W_{hijklm}^{T1}$ for small schools and $A_{lm}^{LS} W_{hijklm}^{T1}$ for large and medium schools.

4.4 Weight Trimming

Extreme variation in sampling weights can cause inflated sampling variances, and offset the precision gained from a well-designed sampling plan. One strategy to compensate for this is to trim extreme weights and distribute the trimmed weight among the untrimmed weights. The method we used⁷ is based on a similar procedure done for the National Assessment of Educational Progress (NAEP). During the weighting task, we will investigate alternative methods based on the distribution of the weights. One specific method considers weights that are multiples of the interquartile range (IQR) away from the median weight as candidates for trimming.

The trimming traditionally used in the NYTS is an iterative procedure. During each iteration, an optimal weight, W_o ⁸ is calculated from the sum of the squared weights in the sample. Then, each weight W_i is marked and trimmed if it exceeds that optimal weight. The trimmed weight is summed within grade and spread out proportionally over the unmarked cases in the grade. This process is repeated for 20 iterations or no weight is being trimmed.

W_{ok} is determined by the following:

$$W_{ok} = c W_i$$

The constant c is arbitrary. Setting it to a low level will generate high levels of trimming; increasing it will reduce the level of trimming. For the current study, c has been set so that approximately 5% of the weight will be trimmed in the first iteration of the trimming algorithm.

Let W_{ik} and W_{ok} be the weight for the i th case and the optimum weight for the k th iteration, respectively, and define t_{ik} as 1 if W_{ik} is greater than or equal to W_{ok} , and zero otherwise.

⁷ Potter, F. (1988). Survey of Procedures to Control Extreme Sampling Weights. *American Statistical Association 1988 Proceedings: Survey Research Methods Section*, pp. 225–230.

⁸ In the following discussion, the subscripts are used to indicate the iterative process used in the trimming algorithm. To avoid overly cumbersome notation, we have omitted the subscripts indexing the sampling stages. W , the initial weight, is taken as the non-response adjusted sampling weight described in the preceding section. The subscripts k and n represent the number of iterations and the number of cases/weights respectively.

Then the trimmed weight for the $k + 1$ iteration is defined as follows:

$$W_{i,k+1} = \begin{cases} W_{ok} & \text{if } W_{ik} \geq W_{ok} \\ W_{ik} & \text{if } W_{ik} < W_{ok} \end{cases}$$

4.5 Post-Stratification to National Student Population Estimates

To obtain accurate counts of students in schools considered eligible for the NYTS by grade, sex, and race for use in post-stratification, we will turn to data available in two NCES data files. Non-public schools enrollments by grade and five racial/ethnic groups will be obtained from the PSS data file, and public school enrollments by grade, sex, and five racial/ethnic categories will be obtained from the CCD data file. These databases will be combined to produce the enrollments for all schools, and to develop population percentages to use as controls in the post-stratification step.

Specifically, population control totals for public school enrollments will be taken from the most recent CCD Public Elementary/Secondary School Universe Survey. Control totals for non-public school enrollments will be taken from the most recent cycle of the PSS.

The post-stratification adjustments will make the sum of the adjusted weights equal to the population control totals within each post-stratum cell. Post-stratification factors will include school type (public versus non-public) grade, sex, and race/ethnicity. One limitation of post-stratification tends to be limited by the number of post-stratum cells and the number of post-stratification factors that can be considered. To circumvent these limitations during the weighting task, we investigate alternative, iterative methods (i.e., raking) that allow for a larger number of post-stratification dimensions.