



Microdata User Guide

National Longitudinal Survey of Children and Youth

Cycle 6

September 2004 to June 2005



Statistics
Canada

Statistique
Canada

Canada

Table of Contents

1.0	Introduction	9
2.0	Background	11
3.0	Objectives	13
4.0	Main Changes at Cycle 6 since Cycle 5	15
4.1	Content Changes	15
4.2	Change to Timing of Collection.....	15
4.3	Methodology Changes	15
4.3.1	Sample	15
4.3.2	Weights	16
4.3.3	Variance	16
4.4	Processing Changes.....	16
4.5	Changes to Released Files	16
5.0	Survey Methodology - Sample.....	19
5.1	Overview of the Sample Design, Cycles 1 to 6.....	19
5.2	The Labour Force Survey	20
5.2.1	Target Population.....	20
5.2.2	Stratification	21
5.2.3	Cluster Delineation and Selection.....	22
5.2.4	Dwelling Selection.....	22
5.2.5	Sample Rotation.....	23
5.2.6	Household Members Eligible for the Labour Force Survey	23
5.3	Details of the National Longitudinal Survey of Children and Youth Sample.....	23
5.3.1	Original Cohort, Cycles 1 to 6.....	23
5.3.2	Early Childhood Development Cohorts Present at Cycle 6	25
5.4	Longitudinal and Cross-sectional Reference Populations for the National Longitudinal Survey of Children and Youth Weights	29
5.4.1	Cohorts and Their Longitudinal Populations.....	29
5.4.2	Cohorts and Their Cross-sectional Populations	31
6.0	Data Collection 33	
6.1	Household Component	33
6.2	Child Component	33
6.3	Adult Component	34
6.4	Youth Component	34
6.5	Sample Sizes at Cycle 6	34
6.6	Direct Assessments	36
6.7	Self-complete Questionnaires – Ages 10 to 17	36
6.7.1	Collection Personnel (Training, Supervision and Control).....	37
7.0	Data Processing.....	39
7.1	Computer Generated Edits	39
7.2	Data Capture.....	40
7.3	Clean Up	40
7.4	Age and Gender Edits.....	41
7.5	Relationship Edits	41
7.6	Pre-edits.....	41
7.7	Flow Edits.....	41
7.8	Coding of Open-ended Questions	41
7.9	Consistency Editing.....	42

7.10	Imputation Flags.....	42
7.11	Creation of Derived Variables and Longitudinal Edits	43
7.12	Standard Coding Structures.....	43
7.13	Naming Convention.....	45
7.14	Examples of Variables Names.....	47
7.15	Acronym Names for the Questionnaire Sections.....	47
7.16	Final Processing Files and Master Files	50
8.0	Content of the Survey	51
8.1	Survey Components.....	51
8.2	Demographic Variables.....	52
8.3	Adult Questionnaire.....	52
8.4	Child Questionnaire.....	54
8.5	Youth Questionnaire (Ages 16 to 21).....	58
8.6	Self-completes (Ages 10 to 17).....	62
9.0	Validation of the Survey Scales	67
9.1	Validation of Scale Data.....	67
9.1.1	Scale Definition	67
9.1.2	Scales and Calculations.....	67
9.1.3	Evaluation of Scale Data.....	67
9.2	Factor Analysis.....	68
9.2.1	Factor Analysis for Scales	68
9.2.2	Data Transformation Using Optimal Scaling.....	69
9.2.3	Factor Analysis Using Weighted Data	69
9.3	Calculation of Scores and Item Imputation	69
9.3.1	Calculation of Scores for Each Factor	69
9.3.2	Example of Factor Score Computation	69
9.3.3	Negative Loading	70
9.3.4	Nonresponse Codes	70
9.3.5	Raw Items	70
9.4	Reliability Measures for Scales.....	71
9.4.1	Cronbach’s Alpha.....	71
9.4.2	Interpretations of Cronbach’s Alpha	71
9.4.3	What is a Satisfactory Level of Reliability?	71
9.5	Parent-reported Scales	72
9.5.1	Depression Rating Scale	72
9.5.2	Home Responsibilities Scale	72
9.5.3	Family Functioning Scale.....	73
9.5.4	Neighbourhood Safety Scale	74
9.5.5	Social Support Scale.....	75
9.5.6	Behaviour Scales	76
9.5.7	Motor and Social Development Scale.....	79
9.5.8	Parenting Scales.....	81
9.5.9	Ages and Stages.....	83
9.6	Youth-reported Scales	83
9.6.1	Depression Scale.....	83
9.6.2	Neighbourhood Structure Scale.....	84
9.6.3	General Self-image Scale	84
9.6.4	Emotional Quotient	84
9.6.5	Social Support Scale.....	89
9.6.6	Friends Scale	89
9.6.7	My Parents and Me Scales.....	90
9.6.8	Conflict Resolution	91
9.6.9	Behaviour Scales	92
9.7	Summary of the Cycle 6 Scales.....	93

9.7.1	Parent-reported Scales	93
9.7.2	Self-completed Scales (Child/Youth Reported)	94
9.7.3	Youth-reported Scales (Self-reported).....	94
10.0	Imputation	95
10.1	Household Income Imputation	95
10.2	Youth Income Imputation – 16 or 17 Years Old.....	98
10.3	Youth Income Imputation – 18 Years Old and Over.....	99
10.4	Motor and Social Development (MSD) Scale Imputation	100
11.0	Weighting and Treatment of Nonresponse	103
11.1	Weights Available at Cycle 6	103
11.1.1	Funnel Weights (Variable FWTCdW1L)	103
11.1.2	Longitudinal Weights (Variable FWTCW01L).....	103
11.1.3	Cross-sectional Weights (Variable FWTCW01C).....	104
11.2	Weighting Method.....	105
11.2.1	The National Longitudinal Survey of Children and Youth Design Weight	105
11.2.2	First Adjustment: Nonresponse Adjustment	105
11.2.3	Second Adjustment: Post-stratification	106
11.2.4	How the Weighting Method at Cycle 6 Differs from the Method at Previous Cycles	106
11.3	Applying the Weighting Method	107
11.3.1	Longitudinal Weighting.....	107
11.3.2	Funnel weighting.....	108
11.3.3	Cross-sectional Weighting	108
12.0	Data Quality, Response Rates and Coverage.....	109
12.1	Sampling Error	109
12.2	Non-sampling Errors	109
12.3	Total Nonresponse and Nonresponse Bias	109
12.3.1	Response Definitions	110
12.3.2	Cross-sectional Response Rates.....	112
12.3.3	Longitudinal Attrition Rates Cycle 1 Cohort.....	114
12.3.4	Longitudinal Response Rate for Children Selected in Cycle 4	115
12.3.5	Longitudinal Response Rate for Children Selected in Cycle 5	116
12.4	Partial Nonresponse.....	116
12.4.1	Child Component	117
12.4.2	Person Most Knowledgeable Component	117
12.4.3	Spouse Component	118
12.4.4	Youth Component	118
12.4.5	PPVT-R	118
12.4.6	Number Knowledge	119
12.4.7	Who Am I?	119
12.4.8	Mathematics Tests.....	119
12.4.9	Problem Solving.....	119
12.4.10	Literacy Assessment.....	120
12.4.11	Numeracy Assessment	120
12.4.12	Self-complete Components.....	120
12.5	Cycle Nonresponse.....	121
12.6	Response Errors Impact for Rare Characteristics	121
12.7	Response Errors Related to Deviant Behaviour or Sensitive Questions.....	122
12.8	Response Errors Due to Approximations	122
12.9	Response Errors Due to Memory Errors.....	122
12.10	Response Errors Due to Collection by Proxy	122
12.11	Response Patterns with Indefinite Response Categories	123
12.12	Language of Interview.....	123

12.13	Conflicting Information	123
12.14	Data Quality for Body Mass Index	124
	12.14.1 Body Mass Index	124
	12.14.2 Body Mass Index – Centers for Disease Control.....	124
	12.14.3 Body Mass Index – Tim Cole, et al.	125
	12.14.4 Body Mass Index – Data Quality.....	125
12.15	Conditioning Bias	126
12.16	Person Most Knowledgeable	126
12.17	Coverage.....	126
	12.17.1 Coverage for the National Longitudinal Survey of Children and Youth	127
	12.17.2 Coverage and Frame Issues.....	127
12.18	Conclusion	132
13.0	Variance Estimation	133
13.1	Terms Related to Sampling Error and Variance	134
13.2	Coefficient of Variation	135
13.3	Importance of Reporting the Sampling Variance	137
13.4	Sampling Variance Calculation	137
	13.4.1 Spreadsheet with Approximate Sampling Variances for Proportions	139
	13.4.2 SAS and SPSS Macros to Calculate the Sampling Variance Using the Bootstrap Weights: Bootvar	139
	13.4.3 Other Computer-based Tools	140
	13.4.4 Taylor Linearization and Other Techniques.....	140
14.0	Direct Assessment	141
14.1	The Peabody Picture Vocabulary Test - Revised	141
	14.1.1 Psychometric Properties of Scores.....	142
14.2	Number Knowledge Assessment.....	143
14.3	Who Am I?	145
14.4	Mathematics Computation Exercise	148
14.5	Problem-Solving Exercise (16- and 17-Year-Olds).....	150
15.0	Survey of Northern Children, Cycle 6	155
15.1	Introduction	155
15.2	Background.....	155
15.3	Establishing the Survey of Northern Children.....	156
16.0	Analytic Issues 161	
16.1	How a Complex Sample Design Affects Analysis.....	161
16.2	Unit of Analysis	162
16.3	Longitudinal versus Cross-sectional Analysis.....	162
16.4	Simple Weighted Estimates (Totals, Averages and Proportions).....	163
	16.4.1 Estimating a Population Total	163
	16.4.2 Estimating a Population Average.....	163
	16.4.3 Estimating a Population Proportion.....	164
	16.4.4 Estimating for Domains of the Population.....	164
16.5	Normalized Weights	165
16.6	Repeated Measures.....	168
16.7	Pooling Data.....	169
16.8	Nonresponse.....	170
16.9	Other Sources of Non-sampling Errors.....	172
16.10	Computing the Variance with Certain Software Applications.....	172
16.11	Coefficients of Variation for Proportions	173
16.12	Standard Deviation Versus Standard Error	173
16.13	Understanding the Difference Between "Not stated", "Don't know", "Refusal" and "Valid skip" 173	

17.0	Guidelines for Tabulation, Analysis and Release	175
17.1	Rounding Guidelines.....	175
17.2	Sample Weighting Guidelines for Tabulation.....	175
17.3	Guidelines for Statistical Modelling.....	176
17.4	Release Guidelines.....	176
	Appendix I.....	177
	Appendix II.....	183
	Appendix III.....	185
	Appendix IV.....	195

1.0 Introduction

The National Longitudinal Survey of Children and Youth (NLSCY) Cycle 6 was conducted from September 2004 to June 2005 by Statistics Canada in partnership with Social Development Canada (now called Human Resources and Social Development Canada).

This manual has been produced to facilitate the manipulation of the microdata files of the survey results and to document data quality and other analytical issues regarding the NLSCY.

Any questions about the data set or its use should be directed to:

Statistics Canada

Client Services
Special Surveys Division
Telephone: (613) 951-3321 or toll free 1 800 461-9050
Fax: (613) 951-4527
E-mail: ssd@statcan.ca

Special Surveys Division, Statistics Canada
150 Tunney's Pasture Driveway
Ottawa, Ontario K1A 0T6

Human Resources and Social Development Canada

Pierre Turcotte
A/Assistant Director
Strategic Analysis, Partnership & Dissemination Directorate
Human Resources and Social Development Canada
Telephone: (613) 957-7483
E-mail: pierre.c.turcotte@sdsc.gc.ca

2.0 Background

The National Longitudinal Survey of Children and Youth (NLSCY) is a long-term study of Canadian children that follows their development and well-being from birth to early adulthood. The NLSCY began in 1994 and is jointly conducted by Statistics Canada and sponsored by Human Resources and Social Development Canada.

The study is designed to collect information about factors influencing a child's social, emotional and behavioural development and to monitor the impact of these factors on the child's development over time.

The survey covers a comprehensive range of topics including the health of children, information on their physical development, learning and behaviour as well as data on their social environment (family, friends, schools and communities).

Information from the NLSCY is being used by a variety of people at all levels of government, in universities, and policy-making organizations.

Survey Population

In Cycle 6, a representative sample of Canadian children aged 0 to 5 years old from each of the provinces was surveyed for longitudinal and cross-sectional purposes. The cohort of children aged 10 to 21 years old was surveyed for longitudinal purposes.

The Northern cohort of children consists of a census of 5-year-old children who are attending kindergarten in the Yukon and Nunavut. Please refer to Chapter 15.0 of this guide for a detailed explanation of the background of the Survey of Northern Children and a description of the children surveyed in Cycle 6.

Target population

The NLSCY objectives are to produce longitudinal and cross-sectional estimates. Therefore, several populations are targeted in the Cycle 6 sample. Please see Chapter 5.0, for more detailed information about the sample.

- Cross-sectionally, the Cycle 6 sample represents all children who were 0 to 5 years old on December 31, 2004.
- Longitudinally, we have three cohorts, i.e., more than one cycle of data:
 - 1) The first cohort represents all children who were 0 to 11 years old as of December 31st, 1994 and who were living in one of the 10 provinces during collection for Cycle 1 in 1994/1995. These children are now 10 to 21 years old in Cycle 6.
 - 2) The second cohort represents all children who were 0 to 1 years old as of December 31st, 2000 and who were living in one of the 10 provinces during collection for Cycle 4 in 2000/2001. These children are now 4 to 5 years old in Cycle 6.
 - 3) The third cohort represents all children who were 0 to 1 years old as of December 31st, 2002 and who were living in one of the 10 provinces during collection for Cycle 5 in 2002/2003. These children are now 2 to 3 years old in Cycle 6.
- Cross-sectionally, the census of children living in the territories is restricted in Cycle 6 to 5-year-old children who attend kindergarten in the Yukon and Nunavut.

Collection Cycles

Data collection occurs at two-year intervals.

Cycle	Collection Start	Collection End
1	December 1994	April 1995
2	December 1996	April 1997
3	October 1998	June 1999
4	September 2000	May 2001
5	September 2002	June 2003
6	September 2004	June 2005

Data Release Strategy

The Cycle 6 data are being released in four files: Longitudinal cohort – child (10 to 17 years old), Longitudinal cohort – youth (18 to 21 years old), Early Child Development cohort (0 to 5 years old), and Northern cohort (5 years old).

- Longitudinal cohort – child: These are the children from the original longitudinal cohort. These children were aged 0 to 7 years old in the first cycle of the NLSCY. For Cycle 6, they are 10 to 17 years old.
- Longitudinal cohort – youth: These are children from the original longitudinal cohort who were 6 to 11 years old in Cycle 1. For Cycle 6, they are 16 to 21 years old. This file contains data collected for youth and their households as well as variables that were brought forward from previous cycles that do not change over time, e.g. birth weight. See Chapter 8.0 for more information about the content of the various questionnaires.
- The Early Childhood Development (ECD) file contains data collected from the person most knowledgeable about the child, for children aged 0 to 5 years old in Cycle 6.
- The Northern cohort includes 5-year-olds from the Yukon and Nunavut. Initially, children from all three territories were to be included in the survey but operational constraints prevented the Northwest Territories from participating in Cycle 6. See Chapter 15.0 for more detail about the Survey of Northern Children.

3.0 Objectives

The objectives of the National Longitudinal Survey of Children and Youth (NLSCY) are:

- to determine the prevalence of various risk and protective factors for children and youth;
- to understand how these factors, as well as life events, influence children's development;
- to make this information available for developing policies and programs that will help children and youth;
- to collect information on a wide variety of topics – biological, social, economic;
- to collect information about the environment in which the child is growing up – family, peers, school, community.

Information comes from different sources (parent and child) and from direct measures (Peabody Picture Vocabulary Test – Revised (PPVT-R), math tests, etc.).

4.0 Main Changes at Cycle 6 since Cycle 5

This chapter outlines the main changes to the survey since Cycle 5. A more detailed explanation will be found in the individual chapters as referenced below.

4.1 Content Changes

Each cycle, there are changes made to the content of the National Longitudinal Survey of Children and Youth (NLSCY). Any new variable or any variable that changed, e.g., wording, response categories, and eligible population, will have an “f” as the fifth character of the variable name. Variable name conventions are described in Chapter 7.0. The survey content is described in detail in Chapter 8.0

The following is a list of the main changes to the content of the survey for Cycle 6:

- The self-complete booklet for 18- to 19-year-olds, Booklet 24, has been dropped. The questions from the booklet are now included in the computer-assisted interviewing (CAI) questionnaire. The concordance table in Appendix IV shows the Cycle 5 and Cycle 6 variable names.
- The oldest respondents in the longitudinal cohort are now 20 and 21 years old. Many of the questions for this age group are the same as those asked in Cycle 5. There are also new questions for these youth. There are new questions about financing post-secondary education (FEDYf165, FEDYf166, FEDYf167, FEDYf168), job quality (FLYYf14A, FLYYf14B, FLYYf14C), dependent children (FDMCfD22), and about voting (FACYfQ16).
- Two new direct assessments were added in Cycle 6. A literacy assessment for 18- to 19-year-olds (FLIYfS01) and a numeracy assessment for 20- to 21-year-olds (FNUYfS01). Assessments are described in Chapter 14.0.
- The Education component (Kindergarten Teacher’s questionnaire) has been dropped.
- Questions were added that allow a respondent to self-identify as an Aboriginal person (FSDCfQ3A, FSDCfQ3BA, FSDCfQ3BB, FSDCfQ3BC).

4.2 Change to Timing of Collection

The allocation of the sample to various waves of collection affects the child’s age in months at the time of interview. This can impact the scores in tests and other direct measures. Some children were interviewed at different times in Cycle 6 than in Cycle 5. For example, 10-year-olds were interviewed in Wave 1 in Cycle 5 but in Wave 2 in Cycle 6. The timing of the Cycle 6 collection waves are described in Chapter 6.0. The variables FMMCbQ1A and FMMCdQ1B give the child’s age at time of collection.

4.3 Methodology Changes

4.3.1 Sample

Cycle 5 had no top-up of 2- to 5-year-olds. At Cycle 6, there was a top-up sample of children aged 2 to 5 years old in provinces other than Quebec and Ontario. The respondents aged 2 to 5 years old from the top-up sample have a cross-sectional weight, but no longitudinal weight.

Between Cycle 5 and Cycle 6 there has been a shift in the sample distribution by age for the Early Childhood Development group as the large cohort introduced at Cycle 3 was dropped at Cycle 6. Note that these children will be re-surveyed as 8- and 9-year-olds at Cycle 7. For more detail, see Chapter 5.0

4.3.2 Weights

The nonresponse adjustment differs from past cycles. For more detail see Chapter 11.0, Weighting and Treatment of Nonresponse.

4.3.3 Variance

A portion of the NLSCY sample was drawn after the Labour Force Survey redesign (re-forming of strata and clusters) that was phased in, beginning with the November 2004 rotation group. This affects how the bootstrap weights were created, yet should be inconsequential for data users.

The methods used to derive the bootstrap weights were modified to correspond with the weighting strategy at Cycle 6. Information about variance estimation can be found in Chapter 13.0.

4.4 Processing Changes

At Cycle 6, the Generalized Processing System was implemented from the Clean-Up stage through to the production of the master files. In Cycle 5, this system was used only after the Clean-up and relationship edits had been completed. As well, for reasons of data quality and efficiency, the longitudinal edits were done at the derived variable stage.

There were 26 youth aged 16 and 17 living independently in Cycle 6 (FDMCfD03=82). All variables in the longitudinal file were set to 'valid skip' for these respondents except for some variables in the Demographic Section, and certain variables in the Sociodemographic and Custody Sections, where some values were carried forward from previous cycles.

Updated classification systems were used to code the Industry and Occupation data. The National Occupational Classification – Statistics (NOC-S) 2001 was used to code the data received for the respondent's occupation(s); the industry questions were coded using the 2002 North American Industry Classification System (NAICS 2002).

4.5 Changes to Released Files

For Cycle 6, the youth variables are released on a separate file. There are now two files for the longitudinal cohort, one for the child (10 to 17 years old) and one for the youth (16 to 21 years old). The youth questions asked of the 16- and 17-year-olds will be on the youth file, while the other data for these respondents will be found on the Child File. The 16- and 17-year-olds who no longer live with a parent or guardian were asked the youth questions only. All data for the 18- to 21-year-olds are included on the Youth File.

Static variables from previous cycles, e.g., country of birth and birth weight, have been added to the Cycle 6 files.

The variables PPERSRUK and SPERSRUK were added to the master file to identify the person most knowledgeable (PMK) and Spouse. These variables will become useful for future cycles as changes to the PMK and/or spouse between cycles will be easily identifiable.

5.0 Survey Methodology - Sample

This chapter provides details on the sample for the National Longitudinal Survey of Children and Youth (NLSCY). Section 5.1 gives an overview of the sample design; Section 5.2 describes the Labour Force Survey (LFS) since all NLSCY children present at Cycle 6 were sampled from the LFS; Section 5.3 provides details on how sampling was performed at each cycle, along with response rates; Section 5.4 describes how the various NLSCY samples at each cycle can be used to make inferences about specific longitudinal and cross-sectional reference populations (ones for which the NLSCY produces survey weights).

5.1 Overview of the Sample Design, Cycles 1 to 6

The NLSCY is a probability survey designed to collect detailed information every two years about the factors influencing a child's cognitive, emotional and physical development and to monitor the impact of these factors over time.

Collection for the first cycle of the NLSCY began in 1994 with one large cohort of 0- to 11-year-olds who lived in one of the ten provinces. This sample is referred to as the original cohort. This cohort is purely longitudinal; it is not topped up to reflect changes that occur in the population over time due to immigration.


At Cycle 2, the scope of the NLSCY expanded to emphasize early childhood development (ECD). The purpose of the ECD component is to collect information on young children and produce some indicators, e.g., motor and social development, emotional problems, hyperactivity, physical aggression, prosocial behaviour, and language skills. As a result, starting at Cycle 2, a new cohort of 0- to 1-year-olds is selected at every cycle and followed for three cycles, i.e., until the children are 4 to 5 years old¹. These samples are designed for both longitudinal and cross-sectional purposes and are referred to as the ECD cohorts.

Typically, children in the NLSCY are selected from households sampled by Statistics Canada's Labour Force Survey (LFS). Exceptions are 1-year-olds in Cycle 3 and some 5-year-olds in Cycles 3 and 4 who were drawn from Birth Registry data. None of these Birth Registry children is present in the Cycle 6 sample since they would be older than 5 at Cycle 6.

At Cycle 6, the NLSCY sample consists of children aged 0 to 5 (ECD children) and 10 to 21 (original cohort). A child's effective age at Cycle 6 is with respect to December 31st, 2004. Thus, 0-year-olds are born in 2004 and 1-year-olds are born in 2003.

The diagram below illustrates the NLSCY sample. The years indicate when collection occurred. The larger arrows represent the original cohort, the smaller arrows represent the ECD cohorts.

1 This is to change at Cycle 7 which will include 6- to 9-year-olds.

Age of children at each cycle, original cohort versus ECD cohorts

For details on how sampling was performed at each cycle, see Section 5.3.

5.2 The Labour Force Survey

The LFS is a monthly survey that collects labour market data from a national sample of about 54,000 dwellings. Following each decennial population census, the LFS is redesigned to reflect changes in the Canadian population and to respond to changes in the information needs of the LFS. Over the life of the NLSCY, there have been two LFS redesigns: one in 1994 and one in 2004/2005 (fully implemented in April 2005). As a result, the sample for the original cohort contains a mixture of pre-1994 LFS design and the 1994 redesign. The vast majority of ECD children surveyed at Cycle 6 were drawn from the 1994 LFS design, with a few children coming from the 2004/2005 LFS redesign. The following sections provide details on the 1994 LFS design.

5.2.1 Target Population

The LFS sample is representative of the civilian, non-institutionalised population 15 years of age or older in Canada's ten provinces. Specifically excluded from the survey's coverage are residents of the Yukon, Nunavut and Northwest Territories, persons living on Indian Reserves, full-time members of the Canadian Armed Forces and inmates of institutions. These groups together represent an exclusion of approximately 2% of the population aged 15 or over.

5.2.2 Stratification

The LFS sample is based upon a stratified, multistage design employing probability sampling at all stages of the design. The design principles of the LFS are the same for each province.

Primary strata

Provinces are divided into economic regions and employment insurance economic regions. Economic Regions (ERs) are geographic areas of more or less homogeneous economic structure formed on the basis of federal provincial agreements. They are relatively stable over time. Employment insurance economic regions (EIERs) are also geographic areas, and are roughly the same in size and number as ERs, but they do not share the same definitions. Labour force estimates are produced for the EIERs for the use of Human Resources and Social Development Canada.

The intersections of the two types of regions form the primary strata for the LFS. Then, sub-stratification takes place within these primary strata (see section 5.2.3). Note that a third set of regions, Census Metropolitan Areas (CMAs), is also respected by stratification in the current LFS design, since each CMA is also an EIER.

Types of areas

The primary strata (ER/EIER intersections) are classified into three types of areas: rural, urban, and remote areas. Urban and rural areas are loosely based on the Census definitions of urban and rural, with some exceptions. Urban areas include the largest CMAs down to the smallest villages categorised by the 1991 Census as urban (1,000 people or more), while rural areas are made up of areas not designated as urban or remote.

All urban areas are further classified into two types: those using an apartment list frame and an area frame, and those using only an area frame.

Approximately 1% of the LFS population is found in remote areas of provinces that are less accessible to LFS interviewers than other areas. For administrative purposes, this portion of the population is sampled separately through the remote area frame. Places with fewer than 10 households or 25 persons and Census Enumeration Areas (EAs) with fewer than 25 households are omitted from the design.

Secondary strata

In urban areas with sufficiently large numbers of apartment buildings, the strata are subdivided into apartment frames and area frames. The apartment list frame is a register maintained for the 18 largest cities across Canada. The purpose of this is to ensure better representation of apartment dwellers in the sample as well as to minimize the effect of growth in clusters, due to construction of new apartment buildings. In the major cities, the apartment strata are further stratified into low income strata and regular strata.

Where it is possible and/or necessary, the urban area frame is further stratified into regular strata, high-income strata, and low population density strata. Most urban areas fall into the regular urban strata, which, in fact, cover the majority of Canada's population. High-income strata are found in major urban areas, while low-density urban strata consist of small towns that are geographically scattered.

In rural areas, the population density can vary greatly from relatively high population density areas to low population density areas, resulting in the formation of strata that reflect these variations. The different stratification strategies for rural areas were based not only on concentration of population, but also on cost-efficiency and interviewer constraints.

The remote area frame is stratified only by province.

5.2.3 Cluster Delineation and Selection

Households in final strata are not selected directly. Instead, each stratum is divided into clusters, and then a sample of clusters is selected within the stratum. Dwellings are then sampled from selected clusters. Different methods are used to define the clusters, depending on the type of stratum.

Within each urban stratum in the urban area frame, a number of geographically contiguous groups of dwellings, or clusters, are formed based upon Census counts. These clusters are generally a set of one or more city blocks or block faces. The selection of a sample of clusters (always six or a multiple of six clusters) from each of these secondary strata represents the first stage of sampling in most urban areas. In some other urban areas, Census EAs are used as clusters. In the low-density urban strata, a three-stage design is followed. Under this design, two towns within a stratum are sampled and then six or 24 clusters within each town are sampled.

For urban apartment strata, instead of defining clusters, the apartment building is the primary sampling unit. Apartment buildings are sampled from the list frame with probability proportional to the number of units in each building.

Other procedures are applied in rural and remote areas. Within each rural stratum, six EAs or two or three groups of EAs are sampled as clusters, whereas remote settlements within each province are sampled proportional to the number of dwellings in the settlement.

5.2.4 Dwelling Selection

In all three types of areas (urban, rural and remote areas) interviewers in the field first visit selected clusters and a listing of all private dwellings in the cluster is prepared. From the listing, a sample of dwellings is then selected. The sample yield depends on the type of stratum. For example, in the urban area frame, sample yields in regular strata within major urban areas are either six or eight dwellings, depending on the size of the city. In the urban apartment frame, each cluster yields five dwellings while in the rural areas and urban EAs, each cluster yields ten dwellings. In all clusters, dwellings are sampled systematically. This represents the final stage of sampling.

5.2.5 Sample Rotation

The LFS employs a panel design whereby the entire monthly sample of dwellings consists of six panels or rotation groups, of approximately equal size. Each of these panels can be considered to be representative of the entire LFS population. Dwellings are in the LFS for six consecutive months. Each month a new panel of dwellings selected from the same or similar clusters replaces the sample dwellings in one of the rotation groups.

This rotation pattern has the statistical advantage of providing a common sample base for month-to-month comparisons of LFS characteristics. It also ensures that the sample of dwellings constantly reflects changes in the current housing stock and helps to minimize the respondent burden and nonresponse that could result if households were to remain in the sample longer than six months. Surveys that use the LFS frame or sample can take advantage of the rotation group feature to use larger or smaller sample sizes than that of the LFS.

5.2.6 Household Members Eligible for the Labour Force Survey

The first month a dwelling is in the LFS, a roster containing information on the household composition is completed. Demographic information including the name, sex, date of birth and education level is obtained for all persons for whom the selected dwelling is the usual place of residence. Labour force information is obtained for all civilian household members 15 years of age or older.

When the dwelling is contacted in subsequent months the roster is updated to reflect changes in household membership from the previous month.

5.3 Details of the National Longitudinal Survey of Children and Youth Sample

At Cycle 6, the NLSCY sample consists of: a new ECD cohort of 0- to 1-year-olds, a top-up sample of new 2- to 5-year-olds, and returning 2- to 5-year-olds and 10- to 21-year-olds from previous cycles. All children belong to LFS households. Cycle 6 was the first time that the LFS was used to top-up existing NLSCY sample.

In addition to explaining how sampling was performed, this section describes which children were surveyed. Some children are sampled but not surveyed because at the previous cycle they were found to be out-of-scope cross-sectionally, e.g., deceased or left the country, or due to too many cycles of nonresponse.

5.3.1 Original Cohort, Cycles 1 to 6

The following describes the composition of the original cohort at each cycle (see also Diagram 1). The original cohort contains a maximum of two children per household.

Cycle 1

The sample of children selected at Cycle 1 was designed to produce reliable -- but not equally reliable -- provincial estimates for children between the ages of 0 and 11, by two-year age groupings, i.e., 0 to 1, 2 to 3, 4 to 5, 6 to 7, 8 to 9, 10 to

11. A maximum of four children per household was selected. Households were sampled from the following sources:

- the old LFS, i.e., prior to 1994,
- the 1994 redesigned LFS,
- the National Population Health Survey (NPHS) (conducted by Statistics Canada).

At the end of Cycle 1, there were 22,831 respondent children in the NLSCY. The household response rate was 86.7% (see Table 1).

For more details, see the User's Guide for Cycle 1.

Cycle 2

At Cycle 2 some children were dropped from the sample for budgeting reasons; all NLSCY households belonging to the NPHS sample were dropped, and to reduce the burden on households, the maximum number of children selected per household was cut from four to two. This resulted in a sample of 16,903 children at the beginning of Cycle 2. The household response rate for collection was 91.7% (see Table 1). The cumulative, longitudinal response rate for households in the original cohort was 79.8% (see Table 2).

For more details, see the User's Guide for Cycle 2.

Cycle 3

At Cycle 3, 185 children were excluded from the sample because at the end of Cycle 2 they were either out-of-scope cross-sectionally (71) or hard refusals (114). Cross-sectionally out-of-scope children include those who died, whose age was not in-scope, who had permanently left the country, or moved to an Indian reserve. Thus, of the 16,903 children sampled for the original cohort, collection was performed on 16,718. The household response rate for collection was 89.6% (see Table 1). The cumulative, longitudinal response rate for households in the original cohort was 77.5% (see Table 2).

For more details, see the User's Guide for Cycle 3.

Cycle 4

At Cycle 4, in order to make collection more efficient, it was decided that households with two or more consecutive cycles of nonresponse would be dropped from collection (along with households with one cycle of nonresponse followed by the status 'temporarily moved'). At Cycle 4, consequently, 1,086 children were dropped from collection because at the end of Cycle 3 they were either out-of-scope cross-sectionally (106) or had two or more cycles of nonresponse (980). A total of 15,632 children were surveyed. The household response rate for collection was 84.8% (see Table 1). The cumulative, longitudinal response rate for households in the original cohort was 69.4% (see Table 2).

For more details, see the User's Guide for Cycle 4.

Cycle 5

At Cycle 5, it was decided that youth who were 18 or 19 years old would be dropped from collection only after three consecutive cycles of nonresponse (versus two for younger children). The reason for this is that at age 18, the youth becomes the sole respondent, while before age 18 the primary respondent is the person most knowledgeable (PMK), who is typically the mother.

At Cycle 5, there were 469 children who were dropped from collection because at the end of Cycle 4 they were either cross-sectionally out-of-scope (34) or had too many consecutive cycles of nonresponse (435). A total of 15,163 children were surveyed at Cycle 5. The household response rate for collection was 83.2% (see Table 1). The cumulative, longitudinal response rate for households in the original cohort was 66.6% (see Table 2).

For more details, see the User's Guide for Cycle 5.

Cycle 6

At Cycle 6, there were 1,506 children who were dropped from collection because at the end of Cycle 5 they were either cross-sectionally out-of-scope (26) or had too many consecutive cycles of nonresponse (1,480). A total of 13,657 children were surveyed at Cycle 6. The household response rate for collection was 83.5% (see Table 1). The cumulative, longitudinal response rate for households in the original cohort was 62.1% (see Table 2).

5.3.2 Early Childhood Development Cohorts Present at Cycle 6

At Cycle 6, the only ECD children present were those introduced as babies in Cycles 4, 5 and 6, and the top-up of new 2- to 5-year-olds selected from the LFS. When the first ECD cohort of babies was selected at Cycle 2, the rule was a maximum of one child per household, except for twins, in which case both were sampled². At Cycle 5, the rule changed to one child per household without exception. Returning twins, however, continued to be surveyed until Cycle 6³.

For the ECD samples, only respondents from the previous cycle are surveyed at the subsequent cycle, unlike the original cohort, where two or three consecutive cycles of nonresponse are required before a child is dropped from collection.

For information on the ECD cohorts introduced in Cycles 2 and 3, please consult the User's Guides for these cycles. For an illustration of the ECD cohorts present at Cycle 6, see Diagram 1.

Cycle 4 ECD Cohort

At Cycle 4, a sample of 0- to 1-year-olds was selected from the LFS along with a top-up sample of 5-year-olds sampled from Birth Registry data. The total sample size was 9,439 households. At the end of Cycle 4 collection, there were 6,960 responding children. The response rate was 75.8% (see Table 1).

² The ECD cohort sampled in Cycle 2 included 0- to 1-year-olds who were younger siblings of children belonging to the original cohort. This was the only cycle in which siblings from the original cohort were selected. There are no Cycle 2 ECD children present in the Cycle 6 sample.

³ For the Cycle 7 sample, it was decided that one of the returning twins would be dropped for returning ECD cohorts. The original cohort continues to have a maximum of two children per household.

At Cycle 5, the 5-year-olds in Cycle 4 were dropped (4,405 households). At the end of Cycle 4, there were 125 households that were cross-sectionally out-of-scope and 1,121 were nonrespondents. Consequently, 3,788 of the returning households with 2- to 3-year-olds from the Cycle 4 ECD cohort were surveyed at Cycle 5. The response rate was 86.9% (see Table 1). The cumulative, longitudinal response rate was 67.1% (see Table 2).

At the end of Cycle 5, of the returning 4- to 5-year-olds, 11 children were cross-sectionally out-of-scope and 507 were nonrespondents. These children were dropped from collection; 3,323 were surveyed at Cycle 6. The response rate was 89.6% (see Table 1). The cumulative longitudinal response rate was 60.1% (see Table 2).

Cycle 5 ECD Cohort

At Cycle 5, a sample of 0- to 1-year-olds was selected from the LFS. The total sample size was 4,492 children and households. At the end of Cycle 5 collection, there were 3,252 responding children. The response rate was 74.0% (see Table 1).

At the end of Cycle 5, 98 children were cross-sectionally out-of-scope and 1,142 were nonrespondents. Consequently, only 3,252 of the 2- to 3-year-olds returning from the Cycle 5 ECD cohort were surveyed at Cycle 6. The response rate was 88.6% (see Table 1). The cumulative longitudinal response rate was 65.6% (see Table 2).

Cycle 6 ECD Cohort

At Cycle 6, a sample of 0- to 1-year-olds along with a top-up sample of new 2- to 5-year-olds was selected from the LFS. The total sample size was 5,795 children and households. At the end of Cycle 6 collection, there were 4,684 responding children. The response rate was 81.3% (see Table 1).

Table 1: Household-level response at collection (number of children in parentheses)

Sampling Cohort	Survey Cycle	Age Range (years)	Collection													
			Sample Size		Sample Reduction		Out-of-scope Dropped from Previous Cycle		Nonrespondents Dropped from Previous Cycle		In-scope Respondents		In-scope Rate (%)	Response Rate (%)		
Original cohort	1	0-11	43,751	(n/a)	n/a		n/a		n/a		15,502	(n/a)	13,439	(22,831)	35.4	86.7
	2	2-13	11,188	(16,903)	25,588	(n/a)	5,345	(n/a)	1,677	n/a	11,140	(16,816)	10,216	(15,391)	99.6	91.7
	3	4-15	11,032	(16,718)	0	(0)	38	(71)	73	(114)	10,937	(16,563)	9,801	(14,777)	99.1	89.6
	4	6-17	10,449	(15,632)	0	(0)	65	(106)	618	(980)	10,417	(15,586)	8,833	(13,174)	99.7	84.8
	5	8-19	10,355	(15,163)	0	(0)	25	(34)	285	(435)	10,320	(15,113)	8,582	(12,280)	99.7	83.2
	6	10-21	9,881	(13,657)	0	(0)	17	(26)	861	(1,480)	9,816	(13,572)	8,201	(11,178)	99.3	83.5
2	2	0-1	5,528	(n/a)	n/a		n/a		n/a		4,865	(n/a)	4,496	(4,634)	88.0	92.4
	3	2-3	3,935	(3,989)	551	(n/a)	598	(n/a)	444	(n/a)	3,893	(3,947)	3,592	(3,640)	98.9	92.3
	4	4-5	3,577	(3,610)	464	(484)	25	(25)	33	(34)	3,552	(3,585)	3,023	(3,052)	99.3	85.1
3	3	0-1 & 5	16,812	(n/a)	n/a		n/a		n/a		15,929	(n/a)	13,256	(13,546)	94.7	83.2
	4	2-3	7,941	(8,118)	6,935	(n/a)	516	(n/a)	1,420	(n/a)	7,896	(8,070)	6,956	(7,111)	99.4	88.1
	5	4-5	6,960	(7,115)	0	(0)	41	(44)	940	(959)	6,919	(7,073)	6,208	(6,340)	99.4	89.7
4	4	0-1 & 5	9,439	(n/a)	n/a		n/a		n/a		9,115	(n/a)	6,907	(6,960)	96.6	75.8
	5	2-3	3,788	(3,841)	4,405	(n/a)	125	(n/a)	1,121	(n/a)	3,776	(3,829)	3,281	(3,324)	99.7	86.9
	6	4-5	3,280	(3,323)	0	(0)	11	(11)	497	(507)	3,270	(3,313)	2,931	(2,964)	99.7	89.6
5	5	0-1	4,492	(4,492)	n/a		n/a		n/a		4,394	(4,394)	3,252	(3,252)	97.8	74.0
	6	2-3	3,252	(3,252)	0	(0)	98	(98)	1,142	(1,142)	3,233	(3,233)	2,866	(2,866)	99.4	88.6
6	6	0-5	5,795	(5,795)	n/a		n/a		n/a		5,763	(5,763)	4,684	(4,684)	99.4	81.3

Table 2: Household-level, longitudinal response (number of children in parentheses)

Sampling Cohort	Survey Cycle	Age Range (years)	Longitudinal							
			Sample Size		In-scope		Respondents		In-scope Rate (%)	Response Rate (%)
Original cohort	1	0-11	18,163	(n/a)	12,818	(n/a)	11,141	(16,903)	70.6	86.9
	2	2-13	18,210	(n/a)	12,863	(n/a)	10,262	(15,472)	70.6	79.8
	3	4-15	18,165	(n/a)	12,818	(n/a)	9,933	(14,997)	70.6	77.5
	4	6-17	18,265	(n/a)	12,914	(n/a)	8,962	(13,361)	70.7	69.4
	5	8-19	18,481	(n/a)	13,134	(n/a)	8,744	(12,535)	71.1	66.6
	6	10-21	18,885	(n/a)	13,532	(n/a)	8,409	(11,483)	71.7	62.1
2	2	0-1	4,977	(n/a)	4,379	(n/a)	4,100	(4,154)	88.0	93.6
	3	2-3	4,977	(n/a)	4,377	(n/a)	3,632	(3,680)	87.9	83.0
	4	4-5	4,513	(n/a)	3,784	(n/a)	2,942	(2,971)	83.8	77.7
3	3	0-1	9,877	(n/a)	9,361	(n/a)	7,949	(8,126)	94.8	84.9
	4	2-3	9,877	(n/a)	9,198	(n/a)	6,838	(6,992)	93.1	74.3
	5	4-5	9,877	(n/a)	9,198	(n/a)	6,141	(6,273)	93.1	66.8
4	4	0-1	5,034	(n/a)	4,909	(n/a)	3,788	(3,841)	97.5	77.2
	5	2-3	5,034	(n/a)	4,907	(n/a)	3,291	(3,334)	97.5	67.1
	6	4-5	5,034	(n/a)	4,911	(n/a)	2,952	(2,985)	97.6	60.1
5	5	0-1	4,492	(4,492)	4,394	(4,394)	3,252	(3,252)	97.8	74.0
	6	2-3	4,492	(4,492)	4,392	(4,392)	2,883	(2,883)	97.8	65.6
6	6	0-1	4,356	(4,356)	4,343	(4,343)	3,521	(3,521)	99.7	81.1

Note: Table 2 contains only longitudinal children i.e., children who are followed through time.

5.4 Longitudinal and Cross-sectional Reference Populations for the National Longitudinal Survey of Children and Youth Weights

In a probability survey, individuals are randomly sampled from a well-defined population such that everyone in the population has a non-zero probability of selection, i.e., anyone may be selected; none is excluded, and this probability can be calculated. For example, if there are 100 children in the population and 10 are selected using simple random sampling, then every sampled child has a probability of selection of $10/100=1/10$.

The child's survey weight is the average number of children in the population that he or she represents. It is calculated as the inverse of the probability of selection (subsequent adjustments are usually made, for example to adjust for nonresponse, to match to demographic counts by age and sex, etc.). Thus, if a child's probability of selection is $1/10$, then the (initial) survey weight is 10, indicating that the child represents 10 children in the population. (For more details on NLSCY weights, see Chapter 11.0).

Survey weights refer to a particular population and they should be used at analysis when making inferences about that population. In a longitudinal survey such as the NLSCY – where children are followed over time – there may be different weights for different populations over time. This is because with time populations change due to deaths, immigration and emigration.

In a longitudinal survey, two types of populations are possible: the longitudinal population and various cross-sectional populations. The longitudinal population is the initial population when the sample was first drawn (there is only one longitudinal population); a cross-sectional population refers to some subsequent time period (there may be many). For example, the longitudinal population for the original cohort is all children whose effective age was 0 to 11 as of December 31st, 1994 and who lived in one of the ten provinces during collection of Cycle 1, in 1994/1995. A cross-sectional population at Cycle 6 could be children aged 0 to 11 as of December 31st, 1994 who lived in one of the ten provinces in 2004, i.e., this population includes immigrants since 1994/1995.

The original cohort can be used to make inferences about the former population, but not the latter since the original cohort has never been topped up for immigrants who arrived after 1994/1995. Cycle 4 was the last cycle for which cross-sectional weights were produced for the original cohort. By Cycle 5, it was felt that the absence of new immigrants was so great that the original cohort should not be used make inferences about the cross-sectional populations after Cycle 4. The ECD cohorts, however, are designed for both cross-sectional and longitudinal purposes, and consequently both cross-sectional and longitudinal weights are produced at each cycle.

The following subsection defines the various longitudinal and cross-sectional populations for which inferences can be made using NLSCY weights. It should be noted that these are not the only populations about which inferences can be made. For example, several ECD cohorts could be pooled to represent a population not listed below. (For more details on how to pool NLSCY samples, see Chapter 16.0). It should also be noted that all final weights are adjusted for nonresponse and to match demographic counts by age, sex and province. (For more details on how the NLSCY weights are calculated, see Chapter 11.0.)

5.4.1 Cohorts and Their Longitudinal Populations

The various original and ECD cohorts represent the following longitudinal populations. Note that for a cohort's first cycle, the longitudinal population is defined by weights that are labelled 'cross-sectional' for that first cycle.

The Original cohort, selected at Cycle 1

- Longitudinal population: children aged 0 to 11 as of December 31st, 1994 who were living in one of the 10 provinces at the time of Cycle 1 collection (1994/1995)

At Cycle 1, a longitudinal sample of children aged 0 to 11 was selected from the LFS. By Cycle 6, these children were aged 10 to 21 (as of December 31st, 2004). Sample reductions were made at Cycle 2. The children dropped between Cycles 1 and 2 can be regarded as Cycle 1 cross-sectional children.

Weights are produced at every cycle for this longitudinal population.

Early Childhood Development cohorts, selected at Cycles 2 to 6

- Longitudinal population of the ECD cohort selected at Cycle 2: children aged 0 to 1 as of December 31st, 1996 who were living in one of the 10 provinces at the time of Cycle 2 collection (1996/1997)

At Cycle 2, a longitudinal sample of children aged 0 to 1 was selected from the LFS. This cohort was followed for only three cycles until ages 4 to 5; they are not present in the Cycle 6 sample.

Weights were produced for this longitudinal population at Cycles 2, 3 and 4.

- Longitudinal population of the ECD cohort selected at Cycle 3: children aged 0 to 1 as of December 1st 1998 who were living in one of the 10 provinces at the time of Cycle 3 collection (1998/1999)

At Cycle 3, a longitudinal sample of children aged 0 to 1 was selected. The 0-year-olds were selected from the LFS while the 1-year-olds were drawn from Birth Registry data. This cohort was followed for only three cycles; none of these children is present at Cycle 6 sample. However, they will be returning at Cycle 7 (aged 8 to 9).

Weights were produced for this longitudinal population at Cycles 3, 4 and 5.

- Longitudinal population of the ECD cohort selected at Cycle 4: children aged 0 to 1 as of December 31st 2000 who were living in one of the 10 provinces at the time of Cycle 4 collection (2000/2001)

At Cycle 4, a longitudinal sample of children aged 0 to 1 was selected from the LFS. At Cycle 6 these children were 4 to 5 years old (as of December 2004). These children will be surveyed at Cycle 7 (aged 6 to 7).

Weights were produced for this longitudinal population at Cycles 4, 5 and 6.

- Longitudinal population of the ECD cohort selected at Cycle 5: children aged 0 to 1 as of December 31st, 2002 who were living in one of the 10 provinces at the time of Cycle 5 collection (2002/2003)

At Cycle 5, a longitudinal sample of children aged 0 to 1 was selected from the LFS. By Cycle 6, these children were 2 and 3 years old. These children will be surveyed at Cycle 7 (aged 4 to 5).

Weights were produced for this longitudinal population at Cycles 5 and 6.

- Longitudinal population of the ECD cohort selected at Cycle 6: children aged 0 to 1 as of December 31st, 2004 who were living in one of the 10 provinces at the time of Cycle 6 collection (2004/2005)

At Cycle 6, a longitudinal sample of children aged 0 to 1 was selected from the LFS. These children will be surveyed at Cycle 7.

Weights were produced for this population at Cycle 6.

5.4.2 Cohorts and Their Cross-sectional Populations

The Original cohort (at Cycle 6, aged 10 to 21 as of December 31st 2004)

Since top-ups for immigrants have never been performed for the original cohort, it is not recommended that the original cohort be used to represent cross-sectional populations after Cycle 4. Thus, the original cohort can be used to make inferences about the longitudinal population defined in 5.4.1, plus the following cross-sectional populations:

- Cycle 2 cross-sectional population: Children aged 2 to 13 as of December 31st, 1996 who were living in one of the 10 provinces at the time of Cycle 2 collection (1996/1997)

Cross-sectional weights were produced for this population at Cycle 2.

- Cycle 3 cross-sectional population: Children aged 4 to 15 as of December 31st, 1998 who were living in one of the 10 provinces at the time of Cycle 3 collection (1998/1999)

Cross-sectional weights were produced for this population at Cycle 3.

- Cycle 4 cross-sectional population: Children aged 6 to 17 as of December 31st, 2000 who were living in one of the 10 provinces at the time of Cycle 4 collection (2000/2001)

Cross-sectional weights were produced for this population at Cycle 4.

The Original cohort and ECD children

Children from both the original cohort and various ECD cohorts can be used to make inferences about the following populations:

- Cycle 2 cross-sectional population: Children aged 0 to 13 as of December 31st, 1996 who were living in one of the 10 provinces at the time of Cycle 2 collection (1996/1997)

This cross-sectional sample consists of:

- The ECD cohort of 0- to 1-year-olds selected at Cycle 2,
- Returning 2- to 13-year-olds belonging to the original cohort.

Cross-sectional weights were produced for this population at Cycle 2.

- Cycle 3 cross-sectional population: Children aged 0 to 15 as of December 31st, 1998 who were living in one of the 10 provinces at the time of Cycle 3 collection (1998/1999)

This cross-sectional sample consists of:

- The ECD cohort of 0- to 1-year-olds selected at Cycle 3,
- Returning 2- to 3-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 2,
- The top-up of 5-year-olds (selected from Birth Registry data),
- Returning 4- to 15-year-olds belonging to the original cohort.

Cross-sectional weights were produced for this population at Cycle 3.

- Cycle 4 cross-sectional population: Children aged 0 to 17 as of December 31st 2000 who were living in one of the 10 provinces at the time of Cycle 4 collection (2000/2001)

This cross-sectional sample consists of:

- The ECD cohort of 0- to 1-year-olds selected at Cycle 4,
- Returning 2- to 3-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 3,
- Returning 4- to 5-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 2,
- Returning 6- to 17-year-olds belonging to the original cohort.

Cross-sectional weights were produced for this population at Cycle 4.

After Cycle 4, inferences about cross-sectional populations should only be made using ECD children.

- Cycle 5 cross-sectional population: Children aged 0 to 5 as of December 31st 2002 who were living in one of the 10 provinces at the time of Cycle 5 collection (2002/2003)

This cross-sectional sample consists of:

- The ECD cohort of 0- to 1-year-olds selected at Cycle 5,
- Returning 2- to 3-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 4,
- Returning 4- to 5-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 3.

Cross-sectional weights were produced for this population at Cycle 5.

- Cycle 6 cross-sectional population: Children aged 0 to 5 as of December 31st 2004 who were living in one of the 10 provinces at the time of Cycle 5 collection (2004/2005)

This cross-sectional sample consists of:

- The ECD cohort of 0- to 1-year-olds selected at Cycle 6,
- Returning 2- to 3-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 5,
- Returning 4- to 5-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 4,
- A new top-up of 2- to 5-year-olds (selected from the LFS).

Cross-sectional weights were produced for this population at Cycle 6.

6.0 Data Collection

Data for Cycle 6 of the National Longitudinal Survey of Children and Youth (NLSCY) were collected between the fall of 2004 and the spring of 2005.

The collection was divided into three periods. The collection periods are described below. The date of the interview is recorded in variable (FMMHfQ06).

Collection period	Age groups
September – January, Wave 1	Ages 0 to 5 and 18 to 21, who do not have selected siblings aged 10 to 17
January – April, Wave 2	Ages 10 to 17 and 18 to 21 not collected in the first collection period.
March – June, Wave 3	Ages 0 to 5.

The survey combines computer-assisted interviewing (CAI) methods and the use of paper questionnaires. There are two types of computer-assisted interviewing applications used in the NLSCY: computer-assisted personal interviewing (CAPI) and computer-assisted telephone interviewing (CATI). For these types of interviews, the interviewer will read the questions on the computer and enter the respondent's answers in the computer. For CAPI the respondent and interviewer complete the questionnaire in-person, whereas for CATI the respondent completes the questionnaire by telephone. The use of CAI allows for complex flows and edits to be built into the questionnaire, which helps with both data quality and ensures that each respondent answers only the questions appropriate to their situation. The questions are identical whether the interview is conducted using CAPI or CATI. Depending on the composition of the household and the nature of the required components, the interview will be conducted partly or completely by telephone and/or field visit. This section provides a brief description of the "Collection tools" or the "Survey instruments", in other words the computer-assisted and paper questionnaire components used in the NLSCY collection. For a more detailed description of the content of the questionnaires, see Chapter 8.0.

6.1 Household Component

The first part of the interview was used to prepare a list of all household members, determine the relationships between them, gather tracing information and record basic demographic characteristics such as sex, date of birth, marital status and relationships between household members.

The person most knowledgeable (PMK) about the child was also identified in this component. The PMK provides the information for all selected children in the household and then gives information about himself/herself and his/her spouse/partner. In some cases it might have been appropriate to label two different people in a household as PMKs. For example, in the case of a step family, it may have been appropriate to label the mother as the PMK for one child and the father for another. However, to simplify the interview procedures, only one PMK is selected per household.

The PMK is selected once the information about the relationships between household members has been collected.

6.2 Child Component

A child component was created for each selected child between 0 and 17 years of age. The person most knowledgeable about the children and youth answered the child component questions. The PMK was usually the child's mother, but it could also be the father, a step-parent

or an adoptive parent who lived in the same dwelling. Only the PMK or his/her spouse was permitted to answer the questions in this component.

For households in which the only child selected was 16 or 17 years old and was living with his/her parents, a shorter version of the child component was asked. If the child was no longer living with his/her parents, the component was not created.

6.3 Adult Component

An adult component was created for the PMK and his/her spouse or partner, if the selected child is 17 years old or younger. Only the PMK or his/her spouse was permitted to answer the questions in this component. Questions in the adult component are asked once per household, even if more than one child was selected in the household.

For households in which the only child selected was 16 or 17 years old and was living with his/her parents, a shorter version of the adult component was asked. If the child was no longer living with his/her parents, the component was not created.

6.4 Youth Component

This component is used for selected respondents aged 16 and above. The youth was the only person permitted to answer the questions in this component, whether he/she was living in the family home or not.

6.5 Sample Sizes at Cycle 6

The number of children and youth sampled in Cycle 6 is shown by age and province in the following tables.

Table 1: Number of Sampled Children and Response Rate, by Age at Cycle 6

Age as of January 1 st , 2005	Sampled	In-scope	Respondents	Cycle 6 Response Rate (%)
0	1,777	1,770	1,457	82.3
1	2,579	2,573	2,064	80.2
2	1,845	1,834	1,627	88.7
3	2,140	2,121	1,834	86.5
4	1,900	1,889	1,662	88.0
5	2,127	2,122	1,870	88.1
6	1	0	0	n/a
7	0	0	0	n/a
8	0	0	0	n/a
9	0	0	0	n/a
10	1,595	1,587	1,378	86.8
11	1,689	1,683	1,447	86.0
12	1,189	1,182	1,015	85.9
13	1,147	1,146	980	85.5
14	1,085	1,077	907	84.2
15	1,026	1,022	884	86.5
16	924	920	790	85.9
17	932	929	795	85.6
18	1,119	1,106	815	73.7
19	1,052	1,044	758	72.6
20	992	978	746	76.3
21	908	898	663	73.8
Total	26,027	25,881	21,692	83.8

Table 2: Number of Sampled Children and Response Rate, by Province of Residence at Cycle 6

Province	Sampled	In-scope	Respondents	Cycle 6 Response Rate (%)
Newfoundland and Labrador	1,408	1,405	1,215	86.5
Prince Edward Island	959	955	815	85.3
Nova Scotia	1,749	1,746	1,493	85.5
New Brunswick	1,659	1,658	1,364	82.3
Quebec	4,385	4,379	3,667	83.7
Ontario	6,827	6,819	5,562	81.6
Manitoba	1,980	1,974	1,696	85.9
Saskatchewan	1,970	1,968	1,697	86.2
Alberta	2,683	2,677	2,263	84.5
British Columbia	2,304	2,300	1,920	83.5
Outside the 10 provinces	103	0	0	n/a
Total	26,027	25,881	21,692	83.8

6.6 Direct Assessments

A variety of direct assessments are administered to the selected respondent. These are summarized in the table below. For detailed information about the assessments, see Chapter 14.0.

Name of assessment	Age group	Method of administration
Peabody Picture Vocabulary Test – Revised (PPVT-R)	4- to 5-year-olds	Computer-assisted interview
Who Am I?	4- to 5-year-olds	Paper questionnaire
Number Knowledge	4- to 5-year-olds	Computer-assisted interview
Mathematics Computation Exercise	10- to 15-year-olds in Grades 4 to 10	Paper questionnaire
Problem Solving Exercise	16- to 17-year-olds	Paper questionnaire
Literacy Assessment	18- to 19-year-olds	Paper questionnaire
Numeracy Assessment	20- to 21-year-olds	Paper questionnaire

6.7 Self-complete Questionnaires – Ages 10 to 17

Respondents between 10 and 17 years of age completed a paper questionnaire on various aspects of their lives. The youth was given the questionnaire during the interview and asked to complete it himself/herself. To ensure confidentiality, the youth placed the completed questionnaire in an envelope, sealed the envelope and gave it to the interviewer.

The Self-complete Questionnaires consisted of a set of four booklets, one for each age group. The table below shows the subjects covered by each age-group section in the booklet. The questions for each subject were different for each age group. The booklets are reproduced in Book 2 of the National Longitudinal Survey of Children and Youth, Cycle 6 Survey Instruments 2004-2005.

Topic	Section			
	10 - 11 Booklet #20 E/F	12 - 13 Booklet #21 E/F	14 - 15 Booklet #22 E/F	16 - 17 Booklet #23 E/F
Friends and Family	A	A	A	A
School	B	B	B	-
About me	C	C	C	B
Feelings and Behaviours	D	D	D	C
My Parent(s)	E	G	G	G
Smoking, Drinking and Drugs	G	F	F	D
Puberty	F	H	H	-
Activities	H	E	E	-
Dating / My Relationships	-	H	H	F
Health	-	H	H	E
Work	-	I	I	-
Thank you	J	J	J	H

In Cycle 5, youth aged 18 and 19 completed a paper questionnaire. In Cycle 6, all the questions for these youth were administered as part of the Youth Component.

6.7.1 Collection Personnel (Training, Supervision and Control)

The NLSCY was conducted by Statistics Canada interviewers. A number of them had worked on one or more previous cycles of the NLSCY. All interviewers report to a staff of senior interviewers who are responsible for ensuring that interviewers are familiar with the survey's concepts and procedures. The senior interviewers ensure that prompt follow-up action is taken for refusal and other nonresponse cases. If necessary, nonresponse cases are transferred to a senior interviewer and re-assigned. The senior interviewers in turn report to the program managers, located at Statistics Canada's regional offices.

For the NLSCY, a combination of classroom training and self-study materials was used to ensure that interviewers and supervisors had a proper understanding of the survey concepts. In the self-study portion, which preceded the classroom training, the program managers, senior interviewers and interviewers read the Interviewer's Manual prepared for the survey and completed a case study exercise. The classroom training was given by a program manager or senior interviewer. There were two sets of training: one for the Computer-Assisted Telephone Interviewing application and one for the Computer-Assisted Personal Interviewing application.

7.0 Data Processing

The main outputs of the National Longitudinal Survey of Children and Youth (NLSCY), Cycle 6 are "clean" data files. This chapter presents a brief summary of some of the processing steps involved in producing these files.

The processing of the NLSCY Cycle 6 data was done using the Generalized Processing Environment. This is a generic system that follows a series of steps to "clean" a file from beginning to end. The main steps were:

- Clean up
- Age and gender edits
- Relationship edits
- Pre-edit
- Flow edits
- Coding
- Consistency edits
- Derived variables
- Final processing file
- Creation of Master File

7.1 Computer Generated Edits

As discussed earlier, all of the information for the household except for the 10 to 17 olds Self-complete Questionnaires was collected in a face-to-face or telephone interview using a computer-assisted interviewing (CAI) application. As such, it was possible to build various edits and checks into the questionnaire for the various household CAI components in order to ensure high quality of the information collected. Below are specific examples of the types of edits used in the NLSCY computer-assisted interviewing application:

Review Screens

Review screens were created for important and complex information. For example, the selection procedures for the person most knowledgeable (PMK), a critical element of the survey, were based on the household roster. The household roster screen showed the demographic information for each household member and his/her relationship to every other household member. The collected information was displayed on the screen for the interviewer to confirm with the respondent before continuing the interview.

Range Edits

Range edits were built into the CAI system for questions asking for numeric values. If values entered were outside the range, the system generated a pop-up window which stated the error and instructed the interviewer to make corrections to the appropriate question. For example, if the value entered into the computer for the child's weight at birth was significantly high or low, a pop-up message would appear asking the interviewer to confirm the answer with the respondent.

Flow Pattern Edits

All flow patterns were automatically built into the CAI system. For example, in the Child Care Section, the PMK is asked if he/she used daycare or babysitting in order that he/she (or a partner/spouse) could work or study. Based on the response given the flow of the questions could be different. If child care was used, the CAI system continued with a series of questions about the specific child care method(s) used for the child. If not, the CAI system automatically skipped this series of questions.

General Consistency Edits

Some consistency edits were included as part of the CAI system which allowed interviewers to "slide back" to previous questions to correct for inconsistencies. Instructions were displayed to interviewers for handling or correcting problems such as incomplete or incorrect data. For example, in the collection of the Labour Force Section, the number of weeks worked, not working, and looking for work should not total more than 52 weeks. If this was the case, the system generated a pop-up window that stated the error and instructed the interviewer to slide back to the appropriate question to confirm the data and make corrections as required.

7.2 Data Capture

Data Capture for Paper Questionnaires

Data capture for the following questionnaires was done in a centralized area at Statistics Canada's Head Office:

- the Self-complete Questionnaires for 10- to 17-year-olds
- the Who Am I? for 4- to 5-year-olds
- the Math tests for 10- to 15-year-olds who are in grades 4 and higher
- the Problem Solving Exercise for 16- and 17-year-olds
- the Literacy Assessment for 18- and 19-year-olds
- the Numeracy Assessment for 20- and 21-year-olds

Any document containing at least one respondent-completed item was captured and a file containing each record was provided to Head Office processing staff for further processing. As part of the capture system, some quality checks were built in to flag unusual entries to warn the operators of potentially incorrect entries.

In cases where more than one response was checked off by the respondent, the operators were instructed to accept the first response. Errors remaining within the questionnaires were then edited at a later stage.

7.3 Clean Up

Defining Requirements

The purpose of this step is to drop full-duplicate records and split-off records with duplicate identification numbers for examination. Then the data is split between response and nonresponse based on pre-determined criteria.

A review was done of the responding and nonresponding questionnaires and specifications were created based on this analysis to determine which records would be dropped due to nonresponse. Essentially, if a record was missing key information or had more than half the questions unanswered, they were dropped from the file.

At the end of this step, records were processed by questionnaire type, i.e., Adult Questionnaire, Child Questionnaire, Youth Questionnaire, Household Questionnaire and Self-complete Questionnaires.

Missing Variables

All missing variables for households were set to "Not stated". If there was not adequate information then the household was dropped from the responding sample and treated as a nonresponse.

The longitudinal file also contains 253 records that were created for some longitudinal children for whom no data was collected in this cycle. These are children who are now deceased or who have moved out of the country, but who will be kept on the longitudinal file for weighting purposes. For these records, all variables except for the longitudinal weight (FWTCW01L) have been set to "Not stated".

7.4 Age and Gender Edits

In this step, verification of all age variables is conducted. A comparison to the previous cycle is done for the date of birth and the reported age. The age of all reporting children and youth is compared with the previous cycle. Also, the age is verified to be consistent with the age cohort. For this cycle the interview date (INTVDATE) was used to edit effective age for all members of the household. INTVDATE now appears on the release. The respondent's sex is also verified to be consistent with the previous cycle.

7.5 Relationship Edits

The relationship edit step establishes the relationship between the members of the household and creates the family derived variables. This step performs a standard set of edits against the relationship information entered for all members of a given household; some inconsistencies are corrected automatically by an application using a set of rules, while others are flagged for manual review and recoding. A related set of derived variables is produced through the relationship edits.

7.6 Pre-edits

For all records where values were missing (blank) from the collection, the value of "9, 99, 999..." was inserted to indicate that no information was collected. The "Don't know" values returned by the CAI application as code "9" are changed to "7" in the pre-edits. As well, the "Mark all that apply" questions were de-strung and values converted to "Yes" ("1") or "No" ("2") responses. Finally, all text answers were removed from the processing file and set aside to be handled separately.

7.7 Flow Edits

The flow edits replicate the flow patterns from the questionnaire. Variables which are skipped based on flows are converted from "Not stated" to "Valid skip" codes (6, 96, 996...).

For skips based on age or based on the answer to certain questions, all skipped questions are set to "Valid skip". For skips based on "Don't know" and "Refusal", all skipped questions are set to "Not stated".

There were 26 youth aged 16 and 17 living independently in Cycle 6 (FDMCfD03=82). All variables in the longitudinal file were set to 'valid skip' for these respondents except for some variables in the Demographic Section, and certain variables in the Sociodemographic and Custody Sections, where some values were carried forward from previous cycles.

7.8 Coding of Open-ended Questions

A few data items on the NLSCY questionnaire were recorded by interviewers in an open-ended format. For example, in the Labour Force Section, a PMK who had worked in the previous 12 months was asked a series of open-ended questions about the current or most recent job:

- What kind of business, service or industry is/was this?
- What kind of work are/were you doing?
- At this work, what are/were your most important duties or activities?

Career aspirations questions were asked in the Youth Questionnaire for 18- to 21-year-olds:

- What kind of career or work would you be interested in having when you are about 30 years old?
- Specify type of career or work.
- Specify type of business.

How they are recorded

The interviewer recorded, in words, the answer provided by the respondent. At Head Office, these written descriptions were converted into industry and occupation codes which describe the nature of the respondent's work.

How they are coded

These open-ended questions were coded using various standard classifications. For Cycle 6 the Occupation questions were coded using the National Occupational Classification – Statistics (NOC-S) 2001, and the industry questions were coded using the 2002 North American Industry Classification System (NAICS 2002)⁴. Grouped versions of these codes are available on the data file (FLFPfD7A and FLFPfD8A for the PMK, and FLFSfD7A and FLFSfD8A for the spouse/partner).

7.9 Consistency Editing

After the flow edits were completed, consistency editing was carried out to verify the relationship between two or more variables. Decision tables are used to specify the consistency edits. The LogiPlus software was used to input the decision tables and generate the SAS code. A report with the “Before” and “After” counts of the variables is generated. Additionally, a report is generated providing the rule counts for each decision table.

For example, in the Socio-demographic Section, for children who were not born in Canada, question FSDCQ2B asks what year they first immigrated to Canada. There was a consistency edit which compared the year of immigration to the child's year of birth. If the year of immigration was before the year of birth then the year of immigration was set to “Not stated” in the edit.

7.10 Imputation Flags

Missing Variables

For various reasons certain variables may be missing for responding households on the NLSCY file. This is usually referred to as item nonresponse or partial response.

Imputation

For a few variables on the NLSCY file, rather than using a special nonresponse code, imputation has been carried out. Imputation is the process whereby missing or inconsistent items are replaced with plausible values. For the NLSCY, imputation was carried out for household income, PMK income, youth income, and motor and social development. See Chapter 10.0 for more details on imputation.

⁴ Information about classification and concordance to previous classifications can be found at www.statcan.ca/english/concepts/index.htm.

Imputation flags have been included on the NLSCY file so that users will have information on the extent of imputation and what specific items have been imputed on what records.

All imputation flag variables on the NLSCY data file have an "I" as the sixth character of the variable name. For example, the imputation flag variable for the income of the PMK would be named FINPc11A.

7.11 Creation of Derived Variables and Longitudinal Edits

For reasons of efficiency and data quality, these steps have been combined for Cycle 6 and future cycles.

Deriving variables from more than one data item

A number of variables have been derived by combining questions on the questionnaire, derived variables or both in order to facilitate data analysis. For example, in the Labour Force Section, a respondent's current working status is derived from questions about his/her current work situation and in the Adult Education Section, there is a question about whether he/she is currently attending a school, college or university. The combination of these two variables forms a derived variable identifying the respondent's current work and study situation (FLFPD51, FLFSD51).

Longitudinal Derived Variables

Longitudinal derived variables were created to indicate changes between data reported in the current and previous cycles for family structure and PMK and spouse changes.

As well, any data inconsistencies between cycles were flagged; the variables that identify these inconsistencies contain a "Z" in the fifth character of the variable name. For example if a respondent said in a previous cycle that they had used child care and in Cycle 6 indicated that they had never used child care, the variable FCRCeZQ6 will have a value of '1'. Inconsistencies between cycles for types of child care arrangements used, as well as for the child's height and weight are also flagged.

At this step, data was also brought forward from previous cycles for variables that are asked of the respondent only once, for example birth weight, breast-feeding information, and socio-demographic information. Data was also brought forward for derived variables from previous cycles that were derived from data that does not change over time.

For Cycle 6, the data brought forward are from Cycles 1 through 5, whereas in Cycle 5, data were brought forward from Cycle 4 only.

Derived Variable Name

All derived variables on the NLSCY data file have a "D" as the fifth character of the variable name. For example, the name of the derived variable for the "Primary care arrangements" is FLFPD51.

7.12 Standard Coding Structures

Some standards have been developed for the coding structure of NLSCY variables in order to explain certain situations in a consistent fashion across all variables. The following describes these various situations and the codes used to describe the situation.

Refusals

During a CAI interview, the respondent may choose to refuse to provide an answer for a particular item. The CAI system has a specific function key that the interviewer presses to

indicate a refusal. This information is recorded for the specific item refused and transmitted back to Head Office.

On the NLSCY data file, an item which was refused is indicated by a code "8". For a variable that is one digit long the code is "8", for a two-digit variable "98", for a three-digit variable "998", etc.

Don't know

The respondent may not know the answer to a particular item. Again the CAI system has a specific function key to describe this situation.

On the NLSCY data file, the code used to indicate that the respondent did not know the answer to an item is "7". For a variable that is one digit long the code is "7", for a two-digit variable "97", for a three-digit variable "997", etc.

Valid skip

In some cases a question was not applicable to the survey respondent. A code "6", "96", "996", etc., was used on the data file to indicate that a question or derived variable is a valid skip.

In some cases a single question or series of questions was not applicable. For example, the question on number of hours per week the child is cared for in a daycare centre (FCRCQ1G1) is only applicable for children for whom this type of care is used (FCRCQ1G = 1). Otherwise there will be a code "996" for this question.

In other cases an entire section of the questionnaire was not applicable or even an entire questionnaire. For example, the Motor and Social Development Section was applicable only to children 0 to 3 years old. For all children outside of this age group, i.e., 4 years old and over, the motor and social development variables have been set to a "Valid skip" ("6", "96", "996").

For cases where the PMK did not have a spouse or common-law partner residing in the household, all spouse variables, e.g., the Labour Force Section and the Education Section for the spouse, have been set to a "Valid skip".

Not stated

In some cases, as part of Head Office processing the answer to an item has been set to "Not stated". The "Not stated" code indicates that the answer to the question is unknown. The "Not stated" codes were assigned for the following reasons:

- As part of the CAI interview, the interviewer was permitted to enter a "Refusal" or "Don't know" code, as described above. When this happened, the CAI system was often programmed to skip out of this particular section of the questionnaire. In the case of a "Refusal", it was assumed that the line of questioning was sensitive and it was likely that the respondent would not answer any more questions on this particular topic area. In the case of a "Don't know", it was assumed that the respondent was not well enough informed to answer further questions. As part of the NLSCY processing system, it was decided that all of these subsequent questions should be assigned a "Not stated" code. A "Not stated" code means that the question was not asked of the respondent. In some cases it is not even known if the question was applicable to the respondent.
- In some cases a questionnaire was not started or it was started but ended prematurely. For example, there may have been some kind of an interruption, or the respondent decided that she/he wished to terminate the interview. If there was enough information collected to establish the household as a responding household, then all remaining unanswered questions on the questionnaire (and on questionnaires that had not yet been started) were set to "Not stated". The one exception was that if it was known that

a certain section or a certain questionnaire was not applicable, then these questions were set to “Not applicable”.

- Some paper questionnaires were mailed back partially complete. If there was enough information collected about the respondent, then all remaining incomplete items on the questionnaire were set to “Not stated”. The one exception was that if it was known that a certain section or a certain questionnaire was not applicable, then these questions were set to “Not applicable”.
- Another situation in which not stated codes were used was as a result of consistency edits. When the relationship between groups of variables was checked for consistency, if there was an error, often one or more of the variables was set to “Not stated”.

For derived variables, if one or more of the input variables (to the derived variable) had a “Refusal”, “Don’t know” or “Not stated” code, then the derived variable was set to “Not stated”.

An item that was coded as “Not stated” is indicated by a code "9". For a variable that is one digit long the code is "9", for a two-digit variable "99", for a three-digit variable "999", etc.

7.13 Naming Convention

The NLSCY microdata file documentation system has employed certain standards to label variable names and values. The intent is to make data interpretation more straightforward for the user.

A naming convention has been used for each variable on the NLSCY data file in order to give users specific information about the variable. All variable names are, at most, eight characters long so that these names can easily be used with analytical software packages such as SAS or SPSS. The “Persruk” and “Fieldruk” identifiers are the exception to this naming convention.

Format for Variable Names

The **first** character of the variable name refers to the NLSCY cycle:

- “**A**” indicates the first cycle,
- “**B**” indicates the second cycle,
- “**C**” indicates the third cycle,
- “**D**” indicates the fourth cycle,
- “**E**” indicates the fifth cycle, and
- “**F**” indicates the sixth cycle.

The **second and third** characters refer to the section of the questionnaire where the question was asked or the section from which the variable was derived. Refer to Section 7.15 for acronym names for each questionnaire sections.

The **fourth** character refers to the collection unit or the unit to which the variable refers. There are nine possibilities⁵.

⁵ It should be noted that while variables do exist for various units of analyses, i.e., the PMK, the spouse/partner and the household, it will only be possible to produce "child estimates" from the NLSCY file. The characteristics of the PMK, spouse/partner and household can be used to describe attributes of the child. For example it will be possible to estimate the number of children living in a household with low income, or the number of children for whom the PMK has scored high on the depression scale etc. However it will **not** be possible to produce estimates of the number of low income households or depressed PMKs.

- “**C**” if the variable refers to the child
- “**P**” if the variable refers to the PMK
- “**S**” if the variable refers to the spouse/partner
- “**H**” if the variable refers to the household
- “**Y**” if the variable refers to youth
- “**W**” if the variable refers to a weight
- “**M**” if the variable refers to the mother

The **fifth, sixth, seventh and eighth** characters of the variable name indicate the cycle in which the variable first appeared (if not Cycle 1), the type of variable and a sequential number assigned to the variable. For example,

- f** the lower case letter refers to the NLSCY cycle in which the variable first appeared on the file or the cycle in which changes to a previously asked question were made.

Some revisions were made to the content of the questionnaire between cycles. If the revision resulted in a change to the meaning or the values of a question in Cycle 6, the variable was treated as new and contains an “f”.

For example, for the variable FLTCfQ1E the coverage was changed from the previous cycle (0- to 2-year-olds) to include children 0-5 years old for Cycle 6. The variable FLTCfQ4B is a new question that was asked in Cycle 6.

- Q** refers to the variable for a question that was asked directly on one of the NLSCY questionnaires.
- S** refers to a score calculated for one of the scales used on the questionnaire.
- D** means the variable was derived from two or more questions that were asked on the questionnaire or coded variables.
- I** means the variable is a flag created to indicate that an item has been imputed.
- Z** means the variable is a flag created to indicate an inconsistency in reported data between the current and previous cycles.
- nnx** refers to the question or variable identification. Generally “nn” is a sequential number assigned to the variable; and “x” is a sequential alphabetic indicator for a series of variables of a similar type.

7.14 Examples of Variables Names

In order to illustrate the naming convention used for variables included on the NLSCY data file the following examples are given.

Variable Name Refers to:

Variable Name	Refers to:
FLFSQ2	Q2 in the Labour Force Section for the spouse/partner
F	Cycle 6 variable
LF	Labour Force Section
S	Spouse/partner
Q	An item asked directly on the questionnaire
2	The second question from the Labour Force Section of the Adult questionnaire

Variable Name	Refers to:
FPRCS03	a positive interaction score on the parenting scale for a 2- to 15-year-old child
F	Cycle 6 variable
PR	Parenting Section
C	Child
S	A score
03	The identification number of the item

7.15 Acronym Names for the Questionnaire Sections

The following table gives the acronyms that were used for each section of the various NLSCY questionnaires. The acronym is embedded in the variable name for all variables on the NLSCY data file. The acronym is the second and third characters of the variable name.

Acronym	Variable	Collected or Derived From:
GE	Geographic	Sample information.
HH	Household	Dwelling characteristics.
MM	Variables collected as part of the household roster.	Basic demographic variables for each household member. These variables are included on the NLSCY data file for the child, youth, the PMK and the spouse/partner.
DM	Demographic - derived to explain the living arrangements of the child or youth.	Information from the household roster and relationship grid.
SD	Socio-demographic	Child on the Child Questionnaire PMK and spouse/partner on the Adult Questionnaire.
HL	Health	PMK and spouse/partner on the Adult Questionnaire Child on the Child Questionnaire Youth 16 and older on the Youth Questionnaire.
CH	Adult Chronic Conditions	PMK and spouse/partner in the Health Section of the Adult Questionnaire.
RS	Restriction of Activities	PMK and spouse/partner in the Health Section of the Adult Questionnaire.
DP	Depression Scale	Parent Questionnaire (this scale was administered to the PMK).
ED	Education	Children aged 4 and 5, and 10 to 15 years old on the Child Questionnaire PMK and spouse/partner on the Adult Questionnaire Youth 16 and older on the Youth Questionnaire.
LF	Labour Force	PMK and spouse/partner on the Adult Questionnaire Youth 16 and older on the Youth Questionnaire.
IN	Income	Household income and personal income of the PMK collected on the Adult Questionnaire Youth 16 and older on the Youth Questionnaire.
FN	Family Functioning	Adult Questionnaire (section asked to the PMK or spouse/partner).
MD	Medical/Biological	Child Questionnaire (0 to 3 years)
TM	Temperament	Child Questionnaire (3 to 35 months).
LT	Literacy	Child Questionnaire (0 to 5 years).
AC	Activities	Child Questionnaire (0 to 5 and 10 to 15 years) Youth 16 years and older on the Youth Questionnaire.
BE	Behaviour	Child Questionnaire (0 to 5 and 10 and 11 years).
MS	Motor and Social Development	Child Questionnaire (0 to 47 months).
WB	Work after Birth	Child Questionnaire (0 to 5 and 10 to 15 years)
RL	Social Relationship	Child Questionnaire (4 to 5 years). Youth 18 and older on the Youth Questionnaire
PR	Parenting Style	Child Questionnaire (0 to 5 and 10 to 15 years).
CR	Child Care	Child Questionnaire (0 to 5 and 10 to 13 years).

Acronym	Variable	Collected or Derived From:
SL	Sleep	Child Questionnaire (0 to 13 years).
PB	Positive Behaviour	Child Questionnaire (3 to 5 years).
CM	Communication Scale	Child Questionnaire (3 to 5 years).
ML	Development Milestones	Child Questionnaire (9 to 47 months).
CS	Custody	Child Questionnaire (0 to 5 and 10 to 17 years old).
AG	Ages and Stages	Child Questionnaire (3 to 47 months).
SF	Neighbourhood Safety	Adult Questionnaire (section asked to the PMK or spouse/partner).
SP	Social Support	Adult Questionnaire (section asked to the PMK or spouse/partner). Youth 18 years and older on the Youth Questionnaire.
AS	Aspirations	Child Questionnaire (16 to 17 years). Youth 18 years and older on the Youth Questionnaire.
EQ	Emotional Quotient	10-17 years Self-complete Questionnaires. Youth 20 and 21 on Youth Questionnaire.
AM	About Me	10 to 17 years Self-complete Questionnaires. Youth 18 and 19 on Youth Questionnaire.
FB	Feelings and Behaviour	10 to 17 years Self-complete Questionnaires. Youth 18 years and older on Youth Questionnaire.
MO	Moving Out of Parental Home	Youth 18 years and older on Youth Questionnaire.
FF	Friends and Family	10 to 17 years Self-complete Questionnaires.
SC	School	10 to 15 years Self-complete Questionnaires.
PM	My Parents and Me	10 to 17 years Self-complete Questionnaires.
PU	Puberty	10 to 17 years Self-complete Questionnaires.
DR	Smoking, Drinking and Drugs	10 to 17 years Self-complete Questionnaire.
AT	Activities	10 to 15 years Self-complete Questionnaires.
HT	Health	12 to 17 years Self-complete Questionnaires.
WK	Work and Sources of Money	12 to 15 years Self-complete Questionnaires.
DA	Dating	12 to 17 years Self-complete Questionnaire.
PP	PPVT-R Test	Aged 4 to 5 years.
WM	Who Am I?	Direct Measure (4 to 5 years).
KN	Number Knowledge	Direct Measure (4 to 5 years).
MA	Math Computation Test	Children 8 to 15 years old in grades 2 to 10 and Problem Solving Exercise for 16- and 17-year-olds.
LI	Literacy Assessment	18- and 19-year-olds.
NU	Numeracy Assessment	20- and 21-year-olds.
WT	Weight	Weight as part of the sample design.

7.16 Final Processing Files and Master Files

The following files were created for the NLSCY, Cycle 6:

- Early Child Development Cohort 0- to 5-year-olds (Adult, Child and Household)
- Longitudinal Cohort 10- to 17-year-olds (Adult, Child and Household)
- Longitudinal Cohort 16- to 21-year-olds (Youth and Household, also static variables brought forward from previous cycles)
- Self-complete Questionnaires
- North (5-year-olds in the Yukon and Nunavut)

8.0 Content of the Survey

The National Longitudinal Survey of Children and Youth (NLSCY) was designed to follow an ecological or holistic approach to measuring child development. The survey captures the diversity and dynamics of the factors affecting children. To ensure that all relevant topic areas affecting child development were adequately addressed by the survey, a multidisciplinary consultation was carried out at the inception of the survey. The selection of specific subject areas, priorities and survey questions was very much a group effort with input and advice from:

- the NLSCY expert advisory group (EAG), which consists of researchers in the area of child development and the social sciences;
- federal departments;
- representatives from the provinces and territories responsible for child development programs.

It was recommended that the NLSCY cover a broad range of characteristics and factors affecting child growth and development. Extensive information was gathered about the child, as well as the child's parent(s), characteristics of the family and the neighbourhood as well as the child's school and school experiences. This section provides an outline of the content for each section of the questionnaire included in the NLSCY data. The different scales used in the NLSCY will be discussed briefly in this chapter but for more information or for a discussion on the validation of the scale scores, please see Chapter 9.0.

8.1 Survey Components

The NLSCY is divided into several components; these are described in Chapter 6.0, Data Collection. Below is a summary of each component.

Household	This is the first part of the interview. The household roster asks for basic demographic information for each household member and his/her relationship to everyone else in the household.
Adult	Questions asked about the person most knowledgeable (PMK) and spouse. For children aged 16 and 17, not all the sections in the adult component are asked. The adult component is completed once even if there are two children in the household. No adult component is generated for youth aged 18 years and older.
Child	Questions about the selected child asked to the PMK. A child component is completed for each selected child aged 0 to 5 and 10 to 17. The only sections of the Child Questionnaire asked about youth aged 16 and 17 are the Aspirations and Expectations section, Custody and the Socio-Demographics section.
Youth	Questions asked about the selected youth, if he/she is aged 16 to 21 years old. In this section respondents answer questions about themselves in a computer-assisted interview (CAI).
Self-completes	Respondents aged 10 to 17 answer questions about themselves in a paper questionnaire.
Direct assessments	Several direct assessments are done with the children and youth; these are described in Chapter 14.0.

8.2 Demographic Variables

The demographic variables are collected in the household questionnaire. As part of the household questionnaire some basic demographic information, e.g., age, gender, and marital status, is collected for all members of the child's household. The relationship grid is also completed as part of this questionnaire, i.e., the relationships of everyone in the household to all the other members of the household. Using this information it is possible to create an extensive set of variables to describe the child's family situation.

It is necessary to perform an extensive series of edits on the collected data. The following are some examples of the types of editing that are carried out.

- a birth parent should be at least 12 years older (and not more than 55 years older) than a birth child
- the difference in age between a husband and wife should be less than 29 years.

8.3 Adult Questionnaire

Education (Parent)

The Education section is completed for both the PMK and spouse/partner. The objective is to gather information on the years of school completed, educational attainment, and current attendance at an educational institution.

Research has indicated a link between maternal educational attainment, the home environment and child development. The questions on full-time and part-time school attendance provide an indicator of the main activities of the PMK and the spouse/partner.

Labour Force

Employment stability impacts the home environment, both in terms of income and stress levels. Research indicates that parental unemployment can adversely impact child outcomes.

The Labour Force section is completed for both the PMK and spouse/partner. The main objective of the section is to determine employment stability as an indicator of the continuity of employment income. Questions include periods of absence from work, reason for the most recent absence, hours worked, and work arrangements, e.g. shifts, during the previous year. A series of questions are asked about the PMK and spouse/partner's current or most recent job held.

A complete description is recorded for the current or most recent job. Industry and occupation coding was carried out using the North American Industry Classification System (NAICS) 2002, the 1991 Standard Occupational Classification codes and the National Occupational Classification for Statistics 2001 (NOC-S).

Labour Force Derived Variables

Several labour force derived variables have been created for the PMK and spouse/partner of the PMK. They include:

FLFPfD5A / FLFSfD5A:	NAICS 2002 code for PMK's/spouse's current job
FLFPfD6A / FLFSfD6A:	NOC-S 2001 code for PMK's/spouse's main job
FLFPfD7A / FLFSfD7A:	Standard industry code for current job (NAICS 2002) – grouped
FLFPfD8A / FLFSfD8A:	Standard occupation code for current job (NOC-S 2001) – grouped

Income

In the Income section of the survey, the sources of income and the income are collected for each household. There are also a few questions that ask for the perceptions of the PMK or the PMK's spouse regarding how well he/she thinks he/she is doing financially. This information provides an indicator of the family's economic situation, an essential component of the child's environment.

As family income is an important part of many studies on child development, we impute a value for household income if the respondent did not answer these questions. See Chapter 10.0 for a detailed explanation of how income is imputed.

A derived variable (FINHD04A) has been created to compare the household income to the low income cut off (LICO)⁶. The LICO is used to distinguish "low income" family units from "other" family units. A family unit is considered "low income" when its income is below the cut off for its family size and its community. The variable FINHD03A gives the value of the LICO by geographic area.

Adult Health

This section asks the PMKs and their spouses about their general health, chronic conditions and restriction of activities, as well as questions on smoking and drinking. The smoking questions have been included because research has indicated that parental smoking behaviours may be predictive of the use of cigarettes by children. Alcohol consumption is covered because of potential impacts on the adult's physical or mental health, the family's economic situation, and family relationships.

Chronic Conditions

PMKs and their spouses are asked whether or not they have any long-term conditions, e.g., allergies, asthma, and high blood pressure. A derived variable (FCHPD01 or FCHSD01) indicates that the respondent answered "Yes", they have at least one of the long-term conditions.

Restriction of Activities

The PMKs and their spouses are asked a series of questions about whether or not their activities are restricted at home, work, school, etc. A derived variable (FRSPdD01 or FRSSdD01) is also created stating whether or not the PMK or spouse reported an activity restriction.

Maternal History

This section is asked to determine pregnancy history. These questions are only asked of those being interviewed for the first time. The questions on pregnancy and birth were provided by Dr. J.-F. Saucier, Ste. Justine Hospital, Montreal, and later modified by the Project Team.

Depression Scale

A Depression Scale (FDPPS01) is administered to the PMK as part of the Adult Questionnaire.

Family Functioning

The objective of the Family Functioning section is to provide a global assessment of family functioning and an indication of the quality of family relationships. This section is asked of the PMK or spouse, if the child is 0-15 years old.

Neighbourhood Safety

This section gathers information about the respondent's satisfaction with his/her neighbourhood as a place to raise children, including perception of the extent of danger and problems, and of social cohesion or "neighbourliness". Two scales are created in this section: Neighbourhood

6 For more information about Statistics Canada's low-income measures, please see Low income cut offs for 2005 and low income measures 2004, Catalogue no. 75F0002MIE2006004.

Safety Score (FSFHS5), indicating the degree of perceived neighbourhood safety and the Neighbours Score (FSFHS6), indicating the degree of neighbour cohesiveness.

Note: This section is not asked for the Northern collection.

Social Support

The purpose of this section is to collect information on the level of support the PMKs feel they have from friends, family members and members of the community. This section is asked of the PMK or the PMK's spouse, if the child is 0 to 15 years old.

Socio-demographic Characteristics

The objective of the Socio-demographic section is to gather information on immigration, ethnic background and the language profile of household members. This will allow for analysis of various components of the Canadian population and will permit identification of visible minorities. As well, there is a question on religious affiliation. These questions are asked of the PMK, spouse and the child.

8.4 Child Questionnaire

Questions in the child questionnaire are administered based on the child's effective age. Instead of using the child's actual age, the NLSCY uses a calculated age called effective age (FMMCQ01). This is done to ensure that the child stays in the age group to which he/she is assigned regardless of whether collection takes place before or after the child's birthday. For Cycle 6, the effective age is calculated as 2004 minus the year of birth. For example, a child born in 2000 would have an effective age of 4 years old (2004-2000). Note that the actual age of the child at the time of the interview is sometimes different from the effective age.

Education (Child)

The objective of this section is to gather basic information about the child's educational experiences. The amount and type of information collected varies depending upon the age of the child, with more information being collected for the older children who have had greater school experience.

Basic information is collected for all age groups, such as: the child's grade level, type of school and language of instruction, whether the child looks forward to school, absenteeism, number of school changes and residential moves.

For children in grade one or higher, additional questions are asked concerning other aspects of the educational experience such as skipping and repeating grades, achievement and special education.

Direct Measures

The purpose of this section is to establish the groundwork for the Direct Measures that will be asked of children aged 4 and 5. If the child does not have the ability to do the direct measures, i.e., does not speak English or French or is colour blind, the measures will not be administered.

Note: This section is not asked for the Northern collection.

Health (Child)

The objective of this section is to provide information on the child's physical health – general health, injuries, limitations and chronic conditions – and use of health services and medications.

For a child 4 or 5 years old, health status information on topics such as hearing, sight, speech and overall mental well-being is also collected. From this information a Health Status Index (HUI3) is calculated (FHLCCD2A). The HUI3 is a generic health status index that is able to

synthesize both quantitative and qualitative aspects of health. The index, developed at McMaster University's Centre for Health Economics and Policy Analysis, is based on the Comprehensive Health Status Measurement System (CHSMS). It provides a description of an individual's overall functional health, based on eight attributes: vision, hearing, speech, mobility (ability to get around), dexterity (use of hands and fingers), cognition (memory and thinking), emotion (feelings), and pain and discomfort.

The scores of the HUI3 embody the views of society concerning health status. Each person's preferences are represented as a numerical value (typically between 0 and 1) for a given health state. (Some of the worse states of health are often given values less than 0, indicating that the individual considers them to be worse than death.) This index is also used by the National Population Health Survey.

Medical/Biological

The Medical/Biological section is completed for children in the 0 to 3 year age group. The major objective is to collect information on factors such as gestational age and birth weight. These factors have been shown to have a direct impact on a child's growth and development. For example, in the long term, underweight babies face higher risks of poor health as well as longer-lasting developmental difficulties.

For each child under 2, the nature of the delivery, general health of the child at birth and the use of specialized services following the birth are collected in this section. The NLSCY also investigates the biological mother's pregnancy and delivery history, including policy-relevant topics such as the mother's breast-feeding experiences and prenatal lifestyle.

There are derived variables created for this section that should be noted. Two variables were derived to indicate the gestational age of the child. FMDCD06 gives the gestational age in days and FMDCD07 indicates if the child was born prematurely (gestational age 258 days or less), in the normal range (gestational age 259 to 293 days) or late (gestational age 294 days or later).

A variable was derived (FMDCD08) to indicate if the child was of normal birth weight (2,500 grams), moderately low birth weight (1,500 to 2,499 grams) or very low birth weight (< 1,500 grams).

Work After Birth

These questions are asked to determine the time interval after which mothers returned to work following the birth of a child and the extent to which these mothers participate in the labour force upon their return.

Ages and Stages Questionnaires

The Ages and Stages Questionnaires (ASQ) are parent-report instruments, developed by Jane Squires, LaWanda Potter, and Diane Bricker, at the University of Oregon, designed to identify infants and young children who show potential developmental problems. There are 19 questionnaires that cover the age range from 4 to 60 months. Each questionnaire includes roughly 30 items covering five domains of development.

- 1) Communication: babbling, vocalizing, listening, and understanding
- 2) Gross Motor: arm, body, and leg coordination
- 3) Fine Motor: hand and finger coordination
- 4) Problem Solving: doing different activities with objects, drawing
- 5) Personal-Social: solitary and social play, dressing and feeding self

The questionnaires also include an overall section that asks about general parental concerns but is not used in the NLSCY, as these questions are similar to those already included in the survey.

The NLSCY is using the ASQs for children aged 3 to 47 months, inclusive. In consultation with the publisher, Statistics Canada has converted the questionnaires so that they could be asked as part of the CAI application. The gross motor portion of the ASQs is not included as this concept is covered in other portions of the survey.

Milestones

These questions are included in order to provide a better measure of early child development. Taken as a package, developmental milestones, such as when the child first said words or took first steps, provide a general sense of a child's development. Experts with the Dunedin study in New Zealand recommended to the Project Team that developmental milestones be used as a measure of development. The items are from the draft questionnaires for the Early Childhood Longitudinal Study Program (Birth Cohort) of the National Center for Education Statistics in the United States.

Temperament

This section measures the temperament of young children by asking the parent about the degree of difficulty their child presents. This measure is based on the assumption that a child's temperament is influenced by the parent's perception of the difficulty of the child, and that temperament is not solely based on biological origins.

Literacy

This section measures children's exposure to books and their interest in reading and learning-related activities that parents do with their children. The focus of this section is the stimulation young children receive at home.

For children aged 0 to 2, several questions are asked to measure how often the parents do certain activities with their children, such as tell stories, sing songs and teach new words. These questions were adapted from the Early Childhood Longitudinal Study in the United States.

Similar questions are asked about children aged 3 to 5, with changes to reflect age appropriate activities. A question has been added on the suggestion of colleagues of Dr. Robbie Case, the developer of the Number Knowledge Test, who state that numeracy is a crucial factor when it comes to learning and literacy. This question was added to provide information on children's numeracy activities that may correlate with results on the Number Knowledge Test.

Communication

The items have been modified from the New Zealand Competent Children Study. They cover a child's ability to understand oral messages and to pass a message on to someone else, as well as to communicate verbally. The final question, about speech being easily understood, is only asked of 3-year-olds. Four and 5-year-olds are asked a similar question as part of the Health Status Index in the Health Section.

Activities

This section measures the child's participation in various non-school activities and the amount of household responsibility taken on by 10- to 13-year-olds at home. The latter questions are used to create the Home Responsibilities Score (FACCS06), indicating the degree of home responsibilities. The section will give some sense of how the child spends his/her time, of personal interests, as well as the degree of interaction with peers.

Several questions are included for children 4 and 5 years old to determine how often parents get to do certain activities with their children, such as eating a meal, playing a game and doing chores together. When there is a spouse/partner in the household, these questions are asked about both the PMK and his/her spouse/partner.

Behaviour

The objective of this section is to assess aspects of the behaviour of children 2 years of age and older and of feeding patterns for 1- to 3-year-olds.

The questions in this section are used to measure the prevalence of behaviours such as hyperactivity and physical aggression. The scales derived from these questions are described in detail in Chapter 9.0.

Positive Behaviour

The objective of this section is to assess positive behaviour of children aged 3 to 5, including perseverance and independence. The New Zealand Competent Children Study has found that perseverance and independence were among a cluster of competencies that are good indicators of a child's overall performance.

Questions have been adapted from the New Zealand study and the behaviour questions used for other ages in the NLSCY.

Sleep

Research suggests that sleeping difficulties are predictive of a child's potential difficulties. Conversely, absence of such difficulties has been correlated with easy temperament and positive outcomes.

The questions in this section ask about hours of sleep, hours of uninterrupted sleep at night, how often the parents sleep was disturbed by the child and so on.

Motor and Social Development

The Motor and Social Development (MSD) Scale measures dimensions of the motor, social and cognitive development of children from birth to age 3; the questions vary by the age of the child. Three scores (FMSCS01, FMSCS02 and FMSCdS03) are derived from these questions.

Relationships

The objective of this section is to provide information about the child's relationships with others. Positive relationships with other children and adults may help to counteract other factors that place a child at risk.

Questions about doing things with friends and getting along with parents, teachers and friends are based on those in the Ontario Child Health Study.

Parenting

Parenting style is considered to have an important influence on child behaviour and development. The objective of this section is to measure certain parenting behaviours. Scales are created from the questions in this section.

The PMKs who have a spouse/partner in the house are asked how often the PMK and spouse/partner agree with each other about parenting decisions. This question was developed by the Project Team and is similar to questions in the Strayhorn and Weidham scale, from which the other parenting questions have been adapted.

Custody

This section was designed to provide information on the child's family arrangements; whether or not his/her parents are married, separated or divorced, the age of the child when parents separated/divorced and so on.

Child Care

This section provides basic information about the methods of care currently provided for the child while the parents are working or studying, plus some information on previous care. Concepts measured include both the amount of time spent by the child in child care and the methods of care used for each child. In addition, information is obtained on the number of changes in child care arrangements that the child has experienced and the reason(s) for changes in the past 12 months. The section also identifies whether or not a child care centre is profit or non-profit, whether home care is licensed or unlicensed, and the ratio of caregivers to children.

Aspirations and Expectations (Ages 16 and 17)

These questions are included to assess parental aspirations and expectations for their youth, and parental views on their youth's school experiences. Providing help with school work, discussing school experiences and future educational plans has been linked to school success.

These questions were developed by the Centre for Education Statistics at Statistics Canada, using NLSCY questions and questions from other education surveys, such as, the Youth in Transition Survey and the School Leavers Survey.

Socio-demographics

These questions gather socio-demographic information on the selected child. Such information as ethnicity and country of origin is collected. In Cycle 6, two questions were included regarding Aboriginal identity.

8.5 Youth Questionnaire (Ages 16 to 21)

Moving Out of the Parental Home (Ages 18 to 21)

There are numerous transitions that a youth goes through from adolescence to adulthood. Undoubtedly, one of these major transitions is when the youth leaves the parental home for the first time to live independently in their own residence.

This section is designed to gather information on how many times the 18- to 21-year-old respondents have left home. Some of these youth will be living away from home for school or work either permanently or temporarily. It was felt that information should be collected on this transition because of its importance in the movement from childhood to adulthood.

The questions in this section will vary depending on already-collected information and information collected during the household component.

The questions were designed on the advice of Dr. Dianne Looker, Chair, Department of Sociology, Acadia University. Dr. Looker supplied us with questions she used in her longitudinal study, "The Transition from Education to Employment: A Longitudinal and Cohort Analysis of Canadian Youth."

Youth Education (Ages 16 to 21)

This section collects information on the youth's education experience. The first few questions establish the current educational status. Based on this information, the respondent is then streamed to the questions applicable to his/her situation. There are four possible streams:

- 1) school leavers (those who are not in school and have not graduated from high school),
- 2) school finishers (those who are not in school and have graduated),
- 3) currently in school (for youth still in high school), and

- 4) currently in post-secondary (for youth who are attending a post-secondary education institution).

The questions were developed by the Centre for Education Statistics at Statistics Canada using NLSCY questions and questions from other education surveys, such as, the Youth in Transition Survey and the School Leavers Survey.

One of the objectives of this section is to help determine the factors involved in youth choosing to continue their schooling or to leave school.

In Cycle 5, an integrated education section was developed. As there are many respondents of different ages and similar education statuses, the youth are streamed to the correct questions based not on age but previous and current education statuses.

Youth Labour Force (Ages 16 and 17)

The Youth Labour Force section is intended to measure youth experience in the labour market. Some youth may be working part-time while attending school, while others may have made the transition to the workforce. These questions are a mix of NLSCY questions from the Self-completes for 14- and 15-year-olds and of the adult labour force questions.

Youth are asked to report about current work, work during the current school year and work last summer.

Labour Force Derived Variables

Several labour force derived variables have been created for the youth (ages 16 and 17). They include: FLFYfD5A, FLFYfD6A, FLFYfD7A, FLFYfD8A, FLFYeD02.

Youth Labour Force (Ages 18 to 21)

The questions in this section are similar to those asked of the youth aged 16 and 17. However, there are more questions taken from the adult labour force section due to the increased age of the respondents. Some of the youth may be working as their main activity and the questions need to reflect this possibility.

These questions collect information that will help to paint a broad picture of youth labour force participation, touching mainly on employment status, job characteristics, number of hours worked, job stability, and the link between work and educational goals and achievements.

Similar to the adult labour force section, a complete description is recorded for the current or most recent job. Industry and occupation coding was carried out using the North American Industry Classification System (NAICS) 2002 and the National Occupational Classification for Statistics 2001 (NOC-S).

Labour Force Derived Variables

Several labour force derived variables have been created for the youth (ages 18 to 21). They include: FLYYfD5A, FLYYfD6A, FLYYfD7A, FLYYfD8A.

Youth Career Aspirations (Ages 18 to 21)

This section collects information on the types of information that the youth has gathered about different career paths. It also identifies whether or not the youth has decided on a future career. The questions vary depending on the age of the respondent.

Career aspirations are thought to provide realistic direction, enabling individuals to find suitable and satisfying jobs. It is important to collect information on future work expectations in order to gain insight into the degree to which young people plan for their future careers.

Some of the questions that appear in this section were developed in consultation with Dr. Dianne Looker from Acadia University. Other items were included that had been used in her own work, “The Transition from Education to Employment: A Longitudinal and Cohort Analysis of Canadian Youth.”

For youth reporting a desired future career, occupation coding was carried out using the National Occupational Classification for Statistics 2001 (NOC-S). From this information, the variable FASYfD03 was created.

Youth Income (Ages 16 and 17)

The Youth Income section asks the youth about his/her income from various sources in the last 12 months. The degree to which youth make autonomous decisions may be measured through their spending behaviours.

Youth Income (Ages 18 to 21)

This is a new section included for Cycle 6 for this age group. These questions are similar to those asked on the adult questionnaire. The youth is asked about his/her personal income if single and about household income if married or living common law.

There is a series of questions concerning payment of housing/shelter expenses. At this age, many youth may be moving out of the parental home for the first time. Determining whether they are paying for shelter is important data to collect in order to assess how youth adjust to financial responsibilities.

Youth Health (Ages 16 to 21)

This section asks about the youth’s general health, injuries, chronic conditions and restriction of activities. These questions are similar to the child and adult health questions. There are also some questions relating to the sleep patterns of the youth. Sleep is an important indicator of the youth’s attitudes towards his/her body and how he/she takes care of him or herself. The amount of sleep reported can be used to help understand if youth are successfully balancing the demands of work, school, volunteering, sports, etc.

Youth Health (Ages 18 to 21)

This section includes questions that ask about height, weight, risky behaviours such as smoking, drinking alcohol and drug use. In Cycle 6, these questions were moved to the CAI application from the Self-Completes.

Furthermore, the depression scale taken from the self-complete questionnaires was included in this section (FHTYfS01). Considering the number of transitions taking place during early adulthood, it can be an extremely stressful time in their lives and certain emotions maybe evoked. There is research that suggests that many people suffer from depression. Good mental health is as important as good physical health and therefore it is important to gather information on aspects of both.

Feelings and Behaviours (Ages 18 and 19)

The objective of this section is to establish whether the youth knows anyone who has committed suicide and whether they have seriously considered or attempted suicide. These questions were adapted from the 1992 British Columbia Adolescent Health Survey and are similar to questions asked of the younger respondents on the self-complete questionnaires.

There are also some questions that ask about engaging in risky behaviour such as stealing, fighting, drinking and driving, and gang membership.

Youth Activities (Ages 16 to 21)

Young adulthood can be a time of high involvement in a variety of activities that are not school related. It is important to measure these activities to understand how this involvement can contribute to good outcomes. The degree to which youth engages in life-long learning and establish value autonomy are outcome areas that are of particular interest to researchers.

This section includes questions about physical activities, literacy activities, television watching, computer use, community involvement and spirituality. These questions have been adapted from the questions asked of younger adolescents. The youth received different questions depending on their age.

Questions are included for the youth aged 16 and 17 about the youth's access to a vehicle and whether or not he or she has a driver's license. Driving is an important "coming of age" activity for this age group and it is important to collect data on this topic.

There are also questions related to the extent to which the youth volunteers in his/her community. In addition, there was also a new question in Cycle 6 asking whether the youth voted in the last municipal, provincial or federal election. These questions attempt to establish the degree to which the youth is civically engaged.

Relationships (Ages 18 to 21)

As youth enter into young adulthood, the nature of relationships with a partner/spouse is an important determinant of overall happiness and quality of life. This information is important to collect in order to determine the impact these relationships have on the youth.

Questions asking about sexual health, pregnancy and the number of romantic partners are asked of youth aged 18 to 21 depending on their age and marital status.

Neighbourhood (Ages 16 and 17)

Neighbourhood factors have been shown to influence child and youth outcomes in a variety of domains, e.g., school achievement, behaviour, emotional and social functioning, motor and social development. These effects increase as children move through the life course, increasing their interactions and exposure to extra-familial environments. This has been evidenced in the academic literature, as well as by research conducted using NLSCY data (Boyle and Lipman, Kohen et. al, Offord and Lipman).

About Me (Ages 18 and 19)

The objectives of the questions at the beginning of this section attempt to establish the level of self esteem the youth experiences. There is an overall self-esteem score derived from the responses (FAMYfS01).

Furthermore, there are additional questions that ask about any painful events the youth may have experienced within the past two years. Included are events such as, a painful break-up with a boyfriend /girlfriend, a serious problem in school or at work, the death or illness of someone close to them, the divorce or separation of their parents, a serious money problem or any other difficult event the youth may have experienced.

The questions found in this section are similar to those found in the self-complete questionnaires given to the younger children.

Emotional Quotient (ages 20 and 21)

The emotional quotient scale was developed by Dr. Reuven BarOn and Dr. James D.A. Parker for the youth aged 20 and 21 in Cycle 6. This scale measures the degree to which the youth relates to other people at home, school and at work. Emotional intelligence involves

the ability to monitor and discriminate feelings and emotions of self and others. The respondents were asked 20 questions related to their feelings, emotions and perceptions. This young adult version is similar to the 15-item measure asked of the 10- to 17-year-olds on the self-complete questionnaires.

Youth Social Support (Ages 18 to 21)

This section is asked of the youth aged 18 years and older as part of the CAI application. The questions vary depending on the age of the youth. These questions collect information on the youth's social support network by administering the Social Support Scale taken from the adult component. This measure was added to ensure age appropriateness. These questions establish the perceived amount of support that the youth feels he or she receives from family and friends.

This section also contains some questions that ask the 20- and 21-year-olds about the number of dependent children they have and their relationship with their mother/father. These questions collect information on the family network and the role it plays in their lives.

8.6 Self-completes (Ages 10 to 17)

The objective of these questionnaires is to collect information directly from the youth on a variety of aspects of his/her life to supplement information obtained from the parent. The questionnaire also collects information from the youth on subjects about which only the youth could reliably report. For 16- and 17-year-olds, some information is still collected on the Self-completes even though these youth are reporting their own information in the CAI portion of the interview. It was felt that youth may be more comfortable answering sensitive questions on a paper questionnaire rather than face-to-face with an interviewer.

Friends and Family (Ages 10 to 17)

The objective of this section is to determine how well the youth feels he/she gets along with others.

The section collects information on the extent and quality of the youth's social support network, such as number of close friends, time spent with friends and presence of someone the youth can confide in. The questions vary depending on the age of the youth. The questions were adapted from the Ontario Child Health Study and the NLSCY Child Questionnaire.

The Friends Scale (FFFC01) is constructed from these questions.

This section also contains a measure of intimacy for the 14- and 15-year-olds. This question, about how often the youth shared secrets and private feelings with close friends, was adapted from Furman and Buhrmester's Network of Relationships Inventory.

School (Ages 10 to 15)

This section asks about the youth's attitude towards school, how well he/she is doing at school, the importance of good grades, feelings of safety and acceptance at school, perception of the teacher with respect to fairness and providing extra help. For 14- and 15-year-olds, there is a series of questions about school based extra-curricular activities, such as sports or drama. These questions have been modified by the Project Team from the Western Australia Child Health Survey, Northwest Territories Health Attitudes, Knowledge and Behaviours Study, Marsh Self-Description Questionnaire, and the World Health Organization (WHO) Survey on Health Behaviours in School Children.

Attitudes about school may be an important influence on a youth's educational accomplishments. Research shows that a negative attitude towards school may be associated with poor school performance.

About Me (Ages 10 to 17)

These questions are used to determine the youth's overall self-esteem. A score is calculated (FAMcS02) based on the answers to these questions.

For youth aged 12 to 17, additional questions are asked about youth's feelings about life now and in the future. These questions are from the Western Australia Child Health Survey.

Also included is a series of questions designed to measure "emotional intelligence". These 15 questions were designed by Dr. Reuven BarOn and Dr. James D.A. Parker. This measure is the youth version of the young adult version given to the youth aged 20 and 21. This measure was selected because it assesses the respondent's social, personal, and emotional abilities, as opposed to their behaviours.

Youth aged 14 to 17 are also asked about painful events, such as a break-up with a boyfriend/girlfriend or death of someone close to them.

Feelings and Behaviours (Ages 10 to 17)

Behaviour Checklist (Ages 10 to 15)

This section replicates the behaviour checklist used in the parent-report CAI Child Questionnaire. It provides indicators of the following behaviours: conduct disorder, hyperactivity, inattention, physical aggression, indirect aggression, emotional disorder, anxiety, and prosocial behaviours. Scores for these behaviours are also created.

Risky Behaviours (Ages 10 to 17)

These questions about risky behaviours, such as staying out all night without permission, are also replicated from the Child Questionnaire. The questions are expanded for the older age groups to capture behaviours that may become more common as the youth get older. These questions were adapted by the Project Team from the National Longitudinal Survey of Youth at Ohio State University, Western Australia Child Health Survey and from questions provided by Dr. Richard Tremblay from the University of Montreal.

Suicide (Ages 12 to 17)

This section includes questions about suicide, including whether the youth knows anyone who has committed suicide and whether they have seriously considered or attempted suicide. These questions were adapted from the 1992 British Columbia Adolescent Health Survey.

Depression (Ages 16 and 17)

Youth are asked about feelings of depression, using the same questions asked of the PMKs and the older youth aged 18 to 21. A score (FHTCbS1b) is calculated based on these questions.

My Parent(s) (Ages 10 to 17)

This section aims to capture the youth's relationship with his/her parents/guardians from several different angles. Questions are geared to uncover the amounts of understanding, fairness and affection received from each parent/guardian as well as conflict resolution practices and parental supervision. The youth's impression of his/her parents'/guardians' relationships and conflict resolution skills are also addressed.

My Parents and Me (Ages 10 to 15)

Three scales are created using these questions:

- 1) Parental Nurturance (FPMcC1),
- 2) Parental Rejection (FPMcS2b) and
- 3) Parental Monitoring (FPMcC3).

Conflict Resolution Scale (Ages 16 and 17)

These questions replicate those asked of parents of 12- to 15-year-olds. For 16- and 17-year-olds, the questions are asked separately about the youth's mother and father. Two scores are derived from these questions:

- 1) Conflict Resolution Scale – Mother (FPMCdS4), and
- 2) Conflict Resolution Scale – Father (FPMCdS5).

Puberty (Ages 10 to 17)

Puberty is an important marker of physical development. This section asks the youth about key physiological indicators and their perceptions of their own puberty. These questions were provided by Dr. Richard Tremblay from the University of Montreal.

Note: For youth aged 12 to 17, these questions are included in the Health Section.

Smoking, Drinking and Drugs (Ages 10 to 17)

This section asks questions to determine if the youth has used cigarettes, alcohol or drugs and the extent of usage. The behaviours have been correlated with negative behaviours and outcomes, such as delinquent behaviours and poor school performance. The questions vary by age.

The smoking questions are adapted from the Youth Smoking Survey, the WHO Survey on Health Behaviours in School Children and the Western Australia Child Health Survey.

The questions on alcohol were adapted from the Western Australia Child Health Survey and from questions provided by Dr. Richard Tremblay from the University of Montreal.

The questions on the use of drugs and addictive substances were adapted from the Northwest Territories Health Attitudes, Knowledge and Behaviours Study.

Questions on driving under the influence of drugs and alcohol and being a passenger in a car with a drunk driver are included for 16- and 17-year-olds. These are important risk-taking behaviours in this age group. The questions have been adapted from the North Carolina Evaluation of School-Based Health Centers.

Activities (Ages 10 to 15)

The objective is to determine the youth's extent of participation in activities outside of school hours and use of free time. Activities include sports, arts, dance or music, Guides or Scouts, jobs and volunteering. Reading for pleasure, using a computer and watching television, are also covered. Generally, the activities are also covered on the CAI parent-report Child Questionnaire for children under 10.

Literacy Activities (Ages 14 and 15)

These questions ask about how often youth engage in literacy activities outside of school, such as using a library or reading. These questions are similar to those asked of the PMK for younger children and of the 16- to 21-year-olds in the CAI questionnaire.

Health (Ages 12 to 17)

Youth are asked to report on their height and weight, symptoms of stress, use of seatbelts and helmets, healthy eating and dating. The questions vary with age.

The questions on physical indicators of stress were adapted from the WHO Survey on Health Behaviours in School Children.

Use of seatbelts and helmets questions were modified from the United States Youth Risk Behaviour Survey, which were also used in the 1992 British Columbia Adolescent Health Survey.

Questions on dating and sexual activity were adapted by the Project Team from various adolescent questionnaires such as the Minnesota Adolescent Health Survey and the 1992 British Columbia Adolescent Health Survey.

Work and Money (Ages 12 to 15)

Youth are asked about their work during the school year and those aged 14 and 15 are asked about work last summer. The 14- and 15-year-olds are asked more detailed questions about their job(s), such as hours worked and pay. They are also asked about whether work reduces the amount of time they spend studying.

Youth are also asked about how much money they received from various sources, such as parents and work.

These questions were developed by the Project Team after reviewing several other surveys.

Dating (Ages 12 to 17)

This section asks youth about their experiences with a boyfriend/girlfriend and their sexual activity. The question about sexual behaviour on the 12- and 13-year-old questionnaire was modified from the Youth and Aids Survey. Questions are also asked about contraceptive use and, for the 16- to 17-year-olds, reasons for abstaining from sex or reasons for not using birth control. These questions were designed by the Project Team in consultation with experts from youth surveys such as the 1992 British Columbia Adolescent Health Survey and the Minnesota Adolescent Health Survey.

9.0 Validation of the Survey Scales

9.1 Validation of Scale Data

9.1.1 Scale Definition

For some of the concepts deemed important to measure in the National Longitudinal Survey of Children and Youth (NLSCY) it was decided that the concept would most appropriately be measured through the use of a scale. A scale is simply a group of questions or items that measures a certain concept when the answers to the items are put together.

For example, on the Child's Questionnaire, it was determined that it was important to have an assessment of certain parenting behaviours. The scale is intended to measure three different constructs or factors related to parenting; positive interaction, ineffective parenting and consistent parenting.

9.1.2 Scales and Calculations

For each factor measured by a scale, a score is calculated. The score for a particular factor can be used to give an ordering of individuals. For example, for the Parenting Scales, for children with higher scores for the “positive interaction” factor, the person most knowledgeable (PMK) reported having more positive encounters with the child, e.g., laughed with them more, praised them more. The score for a particular factor is usually based on a series of items, since one single item usually cannot measure the construct or factor with adequate precision.

During the development of the NLSCY, when consideration was being given to what scales should be used to measure a particular concept, an attempt was made to select scales that had been used in other studies. In this way, the psychometric properties of the measures produced by each scale were available with complete references.

9.1.3 Evaluation of Scale Data

In many instances, the wording of certain questions in the original scale was modified and in some cases new questions were added. Sometimes the scale that was used had not previously been used for children in Canada, or had only been used for very small samples. Given these concerns and further concerns regarding interviewing conditions, it was felt that the factor structures of the scales used in the NLSCY could be different from the ones given in the literature. Therefore the Project Team felt the need to carry out an extensive evaluation of the scale data to ensure that the psychometric properties found in other studies also held true for the NLSCY experience.

There were three major steps in the analysis of the scale data. First a new factor analysis was performed on all scales to determine the constructs or factors inherent in each scale. Then scale scores were calculated based on this factor structure. Finally reliability measures were produced. The general procedures followed for each of these steps are described in detail in the following pages.

Note: Many of the scales were developed and validated in Cycle 1. In subsequent cycles, the same factor structure which emerged from the Cycle 1 analysis was imposed. Imposing the same factor structure ensures that the scales are

consistent across time to allow for longitudinal analysis and cross-sectional comparisons. Each scale has a note indicating in which cycle the validation was performed.

9.2 Factor Analysis

9.2.1 Factor Analysis for Scales

The factor structure of each scale was determined based on data from the first cycle. The factor structure imposed on the scales already used in the first cycle and repeatedly utilized in subsequent cycles of the survey was the result of analysis of data from the first cycle. For detailed results from the Cycle 1 factor analysis, please refer to the Cycle 1 Microdata User Guide.

1. The sample of respondents for each scale (and age group, if the scale used different questions for different age groups), was randomly divided into two half-samples. This was done to find out whether different samples would yield the same results.
2. Principal component analysis was carried out separately on each half-sample to find out how many factors should be extracted in the subsequent factor analysis. In principle, the same number of factors as was found in the literature was expected. In practice, however, some scales showed a different number of factors because in some cases factors combined while in others new factors emerged.
3. Factor analysis was done on each half-sample and the factor structure and loading of each factor were compared across the half-samples.
4. In the factor analysis, the items for each child in the appropriate age group were used, multiplied by the child's normalized weight. An individual's statistical weight is normalized by dividing his/her weight (`_WTCW01C`)⁷ by the average weight for all individuals. Thus, the sum of the normalized weights is equal to the sample size.
5. Once the factor structures were analysed and the items included in each factor were determined, scores were calculated. To produce the scores, one was subtracted from each item so that the lowest possible score would be zero (0). A score of zero indicates that the child has no problems for all factors in the Behaviour Scale except for the prosocial factor, where a score of zero indicates the absence of prosocial behaviour. Some items were imputed. The imputed values were computed by a procedure (the SAS PRINQUAL procedure) that determines which of the possible values for an item is the most plausible for an individual in view of his/her response profile, the response profiles of others in the sample, and the number of factors included in the analysis.
6. The score for each factor on the scale was derived by totaling the values of the items that made up that factor (including imputed values). The score was set to "missing" if too many of the values of any items included in the factor were unreported. A value may be missing if the parent refused to answer or did not know the answer to the item.

⁷ In this chapter, an underscore “_” is used at the beginning of each variable name rather than a letter indicating a specific cycle. For example, the variable name FPRCS01 in Cycle 6 begins with the letter "F" on the microdata file and here is referred to as `_PRCS01`.

9.2.2 Data Transformation Using Optimal Scaling

Before performing the factor analysis for each of the NLSCY scales, the data were transformed using optimal scaling. The method used was one proposed by Young and several associates (Young, 1981) which is a variant of Fisher's optimal scaling technique. The method is presented as a means of transforming data that are fundamentally nominal or ordinal in nature to interval or ratio level data so that statistical techniques which are appropriately applied only to interval and ratio data may be utilized.

9.2.3 Factor Analysis Using Weighted Data

Factor analysis requires that the data have the property of interval or ratio data, meaning that the distance between each answer category of the question should be the same. For example, in scales where the answer choices are: "Never", "Sometimes", "Often" and "Always", one must assume that the distance between "Never" and "Sometimes" is the same as that between "Sometimes" and "Often" in the respondent's perception. It was felt that this was not necessarily true in the case for the scales used in the NLSCY.

9.3 Calculation of Scores and Item Imputation

9.3.1 Calculation of Scores for Each Factor

The results of the factor analysis were used to determine which items "loaded" into each factor, i.e., were a part of each factor. The next step was to calculate a score for each factor. This was done by summing the values for each individual item that made up the factor. In some cases some rescaling of values was done before the final score was calculated. The following example illustrates how factor scores were computed.

9.3.2 Example of Factor Score Computation

One of the constructs that emerged in the factor analysis for the Parenting Scale on the Child's Questionnaire was the ineffective parenting factor (aged 2 to 11 years). In the factor analysis on Cycle 1 data seven items were found to load into this factor.

_PRCQ04	How often do you get annoyed with your child for saying or doing something he/she is not supposed to?
_PRCQ08	Of all the times you talk to your child about his/her behaviour, what proportion is praise?
_PRCQ09	Of all the times you talk to your child about his/her behaviour, what proportion is disapproval?
_PRCQ13	How often do you get angry when you punish your child?
_PRCQ14	How often do you think the kind of punishment you give your child depends on your mood?
_PRCQ15	How often do you feel you have problems managing your child in general?

_PRCQ18 How often do you have to discipline your child repeatedly for the same thing?

The answer categories for these items were of two types:

- 1 Never
- 2 About once a week or less
- 3 A few times a week
- 4 One or two times a day
- 5 Many times each day

- 1 Never
- 2 Less than half the time
- 3 About half the time
- 4 More than half the time
- 5 All the time

In the calculation of the score for this ineffective parenting factor, the categories were rescaled to 0 to 4, i.e., the category "Never" was scored as 0, the category "About once a week or less/Less than half the time" was scored as 1, ... and the category "Many times each day/All the time" was scored as 4. In order to compute the score, these values were summed across the seven items involved in the factor resulting in an ineffective parenting score in the range 0 to 28. A low score of zero represents the absence of a problem and a high score of 28 indicates a high degree of problems. For most of the scores calculated for the NLSCY, a score of zero represents the absence of a problem. However there are exceptions to this which are noted in the documentation for each particular scale.

9.3.3 Negative Loading

Note that the second item that loaded into the ineffective parenting factor, _PRCQ08 (Of all the times you talk to your child about his/her behaviour, what proportion is praise?) is in the opposite direction compared to the other items. In fact, the item loaded "negatively" into the factor. Therefore, when computing the score the values for this item were reversed, i.e., "All the time" was scored as 0, "More than half the time" as 1, ... and "Never" as 4. In the documentation for each scale any item that was reversed for the scoring algorithm due to a negative loading is indicated.

9.3.4 Nonresponse Codes

The score for the ineffective parenting factor is labelled as _PRCS04 on the record layout for the microdata file. An "S" in the fifth position of the variable name indicates a score.

When the score was being calculated for each factor there was a possibility that one or more of the items making up the score had a nonresponse code ("Don't know", "Refusal" or "Not stated"). If any of the items had a nonresponse code, the factor score was set to "Not stated".

9.3.5 Raw Items

It should be noted that in addition to the scores, the raw items for each scale are included on the microdata file. This will allow researchers to consider alternate factor

structures if desired. For the raw items the original values (in the 1 to 5 range for the Parenting Scale) have been retained before any rescaling or reversal of values took place.

9.4 Reliability Measures for Scales

Reliability refers to the accuracy, dependability, consistency or ability to replicate a particular scale. In more technical terms, reliability refers to the degree to which the scale scores are free of measurement error. There are many ways to measure reliability.

9.4.1 Cronbach's Alpha

One of the most commonly used reliability coefficients is Cronbach's alpha (Cronbach, 1951). Alpha is a measure of the internal consistency of the items within the factor. It is based on the average covariance of items within the factor. It is assumed that items within a factor are positively correlated with each other because they are attempting to measure, to a certain extent, a common entity or construct.

9.4.2 Interpretations of Cronbach's Alpha

Cronbach's alpha has several interpretations. It can be viewed as the correlation between the scale or factor and all other possible scales containing the same number of items, which could be constructed from a hypothetical universe of items that measure the characteristic of interest. For example, in the ineffective parenting factor, the seven questions included in the scale can be viewed as a sample from the universe of all possible items. Parents could also have been asked: "How often do you raise your voice when you discipline your child?" or "How often do you threaten punishment more often than you use it?" Cronbach's alpha indicates how much correlation can be expected between the scale which was used and all other possible seven-item scales measuring the same thing.

Another interpretation of Cronbach's alpha is the squared correlation between the score an individual obtains on a particular factor (the observed score) and the score he/she would have obtained if questioned on all possible items in the universe (the true score). Since alpha is interpreted as a correlation coefficient, it ranges from 0 to 1.

Generally, it has been shown that alpha is a lower bound to the reliability of a scale of n items (Novick and Lewis, 1967). In other words, in most situations alpha provides a conservative estimate of a score's reliability.

9.4.3 What is a Satisfactory Level of Reliability?

It is difficult to specify a single level that should apply in all situations. Some researchers believe that reliabilities should not be below 0.8 for widely used scales. At that level, correlations are affected very little by random measurement error. At the same time, it is often very costly in terms of time and money to obtain a higher reliability coefficient. It should be noted that for some of the factors for which scores were computed for the NLSCY, the reliabilities are below this level. The Cronbach's alpha is given in the documentation for each score that has been calculated. Researchers can determine for themselves whether or not the score has adequate reliability for their specific purposes.

Finally, it should be mentioned that for the NLSCY the Cronbach's alpha for each factor score was computed using SAS. Typically, the alpha coefficients calculated using SAS are lower than those calculated using SPSS.

9.5 Parent-reported Scales

The remainder of this chapter provides an in-depth description of the sources of the NLSCY scales and all analytical results of factor and reliability analysis. Changes made to the scales across cycles are also described. The scales are listed in the order they appear in the questionnaire.

9.5.1 Depression Rating Scale

Objectives and Overview

The Depression Rating Scale was administered to the PMK as part of the Parent Questionnaire. Questions for this scale (_DPPQ12A to _DPPQ12L) are a shorter version of the Depression Rating Scale (CES-D), comprising 20 questions, developed by L. S. Radloff of the Epidemiology Study Center of the National Institute of Mental Health in the United States. This rating scale is used to measure the frequency of symptoms in the public at large. The occurrence and severity of symptoms associated with depression during the previous week are measured. The rating scale was reduced to 12 questions by Dr. M. Boyle of the Chedoke-McMaster Hospital, McMaster University.

This rating scale is aimed at gathering information about the mental health of respondents, with particular emphasis on symptoms of depression. Several members of the NLSCY advisory group of experts pointed out that the best way of proceeding was to measure one particular aspect of the PMK's mental health instead of trying to measure overall mental health. It was proposed that this section focus on depression for the following reasons: depression is a prevalent condition; it has been demonstrated that depression in a parent affects the children; present research on this subject is generally based on demonstration groups and not on population samples; and it is felt that introducing policies in this area could make a difference.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
PMK Depression (_DPPS01)	_DPPQ12A, _DPPQ12B,	3,178	0 to 1	7.2 to 8.0%	0.812
	_DPPQ12C, _DPPQ12D,	3,312	2 to 3	3.8 to 4.0%	0.835
	_DPPQ12E, _DPPQ12F*,	3,364	4 to 5	4.0 to 4.2%	0.827
	_DPPQ12G, _DPPQ12H*,	2,764	10 to 11	1.7 to 1.8%	0.857
	_DPPQ12I, _DPPQ12J*,	1,950	12 to 13	1.7 to 2.0%	0.859
	_DPPQ12K, _DPPQ12L	1,756	14 to 15	1.5 to 1.6%	0.862

**Indicates that the values have been reversed.*

9.5.2 Home Responsibilities Scale

Objectives

The objective of the activities scale is to measure the child's participation in home responsibilities. This set of questions is from the Home Observation for Measurement of

the Environment-Short Form questionnaire in the National Longitudinal Survey of Youth, Ohio State University.

In Cycle 1, these questions were only asked of 10- and 11-year-olds as they were the eldest age group. In subsequent cycles these questions were asked of all children aged 10 to 13.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Home Responsibilities (_ACCS6)	_ACCQ6A, _ACCQ6B, _ACCQ6C,	2,811	10 to 11	0.4%	0.794
	_ACCQ6D, _ACCQ6E	1,988	12 to 13	0.4%	0.794

9.5.3 Family Functioning Scale

Objectives and Overview

Questions related to family functioning, i.e., _FNHQ01A to _FNHQ01L, were developed by researchers at the Chedoke-McMaster Hospital, McMaster University and have been used widely both in Canada and abroad. This scale is used to measure various aspects of family functioning, e.g., problem solving, communications, roles, affective involvement, affective responsiveness and behaviour control.

Question _FNHQ01M, drawn from the Follow-up to the Ontario Child Health Study, was added to the original scale to determine whether alcohol consumption had an effect on global family dynamics. However, it was not used in the analysis of the scale.

This scale is aimed at providing a global assessment of family functioning and an indication of the quality of the relationships between parents or partners. For this reason and because of the small number of questions, no attempt was made to measure the various aspects of family functioning.

Other surveys have shown that the relationship between family members has a considerable effect on children. The results of the Ontario Child Health Study have shown, for example, that there is an important link between family dysfunction and certain mental conditions in children.

Administering the Family Functioning Scale

The Family Functioning Scale was administered to either the PMK or the spouse/partner as part of the Parent Questionnaire. The scale includes 12 questions, each of which contains four response categories. In order for the lowest score value to be zero, the value of the categories was reduced by one in calculating the score. The order of the categories was reversed for questions having a negative loading (_FNHQ01A, _FNHQ01C, _FNHQ01E, _FNHQ01G, _FNHQ01I, and _FNHQ01K). The total score (_FNHS01) may therefore vary between 0 and 36, a high score indicating family dysfunction.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Family Functioning (_FNHS01)	_FNHQ01A*,	3,232	0 to 1	4.7 to 5.8%	0.915
	_FNHQ01B,	3,353	2 to 3	2.1 to 2.3%	0.909
	_FNHQ01C*,	3,439	4 to 5	1.5 to 1.8%	0.910
	_FNHQ01D,	2,775	10 to 11	1.0 to 1.1%	0.916
	_FNHQ01E*,_FNHQ01F,	1,968	12 to 13	0.3 to 0.6%	0.907
	_FNHQ01G*,	1,764	14 to 15	0.6 to 0.7%	0.906
	_FNHQ01H,_FNHQ01I*, _FNHQ01J,_FNHQ01K*				

9.5.4 Neighbourhood Safety Scale

Objectives and Overview

To gather information on the respondent's satisfaction with his/her neighbourhood as a place to raise children, including perception of the extent of danger and problems, and of social cohesion or "neighbourliness". Research by Dr. Jacqueline Barnes at the Judge Baker Children's Centre, Harvard University in Boston has found that parents' fear of danger and perception of social disorder in the neighbourhood affected their sense of attachment to the neighbourhood and their disciplinary strategies.

Questions _SFHQ01, _SFHQ02 and _SFHQ05A to _SFHQ06E cover the length of residency in the neighbourhood, satisfaction with the neighbourhood as a place to bring up children, safety, social cohesion and neighbourhood problems. They represent a revised version of specific sections of the Simcha-Fagan Neighbourhood Questionnaire used by Dr. Jacqueline Barnes in her studies of neighbourhoods in Boston and Chicago. Revisions were made based on the factor analysis of the sections, in consultation with Dr. Barnes. Question _SFHQ03 on volunteer involvement is based on a question in the National Population Health Survey.

Changes to Neighbourhood Section Across Cycles

These scales have been used intermittently over the four cycles of the NLSCY. In Cycle 1, three scales were created: neighbourhood safety (_SFHQ05A to _SFHQ05C), neighbours (_SFHQ06A to _SFHQ06E) and neighbourhood problems. The entire Neighbourhood section was not asked of survey participants in Cycle 2. In Cycle 3, the Neighbourhood section was reintroduced without questions _SFHQ05A to _SFHQ05C and without questions ASFHQ07A to ASFHQ07F.

The Cycle 6 scale questions are the same as the Cycle 1 questions with the exception of _SFHQ05C where there has been a small wording change. Also, the questions that made up the neighbourhood problems scale in Cycle 1 (ASFHQ07A to ASFHQ07F) were dropped after Cycle 4.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Safety Score (_SFHS5)		3,251	0 to 1	4.8 to 5.0%	0.673
		3,354	2 to 3	2.0 to 2.1%	0.672
	_SFHQ05A,	3,436	4 to 5	1.5 to 1.5%	0.682
	_SFHQ05B,	2,759	10 to 11	0.8 to 0.9%	0.704
	_SFHQ05C	1,973	12 to 13	0.3 to 0.3%	0.702
		1,765	14 to 15	0.5 to 0.6%	0.719
Neighbours Score (_SFHS6)		2,674	0 to 1	5.2 to 5.6%	0.889
	_SFHQ06A,	2,918	2 to 3	2.1 to 2.5%	0.897
	_SFHQ06B,	3,021	4 to 5	1.6 to 1.9%	0.901
	_SFHQ06C,	2,498	10 to 11	0.9 to 1.4%	0.910
	_SFHQ06D,	1,768	12 to 13	0.3 to 0.8%	0.906
	_SFHQ06E	1,586	14 to 15	0.6 to 0.9%	0.900

9.5.5 Social Support Scale

Objectives and Overview

This section is asked of all PMKs with children/youth less than 16 years of age. The original scale contains 24 items from Robert Weiss's Social Provisions Model that describes six different social functions or "provisions" that may be acquired from relationships with others. Due to the length of the scale, and on the advice of Dr. M. Boyle at Chedoke-McMaster Hospital, McMaster University, the survey uses the shortened version (containing six items) that was derived for the Government of Ontario's, Better Beginnings, Better Futures Project. This measures guidance (two questions), reliable alliance (two questions) and attachment (two questions). Furthermore, in Cycle 1, four additional questions on different types of social support, i.e., religious and community services, were added as suggested by Dr. Tom Hay. These questions were not included for Cycle 3, however, due to a lack of variability in response. Questions similar to those suggested by Dr. Hay were taken from the Family Crisis Oriented Personal Evaluation Scales (F-COPES) and included in Cycle 4 and Cycle 5. F-COPES draws upon the coping dimensions of the Resiliency Model of Family Adjustment and Adaptation (McCubbin, Olson & Larsen: 1981). The total social support measurement includes six questions and not only focuses on the quantity of social support but on the quality of social supports as well.

In Cycle 2, the entire Social Support section was dropped due to a belief that there would be little temporal variation in the amount individuals received and concerns regarding response burden.

Changes to Social Support Section Across Cycles

In Cycle 4, the following changes were made to the Social Support section and these changes were kept for subsequent cycles:

- The original six items used in Cycle 1 and Cycle 3 were kept; however, items _SPHQ02A to _SPHQ02D used in Cycle 1 were replaced by the F-COPES items.
- Two additional questions from the above mentioned social integration sub-scale (items _SPHQ01H and _SPHQ01I) were also added. The questions on social integration are significant because they assess one's feeling of belonging to a

group that shares similar interests, concerns, and activities thus measuring another factor of social support.

- Four supplementary questions from the F-COPES were added, as well as one question based on the F-COPES framework that all centre on the same reasoning as those questions used in Cycle 1. However, the suggested questions steer away from the simple "Yes" and "No" responses that fail to indicate variability and instead use the response categories of "Strongly disagree", "Disagree", "Agree" and "Strongly Agree".

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Social Support (_SPHS01)	_SPHQ01A*	3,270	0 to 1	4.7 to 5.7%	0.906
	_SPHQ01B,	3,362	2 to 3	2.1 to 2.2%	0.898
	_SPHQ01C,	3,450	4 to 5	1.6 to 1.8%	0.901
	_SPHQ01D*,	2,785	10 to 11	0.8 to 0.9%	0.909
	_SPHQ01E*,	1,979	12 to 13	0.3 to 0.4%	0.912
	_SPHQ01F,	1,777	14 to 15	0.5 to 0.5%	0.915
	_SPHQ01H, _SPHQ01I*.				

*Indicates that the values have been reversed.

9.5.6 Behaviour Scales

Objectives

The objective of the Behaviour Scales is to assess aspects of the behaviour of children 2 to 11 years of age.

Separation Anxiety (Aged 2 to 3 years)

Includes items _BECQ6CC, _BEC6QDD1, _BECQ8LL1, _BECQ8PP1 and _BEC8TT1 from Achenbach's Child Behaviour checklist (CBCL).

Opposition (Aged 2 to 3 years)

Includes items _BECQ6G, _BECQ6R1, _BECQ8E1, _BECQ8T1, _BECQ8Z1 and _BECQ8J1 drawn from Achenbach's CBCL.

Conduct Disorder (Aged 10 to 11 years)

Includes items _BECQ6AA, _BECQ6FF, _BECQ6JJ and _BECQ6NN from the Ontario Child Health Study (OCHS).

Hyperactivity (Aged 2 to 5 and 10 to 11 years)

Includes items _BECQ6B, _BECQ6I, _BECQ6P and _BECQ6W from the OCHS and _BECQ6QQ and _BECQ8HH from the Montreal Longitudinal Survey. In previous cycles, item _BECQ6N was included in this construct. A decision was made to drop this item from Cycle 4 and all future cycles as respondents found it to be too repetitive.

Emotional Disorder and Anxiety (Aged 2 to 5 and 10 to 11 years)

Includes items _BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM and _BECQ6RR from the OCHS. Anxiety includes NLSCY items taken from OCHS emotional disorder items (_BECdQ6F, _BECQ6Q, _BECQ6V and _BECQ6CC). In previous cycles the items _BECQ6Y and _BECQ6II were included. A decision was made to remove both items from Cycle 4 and all future cycles.

Indirect Aggression (Aged 10 to 11 years)

Includes items _BECQ6J, _BECQ6R, _BECQ6Z, _BECQ6LL and _BECQ6TT from Lagerspetz, Bjornqvist and Peltonen of Finland.

Physical Aggression (Aged 2 to 5 and 10 to 11 years)

Includes items _BECQ6X from the Montreal Longitudinal Survey and _BECQ6G, _BECQ6AA and _BECQ6NN from the OCHS.

Inattention (Aged 2 to 5 and 10 to 11 years)

Includes items _BECQ6P from the OCHS and _BECQ6QQ from the Montreal Longitudinal Survey.

Prosocial behaviour (Aged 10 to 11 years)

Includes items _BECQ6A, _BECQ6H, _BECQ6M, _BECQ6GG and _BECQ6OO from the OCHS and _BECQ6D, _BECQ6U, _BECQ6BB, _BECQ6SS and _BECc6UU from the Montreal Longitudinal Survey; the last four items are from a scale devised by K. Weir and G. Duveen. In Cycles 1 through 3, these items were asked of all children aged 4 to 11. In Cycle 4, all 4- and 5-year-olds were excluded from this scale and were asked the questions in the positive behaviour section instead.

Overview of Behaviour Scales for 2- and 3-year-olds

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Hyperactivity - Inattention (_BECdS01)	_BECQ6B, _BECQ6I, _BECQ6P, _BECQ6W, _BECQ6QQ, _BECQ8HH	3,391	2 to 3	0.55 to 0.84%	0.730
Emotional Disorder - Anxiety (_BECdS03)	_BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM, _BECQ6RR	3,407	2 to 3	0.61 to 0.87%	0.602
Physical Aggression - Opposition (_BECs04)	_BECQ6G, _BECQ6X, _BECQ6NN, _BECQ6R1, _BECQ8E1, _BECQ8T1, _BECQ8Z1, _BECQ8J1	3,376	2 to 3	0.61 to 0.81%	0.727
Separation Anxiety (_BECs05)	_BECQ6CC, _BECQ6DD1, _BECQ8LL1, _BECQ8PP1, _BEC8TT1	3,421	2 to 3	0.78 to 0.89%	0.566

Overview of Behaviour Scales for 4- and 5-year-olds

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Hyperactivity - Inattention (_BECDS06)	_BECQ6B, _BECQ6I, _BECQ6P, _BECQ6S, _BECQ6W, _BECQ6QQ, _BECQ8HH	3,471	4 to 5	0.96 to 1.19%	0.787
Emotional Disorder - Anxiety (_BECDS08)	_BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM, _BECQ6RR, _BECQ6CC	3,475	4 to 5	0.88 to 1.28%	0.665
Physical Aggression - Conduct Disorder (_BECDS09)	_BECQ6G, _BECQ6X, _BECQ6AA, _BECQ6FF, _BECQ6JJ, _BECQ6NN	3,483	4 to 5	0.88 to 0.93%	0.764
Indirect Aggression (_BECs10)	_BECQ6J, _BECQ6R, _BECQ6Z, _BECQ6LL, _BECQ6TT	3,409	4 to 5	1.56 to 2.04%	0.640

Overview of Behaviour Scales for 10- and 11-year-olds

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Prosocial Behaviour (_BECdS07)	_BECQ6A, _BECQ6D, _BECQ6H, _BECQ6M, _BECQ6U, _BECQ6BB, _BECQ6GG, _BECQ6OO, _BECQ6SS, _BECc6UU	2,647	10 to 11	0.9 to 2.9%	0.841
Property Offences (_BECdS11)	_BECQ6C, _BECQ6E, _BECQ6L, _BECQ6T, _BECQ6DD, _BECQ6PP	2,784	10 to 11	0.6 to 1.1%	0.608
Hyperactivity - Inattention (_BECdS06)	_BECQ6B, _BECQ6I, _BECQ6P, _BECQ6S, _BECQ6W, _BECQ6QQ, _BECQ8HH	2,782	10 to 11	0.6 to 1.1%	0.847
Emotional Disorder - Anxiety (_BECdS08)	_BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM, _BECQ6RR, _BECQ6CC	2,787	10 to 11	0.6 to 1.0%	0.776
Physical Aggression - Conduct Disorder (_BECdS09)	_BECQ6G, _BECQ6X, _BECQ6AA, _BECQ6FF, _BECQ6JJ, _BECQ6NN	2,779	10 to 11	0.7 to 1.2%	0.795
Indirect Aggression (_BECs10)	_BECQ6J, _BECQ6R, _BECQ6Z, _BECQ6LL, _BECQ6TT	2,565	10 to 11	2.0 to 6.1%	0.798

9.5.7 Motor and Social Development Scale

Motor and Social Development Section

The Motor and Social Development Section of the Child's Questionnaire was completed for children in the 0 to 3 age group. The objective was to measure motor, social and cognitive development of young children. A scale made up of 48 questions (_MSCQ01 to _MSCQ48), was used to assess these concepts. According to the age in months, 15 questions were asked of each child.

The Motor and Social Development Scale

The Motor and Social Development (MSD) Scale was developed by Dr. Gail Poe of the United States National Center for Health Statistics. The MSD Scale consists of a set of 15 questions that vary by the age of the child, which measure dimensions of the motor, social and cognitive development of young children from birth to age 3. Each item asks whether or not a child is able to perform a specific task. The scale has been used in collections of the National Longitudinal Survey of Youth in the United States and in the National Child Development Survey in England.

The following table shows which questions were asked to each age group.

Age in Months	MSD Questions
0 to 3	_MSCQ01 to _MSCQ15
4 to 6	_MSCQ08 to _MSCQ22
7 to 9	_MSCQ12 to _MSCQ26
10 to 12	_MSCQ18 to _MSCQ32
13 to 15	_MSCQ22 to _MSCQ36
16 to 18	_MSCQ26 to _MSCQ40
19 to 21	_MSCQ29 to _MSCQ43
22 to 47	_MSCQ34 to _MSCQ48

Raw Scores

A Raw Score was calculated for each child by summing the number of "Yes" answers to each item in the scale (_MSCS01). Due to a problem with the application question 26 (_MSCQ26) was not asked of the children aged 7 and 8 months. As a result these children have a Raw Score that has a maximum of 14. Using data from previous cycles it was noted that in at least 93% of cases children of these ages responded "No" to this question. As well, the children who would have responded "Yes" would still have the highest scores for this age group even without taking that question into consideration. Therefore no adjustment was done to compensate for this error.

Although there were different sets of questions depending on the age in months of the child, differences were observed when comparing score within these age bands. For example, there was a specific set of questions for children aged 4 to 6 months. It was found that children who were 6 months old had scores that were on average higher than those 4 months old. Therefore a decision was made to produce standardized scores. These scores, calculated for each age in months, would make it possible to compare scores across ages. All children, aged 3 years old or less, received a standardized score based on Cycle 1 data and a standardized score based on the Cycle 6 data.

Standardized Scores based on Cycle 6 norms

Each child aged 3 to 47 months was assigned a standard score. This standardization was done by one month age groups. For each month age group the mean and standard deviation of the raw score was found and was used to produce a normalized score with a mean of zero and a standard deviation of one. This score was adjusted such that the mean MSD score was 100 and the standard deviation was 15. Therefore, children who are 3 months old have an average MSD score of 100, children who are aged 4 months have an average MSD score of 100, and children aged 47 months have an average MSD score of 100.

Once these scores were calculated, children who were more than three standard deviations away from the mean (scores smaller than 55 or greater than 145) were identified, and the norms were recalculated not including these children. These children were considered outliers and are not representative of other children their age. Therefore the average of the MSD scores on the data file by age in months may not be exactly 100. Using this standardized score (_MSCdS03) makes it possible to compare scores of children across the 3- to 47-month-old age group, without having to control for age.

This score was not calculated for children aged 0 to 2 months old as there were not enough respondent children by age in months to establish a norm.

Standardized Scores based on Cycle 1 norms

A second standardized score (_MSCS02) was calculated for all children 0 to 47 months. This score was calculated in the same way as mentioned above, except that the norms were derived using the data from Cycle 1, and then applied to the Cycle 6 data.

Overall there are no major differences found when comparing the scores found using Cycle 6 norms and the scores found using Cycle 1 norms. The score calculated using Cycle 1 norms should be used in order to compare scores over cycles. This score is available for all cycles of data.

The Motor and Social Development Scale questions have remained unchanged throughout the six cycles of the National Longitudinal Survey of Children and Youth, but there have been changes to the calculation of the final scores. For more information on these changes please refer to the Appendix on Revisions to Previous Cycles, in the Cycle 4 Microdata User Guide.

9.5.8 Parenting Scales

Objectives and Overview

The objective of this scale is to measure certain parenting practices. Specifically, two scales were used. The first was designed to measure the positive interaction, ineffectiveness and consistency of the parenting of the child. The second scale was designed to measure parental practices that may or may not provoke aversion.

The questions from the Child's Questionnaire used to measure these aspects of parenting are identified in the following paragraphs. A complete factor analysis was done on the parenting scales to evaluate the psychometric properties of these scales for the NLSCY population.

Questions _PRCQ01 to _PRCQ18 and _PRCQ21 to _PRCQ24 on positive interaction, ineffectiveness and on coherence were provided by Dr. M. Boyle of the Chedoke-McMaster Hospital, McMaster University, based on the work of Dr. Ken Dodge (Vanderbilt University) and an adaptation of the Parent Practices Scale of Strayhorn and Weidman.

Analytical Results for 0- to 1-year-olds for the Parenting Scales

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Positive Interaction (_PRCS01)	_PRCQ01, _PRCQ02, _PRCQ03, _PRCQ06, _PRCQ07	3,393	0 to 1	2.1 to 3.4%	0.638
Ineffective Parenting (_PRCS02)	_PRCQ04, _PRCQ05	3,409	0 to 1	2.5 to 2.9%	0.367

Analytical Results for 2- to 5-year-olds for the Parenting Scales

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Positive Interaction (_PRCS03)	_PRCQ01, _PRCQ02, _PRCQ03, _PRCQ06, _PRCQ07	3,375	2 to 3	2.4 to 2.5%	0.703
		3,454	4 to 5	2.1 to 2.1%	0.703
Ineffective Parenting (_PRCS04)	_PRCQ04, _PRCQ08*, _PRCQ09, _PRCQ13, _PRCQ14, _PRCQ15, _PRCQ18	3,323	2 to 3	2.5 to 3.2%	0.680
		3,416	4 to 5	2.2 to 2.6%	0.678
Consistent Parenting (_PRCS05)	_PRCQ10, _PRCQ11, _PRCQ12*, _PRCQ16*, _PRCQ17*	3,260	2 to 3	2.7 to 3.3%	0.668
		3,374	4 to 5	2.3 to 2.6%	0.665
Rational Parenting (_PRCS06)	_PRCQ21, _PRCQ22*, _PRCQ23, _PRCQ24*	3,354	2 to 3	2.8 to 3.1%	0.545
		3,442	4 to 5	2.4 to 2.4%	0.566

*Indicates items were reversed.

Analytical Results for 10- to 11-year-olds for the Parenting Scales

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Positive Interaction (_PRCS03)	_PRCQ01, _PRCQ02, _PRCQ03, _PRCQ06, _PRCQ07	2,769	10 to 11	1.8 to 1.9%	0.721
Ineffective Parenting (_PRCS04)	_PRCQ04, _PRCQ08*, _PRCQ09, _PRCQ13, _PRCQ14, _PRCQ15, _PRCQ18	2,750	10 to 11	1.8 to 2.1%	0.733
Consistent Parenting (_PRCS05)	_PRCQ10, _PRCQ11, _PRCQ12, _PRCQ16, _PRCQ17	2,658	10 to 11	2.2 to 5.2%	0.686
Rational Parenting (_PRCS06)	_PRCQ21, _PRCQ22*, _PRCQ23, _PRCQ24*	2,765	10 to 11	2.1 to 2.1%	0.536
Conflict Resolution (_PRCS09)	_PRCB30A*, _PRCB30B, _PRCB30C, _PRCB30D, _PRCB30E, _PRCB30F, _PRCB30G, _PRCB30H*	1,957	12 to 13	1.5 to 1.9%	0.747
		1,751	14 to 15	1.6 to 1.8%	0.730

*Indicates items were reversed.

9.5.9 Ages and Stages

The Ages and Stages component was generated for all children 3 to 47 months. The questions were grouped into the four categories listed below with each respondent receiving a score in the range of 0 to 60. For this measure, a high score indicates that the child is at or above the normal range for their age group. For more information about this measure please refer to Chapter 8.0 or contact Brookes Publishing Co. and Health Professions Press for a copy of the individual items.

Factor	Score	Range of Scores
Problem solving score	_AGCdS01	0 to 60
Personal score	_AGCdS02	0 to 60
Communication score	_AGCdS03	0 to 60
Fine motor score	_AGCdS04	0 to 60

9.6 Youth-reported Scales

9.6.1 Depression Scale

Overview

The depression scale used to measure PMK depression was also used for youth 16 years of age and older. For 16- and 17-year-olds the questions were asked in the self-completed paper questionnaire and for the youth 18 and older the questions were asked as part of their computer-assisted interview (CAI) questionnaire. The factor structure that was used for the PMK scale was also imposed on the youth scale.

The total score (_HTCbS1B and _HTYfS01) may therefore vary between 0 and 36, a high score indicating the presence of depression symptoms.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Depression (_HTCbS1B)	_FBCd10A, _FBCd10B, _FBCd10C, _FBCd10D, _FBCd10E, _FBCd10F*, _FBCd10G, _FBCd10H*, _FBCd10I, _FBCd10J*, _FBCd10K, _FBCd10L	1,176	16 to 17	0.5 to 0.5%	0.602
Depression (_HTYfS01)	_HTYf14A, _HTYf14B, _HTYf14C, _HTYf14D, _HTYf14E, _HTYf14F*, _HTYf14G, _HTYf14H*, _HTYf14J, _HTYf14K*, _HTYf14M, _HTYf14N	1,567	18 to 19	0.0 to 0.3%	0.811
		1,400	20 to 21	0.0 to 0.3%	0.811

*Indicates items were reversed.

9.6.2 Neighbourhood Structure Scale

Objectives and Overview

To gather information on the respondent's satisfaction with his/her neighbourhood, including perception of the extent of danger and problems, and of social cohesion or "neighbourliness". These questions are asked of 16- and 17-year-olds in the Youth Questionnaire.

The items included in the score represent a revised version of specific sections of the Simcha-Fagan Neighbourhood Questionnaire used by Dr. Jacqueline Barnes in her studies of neighbourhoods in Boston and Chicago.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Neighbourhood Structure (_ACYDS01)	_ACYD13A, _ACYD13B, _ACYD13C, _ACYD13D, _ACYD13F, _ACYD13G	1,359	16 to 17	9.0 to 13.4%	0.829

9.6.3 General Self-image Scale

The objective of the General Self-image Scale is to measure the youth's overall self-esteem. The self-esteem scale is asked of youth 10 to 17 in the self-completed paper questionnaire and of youth 18 to 19 years of age in the computer-assisted personal interview (CAPI) questionnaire. These questions on overall self-esteem were taken from the General Self-image Scale of the Marsh Self-description Questionnaire developed by H.W Marsh.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
General Self-Image (_AMCS02)	_AMCQ01A, _AMCQ01B, _AMCQ01C, _AMCQ01D	2,459	10 to 11	0.2 to 0.2%	0.792
		1,707	12 to 13	0.2 to 0.3%	0.823
		1,507	14 to 15	0.3 to 0.3%	0.844
		1,199	16 to 17	0.4 to 0.4%	0.862
General Self-Image (_AMYfS01)	_FAMYfQ01, _FAMYfQ02, _FAMYfQ03, _FAMYfQ04	1,568	18 to 19	0.3 to 0.3%	0.746

9.6.4 Emotional Quotient

Developed by Dr. Reuven BarOn and Dr. James D.A. Parker, the Emotional Quotient Adult Version (EQ-i) and the Emotional Quotient Inventory Youth Version (EQ-i:YV) are measures of emotional intelligence. These measures are comprised of five major dimensions: intrapersonal, interpersonal, adaptability, stress management, and general mood. When compared to other possible measures, there are several reasons why this

instrument was preferred and, eventually chosen. First, the questions are generally very positive and are short and simple. They also address the respondent’s social, personal, and emotional “abilities”, as opposed to their behaviours.

The EQ-i:YV measure was asked of youth aged 10 to 17 and EQ-i was asked of youth 20 to 21. Youth 10 to 17 were asked in the self-completed paper questionnaire while the older youth were asked the questions through their CAPI questionnaire.

Prior to calculating the EQ-i score, the response category values were reversed for the Intrapersonal and Stress Management items. For the calculation of the EQ-i:YV score only the Stress Management items were reversed. After reversing the values, 1 was subtracted from each of the items to permit a score of zero. Once these two steps had been completed, the values were summed for each of the dimensions and five scores were created.

The lowest scores for a particular scale represent the negative end of the Emotional Intelligence continuum, while the higher scores represent the positive end. For example, 33 on total EQ would mean that the individual is much more emotionally intelligent than an individual who receives a score of 12 on the same scale. The dividing line between (a) extremely high, (b) high, (c) average, (d) low and (e) very low scores is essentially +/- 1 standard deviation from the mean value for the particular scale involved.

The standard scores are not provided on the master file however, the table below presents interpretive guidelines should data users decide to create the scores on their own. Standard scores for the EQ-i and EQ-i:YV set the mean values at 100 and each standard deviation at 15; however, you will notice there are 10 points around the mean values for differentiating between the descriptors in the table below.

Interpretative Guidelines for Standardized EQ-i and EQ-i:YV Scales Scores	
130 and above	Markedly high (atypically well-developed emotional/social capacity)
120 to 129	Very high (extremely well-developed emotional/social capacity)
110 to 119	High (well-developed emotional/social capacity)
90 to 109	Average (adequate emotional/social capacity)
80 to 89	Low (under-developed emotional/social capacity)
70 to 79	Very low (extremely under-developed emotional/social capacity)
under 70	Markedly low (atypically impaired emotional/social capacity)

Analytical Results for the EQ-i:YV

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Intrapersonal (_EQYES06)	_AMCE25A, _AMCE25F, _AMCE25K	2,427	10 to 11	2.0 to 2.9%	0.810
		1,703	12 to 13	2.2 to 2.7%	0.825
		1,499	14 to 15	3.1 to 3.3%	0.855
		1,197	16 to 17	4.2 to 4.5%	0.855
Interpersonal (_EQYES07)	_AMCE25B, _AMCE25G, _AMCE25L	2,415	10 to 11	2.1 to 3.0%	0.631
		1,699	12 to 13	2.3 to 2.8%	0.635
		1,504	14 to 15	3.0 to 3.2%	0.666
		1,190	16 to 17	4.4 to 4.7%	0.618
Stress Management (_EQYES08)	_AMCE25C*, _AMCE25H*, _AMCE25M*	2,403	10 to 11	2.9 to 3.3%	0.631
		1,686	12 to 13	2.8 to 2.9%	0.649
		1,494	14 to 15	3.2 to 3.5%	0.658
		1,198	16 to 17	4.2 to 4.4%	0.678
Adaptability (_EQYES09)	_AMCE25D, _AMCE25I, _AMCE25N	2,439	10 to 11	2.3 to 2.6%	0.679
		1,710	12 to 13	2.2 to 2.4%	0.734
		1,503	14 to 15	3.1 to 3.2%	0.778
		1,197	16 to 17	4.2 to 4.4%	0.774
General Mood (_EQYES10)	_AMCE25E, _AMCE25J, _AMCE25O	2,437	10 to 11	2.2 to 2.9%	0.689
		1,705	12 to 13	2.3 to 2.6%	0.704
		1,501	14 to 15	3.0 to 3.5%	0.676
		1,197	16 to 17	4.3 to 4.4%	0.673
Emotional Quotient (EQ4) (_EQYES04)	_AMCE25A, _AMCE25F, _AMCE25K, _AMCE25B, _AMCE25G, _AMCE25L, _AMCE25C*, _AMCE25H*, _AMCE25M*, _AMCE25D, _AMCE25I, _AMCE25N	2,297	10 to 11	2.0 to 3.3%	0.563
		1,655	12 to 13	2.2 to 2.9%	0.526
		1,477	14 to 15	3.1 to 3.5%	0.481
		1,182	16 to 17	4.2 to 4.7%	0.42
Emotional Quotient (EQ5) (_EQYES05)	_AMCE25A, _AMCE25F, _AMCE25K, _AMCE25B, _AMCE25G, _AMCE25L, _AMCE25C*, _AMCE25H*, _AMCE25M*, _AMCE25D, _AMCE25I, _AMCE25N, _AMCE25E, _AMCE25J, _AMCE25O	2,270	10 to 11	2.0 to 3.3%	0.691
		1,645	12 to 13	2.2 to 2.9%	0.667
		1,470	14 to 15	3.0 to 3.5%	0.636
		1,181	16 to 17	4.2 to 4.7%	0.585

Analytical Results for the EQ-i

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Intrapersonal (_EQYFS11)	_EQYfQ02*, _EQYfQ07*, _EQYfQ12*, _EQYfQ17*	1,395	20 to 21	0.0 to 0.1%	0.627
Interpersonal (_EQYFS12)	_EQYfQ01, _EQYfQ06, _EQYfQ11, _EQYfQ16	1,407	20 to 21	0.1 to 0.9%	0.648
Stress Management (_EQYFS13)	_EQYfQ03*, _EQYfQ08*, _EQYfQ13*, _EQYfQ18*	1,403	20 to 21	0.0 to 0.4%	0.642
Adaptability (_EQYFS14)	_EQYfQ04, _EQYfQ09, _EQYfQ14, _EQYfQ19	1,403	20 to 21	0.1 to 0.2%	0.678
General Mood (_EQYFS15)	_EQYfQ05, _EQYfQ10, _EQYfQ15, _EQYfQ20	1,404	20 to 21	0.0 to 0.2%	0.689
Emotional Quotient (EQ4) (_EQYFS16)	_EQYfQ02*, _EQYfQ07*, _EQYfQ12*, _EQYfQ17*, _EQYfQ01, _EQYfQ06, _EQYfQ11, _EQYfQ16, _EQYfQ03*, _EQYfQ08*, _EQYfQ13*, _EQYfQ18*, _EQYfQ04, _EQYfQ09, _EQYfQ14, _EQYfQ19	1,388	20 to 21	0.0 to 0.9%	0.598
Emotional Quotient (EQ5) (_EQYFS17)	_EQYfQ02*, _EQYfQ07*, _EQYfQ12*, _EQYfQ17*, _EQYfQ01, _EQYfQ06, _EQYfQ11, _EQYfQ16, _EQYfQ03*, _EQYfQ08*, _EQYfQ13*, _EQYfQ18*, _EQYfQ04, _EQYfQ09, _EQYfQ14, _EQYfQ19, _EQYfQ05, _EQYfQ10, _EQYfQ15, _EQYfQ20	1,386	20 to 21	0.0 to 0.9%	0.709

Below are brief definitions of what is measured by the five composite scales and the 15 subscales (only the five composite scales appear in the EQ-i:YV while these and the 15 subscales appear in the EQ-i which are bulleted below under each of the composite scales):

1) Intrapersonal competencies - self-awareness and self-expression:

These competencies include the following sub-components that govern our ability to be aware of ourselves, to understand our strengths and weaknesses, and to express our thoughts and feelings nondestructively.

- Self-regard: The ability to be aware of, understand and accept ourselves.

- Emotional self-awareness: The ability to be aware of and understand our
- Emotions.
- Assertiveness: The ability to express our feelings and ourselves nondestructively.
- Independence: The ability to be self-reliant and free of emotional dependency on others.
- Self-actualization: The ability to set goals and the drive to achieve them.

2) Interpersonal competencies - social awareness and interpersonal relationship:

These competencies include the following sub-components that govern our ability to be aware of others' emotions, feelings and needs, and to be able to establish and maintain co-operative, constructive and mutually satisfying relationships.

- Empathy: The ability to be aware of and understand how others feel.
- Social responsibility: The ability to identify with and feel part of our social group.
- Interpersonal relationship: The ability to establish and maintain mutually satisfying relationships with others.

3) Stress management competencies - emotional management and regulation:

These competencies include the following sub-components that govern our ability to manage emotions so that they work for us and not against us.

- Stress tolerance: The ability to effectively and constructively manage our emotions.
- Impulse control: The ability to effectively and constructively control our emotions.

4) Adaptability competencies - change management:

These competencies include the following sub-components that govern our ability to manage change, by realistically and flexibly coping with the immediate situation and effectively solving problems as they arise.

- Reality-testing: The ability to validate our feelings and thinking with external reality.
- Flexibility: The ability to cope with and adapt to change in our daily life.
- Problem-solving: The ability to generate effective solutions to problems of a personal and social nature.

5) General mood - self-motivation:

General mood is a facilitator of emotionally and socially intelligent behavior and includes the following sub-components that govern our ability to be optimistic, positive and sufficiently self-motivated to set and pursue our goals.

- Optimism: The ability to have a positive outlook and look at the brighter side of life.
- Happiness: The ability to feel content with ourselves, others and life in general.

For further information, see the following:

- 1) BarOn, R. (2004). The Bar-On Emotional Quotient Inventory (EQ-i): Rationale, description, and summary of psychometric properties. In Glenn Geher (Ed.),

Measurement of emotional intelligence: Common ground and controversy.
Hauppauge, NY: Nova Science Publishers, pp. 111-142.

- 2) BarOn, R., & Parker, J.D.A. (2000). *Emotional Quotient Inventory: Youth Version (EQ-i:YV): Technical manual.* Toronto, Canada: Multi-Health Systems.

9.6.5 Social Support Scale

The original scale contains 24 items from Robert Weiss's Social Provisions Model that describes six different social functions or "provisions" that may be acquired from relationships with others. Due to the length of the scale, and on the advice of Dr. M. Boyle at Chedoke-McMaster Hospital, McMaster University, the survey uses the shortened version (containing six items) that was derived for the Government of Ontario's, Better Beginnings, Better Futures Project. This measures guidance (two questions), reliable alliance (two questions) and attachment (two questions). Four additional questions on different types of social support, i.e., religious and community services, were added as suggested by Dr. Tom Hay. Questions similar to those suggested by Dr. Hay were taken from the Family Crisis Oriented Personal Evaluation Scales (F-COPES). F-COPES draws upon the coping dimensions of the Resiliency Model of Family Adjustment and Adaptation (McCubbin, Olson & Larsen: 1981). The total social support measurement includes eight questions and not only focuses on the quantity of social support but on the quality of social supports as well. The questions are asked of 18- and 19-year-olds in the Youth Questionnaire.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Social Support (_SPYES01)	_SPYEQ1A*, _SPYEQ1B, _SPYEQ1C, _SPYEQ1D*, _SPYEQ1E*, _SPYEQ1F, _SPYEQ1G, _SPYEQ1H*	1,562	18 to 19	0.5 to 1.6%	0.847

*Indicates items were reversed.

9.6.6 Friends Scale

The Friends Scale is intended to measure how well the youth feels he/she gets along with his/her peers. This information is important in identifying the extent and quality of the child's social support network. These questions form the Peer Relations Subscale in the Marsh Self-Descriptive Questionnaire, developed by H.W. Marsh.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Friends Score (_FFCS01)	_FFCQ01, _FFCQ02, _FFCQ03, _FFCQ04	2,461	10 to 11	1.5 to 2.1%	0.770
		1,702	12 to 13	1.9 to 2.8%	0.809
		1,473	14 to 15	4.3 to 5.1%	0.835
		1,191	16 to 17	4.3 to 5.0%	0.874

9.6.7 My Parents and Me Scales

The Parenting Scales are intended to complement the parenting section in the parent-reported Children's Questionnaire by gathering information from the child regarding his/her perception of his/her relationship with parents. For the self-completed questionnaire, it was also considered important to obtain a measure of parental supervision, i.e., monitoring, as this has been shown to be linked to child outcomes – there is a correlation between a lack of supervision and negative outcomes, such as juvenile delinquency and other risk-taking behaviours.

This scale is used in the Western Australia Child Health Survey. It was developed by Lempers et al (1989) based on the work of Schaefer (1965) and Roberts et al (1984) and measures parental nurturance, rejection and monitoring.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Parental Nurturance (_PMCCS1)	_PMCCQ1A, _PMCCQ1D, _PMCCQ1K, _PMCCQ1H, _PMCCQ1L, _PMCCQ1M, _PMCCQ1Q	2,269	10 to 11	3.2 to 5.9%	0.860
		1,581	12 to 13	3.1 to 6.6%	0.904
		1,466	14 to 15	4.2 to 4.6%	0.920
Parental Rejection (_PMCBS2B)	_PMCCQ1C, _PMCCQ1G, _PMCCQ1J, _PMCCQ1L, _PMCCQ1O, _PMCCQ1P, _PMCCQ1R	2,197	10 to 11	3.9 to 8.0%	0.664
		1,583	12 to 13	3.6 to 5.8%	0.739
		1,456	14 to 15	4.3 to 5.1%	0.771
Parental Monitoring (_PMCCS3)	_PMCCQ1B, _PMCCQ1F, _PMCCQ1N, _PMCCQ1E, _PMCCdQ1T	2,321	10 to 11	3.4 to 6.0%	0.482
		1,648	12 to 13	3.6 to 4.5%	0.505
		1,469	14 to 15	4.2 to 4.9%	0.481

9.6.8 Conflict Resolution

Two conflict resolution scores were created for youth aged 16 and 17 based on questions asked in the self-completed questionnaire. One score relates to the relationship between the youth and the mother and the other score refers to the relationship between the youth and his/her father. A high score indicates an elevated number of disagreements between the youth and his/her parent.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Conflict Resolution - Mother (_PMCD54)	_PMCDQ6C, _PMCDQ6D, _PMCDQ6E, _PMCDQ6F, _PMCDQ6G, _PMCDQ6H, _PMCDQ6I, _PMCDQ6J, _PMCDQ6K, _PMCDQ6L	1,162	16 to 17	4.0 to 4.4%	0.617
Conflict Resolution - Father (_PMCD55)	_PMCDQ9C, _PMCDQ9D, _PMCDQ9E, _PMCDQ9F, _PMCDQ9G, _PMCDQ9H, _PMCDQ9I, _PMCDQ9J, _PMCDQ9K, _PMCDQ9L	1,110	16 to 17	4.9 to 5.3%	0.669

9.6.9 Behaviour Scales

This section replicates the behaviour checklist included in the Children's Questionnaire completed by the PMK. All youth aged 10 to 15 answer these questions in the self-complete portion of the survey. It is intended to provide indicators of the following behaviours: conduct disorder, hyperactivity, inattention, physical aggression, indirect aggression, emotional disorder, anxiety, prosocial behaviours and property-offence related behaviour.

Score	Items Included	Sample Size	Universe (Age in years)	Item Nonresponse	Cronbach's Alpha (Standardized)
Indirect Aggression (_FBCS01)	_FBCQ01J, _FBCQ01R, _FBCQ01Z, _FBCQ1LL, _FBCQ1TT	2,346	10 to 11	3.7 to 4.3%	0.730
		1,677	12 to 13	2.4 to 3.1%	0.747
		1,495	14 to 15	3.3 to 3.7%	0.714
Emotional Disorder - Anxiety (_FBCdS02)	_FBCQ01F, _FBCQ01K, _FBCQ01Q, _FBCQ01V, _FBCQ1CC, _FBCQ1MM, _FBCQ1RR	2,295	10 to 11	3.9 to 4.8%	0.730
		1,665	12 to 13	2.4 to 3.3%	0.775
		1,485	14 to 15	3.3 to 3.9%	0.784
Physical Aggression - Conduct Disorder (_FBCS03)	_FBCQ01G, _FBCQ01X, _FBCQ1AA, _FBCQ1FF, _FBCQ1JJ, _FBCQ1NN	2,334	10 to 11	3.8 to 4.3%	0.760
		1,681	12 to 13	2.4 to 3.0%	0.803
		1,489	14 to 15	3.3 to 3.7%	0.801
Hyperactivity - Inattention (_FBCdS04)	_FBCC01B, _FBCQ01I, _FBCQ01P, _FBCQ01S, _FBCQ01W, _FBCQ1HH, _FBCQ1QQ	2,294	10 to 11	3.8 to 5.1%	0.755
		1,664	12 to 13	2.5 to 4.1%	0.769
		1,474	14 to 15	3.3 to 4.0%	0.796
Prosocial Behaviour (_FBCS05)	_FBCQ01A, _FBCQ01D, _FBCQ01H, _FBCQ01M, _FBCQ01U, _FBCQ1BB, _FBCQ1GG, _FBCQ1OO, _FBCQ1SS, _FBCC1UU	2,296	10 to 11	3.5 to 5.0%	0.817
		1,658	12 to 13	2.4 to 3.3%	0.857
		1,490	14 to 15	3.1 to 3.7%	0.874
Property Offences (_FBCS07)	_FBCQ01C, _FBCQ01E, _FBCQ01L, _FBCQ01T, _FBCQ1DD, _FBCQ1PP	2,314	10 to 11	3.6 to 5.7%	0.633
		1,680	12 to 13	2.4 to 3.3%	0.704
		1,483	14 to 15	3.1 to 4.0%	0.712

9.7 Summary of the Cycle 6 Scales

9.7.1 Parent-reported Scales

Score Variable	Scale Name	Universe
_DPSS01	Depression Rating Scale	PMK of children 0 to 15 years
_FNHS01	Family Functioning Scale	PMK or spouse of children 0 to 15 years
_SFHS5	Neighbourhood Safety Scale	PMK or spouse of children 0 to 15 years
_SFHS6	Neighbours Scale	PMK or spouse of children 0 to 15 years
_SPHS01	Social Support Scale	PMK or spouse of children 0 to 15 years
_ACCS6	Home Responsibilities Scale	PMK of children 10 to 13 years
_BECdS01	Hyperactivity - Inattention	PMK of children 2 to 3 years
_BECdS03	Emotional Disorder - Anxiety	PMK of children 2 to 3 years
_BECS04	Physical Aggression - Opposition	PMK of children 2 to 3 years
_BECS05	Separation Anxiety	PMK of children 2 to 3 years
_BECdS06	Hyperactivity - Inattention	PMK of children 4 to 11 years
_BECdS07	Prosocial Behaviour	PMK of children 8 to 11 years
_BECdS08	Emotional Disorder - Anxiety	PMK of children 4 to 11 years
_BECdS09	Physical Aggression - Conduct Disorder	PMK of children 4 to 11 years
_BECS10	Indirect Aggression	PMK of children 4 to 11 years
_BECdS11	Property Offences	PMK of children 8 to 11 years
_MSCS01	MSD Raw Score	PMK of children 0 to 47 months
_MSCS02	MSD Standardized Score - based on Cycle 1 norms	PMK of children 0 to 47 months
_MSCdS03	MSD Standardized Score - based on Cycle 6 norms	PMK of children 3 to 47 months
_PRCS01	Positive Interaction	PMK of children 0 to 23 months
_PRCS02	Ineffective Parenting	PMK of children 0 to 23 months
_PRCS03	Positive Interaction	PMK of children 2 to 11 years
_PRCS04	Ineffective Parenting	PMK of children 2 to 11 years
_PRCS05	Consistent Parenting	PMK of children 2 to 11 years
_PRCS06	Rational Parenting	PMK of children 2 to 11 years
_PRCbS09	Conflict Resolution Scale	PMK of children 12 to 15 years

9.7.2 Self-completed Scales (Child/Youth Reported)

Score Variable	Scale Name	Universe
_FFCS01	Friends Scale	Children/Youth 10 to 17 years
_AMCS02	General Self-image Scale	Children/Youth 10 to 17 years
_FBCS01	Indirect Aggression	Children/Youth 10 to 15 years
_FBCdS02	Emotional Disorder - Anxiety	Children/Youth 10 to 15 years
_FBCS03	Physical Aggression - Conduct Disorder	Children/Youth 10 to 15 years
_FBCdS04	Hyperactivity - Inattention	Children/Youth 10 to 15 years
_FBCS05	Prosocial Behaviour	Children/Youth 10 to 15 years
_FBCS07	Property Offences	Children/Youth 10 to 15 years
_PMCcS1	Parental Nurturance	Children/Youth 10 to 15 years
_PMCbS2b	Parental Rejection	Children/Youth 10 to 15 years
_PMCcS3	Parental Monitoring	Children/Youth 10 to 15 years
_EQYES04	Emotional Quotient (4 factors)	Children/Youth 10 to 17 years
_EQYES05	Emotional Quotient (5 factors)	Children/Youth 10 to 17 years
_EQYES06	Intrapersonal Skills	Children/Youth 10 to 17 years
_EQYES07	Interpersonal Skills	Children/Youth 10 to 17 years
_EQYES08	Stress Management Skills	Children/Youth 10 to 17 years
_EQYES09	Adaptability Skills	Children/Youth 10 to 17 years
_EQYES10	General Mood	Children/Youth 10 to 17 years
_PMCdS4	Conflict Resolution – Mother	Youth 16 to 17 years
_PMCdS5	Conflict Resolution – Father	Youth 16 to 17 years
_HTCbS1B	Depression Scale	Youth 16 to 17 years

9.7.3 Youth-reported Scales (Self-reported)

Score Variable	Scale Name	Universe
_ACYdS01	Neighbourhood Structure Scale	Youth 16 to 17 years
_HTYfS01	Depression Scale	Youth 18 years and older
_SPYES01	Social Support Scale	Youth 18 to 19 years
_AMYfS01	General Self-image Score	Youth 18 to 19 years

10.0 Imputation

The definition of a respondent is a child or youth, who has at least one of the child, youth or adult components completed. For the respondents, there exist many cases of partial nonresponse. This may be for an entire component or only for certain questions. Imputation is the process whereby missing or inconsistent items are "filled in" with acceptable values. In the National Longitudinal Survey of Children and Youth (NLSCY), imputation is carried out for certain variables in the adult and youth income sections as well as in the child motor and social development section.

Imputation flags have been included on the NLSCY file so that users will have information on the extent of imputation and what specific items have been imputed on what records. All imputation flags on the NLSCY data file have an "I" as the sixth character of the variable name. For example, the name of the imputation flag for household income (FINHeQ03) is FINHeI02.

10.1 Household Income Imputation

Several income questions were asked during the NLSCY household interview. Information on income, broken down into three sources, was asked for the person most knowledgeable (PMK) and his or her spouse. Those three income sources are: wages and salary, self-employment net income, and Employment Insurance benefits. Information on income, broken down into four sources was also asked at the household level. Those four income sources are: Child Tax Benefit/National Child Benefit, social assistance, child and spousal support and other sources. The total household income represents the sum of these ten sources of income.

Income is a sensitive topic. As a result, some respondents refused to provide answers to the detailed income questions. Among those, some provided an estimate of their total household income or an estimate of their income using ranges. Finally, for those who responded, amounts declared in the income section were sometimes incoherent with answers provided in the labour force section (for example, an individual might have reported working in the past 12 months according to answers provided in the labour force section but no wages or self-employment income were reported in the income section). Income imputation was carried out to fill in the holes resulting from partial nonresponse as well as to rectify, when possible, these incoherencies. Imputation was also done for households whose total reported income was less than \$6,000.

Imputation of the household income was done only for those households that were eligible for an adult component. This includes all households with the exception of those that only have selected youth aged 18 and over and those that only have youth aged 16 or 17 who live independently. Of the 16,491 eligible households at least one income variable was imputed for a total of 3,160 households. The 16,491 eligible households actually correspond to 18,684 children or youth who are split in two files depending on their effective age: i) the longitudinal file for children or youth part of the original cohort selected in Cycle 1 who are aged from 10 to 21 years old in Cycle 6 and ii) the Early Childhood Development (ECD) file for children or youth who are part of the birth cohorts selected in Cycles 4, 5 and 6 and who are aged from 0 to 5 years old in Cycle 6. The longitudinal file contains 10,514 children or youth of which 2,291 were imputed while the ECD file contains 8,170 children of which 1,159 were imputed. Overall, 3,450 of the 18,684 children or youth were imputed.

The most critical piece of information in the adult income section is the total household income. Our imputation strategy for Cycle 6 was designed to determine the best possible total household income value, occasionally at the expense of the reported sources of income. Imputation was carried out using various forms of nearest neighbour imputation. This method first identifies a respondent to the income section (a donor) who has similar characteristics to the respondent with incomplete income data (the recipient). The donor record is then used to compute imputed values for the recipient record. Imputation was done in four steps: a) households that provided an estimated income, b) households that responded in Cycle 5, c) households that reported their income in ranges and d) remaining households.

a) Imputation of households that provided an estimated income

For these households, we considered the estimated income as the total household income. If only one source of income was missing, it was imputed deterministically. If more than one source of income was missing, the ten sources of income were imputed using the distribution of the income sources from a donor household (donor ratio imputation). Each source of income was actually imputed if it was not reported or if the calculated value based on the donor differed from the reported value by more than 10%. In the end, the actual total household income obtained by summing up the ten sources of income may vary slightly from the provided estimated income.

b) Imputation of households that responded in Cycle 5

In order to preserve longitudinal coherence through time, the imputation of the total income of households that responded in Cycle 5 was performed by nearest neighbour trend imputation, excluding from the donor pool households with extreme income trends from Cycle 5 to Cycle 6. When a recipient household reported its total income in ranges, we made sure that the imputed total income respected the specified range. Similar to a), the donor household's distribution of income sources was used to impute the ten sources of income for the recipient household (donor ratio imputation).

c) Imputation of the households that reported their income in ranges

Imputation was carried out for each missing source of income using a classic nearest neighbour approach. The sources of income reported by the donor were directly used to impute the missing income sources for the recipient. The missing sources of income were imputed at once, using the same donor household, making sure that the total household income respected the specified range.

d) Imputation of the remaining households

Not a lot of information on the total income was available for the remaining households. As in c), imputation was carried out for each missing source of income using a classic nearest neighbour approach. The sources of income reported by the donor were directly used to impute the missing income sources for the recipient. However, the imputation was performed in three steps this time: i) imputation of the PMK's income sources, ii) imputation of the spouse's income sources and iii) imputation of the household level income sources.

The breakdown of imputed households by imputation step can be found in the following table.

Imputation Step	Number of Households	Percentage of Imputed Households
a) Households that provided an estimated income	698	22,1 %
b) Households that responded in Cycle 5	1 314	41,6 %
c) Households that reported their income in ranges	388	12,3 %
d) Remaining households	760	24,0 %
Total	3 160	100 %

The imputation flags provide information on how the imputation was done. The description of the flag values follow.

Total Household Income Flag (FINHeI02):

Imputation flag = 0	No imputation
Imputation flag = 1	Estimated income imputation
Imputation flag = 2	Donor trend imputation by income range
Imputation flag = 3	Donor trend imputation
Imputation flag = 4	Donor imputation by income range
Imputation flag = 5	Donor imputation

Subtotal Income Flags (FINPeI03 and FINSeI03):

Imputation flag = 0	No imputation
Imputation flag = 1	At least one source of income imputed
Imputation flag = 6	Valid skip (no spouse in the household)

Income Source Flags (FINPI1AA, FINPI1AB, FINPI1AC, FINSI1AA, FINSI1AB, FINSI1AC, FINHI1AD, FINHI1AF, FINHI1AE and FINHI1AG):

Imputation flag = 0	No imputation
Imputation flag = 1	Deterministic imputation
Imputation flag = 2	Donor ratio imputation
Imputation flag = 3	Donor imputation
Imputation flag = 6	Valid skip (no spouse in the household)

Child-level imputation rates for the income variables can be found in the following table.

Variable	Imputation Flag	Overall Imputation Rate	Imputation Rate for Longitudinal File (Ages 10 to 21)	Imputation Rate for ECD File (Ages 0 to 5)
FINPc1AA PMK income from wages and salaries	FINPI1AA	11.1%	9.6%	12.2%
FINPc1AB PMK income from self-employment	FINPI1AB	9.2%	7.8%	10.2%
FINPc1AC PMK income from Employment Insurance benefits	FINPI1AC	9.8%	7.6%	11.5%
FINPeD04 Total personal income for PMK	FINPeI03	14.0%	11.6%	15.9%
FINS1AA Spouse income from wages and salaries	FINSI1AA	14.3% *	12.0% *	16.0% *
FINS1AB Spouse income from self-employment	FINSI1AB	11.9% *	10.0% *	13.3% *
FINS1AC Spouse income from Employment Insurance benefits	FINSI1AC	10.5% *	8.4% *	11.9% *
FINS1D04 Total personal income for spouse	FINS1E03	16.4% *	13.2% *	18.8% *
FINHe1AD Household income from child tax benefits	FINHI1AD	14.5%	12.4%	16.1%
FINHe1AE Household income from social assistance	FINHI1AE	8.8%	7.1%	10.1%
FINHe1AF Household income from child and spousal support	FINHI1AF	8.9%	7.5%	10.0%
FINHe1AG Household income from other sources	FINHI1AG	9.4%	7.8%	10.6%
FINHeQ03 Total household income	FINHeI02	18.5%	14.2%	21.8%

* Households where there was no spouse were not included in the calculation of the imputation rate.

10.2 Youth Income Imputation – 16 or 17 Years Old

Information on income, broken down in five sources, was asked as part of the youth component for those who were 16 or 17 years old in Cycle 6. The youth were asked their income from odd jobs, employers, scholarships or bursaries, parents and other sources. The total youth income represents the sum of these five sources of income.

Income is a sensitive topic. As a result, some respondents refused to provide answers to the detailed income questions. Among those, some provided an estimate of their income using ranges. Finally, for those who responded, amounts declared in the income section were sometimes incoherent with answers provided in the labour force section (for example, a youth might have worked for odd jobs for pay according to answers provided in the labour force section but no income from odd jobs was reported in the income section). Income imputation was carried out to fill out the holes resulting from partial nonresponse as well as to rectify, when possible, these incoherencies.

Imputation of income was considered for all respondents who were 16 or 17 years old and have completed the youth component. Of the 1,585 respondent youth aged 16 or 17 years old in Cycle 6, 1,436 completed the youth component. The remaining 149 youth only completed the child and/or the adult component. There was not enough information about the youth who did not complete the youth component to be able to impute them. From those who completed the youth component a total of 358 had at least one source of income imputed.

Imputation was carried out for each of the five sources of income. Imputation for most cases was done using a nearest neighbour approach. This method first identifies a respondent to the income section (a donor) who has similar characteristics to the individual with incomplete income data (the recipient). Once the nearest neighbour has been identified, the missing sources of income are copied to the recipient record. When provided, the total income range was taken into consideration in the donor selection so that the imputed total income would respect the specified range. When only one of the four income sources required imputation and there was a total income provided in ranges, a plausible value randomly chosen from a uniform distribution of possible values was imputed.

The imputation flags provide information on how the imputation was done. The description of the flag values follow.

Imputation flag = 0	No imputation
Imputation flag = 1	Donor imputation by income range
Imputation flag = 2	Plausible value imputation
Imputation flag = 3	Donor imputation

The following table shows the breakdown of imputed youth, that is respondents who had at least one of their sources of income imputed, by imputation method.

Imputation Method (Imputation flag FINYeI02)	Number of Youth	Percentage of Imputed Youth
Plausible value imputation	61	17,0 %
Donor imputation by income range	79	22,1 %
Donor imputation	218	60,9 %
Total	358	100 %

Imputation rates for the income variables can be found in the following table.

Variable	Imputation Flag	Imputation Rate
FINYeQ1A Income from odd jobs	FINYeI1A	12,4 %
FINYdQ1B Income from employers	FINYeI1B	12,3 %
FINYeQ1E Income from scholarships or bursaries	FINYeI1E	2,9 %
FINYeQ1C Income from parents	FINYeI1C	8,6 %
FINYdQ1D Income from other sources	FINYeI1D	5,4 %
FINYeD01 Total youth income	FINYeI02	24,9 %

10.3 Youth Income Imputation – 18 Years Old and Over

Information on income was asked as part of the youth component for those who were 18 years old and over. The youth were asked their total income amount as well as their sources of income without the detailed amounts. The list of income sources that they had to choose from follows.

1. Wages and salaries, (including commissions, tips and bonuses)
2. Income from self-employment
3. Scholarships
4. Government Student Loans
5. Registered Education Savings Plan (RESP)
6. Interest, dividends, capital gains, or other investment income
7. Employment Insurance (EI benefits)
8. Worker's compensation
9. Benefits from Canada or Québec Pension Plan, Guaranteed Income Supplement or Spouse's Allowance
10. Child Tax Benefit
11. Provincial or municipal social assistance or welfare
12. Child Support
13. Alimony
14. Other
15. No income

Income is a sensitive topic. As a result, some respondents refused to provide answers to the income questions. Among those, some provided an estimate of their income using ranges. Finally, for those who responded, amounts declared in the income section were sometimes incoherent with answers provided in the labour force section (for example, a youth might have worked for pay according to answers provided in the labour force section but reported no income in the income section). Income imputation was carried out to fill the holes resulting from partial nonresponse as well as to rectify, when possible, these incoherencies.

Imputation of income was considered for all respondents who were 18 years old and over. Of the 2,982 respondent youth aged 18 years old and over in Cycle 6, 300 had their total income imputed while 11 of these 300 respondents had their list of income sources imputed.

Imputation was done using a nearest neighbour approach. This method first identifies a respondent to the income section (a donor) who has similar characteristics to the individual with incomplete income data (the recipient). Once the nearest neighbour has been identified, the missing data are copied to the recipient record. When provided, the total income range was taken into consideration in the donor selection so that the imputed total income would respect the specified range.

The imputation flags provide information on how the imputation was done. The description of the flag values follow.

Imputation flag = 0	No imputation
Imputation flag = 1	Donor imputation by income range
Imputation flag = 2	Donor imputation

The following table shows the breakdown of imputed youth, that is respondents who had their total income imputed, by imputation method.

Imputation Method (Imputation flag FIYYfI2A)	Number of Youth	Percentage of Imputed Youth
Donor imputation by income range	236	78,7 %
Donor imputation	64	21,3 %
Total	300	100 %

Imputation rates for the income variables can be found in the following table.

Variable	Imputation Flag	Imputation Rate
FIYYfQ1A to FIYYfQ1O Income sources	FIYYfI1	0.4%
FIYYfD01 Total youth income	FIYYfI2A	10.1%

10.4 Motor and Social Development (MSD) Scale Imputation

The imputation of the motor and social development (MSD) scale in Cycle 6 uses the same methodology as the one used in Cycle 5. Prior to Cycle 5, to obtain the raw MSD score for a child (variable **FMSCS01**), all 15 applicable questions had to be answered either "Yes" or "No". However, it was noted that a large proportion of the records with incalculable raw scores had only one or two missing responses among the 15 questions. We therefore decided, starting with Cycle 5, that we could obtain a reasonably accurate score making use of 13 or 14 valid responses and imputing for the missing items.

Specifically, if a child had 13 or 14 valid responses, a donor record was randomly chosen from among the children having complete response and the same response pattern to the common questions. The "Yes" or "No" from the selected donor replaced the original missing value. When two items were imputed, these were done independently. Consequently, there could be two different donors for the two missing values.

A donor matching the exact response pattern for the common questions could not always be found. These situations were handled by then choosing a donor among the children having complete response and the same partial score for the common questions.

Naturally, to have 13 or 14 questions in common, all potential donors had to be in the same age range in months as the child to be imputed. For example, an 8-month-old child missing FMSCQ21 had potential donors age 7 to 9 months who were asked the same 15 questions (FMSCQ12 – FMSCQ26) and had the same pattern of "Yes" and "No" responses for FMSCQ12 – FMSCQ20 and FMSCQ22-FMSCQ26.

Through this process, a valid response was never changed from "Yes" to "No" or vice versa. Only missing values were overwritten with a "Yes" or "No".

In total, 525 additional MSD scores were obtained by having at least one response imputed; 472 had exactly one response imputed and 53 had exactly two responses imputed. This represents 7.5% of all eligible children.

The imputation flag variables FMSCIS1A and FMSCIS1B identify which MSD questions were imputed. A value of zero for both of these flags means that no imputation was done for the MSD questions.

11.0 Weighting and Treatment of Nonresponse

The National Longitudinal Survey of Children and Youth (NLSCY) is a probability survey. As is the case with any probability survey, the sample is selected so as to be able to produce estimates for a reference population. Therefore, each unit in the sample represents a number of units in the population. In the NLSCY, several populations are represented. The total sample for Cycle 6 is a combination of samples selected in Cycles 1, 4 and 5 (1994, 2000 and 2002) and a new sample selected in Cycle 6 (2004). For details on the NLSCY's sample design and the composition of the sample at each cycle, see Chapter 5.0.

Recall from Chapter 5.0 that in a longitudinal survey such as the NLSCY, two types of populations are possible: longitudinal and cross-sectional. The longitudinal population is the initial population when the sample was first drawn; cross-sectional populations refer to subsequent time periods. Differences between the longitudinal population and subsequent cross-sectional populations are due to births, deaths, immigration and emigration.

The NLSCY produces three sets of weights at each cycle: two longitudinal, one cross-sectional. The difference between the two longitudinal weights is that the funnel weights are only for those in the original cohort who have responded to every cycle, while the non-funnel longitudinal weights are for longitudinal respondents who may not have responded at every cycle (either original cohort or Early Childhood Development (ECD) cohorts).

Aside from the original cohort, longitudinal weights are only assigned to returning individuals who started the NLSCY as 0- to 1-year-old children. The first cycle an individual is introduced he/she only receives cross-sectional weights.

11.1 Weights Available at Cycle 6

The following describes the reference populations for the various weights produced at Cycle 6.

11.1.1 Funnel Weights (Variable FWTCdW1L)

The funnel weights are available only for longitudinal members of the original cohort who responded at every cycle.

Longitudinal population: children aged 0 to 11 years old as of December 31st, 1994 who were living in one of the 10 provinces at the time of Cycle 1 collection (1994/1995)

At Cycle 1, a longitudinal sample of children aged 0 to 11 was selected from the LFS. By Cycle 6, these children were 10 to 21 years old (as of December 31st, 2004). Sample reductions were made at Cycle 2. The children dropped between Cycles 1 and 2 can be regarded as Cycle 1 cross-sectional children.

Funnel weights are produced at every cycle, starting at Cycle 4, for this longitudinal population.

11.1.2 Longitudinal Weights (Variable FWTCW01L)

Longitudinal (non-funnel) weights are available for returning, longitudinal members of the original cohort and ECD cohorts who were respondents at Cycle 6 but who may not have been respondents at every previous cycle.

The original cohort

Longitudinal population (same as for the funnel weights): children aged 0 to 11 years old as of December 31st, 1994 who were living in one of the 10 provinces at the time of Cycle 1 collection (1994/1995)

At Cycle 1, a longitudinal sample of children aged 0 to 11 was selected from the LFS. By Cycle 6, these children were 10 to 21 years old (as of December 31st, 2004). Sample reductions were made at Cycle 2. The children dropped between Cycles 1 and 2 can be regarded as Cycle 1 cross-sectional children.

Longitudinal weights are produced at every cycle for this longitudinal population.

These individuals can be identified on the data files by the condition: MEMCYCLE = 01. The variable MEMCYCLE indicates the cycle in which the child entered the survey.

ECD Cohort introduced in Cycle 4

Longitudinal population of the ECD cohort selected at Cycle 4: children aged 0 to 1 year old as of December 31st 2000 who were living in one of the 10 provinces at the time of Cycle 4 collection (2000/2001)

At Cycle 4, a longitudinal sample of children aged 0 to 1 was selected from the LFS. At Cycle 6 these children were 4 to 5 years old (as of December 2004).

Longitudinal weights were produced for this population at Cycles 5 and 6.

These individuals can be identified on the data files by the condition: MEMCYCLE = 04. The variable MEMCYCLE indicates the cycle in which the child entered the survey.

ECD Cohort introduced in Cycle 5

Longitudinal population of the ECD cohort selected at Cycle 5: children aged 0 to 1 year old as of December 31st 2002 who were living in one of the 10 provinces at the time of Cycle 5 collection (2002/2003)

At Cycle 5, a longitudinal sample of children aged 0 to 1 year old was selected from the LFS. By Cycle 6, these children were 2 to 3 years old (as of December 2004).

Longitudinal weights were produced for this longitudinal population at Cycle 6.

These individuals can be identified on the data files by the condition: MEMCYCLE = 05. The variable MEMCYCLE indicates the cycle in which the child entered the survey.

11.1.3 Cross-sectional Weights (Variable FWTCW01C)

Cross-sectional weights (variable FWTCW01C): only for members of the NLSCY sample at Cycle 6 who can be used to represent the following cross-sectional population.

Cross-sectional population at Cycle 6: Children aged 0 to 5 years old as of December 31st 2004 who were living in one of the 10 provinces at the time of Cycle 5 collection (2004/2005)

This cross-sectional sample consists of:

- the ECD cohort of 0- to 1-year-old children selected at Cycle 6,
- returning 2- to 3-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 5,
- returning 4- to 5-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 4,
- a new top-up of 2- to 5-year-old children (selected from the LFS).

11.2 Weighting Method

Many details of the weighting method are described below. However users, who are interested in knowing more details than are presented here, may contact Statistics Canada for more information.

11.2.1 The National Longitudinal Survey of Children and Youth Design Weight

The NLSCY weighting strategy is based on a series of adjustments applied to the NLSCY design weight. Each child's NLSCY design weight is equal to the inverse of his/her probability of selection. All children still in the NLSCY at Cycle 6 were selected, at one time or another, from the LFS. The LFS subweight accounts for all of the sample design information for the LFS sample. For NLSCY, we further select households from LFS and then children within these households. To reflect these additional sample design steps, the LFS subweight is multiplied by several other factors. After multiplying the LFS subweight by these various factors, we arrive at the NLSCY design weight. The final weight is obtained by applying nonresponse and post-stratification adjustments to the NLSCY design weight. In this chapter, we will denote the NLSCY design weight as:

$W_{\text{NLSCY design}}$

11.2.2 First Adjustment: Nonresponse Adjustment

It is a reality of most surveys that not everyone who is sampled, responds. NLSCY is no exception. Because NLSCY suffers from nonresponse, we need to adjust the weights so that the respondents represent the nonrespondents. Otherwise, for example, we would underestimate totals.

More precisely, the goal of the nonresponse adjustment is to inflate the NLSCY design weights of the respondents so that their nonresponse adjusted weights add up to the sum of the NLSCY design weights for everyone in the original sample. In other words, the nonresponse adjustment aims to give the NLSCY design weights of the nonrespondents to the respondents in an intelligent way. Children in the Cycle 1 sample have experienced six cycles of nonresponse (Cycles 1 through 6, inclusive). Those in the Cycle 4 sample have lost respondents over three cycles (Cycles 4, 5 and 6), while the Cycle 5 sample of children has been affected by nonresponse over two cycles (Cycles 5 and 6). The Cycle 6 sample has been subjected to only one cycle of nonresponse (Cycle 6). The nonresponse adjustment aims to adjust for all of this nonresponse in one step.

To decide how to give the weight of the nonrespondents to the respondents, we apply the method of response homogeneous groups (RHGs). The RHG method involves grouping individuals with the same likelihood of response. Then an adjustment factor is computed for each RHG. This factor is defined as follows:

$$\text{Non-response adjustment in an RHG} = \frac{\sum_{\substack{\text{Respondents} \\ \text{and} \\ \text{Non-respondents}}} w_{\text{NLSCY design}}}{\sum_{\text{Respondents}} w_{\text{NLSCY design}}}$$

The likelihood of response at Cycle 6 is determined through logistic regression models. The model output is predicted response probabilities (value between 0 and 1) for everyone. These probabilities are then sorted, and divided into Y RHGs, where Y is determined so that: a monotonic response rate across the groups is achieved and, so that certain constraints (size of the adjustment factor and minimum size of each RHG) are satisfied so that reasonable, reliable adjustment factors can be obtained.

For Cycle 6 weighting, three logistic regression models are built: one for the longitudinal response of the Cycle 1 sample, one for the funnel response of the Cycle 1 sample, and one for the response of the Cycles 4, 5 and 6 samples. Separate models are constructed because the Cycle 1 sample has suffered many more cycles of nonresponse than the ECD cohorts, and there is every reason to believe that nonresponse behaviour varies as the number of times the individual has been interviewed increases. From these models, RHGs are built separately for the funnel, longitudinal and cross-sectional weighting.

11.2.3 Second Adjustment: Post-stratification

The second adjustment factor ensures consistency between the estimates produced by NLSCY and Statistics Canada's population estimates by age, sex and province. This method is called post-stratification. The post-stratification totals depend on the population of reference.

Both the cross-sectional and longitudinal (including funnel) weights are post-stratified. For Cycle 6 cross-sectional weighting, the reference year to calculate a child's effective age is 2004. The post-stratification counts used refer to January 2005, so that we have a reliable count of children of a given age as of December 31st of the previous year. In a similar vein, longitudinal weighting uses the January counts of the year following the reference year for the population of interest.

To find out which post-stratum a given individual belongs to, see the variables PSTRATC (post-stratum for cross-sectional weighting purposes) and PSTRATL (post-stratum for longitudinal weighting purposes).

11.2.4 How the Weighting Method at Cycle 6 Differs from the Method at Previous Cycles

The weighting method described above for Cycle 6 differs from the method employed in previous cycles. We redesigned the nonresponse model for Cycle 6, so that the nonresponse model for longitudinal response is now truly longitudinal. The method for modelling nonresponse was also changed from segmentation modelling in previous cycles to logistic regression modelling at Cycle 6. The result is fewer, more robust and

discriminating RHGs. One innovation is the use of a “co-operation variable” in the new nonresponse model: this variable assigns individuals a score based on how well they participated in the survey in the past, which turns out to be a predictor of nonresponse.

For example, the co-operation score for the Cycle 1 sample is derived using Cycle 1 information. Essentially it measures the extent to which questions asked of the respondent were answered in Cycle 1. If all of the questions posed were answered, then the co-operation score was high. If however, a large number of the questions asked at Cycle 1 were not answered, they are assigned a low co-operation score. The score itself is a categorical variable. In the case of Cycle 1 nonrespondents, the co-operation score was imputed.

11.3 Applying the Weighting Method

11.3.1 Longitudinal Weighting

The longitudinal weight, FWTCW01L, applies to three separate samples of NLSCY children:

- The original cohort of 0- to 11-year-old children selected at Cycle 1
- The ECD cohort of 0- to 1-year-old children selected at Cycle 4
- The ECD cohort of 0- to 1-year-old children selected at Cycle 5

Notice that for each sample, the longitudinal weight relates to a distinct reference population.

Definition of a longitudinal respondent

A longitudinal respondent is a child who was introduced in a previous cycle and whose adult component or child or youth component is complete. For youth 18 years old and above, the youth component must be completed in order to consider the youth a respondent. Children who were introduced in a previous cycle and died or moved outside Canada's 10 provinces are also longitudinal respondents. They represent similar children in the reference population.

First adjustment: nonresponse adjustment

Two nonresponse models were created: one for the original cohort, and another for the ECD children. Both models used LFS variables (such as: is the dwelling rented or owned? and, highest level of education), with the original cohort model having an additional co-operation score as an independent variable. The nonresponse weight adjustment is calculated by RHG using the formula presented earlier.

For the original cohort, nine RHGs were used. For the Cycle 4 sample, five RHGs were used. The Cycle 5 sample also used five RHGs. Although Cycles 4 and 5 were modelled together, their nonresponse adjustments are calculated separately, for the simple reason that they each represent their own distinct reference populations.

Second adjustment: post-stratification adjustment

For the sample of children selected in Cycle 1, the reference population is the population of all children aged 0 to 11 years old as of December 31st, 1994. For the sample of children selected in Cycle 4, the reference population is that of all children aged 0 and 1 year as of December 31st, 2000. Finally, for the sample of children selected in Cycle 5, the reference population is that of all children aged 0 and 1 year old as of December 31st, 2002. Each group was post-stratified to the relevant age-sex-province population counts.

11.3.2 Funnel weighting

We only produce the funnel weight, FWTCdW1L, for:

- The original cohort of 0- to 11-year-old children selected at Cycle 1

Definition of a funnel respondent

A funnel respondent is a longitudinal respondent at Cycle 6 (see definition above) who was also a respondent at all previous cycles; in this case, a respondent at Cycles 1 through 6. Notice that the all-cycle respondents are a subset of the Cycle 6 respondents.

First adjustment: nonresponse adjustment

Once again the nonresponse adjustment is based on the creation of RHGs. The nonresponse model is created using LFS variables and a co-operation score. Using the model output, RHGs are generated. Note that these RHGs are different from the RHGs created for longitudinal weighting since Cycle 6 respondents who did not respond in all of the previous cycles have a different nonresponse mechanism than respondents to all cycles. The adjustment factor is computed for each RHG. In total, 11 RHGs were used.

Second adjustment: post-stratification adjustment

The reference population is again children aged 0 to 11 years old as of December 31st, 1994. The adjustment is computed for each age-sex-province combination.

11.3.3 Cross-sectional Weighting

The cross-sectional weight, FWTCW01C, applies to four separate samples of NLSCY children:

- The ECD cohort of 0- to 1-year-olds selected at Cycle 6,
- Returning 2- to 3-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 5,
- Returning 4- to 5-year-olds from the ECD cohort of 0- to 1-year-olds selected at Cycle 4,
- A new top-up of 2- to 5-year-olds

Cross-sectional weighting for 0- to 5-year-old children involves representing the population at the time of Cycle 6 collection.

Definition of a cross-sectional respondent

A cross-sectional respondent is a child whose adult component or child component is complete. In contrast to longitudinal respondents, children who were introduced in a previous cycle and died or moved outside Canada's 10 provinces are out-of-scope. They are not in the target population.

First adjustment: nonresponse adjustment

Children in the NLSCY aged 0 to 5 years old in 2004 were given a cross-sectional weight. The nonresponse model used LFS variables. In total, 10 RHGs were used to perform the nonresponse adjustment for these children.

Second adjustment: post-stratification

The reference population is children aged 0 to 5 years old as of December 31st, 2004. The adjustment is computed for each age-sex-province combination.

12.0 Data Quality, Response Rates and Coverage

This chapter provides the user with information about the various factors affecting the quality of the survey data. There are two main types of error: sampling error and non-sampling errors. We will pay special attention to non-sampling errors in this chapter.

Also, more general information on survey data quality and quality assurance is available at www.statcan.ca.

12.1 Sampling Error

The estimates derived from this survey are based on a sample of children. If we had done a census of the target population with the same questionnaires, interviewers, supervisors, processing methods and so on, we might have obtained slightly different values. The difference between the estimates produced by a sample and the estimates obtained through complete enumeration under similar conditions is known as the sampling error of the estimates.

Sampling error can be estimated using the sampling variance. For more details on calculating the estimated sampling error, see Chapters 13.0 and 17.0.

12.2 Non-sampling Errors

There are many sources of non-sampling errors in any survey. Interviewers may misunderstand survey instructions; respondents may make mistakes in answering the questions; responses may be recorded in the questionnaire incorrectly; and errors may be made in processing the data. These examples of non-sampling errors are difficult to quantify. Other kinds of error, especially nonresponse and the coverage of the intended population, are more easily quantifiable.

Non-sampling errors can cause bias, defined as a difference between the expected survey estimated value and the true population value. As the true population values are not known, it is very difficult to measure bias.

12.3 Total Nonresponse and Nonresponse Bias

In surveys, nonresponse results from the inability to obtain a set of measurements for a given unit in the sample. Nonresponse can be classified into two types: total (unit) nonresponse, and partial (item) nonresponse. Unit nonresponse arises when none of the survey measurements for a given unit are available. Such a unit is labelled a nonrespondent. Item nonresponse is characterized by the inability to gather *some* measurements, but enough measurements are observed to qualify the unit as a respondent. This section focuses on unit nonresponse and Section 12.4 discusses item nonresponse.

Nonresponse is a situation that can lead to bias in the survey estimates. Biased estimates can result if nonrespondents have significantly different characteristics from respondents. Both the amount of nonresponse and the degree to which the nonrespondents would have reported different answers than the respondents affect the amount of bias in the estimate. We are unable to accurately measure what the nonrespondents “would have reported”, but we can measure the level of nonresponse. Later in this section, cross-sectional response rates and longitudinal attrition rates are given. More details on the weighting procedure and how it attempts to adjust for total nonresponse are given in Chapter 11.0, and a general discussion of bias can be found in the Chapter 13.0.

Our weighting procedures adjust the sampling weights to attempt to reduce the potential bias due to nonresponse. However, this practice is based on certain assumptions, and it does not guarantee that there is no bias due to nonresponse.

For the National Longitudinal Survey of Children and Youth (NLSCY), response homogeneous groups (RHG) were created such that the weights of respondents will compensate for the nonrespondents having similar predicted propensities to respond, where this predicted propensity to respond is based on previously collected characteristics, e.g., education level and type of dwelling. Still, within any given RHG, the nonrespondents may differ from the respondents in important unobserved or unknown ways.

Nonresponse cumulates over time. As we have fewer and fewer participants, the estimated sampling error increases, but the potential for bias also increases. After many cycles, it would be highly improbable that the participants who continue to co-operate are a random sub-sample of the Cycle 1 respondents.

In fact, in extreme cases, certain subsets of the population may no longer be represented by the remaining sample. For a purely hypothetical example, if the initial sample contained 20 girls with autism in some province, yet none of these 20 responds at Cycle 6. Regardless of the weighting procedure, the survey could no longer produce estimates for autistic girls in that province.

12.3.1 Response Definitions

There are two distinct types of response rates: **collection phase response rates** measure the effectiveness of the data collection process and are based on the units actually sent to the field for collection.

Estimation phase response rates are an indicator of the quality of the estimates produced. Estimation phase response rates are in terms of the statistical unit (for the NLSCY, this is the child), and intend to show the degree to which data are missing.

For the NLSCY, the key difference between the two types of response rates is the children who remain in scope for the survey, but who were not part of the Cycle 6 sample. Examples are hard refusals or chronic refusals in previous cycles for the original cohort and any nonresponse in previous cycles for the Early Childhood Development (ECD) cohort. These are removed from the Cycle 6 sample, and their absence affects data quality. Not being sent to the field in Cycle 6, these children are not considered in the collection phase rates. Some child-level collection phase rates are given in Chapter 6.0, and some household level rates are given in Chapter 5.0. The estimation phase rates given in this section do not contradict those collection phase rates, but are intended to complement them and give slightly different information.

Weighted versus unweighted rates

Often, estimation phase response rates are weighted to reflect the idea that nonresponse from certain units, i.e., those with large design weights in social surveys; those with large influence on estimates in business surveys, is more damaging to the survey estimate than nonresponse from less influential units. This is valid reasoning.

There are advantages to using unweighted response rates, too. They are more easily defined and produced (and perhaps interpreted), not affected by revisions to the survey weights, and have been used in past cycles of the NLSCY. Further, within provinces, we observe very little difference between the weighted and unweighted rates for the NLSCY.

The response rates given in this chapter are **unweighted**.

Definitions

The following terms are relevant to understanding the tables provided in this chapter.

Different survey components are administered for children and youth of different age groups. Likewise, the criteria for being considered a respondent varies by the age of the selected respondent.

The child component is a computer-assisted interviewing (CAI) component where questions about the child are asked to the person most knowledgeable (PMK). The PMK component is a CAI component where questions about the PMK are asked to the PMK. The spouse component is a CAI component where questions about the spouse are asked to the spouse of the PMK. The PMK and spouse components are sometimes referred to as adult components. The youth component is a CAI component where questions about the youth are asked to the youth.

A child with Effective Age 0 to 15 is considered a **respondent** if the child component or an adult component (either PMK or spouse) of the survey is completed. A component with a set of key questions answered is considered **completed**. Note that substantial item nonresponse (see section 11.4) can be present within components classified as completed.

A youth with Effective Age 16 or 17 is considered a **respondent** if the child component or youth component or an adult component (either PMK or spouse) of the survey is completed.

A youth with Effective Age 18 or older is considered a **respondent** if the youth component of the survey is completed. There is no adult component administered for youth 18 or older.

An **out-of-scope child** is a child that is selected into the sample, but is not part of the survey population. A child may be out-of-scope either because he/she is deceased, residing outside of Canada, or an inmate of an institution. In contrast, an **in-scope child** is a child who is selected into the sample and is part of the target population. The sum of the number of out-of-scope and in-scope children equals the sample size. Note that it is possible for children to be cross-sectionally out-of-scope but to be longitudinally in-scope. This situation occurs for example with children who are deceased or children who have moved out of the country. Cross-sectionally, these children do not represent anyone in the target population. However, longitudinally, these children represent other children in the same situation who were present in the longitudinal target population when first selected in the survey. On the other hand, it is also possible to have children who are cross-sectionally in-scope but longitudinally out-of-scope, e.g. children in the top-up sample.

A **nonrespondent** is an in-scope child or youth who does not meet our response criteria. Nonresponse can occur because the targeted participants refused to do the survey (or did not answer sufficiently), because the child or youth could not be traced, or because the interviewer was unable to complete the interview for other reasons.

The **response rate** is defined as the number of respondent children or youth over the number of in-scope children or youth.

12.3.2 Cross-sectional Response Rates

The cross-sectional Cycle 6 sample is the aggregation of children introduced in three different cohorts, namely Cycles 4, 5 and 6. Only respondents from the previous cycle were sent to the field in Cycle 6.

Unweighted Cross-sectional Response Rates by Sample Cohort

Effective Age	Initial sample	Out-of-scope	In-scope sample	Respondents	NR C4	NR C5	NR C6	Estimation Response rate (%)
0 and 1	4,356	13	4,343	3,521	na	na	822	81.1
2 and 3 returners	4,492	117	4,375	2,866	na	1,142	367	65.5
2 and 3 top-up	734	12	722	595	na	na	127	82.4
4 and 5 returners	5,087	144	4,943	2,964	1,125	505	349	60.0
4 and 5 top-up	703	5	698	568	na	na	130	81.4
Total	15,372	291	15,081	10,514	1,125	1,647	1,795	69.7

Column definitions for the above table

Effective Age: This is based on year of birth. Those with Effective Age = 0 were born in 2004; those with Effective Age = 1 were born in 2003; etc.

Initial sample: This is the total number of children selected from the frame. This includes children who were classified as out-of-scope at the initial or any subsequent collection.

Out-of-Scope: This is the count of cross-sectionally out-of-scope children at Cycle 6. The survey may have classified the child as out-of-scope at the time of Cycle 6 collection, at Cycle 5 collection, or at Cycle 4 collection. As expected, the proportion of out-of-scope children is small for the sample selected at Cycle 6, and larger as more time passes since the initial selection. As more time passes there is a corresponding longer opportunity for emigration.

In-scope sample: This count forms the denominator of the response rate. See the previous section for a complete definition.

Respondents: This count is the numerator of the response rate. These children meet our response criteria given in the previous section.

NR C4: This is the count of nonrespondents at Cycle 4 collection among those children in-scope at Cycle 6. These children were dropped from the Cycle 5 and Cycle 6 samples.

NR C5: This is the count of nonrespondents at Cycle 5 collection among those children in-scope at Cycle 6. These children were dropped from the Cycle 6 sample.

NR C6: This is the count of nonrespondents at Cycle 6 collection.

This table illustrates the difference between the collection phase and estimation phase response rates. For the returners with Effective Age of 4 or 5, the response rate among the children in the Cycle 6 sample was close to 90%, but when one takes into account that several hundred in-scope children were dropped from the Cycle 6 sample due to nonresponse at a previous cycle, the estimation phase response rate is substantially lower at 60%. With these lower response rates, we grow more concerned about the potential for nonresponse bias discussed earlier in this chapter.

In light of this, we have modified our collection strategy for the Cycle 7 collection. For Cycle 7, nonresponding children from previous cycles were sent to the field for a new attempt. This should improve the estimation phase cross-sectional response rates at Cycle 7.

The following table is a summary of the earlier table *Unweighted Cross-sectional Response Rates by Sample Cohort*, with the returning and top-up samples combined.

Unweighted Cross-sectional Response Rates by Effective Age

Effective Age	In-scope sample	Respondents	Estimation Response rate (%)
0 and 1	4,343	3,521	81.1
2 and 3	5,097	3,461	67.9
4 and 5	5,641	3,532	62.6
Total	15,081	10,514	69.7

The next table shows these rates by province. As in past cycles, Ontario has the lowest response rate.

Unweighted Cross-sectional Response Rates by Province

Province	In-scope sample	Respondents	Estimation Response rate (%)
Newfoundland and Labrador	781	618	79.1
Prince Edward Island	645	494	76.6
Nova Scotia	931	701	75.3
New Brunswick	995	702	70.6
Quebec	2,288	1,589	69.4
Ontario	4,230	2,746	64.9
Manitoba	1,178	847	71.9
Saskatchewan	1,063	819	77.0
Alberta	1,439	1,038	72.1
British Columbia	1,431	960	67.1
Canada	15,081	10,514	69.7

The *In-scope sample* column is based on province of selection – where the child was chosen for the survey. The *Respondents* column is based on the province of residence at Cycle 6 – known only for respondents.

12.3.3 Longitudinal Attrition Rates Cycle 1 Cohort

In a longitudinal survey, the longitudinal response rate shows proportion of respondents remaining in the survey. Normally, this rate is represented by the ratio of the number of longitudinal children who responded in the current cycle to the number of children that were *selected* in the first cycle. However, since the sampling method used in the first two collection cycles differs from the sampling method for subsequent cycles, it is not possible to obtain an accurate longitudinal response rate that incorporates nonresponse at Cycle 1 for children introduced at Cycle 1. Specifically, the actual number of eligible children in nonresponding households is unavailable. Therefore, the denominator needed to determine the response rate is also unknown.

Instead, we define the **attrition rate** as the proportion of proportion of respondents remaining in the survey relative to the number of *respondents* at Cycle 1.

Here are some points related to the table below:

- The denominator for all the percentages shown in these tables is the number of responding children in Cycle 1 who were followed in Cycle 2. Note that since the sample size has been reduced from Cycle 1 to Cycle 2, not all Cycle 1 responding children are considered longitudinal.
- Those receiving a longitudinal weight at Cycle 6 can be classified into two categories: those with reported data and those without reported data. Children who have died or moved outside the ten provinces belong in the second category. These children do not have reported data but still have a longitudinal weight since they represent other children in the longitudinal population in the same situation. More commonly, a child or youth with a longitudinal weight has responded to the survey. Those with reported data appear in the *Collected data for Cycle 6 column* and are a subset of those given a longitudinal weight.
- In order to retain the highest possible number of children from the original cohort, attempts are made to convert children who did not respond in a previous cycle. It is then possible to have children who have not responded in a particular cycle but have responded in the current cycle.
- The *Possess funnel weight at Cycle 6* column shows the number of children who have never been considered nonrespondents. That is, they have received a longitudinal weight on each cycle's master file. Again, the majority of these children reported data in each cycle. A much smaller number, those who have died or moved outside the ten provinces, retain a weight without reported data. See the Weighting Chapter for more details about the funnel weight, variable FWTCWd1L.
- For those in the first row of the table, the primary respondent is the PMK. Youth 18 or older respond for themselves.

Unweighted Longitudinal Attrition Rates for Cycle 1 Children by Effective Age

Effective Age at Cycle 6	Number of Respondents in Cycle 1	Collected data for Cycle 6		Possess longitudinal weight at Cycle 6		Possess funnel weight at Cycle 6	
		Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1
10 to 17	12,027	8,196	68.1	8,399	69.8	7,364	61.2
18 to 21	4,876	2,982	61.2	3,084	63.2	2,490	51.1
Total	16,903	11,178	66.1	11,483	67.9	9,854	58.3

The following table shows attrition up to Cycle 5. In the table below, “Respondents” are those with a longitudinal weight including those without reported data. At Cycle 5, 74.1% of the Cycle 1 respondents had a longitudinal weight. At Cycle 6, this drops to 67.9%, and is lower for those 18 or older.

Unweighted Longitudinal Attrition Rates (Cycles 2 to 5) for Cycle 1 Children, by Province

Province in Cycle 1	Number of Respondents in Cycle 1	Respondents in Cycle 2		Respondents in Cycle 3		Respondents in Cycle 4		Respondents in Cycle 5		Respondents in all cycles	
		Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1
Newfoundland and Labrador	950	892	93.9	845	88.9	777	81.8	755	79.5	689	72.5
Prince Edward Island	467	443	94.9	434	92.9	392	83.9	364	77.9	330	70.7
Nova Scotia	1,191	1,068	89.7	1,085	91.1	988	83.0	903	75.8	811	68.1
New Brunswick	1,070	958	89.5	958	89.5	836	78.1	792	74.0	691	64.6
Quebec	3,182	2,944	92.5	2,844	89.4	2,522	79.3	2,361	74.2	2,108	66.2
Ontario	4,342	3,899	89.8	3,760	86.6	3,318	76.4	3,104	71.5	2,714	62.5
Manitoba	1,232	1,161	94.2	1,112	90.3	1,019	82.7	1,004	81.5	891	72.3
Saskatchewan	1,413	1,305	92.4	1,257	89.0	1,073	75.9	1,002	70.9	893	63.2
Alberta	1,599	1,465	91.6	1,420	88.8	1,242	77.7	1,162	72.7	1,031	64.5
British Columbia	1,457	1,333	91.5	1,282	88.0	1,143	78.4	1,076	73.9	978	67.1
Canada	16,903	15,468	91.5	14,997	88.7	13,310	78.7	12,523	74.1	11,136	65.9

12.3.4 Longitudinal Response Rate for Children Selected in Cycle 4

The response rate given in the Section 12.3.2 table *Cross-sectional Response Rates* in the 4- and 5-year-old returners row, gives a good indication of the longitudinal response rate as well.

12.3.5 Longitudinal Response Rate for Children Selected in Cycle 5

The response rate given in the Section 12.3.2 table *Cross-sectional Response Rates* in the *2- and 3-year-old returners* row, gives a good indication of the longitudinal response rate as well.

12.4 Partial Nonresponse

The previous section on total nonresponse dealt with the issues of representation of the sample for cross-sectional or longitudinal estimation. These types of non-sampling errors can usually be dealt with effectively by adjusting the survey weight to reflect the survey population. Other types of nonresponse are also measured in this survey and are usually not corrected through an adjustment of survey weights.

Although a person may provide enough information to qualify as a respondent, some of the questions (variables) may still be not answered, resulting in partial nonresponse. Some reasons for this are: (in no particular order) co-operation from some, but not all, of the targeted participants within the household, unwillingness to answer sensitive questions, respondent fatigue, accidental skipping of parts of the questionnaire, or operational difficulties.

Component nonresponse can happen when one individual participates, but others do not. For example, in the household of a selected 17-year-old, the PMK may co-operate and answer the Child and PMK components, but her spouse may refuse to do his spouse component, and the selected youth may refuse to complete the youth component. By our definitions, this youth is considered a respondent and a record exists for him on the master file, but we have partial nonresponse. Another cause of partial nonresponse is when the telephone portion is collected but the other components are missing.

Usually, the nature of partial nonresponse depends on the subject matter. For instance, the Motor and Social Development module, for children aged 0 to 3, is thoroughly answered since parents have a greater interest in this topic, whereas the questions on income may be considered too personal by some respondents, resulting in some partial nonresponse.

Item nonresponse is measured at the variable level and represents information that was not collected from the respondent at the time of the interview. This type of nonresponse is left uncorrected except where specifically noted by imputation flags. Item nonresponse is detailed in the code book with categories such as “Refusal” or “Not stated”. The “Don’t know” category is regarded as a nonresponse during analysis, but some analysts may consider it a valid response depending on the information sought and the interpretation of specific variables. For analytical purposes, researchers should remember that the “Refusal” and “Don’t know” categories are used when the respondent was questioned about this particular piece of information, while the “Not stated” category usually indicates that the respondent was not asked for the information. This is true for computer-assisted response capture but not for self-completed paper questionnaires. For the latter, “blank responses” are categorized as “Not stated” even though the respondent may have seen the question.

Note that the “Valid skip” category is not a nonresponse but a valid skip of a particular piece of information for a particular respondent. For example, many questions are age specific and children outside the targeted age group have “Valid skip” for those variables.

For item level details about item nonresponse, consult the code book that accompanies the microdata file. Some of those classified as “component respondents” may have answered only a portion of the component.

Analysts using NLSCY data should be aware of how partial nonresponse affects the data they are attempting to analyze. As in the case of total nonresponse, partial nonresponse may be higher for respondents with a particular characteristic, e.g., teenagers doing poorly in school may be more reluctant to fill out the Self-complete Questionnaire. This leads to bias, and if severe, can cast analytical results into question. There are techniques available to deal with partial nonresponse, for instance, re-weighting and imputation. Users are strongly encouraged to make themselves aware of the extent of the nonresponse in the analysis they are doing and, if appropriate, take corrective measures to compensate for the nonresponse. At minimum, they should detail the impact of component or item nonresponse in their findings. This is also discussed later in the Analysis Chapter.

The following sections will explore the issue of component nonresponse for the NLSCY. This is intended to inform researchers who use these variables in their analysis of possible sources of error not remedied by the survey weights. All rates in the following tables are unweighted and the denominator is the count of children eligible for that component among those who are considered respondents.

Further, this has been written before the release of the master file, and small differences may exist between what is stated here and what is eventually seen on the final master files.

12.4.1 Child Component

The child component is a computer-assisted interviewing (CAI) component where questions about the child are asked to the PMK. All children with Effective Age 0 to 17 are eligible except those 16- and 17-year-olds who are living independently.

Note that the *Number eligible* is based on the number of children with reported data – longitudinal in-scope children who did not report in Cycle 6, who nonetheless appear on the master file with a longitudinal weight, are not included.

The *Number answered* column includes fully completed and partially completed components.

Number eligible	Number answered	Component response rate
18,684	18,561	99.3%

Among survey respondents, it is exceedingly rare that we are missing the entire child component. The children without an answered child component have an answered adult component or youth component.

12.4.2 Person Most Knowledgeable Component

The PMK component is a CAI component where questions about the PMK are asked to the PMK. All children with Effective Age 0 to 17 are eligible except those 16- and 17-year-olds who are living independently.

This table is in terms of children, not adults. In households with two selected children, there is a single PMK. We desire PMK information from 16,491 individuals. In terms of children, we desire 18,684 children to have PMK information.

Number eligible	Number answered	Component response rate
18,684	18,349	98.2%

Among survey respondents, it is rare that we are missing the entire PMK component.

12.4.3 Spouse Component

The spouse component is a CAI component where questions about the spouse are asked to the spouse. All children with Effective Age 0 to 17 whose PMK has a partner are eligible. Those 16- and 17-year-olds who are living independently have no PMK and, of course, there is no spouse of the PMK.

This table is in terms of children, not adults. In households with two selected children, there is one spouse answering. We desire spouse information from 14,017 individuals. In terms of children, we desire 15,867 children to have spouse information. Children living with a single parent are not eligible for the spouse component.

Number eligible	Number answered	Component Response rate
15,867	15,411	97.1%

Among survey respondents, it is rare that we are missing the entire spouse component. Note that proxy responses, where the PMK could answer on behalf of the spouse, were permitted for this component. See Section 12.10 for more details on proxy responses.

12.4.4 Youth Component

The youth component is a CAI component where questions about the youth are asked to the youth.

The component response rate is only meaningful for youth aged 16 to 17. Youth 18 and older must complete the youth component to be considered a response.

Number eligible	Number answered	Component Response rate
1,585	1,457	91.9%

Note also that of the 1,457 who answered at least partially, 1,436 answered the component enough to proceed with imputing missing income.

Among the 16- and 17-year-old survey respondents, it was relatively common to have participation of an adult, but no participation from the youth.

12.4.5 PPVT-R

The Peabody Picture Vocabulary Test – Revised (PPVT-R) was administered to children aged 4 or 5 years old.

Number eligible	Number answered (enough to get score)	Component Response rate
3,532	3,184	90.1%

Recall that from the total nonresponse section, there were 5,641 in-scope children in this age group. The PPVT was completed by 3,184 children. The total and partial nonresponse together mean that 56.4% of children for which a PPVT test was desired completed the test. The total nonresponse is compensated for by the survey weights, but the component non response is not.

12.4.6 Number Knowledge

The Number Knowledge assessment was administered to children aged 4 or 5 years old.

Number eligible	Number answered (enough to get score)	Component Response rate
3,532	3,167	89.7

12.4.7 Who Am I?

The Who Am I? assessment was administered to children aged 4 or 5 years old.

Number eligible	Number answered (enough to get score)	Component Response rate
3,532	2,798	79.2

One can speculate why the component response is worse for this assessment compared to the PPVT and Number Knowledge. Perhaps this assessment was too difficult for some eligible children. This rate is slightly lower than the component response rate observed for Cycle 5.

12.4.8 Mathematics Tests

The NLSCY mathematics tests are made of 20 computational questions answered in the home by respondents aged 10 to 15. The level of test (ranging from 4 to 10) was determined by the child's grade. If the grade was not known, the child's Effective Age determined which level of test was administered.

Number eligible	Number answered (enough to get score)	Component Response rate
6,611	5,825	88.1

12.4.9 Problem Solving

The Problem Solving Assessment for 16- and 17-year-olds in the NLSCY consisted of 20 questions to assess strengths in reading comprehension, problem solving and decision making. It also tests some mathematical skills.

Number eligible	Number answered (enough to get score)	Component Response rate
1,585	1,290	81.4

12.4.10 Literacy Assessment

The Literacy Assessment for 18- and 19-year-olds consisted of 36 questions with an emphasis on extracting information from texts, tables and graphs. The test required a personal visit while the youth component could be completed by phone.

Number eligible	Number answered (enough to get score)	Component Response rate
1,573	1,329	84.5

12.4.11 Numeracy Assessment

The Numeracy Assessment for 20- and 21-year-olds consisted of 32 questions. It aims to test the ability of young adults to function in society and manage mathematical demands in diverse situations. The test required a personal visit while the youth component could be completed by phone.

Number eligible	Number answered (enough to get score)	Component Response rate
1,409	1,145	81.3

12.4.12 Self-complete Components

The Self-complete Component in the NLSCY is a short booklet comprising questions mostly of a private nature on topics such as misbehaviour, feelings, parents and puberty. The specific topics covered vary by age group. These are self-administered questionnaires that the child completes in private, away from both parents and interviewer. Questionnaires are returned in a sealed envelope to the interviewer during the visit.

Effective Age	Number eligible	Number answered	Component Response rate
10 to 11	2,825	2,544	90.1
12 to 13	1,995	1,782	89.3
14 to 15	1,791	1,562	87.2
16 to 17	1,585	1,281	80.8
Total	8,196	7,169	87.5

12.5 Cycle Nonresponse

Certain longitudinal respondents do not participate in every cycle. This is cycle nonresponse. When dealing with the longitudinal data for a respondent, data from every cycle are not necessarily available. For example, a child may be a respondent in Cycles, 1, 3, 4, and 6, but not Cycle 2 and Cycle 5.

If data from every cycle are crucial, the analyst can limit himself to children without cycle nonresponse and use the funnel longitudinal weights for this group, variable FWTCWd1L. This weight is available for children introduced in Cycle 1.

12.6 Response Errors Impact for Rare Characteristics

General population surveys are not well suited to measuring rare characteristics.

Survey response or recording errors do occur in the course of collection. As one simple example, of the several thousand interviews conducted, we expect that some percentage of respondents will not answer every question honestly. Other times, the interviewer may simply hit the wrong button. For most purposes, the effect of this type of misreporting is not large. For many variables, the errors “even out”, and the overall impact is minimal. However, if you are using the survey to make inferences about rare characteristics, events, or behaviours, these response errors can become relatively more important and influential. The errors are no longer expected to “even out”; instead, if response errors occur randomly, there is a systematic overestimation of the rare characteristic. Imagine a general survey where highest level of education is asked of 1,000 adults -- 995 without a PhD and five with a PhD. There are many more chances for a non-PhD to falsely report having a PhD than the other way around. Suppose that there is response error to this question at a rate of 0.2 % (0.2% of 995 is about 2 and 0.2% of 5 is very close to 0). The survey would estimate the proportion of PhDs to be 7/1,000 rather than 5/1,000. The difference is not large, but in relative terms, it is a substantial and worrisome 40% overestimation. There are techniques, like asking a series of questions instead of one question, that can reduce this effect, but these add length and complexity to the survey. With the broad content of the National Longitudinal Survey of Children and Youth (NLSCY), it was not practical or possible to devote this level of attention to every item collected.

Users of the NLSCY data wishing to study rare behaviours like heavy drug use or violent behaviour should keep this limitation in mind.

Also, for many variables, the assumption of random response error may not hold, particularly for responses seen as socially undesirable. This is discussed in Section 12.7. For example, the

chance that a non-violent youth falsely reports violent behaviour may differ from the chance that a violent youth falsely reports no violent behaviour.

12.7 Response Errors Related to Deviant Behaviour or Sensitive Questions

In an interview, respondents will not always be truthful about behaviours that are considered negative or abnormal. This is called social desirability bias. For example, parents who frequently use physical punishment may not respond truthfully when asked about this. Likewise, respondents may lie, and portray themselves and their children in an unrealistically positive way. For example, some parents may not answer honestly when asked about reading to the child, recognizing that they *should* do this frequently.

Since much of the survey data are reported by the respondents, rather than physically observed or measured, statements of survey results should make clear this distinction. For example, one cannot conclude from the NLSCY that “X % of children in Canada sometimes receive physical punishment”. In fact, the survey allows only statements like “X % of children in Canada are **reported** to sometimes receive a physical punishment”.

12.8 Response Errors Due to Approximations

It is perhaps obvious, but bears mentioning that certain collected values are often approximated by the respondent. Data users should be aware that variables measuring concepts, like income or height, which can properly be considered continuous in the population, do not necessarily retain these properties on the survey file. For example, we see many incomes reported as exact multiples of \$10,000, and many heights reported in exact feet (see chart in Section 12.10.4). In the population, the number of households with income \$19,501 to \$20,500 is probably comparable in size to the number of households with income \$20,501 to \$21,500. The survey results would show a very different picture with the first group many times larger than the second due to respondent approximation of income.

This phenomenon is also seen when asking about the child’s age at the time of some event. For example, we ask for the child’s age in years and months at the time of parental separation, but for the month component, zero months is by far the most frequently reported.

12.9 Response Errors Due to Memory Errors

Another type of response error occurs when the respondent cannot accurately recall the information, particularly when the reference period is long. For example, the respondent may not know exactly how many times the child visited a doctor in the past 12 months. Minor illnesses several months in the past may be forgotten. On the other hand, respondents may “telescope” major events and report them as occurring within the reference period, even when the event actually occurred before the reference period.

12.10 Response Errors Due to Collection by Proxy

The NLSCY allows proxy response for the adult components. This means that information about one person is given by another person, e.g., the child’s mother answers her own PMK component and the spouse component on behalf of her husband. One member of the couple usually is sufficiently knowledgeable about the other person to answer the questions appropriately. However, it is possible that the targeted person would have given different answers from those given by the proxy respondent.

Proxy rates are monitored by the NLSCY, but no detailed studies on proxy response patterns have been undertaken. Of course, for any given case, it is impossible to know what the non-proxy respondent would have reported.

Note the following table is in terms of children, not adults.

Component	Number of eligible children	Children with proxy responses	Proxy rate
PMK	18,684	309	1.7%
Spouse	15,867	9,644	60.8%

It is rare that the spouse responds to the PMK component, but it is quite common that the PMK will answer both adult components.

12.11 Response Patterns with Indefinite Response Categories

For many items on the NLSCY questionnaire, the response categories available are indefinite or not concretely and precisely defined, e.g., Never, Sometimes, or Often. One person's threshold between "Sometimes" and "Often" may be very different from another person's. The same is true for "Strongly agree" and "Agree". For this reason, we have the undesirable consequence that respondents with the same behaviour patterns will not necessarily have identical survey data. Generally, this does not mean that the data based on indefinite response categories are incorrect or unreliable, but caution is warranted when comparing different groups. One should be aware that differences in response patterns by region or ethnicity may not necessarily be due to true differences in the children. For example, there may be cultural patterns in the propensity to respond "Often" rather than "Sometimes".

12.12 Language of Interview

Due to the nuances of language, exact translation of some phrases and questions is not possible. This can introduce artificial differences in the survey results when there is no true difference in the populations. Also note that interviewers can switch between English and French during an interview. The language variable gives the primary language of the interview, but some questions could have been posed in the other language.

Also, a small number of interviews are conducted in languages other than French or English with the interviewer translating the questions into the respondent's preferred language.

12.13 Conflicting Information

Occasionally, respondents give conflicting information. In some cases, the inconsistency can be resolved through deterministic edit rules. For example, if a respondent reports year of immigration less than her year of birth, the year of immigration is set to the year of birth.

In other cases the inconsistency cannot be easily resolved. For example, a respondent may answer "Yes" to "Does your child say eight or more words in addition to 'Mama' and 'Dada'?" in the Ages and Stages module, and also answer "No" to "Has he/she said two recognizable words besides 'Mama' or 'Dada'?" in the Motor and Social Development module. Clearly, these responses are inconsistent, but such situations are left unchanged.

It is frustrating that the collected information is inconsistent, but since we cannot confidently render it consistent and accurate, inconsistencies remain on the final survey files.

The data from the current cycle can also conflict with what has been collected in past cycles. For example, for some children, a parental separation was reported at Cycle 1, but at a later cycle the person most knowledgeable (PMK) reports that the parents have lived together continuously since the child's birth. There are also instances where, over the course of the survey, more than one person has reported being the biological mother or father of the child. In such cases, we accept what has been reported in the current cycle.

The results from the NLSCY can also conflict with other sources. Definitions and concepts may not be exactly compatible, or different practices may have been used in collection. It is also possible that an error has occurred in the processing of the microdata file.

12.14 Data Quality for Body Mass Index

12.14.1 Body Mass Index

Body Mass Index (BMI) is a standardized scale to measure body mass. A BMI score is calculated by dividing weight by height squared:

$$\text{BMI} = \frac{\text{Weight in Kilograms}}{(\text{Height in Metres}) \times (\text{Height in Metres})}$$

The height and weight variables used to derive BMI from NLSCY data are:

- PMK reported for 2- to 11-year-olds - FHLCQ03B (height) and FHLCQ04A (weight) yield the respondent's BMI score, FHLCEs01.
- Self-complete for 12- to 17-year-olds - FHTCbQ01 (height) and FHTCbQ02 (weight) yield the respondent's BMI score, FHLCEs01.
- Self reported by the youth in the youth component for 18- to 21-year-olds - FHTYeD01 (height) and FHTYeD02 (weight) yield the respondent's BMI score, FHTYeS03.

By calculating a BMI score, this score can then be compared with others to see into which percentile it falls. Differing cutoffs or percentile ranges have been proposed to help identify whether one's BMI score is classified as underweight, normal, at risk of overweight, overweight, or obese. The United States Centers for Disease Control (CDC) has proposed cutoffs for children, youth, and adults. Similarly, Tim Cole et al. have proposed international cutoffs for children and youth using a different methodology.

12.14.2 Body Mass Index – Centers for Disease Control

The CDC have proposed a set of percentile ranges to classify BMI scores as either: underweight, normal, at risk of overweight or overweight. These percentile ranges are age-specific by sex, and are based on American height and weight data. The CDC cutoffs are based on the person's age broken down into one month intervals. Consequently, in processing the NLSCY data, the age in months variable (FMMCdQ1B) was used to derive the cutoffs. The percentile ranges proposed by the CDC can potentially be used for 0- to 20-year-olds. The release name for this variable is FHLCEd03 for those 17 or younger and FHTYeD05 for those 18 or older.

More information on the CDC BMI cutoffs for children and youth can be obtained at the following website: <http://www.cdc.gov/nccdphp/dnpa/bmi/bmi-for-age.htm>.

12.14.3 Body Mass Index – Tim Cole, et al.

A set of international BMI cutoffs for 2- to 18-year-olds were proposed by Tim Cole, Mary Bellizzi, Katherine Flegal, and William Dietz in the British Medical Journal (Establishing a standard definition for child overweight and obesity worldwide: international survey. *BMJ* 2000;320:1–6). These cutoffs classify BMI scores into three categories: normal, overweight or obese. Similar to the CDC cutoffs, these categories are age and sex specific. In contrast to the CDC cutoffs, the international cutoffs are in half-year intervals for age and were based on studies from six nationally representative datasets of body mass indices in childhood (United States, Brazil, Great Britain, Hong Kong, Netherlands, Singapore). The release name for this variable is FHLCeD02 for those 17 or younger and FHTYeD04 for those 18 or older.


More information on the BMI cutoffs proposed by Tim Cole et al. can be obtained at the following website: <http://www.bmj.com>

12.14.4 Body Mass Index – Data Quality

There are several issues that affect the quality of the BMI scores. First, there is a higher rate of nonresponse for the BMI variables as it is necessary that both the height and weight variables contain valid responses in order to calculate a score. Second, the data collected for height and weight are based solely on estimates provided by the parent or the youth rather than on accurate clinical measurements. The result of this method of collection, particularly parent reporting, is less accurate for height and weight and correspondingly less accurate for the BMI. Typically, a respondent will round the values of height or weight that they report, which leads to different BMI values than would otherwise have been calculated based on clinical measurements. For example, a PMK will report the child as being 5' (feet) tall instead of 5'2"(inches) or 4'10", or maybe they will report that the child's weight is 110 pounds instead of 113 or 108. These small errors in estimated height and weight can translate into a much larger degree of error in the BMI resulting in a change in classification from 'overweight' to 'normal' or 'obese' depending on how height and/or weight was rounded.

One can see in the chart below that many more children were reported as being exactly three feet tall than were reported as close to three feet tall.

Cycle 6: Unweighted counts of some reported heights in inches



12.15 Conditioning Bias

Participants in a longitudinal survey may act differently because they know that they are in the study. Further, the process of answering the questionnaire has the potential to affect the behaviour of respondents. For example, after being asked about frequency of reading to the child at Cycle 1, the parent may decide to read more frequently to the child. This parent is no longer representative of other Canadian parents who have not participated in the survey – participating in the survey has affected her behaviour.

There is also the possibility that respondents may answer in ways known to reduce the interview length. Respondents may realize that answering “Yes” to certain questions triggers a series of detailed follow-up questions and may not answer such questions truthfully.

Though expected to be negligible, it is impossible to precisely measure these biases.

12.16 Person Most Knowledgeable

At each cycle, one adult in the household is identified as the person most knowledgeable about the child. The PMK answers the Child component, giving information about the child’s health, education, behaviour, etc. The child’s characteristics are measured indirectly as reported by the PMK. From cycle to cycle, the PMK can change. For a given child, perhaps the mother was the PMK at Cycle 1, then the father at Cycle 2 and Cycle 3, and then the mother again at Cycle 4, Cycle 5, and Cycle 6. Observed changes in the child’s characteristics may be due, in part, to having different people answer these questions. The child’s environment is not truly different, but the person answering the questions has changed and, naturally, has a different perspective. The variable FDMCD06 shows the relationship of the PMK to the child.

12.17 Coverage

To speak about coverage, some definitions are required.

The **target population** is the population for which information is desired.

The **survey population** is the population actually covered by the survey. Various survey populations for the NLSCY are described in Chapter 5.0.

The **survey frame** provides the means of identifying and contacting the units of the survey population.

Differences between the survey population and target population often arise from limitations of the survey frame. Differences can also arise from the operational realities of collecting the survey information. For example, a “snapshot” survey cannot be collected instantaneously; the passage of time between the selection of the sample and the collection, as well as the passage of time during collection can affect the survey population. The extent to which the survey population agrees with the target population is the **coverage** of the survey.

For example, imagine the target population for some opinion survey is all eligible voters in Canada. The survey frame is a random digit dialling (RDD) procedure. The survey population becomes all eligible voters in Canada who live in a household with a telephone. The survey frame imposes some exclusions to the original target population.

Frame errors (or coverage errors) consist of omissions, erroneous inclusions, duplications, and misclassification of units in the survey frame. After stating the coverage and defining the survey frame and survey population, the frame may, in practice, have certain flaws. For example, the survey frame of the list of registered births in Canada theoretically includes all births, but, in practice, this list may not be up to date. The missing births constitute a frame error.

12.17.1 Coverage for the National Longitudinal Survey of Children and Youth

The target population for the original cohort of the NLSCY can be precisely defined as children born from 1983 to 1994 and living in the ten provinces of Canada on December 31, 1994.

The cross-sectional target population for the Early Childhood Development (ECD) cohort can be defined as children born from 1999 to 2004 and living in the ten provinces of Canada on December 31, 2004.

The survey frame for the NLSCY is households who have participated in the Labour Force Survey (LFS). This frame imposes certain restrictions on our survey population. Specifically, as described in Chapter 5.0, the LFS excludes persons living on Indian Reserves, full-time members of the Canadian Armed Forces and inmates of institutions.

One obvious difference between the target population and the survey population for the NLSCY is that children living on reserves are not part of the survey population. This and other issues relating to frame errors and coverage are discussed in the following section.

12.17.2 Coverage and Frame Issues

1) Labour Force Survey exclusions

Note that children living in the Territories are excluded from our target population and survey population. Further, the Labour Force Survey excludes persons living on Indian Reserves, full-time members of the Canadian Armed Forces and inmates of institutions. Of these, the exclusion of reserves is most pertinent for the NLSCY. Though, in total, children living on reserves are not a large proportion of the population (around 2%), the impact varies considerably by province. According to the 2001 Census, in Saskatchewan and Manitoba, close to 10% of children 0 to 5 years of age live on reserves.

2) Eligibility for the NLSCY is based on the values of Labour Force Survey variables

Frame errors do exist on the NLSCY frame. For example, on the LFS file, children can be misclassified by year of birth. Each cycle, a small number of children who are not truly eligible for the NLSCY appear to be eligible based on the LFS information. Likewise, we assume there are a small number of children who are eligible for the NLSCY, yet appear to be ineligible according to the LFS information.

To be eligible for the NLSCY, the child must live in a household that responds to the LFS. NLSCY eligible children living within nonresponding households selected by the LFS are not surveyed for the NLSCY – primarily because we cannot identify children unless the household responds to the LFS. The survey weights account for this nonresponse, but a potential for bias exists. Not surprisingly, LFS nonresponse can affect the NLSCY quality.

The response rate for the LFS is consistently high (around 90% to 93%) and varies slightly month to month, which reduces the likelihood of a substantial bias associated with LFS nonresponse.

3) Immigration

For part of the ECD file, between the time when the sample was selected and when the Cycle 6 collection started, a period of two or four years has elapsed. Immigrant children who arrived in Canada during that period were not eligible for selection.

For example, at Cycle 6, the entire 4- and 5-year-old sample in Ontario was selected as 0- and 1-year-olds in Cycle 4 (selected in 2000). Any child born in 1999 or 2000, who arrived in Canada to settle in Ontario in 2001 or later, had no chance of being selected for the NLSCY.

The following table attempts to show the differences between the cross-sectional and longitudinal target populations. The *Total Number of Children* column is the total number of children aged 2 to 5 in January 2005. In other words, this is the size of the cross-sectional target population for this age group. The *Longitudinal Target Size* column is the sum of the number of 0- and 1-year-old children present in January 2001 and the number of 0- and 1-year-old children in January 2003. Basically, this is the size of our longitudinal target population.

The source of both columns of control totals are Census based projections of the LFS eligible population. Totals from this same source, at a different level of detail, are used to post-stratify the Labour Force Survey.

Difference Between Estimates of Total Population for the 2- to 5-Year-old Age Group by Province for Cycle 6

Province	Total Number of Children	Longitudinal Target Size	Difference	Relative Difference (%)
Ontario	557,463	527,435	30,028	5.39
Quebec	294,803	283,481	11,322	3.84
Alberta	151,323	145,798	5,525	3.65
British Columbia	163,351	158,153	5,198	3.18
Newfoundland and Labrador	19,664	19,106	558	2.84
Manitoba	51,610	50,230	1,380	2.67
Prince Edward Island	5,758	5,609	149	2.59
New Brunswick	29,056	28,474	582	2.00
Saskatchewan	43,786	43,191	595	1.36
Nova Scotia	35,709	35,258	451	1.26
Canada	1,352,523	1,296,735	55,788	4.12

The differences are fairly large for Ontario (over 5%) and overall (4.1%). In every province there are more children aged 2 to 5 in January 2005 than the sum of children aged 0 to 1 in January 2001 plus children aged 0 to 1 in January 2003.

Since recent immigrants and children born in Canada may have different characteristics, we run the risk of bias. This gap in coverage is due to the long period between sample selection and data collection for the returning cohorts. For children aged 2 to 5 at Cycle 6, we still consider the sample to be cross-sectionally representative, but urge users to be aware that recent immigrants are not truly covered by the sample.

Also, this table does not tell the whole story. A certain number of children aged 2 to 5 die or leave the ten provinces between when they are selected and the Cycle 6 collection. These children cannot contribute to the cross-sectional estimates at Cycle 6 and such children are not counted in the first column. Our cross-sectional under-coverage is at least 4.1% and, in reality, is slightly larger because 55,788 is a count of net immigration, and ideally, we would use a count of gross immigration. Gross immigration is calculated as the cross-sectional target population size minus the longitudinal target population size plus the death and emigration count.

There is also top-up sample in certain provinces at Cycle 6. In provinces with a top-up sample, there is partial coverage of recent immigrants. Practically, only a small proportion of recent immigrants aged 2 to 5 could be selected by the NLSCY because the provinces that draw the most immigrants (Ontario and Quebec) did not have any top-up sample.

4) Inter-provincial migration

Another consequence of the long lag between sample selection and data collection is that inter-provincial migration occurs. The cross-sectional sample is intended to represent the population of children for each province at the time of collection. If the child has moved, the province of residence may be different from the province at the time of selection. The initial weight is based on the province at time of selection and can potentially be much larger or smaller than the weights of the other children in the province of residence at the time of collection. This can lead to unstable cross-sectional estimates.

The problem can be described using this scenario. Children, who were selected in provinces with low probability of selection, such as Ontario and Quebec, moved to a small province with a much higher probability of selection. Such children might dominate the estimation because of their excessively large sampling weight. To address this situation, outlier detection techniques have been used to determine which cases required a modification to the initial weight.

Basically, in some instances, the initial weights from an outdated frame (either 2 or 4 years old) are no longer cross-sectionally appropriate. In creating the cross-sectional weights, steps are taken to address this frame problem.

5) Non-uniform coverage of month of birth


In the LFS, each month a new panel or rotation group is introduced while another rotates out or leaves the LFS. To meet the targeted sample sizes for the NLSCY, several rotation groups are considered, and all households with an NLSCY-eligible child are chosen for the NLSCY.

One difficulty is, when trying to target 0-year-old children, some of the rotation groups used can identify only a portion of the children born in the reference year.

For example, for our sample of 0-year-olds (children born in 2004), households interviewed for the LFS in January 2005 can report children born at any time of 2004. These rotation groups fully cover 2004 births. We do use all these households, but we desire to include more 0-year-old children in the sample and we are forced to use rotation groups that leave the LFS during 2004 and have only partial coverage of 2004 births. For example, those in the rotation group that leaves the LFS in August 2004 can report only births occurring in January to August 2004. We are unaware of any births in these households occurring in September to December 2004.

For this reason, we do not have uniform coverage by month of birth for children introduced as 0-year-olds – that is, children born in 2000, 2002, or 2004. For these cohorts we have more children born in the first months and fewer in the last months. The survey weights are adjusted for the partial coverage of births, but we did not attempt to erase the distortion in terms of month of birth. The extent of this effect is shown in the chart below.

Cycle 6 Count of Responding Children Born in 2004 by Month of Birth



6) Non-uniform coverage by age in months

There is a slight overrepresentation of children aged 24 to 28 months at the time of the Cycle 6 interview (variable FMMCdQ1B). This is a consequence of how the sample was collected in Cycle 6. At the first wave of collection, starting in October 2004, Cycle 5 introduced children born in summer 2002 were collected having age in months of 24 to 28 months. In the final wave in April 2005, Cycle 6 introduced children born in winter 2003 were interviewed. These children, too, had an age in months at the time of interview of 24 to 28 months.

7) Original cohort lag between selection and collection

At Cycle 1, dwellings containing NLSCY eligible children from the Labour Force Survey formed the NLSCY sample. There was a lag of several weeks between the sample selection and the start of collection. Once collection started, some cases were not contacted for a few additional weeks.

If the children changed address before the NLSCY interviewer reached that dwelling, the case was considered out of scope. As the time lag was relatively short, the count of such cases was fairly small. Still, it is possible that children living in families that move frequently could be slightly underrepresented among the Cycle 1 respondents.

This is not a factor for children introduced in Cycle 5 or Cycle 6 as these children were pre-selected and we attempted to contact them if they had moved since sample selection. For children introduced in Cycle 4, differing field protocols were used, with some movers traced and others not.

8) Effect of sample cut at Cycle 2 on coverage

At Cycle 2, the longitudinal cohort was reduced. Some households were dropped and within some households, the number of selected children was reduced to a maximum of two children from the Cycle 1 maximum of four children.

For the great majority of households, the choice of children retained was random, and the retained children accurately reflect the survey population.

For a portion of the sample, namely for the households with at least one 0- or 1-year-old child selected from the LFS after the 1994 re-design plus all households in New Brunswick, the choice of which children to retain to respect the new maximum of two children per household was not random. Instead, for this portion of the sample, the youngest two children were retained.

The result is a slight distortion by birth order. For example, in our Cycle 1 longitudinal sample, we have a slightly elevated proportion of 0- and 1-year-old children (age at Cycle 1) with two or more older siblings when compared to the overall population. In our Cycle 1 longitudinal sample, there is a corresponding slightly decreased proportion of children aged 2 to 11 with two or more younger siblings.

The effect is not minute only in New Brunswick, where all 2- to 11-year-olds with at least two younger siblings were cut from the sample. For New Brunswick, the survey weight adjustment is valid only if we believe that children who are not the first born children in a large family can represent those who are first born children in large families.

9) LFS design and coverage

The LFS underwent its last redesign in 2004. Over time, the number of occupied dwellings in certain areas changes and the efficiency of the survey design gradually deteriorates. This has a small impact on all surveys using the LFS as a frame, including the NLSCY.

Also, the LFS, though generally a good vehicle for social surveys, was designed to measure labour force characteristics, not necessarily to measure characteristics of Canadian children. There is not any real problem per se. We are only noting that the sample design, though near-optimal and very cost efficient, is not optimal for a survey targeting children.

12.18 Conclusion

Data quality is affected by various sources of error. Efforts are made at all steps (interviewer training, collection monitoring, processing, weighting, etc.) to reduce the potential for errors.

Data users are encouraged to consider how sampling and non-sampling errors may affect the variables they are attempting to analyze.


13.0 Variance Estimation

The National Longitudinal Survey of Children and Youth (NLSCY) is a probabilistic survey for which samples have been selected to represent various target populations. A quantity of interest about a given target population is called a parameter. The average height of children of a same age is an example of a parameter. The value of a parameter is unknown and we conduct a survey like the NLSCY to obtain valuable information from a sample of all the individuals composing the target population of interest. The relevant information contained in this observed sample about the parameter is extracted from the data using a mathematical tool called an **estimator**. The value which represents a reasonable guess about the parameter one can make from the observed information is called an **estimate**; it is simply the output of the estimator when the observed sample is fed into it. Feeding different samples through the estimator result in different numerical guesses, i.e., the estimates, being made about the parameter. The extent to which these estimates would differ as a whole from the value of the parameter is the sampling error. A key feature of survey sampling is to measure mathematically the magnitude of the sampling error. By definition, a census has no sampling error since the only possible sample is actually made of the whole population (and the only sensible estimate we should get corresponds exactly then to the parameter's value we are after).


Even if sampling error could not be measured in a given context, it would still be possible to say a few basic things about it. For example, the larger the sample size the smaller the sampling error generally; this is because a larger sample contains valuable information about a greater part of the population. Some of the other factors influencing the magnitude of the sampling error are the size of the population, how the sample is drawn (this is specified through what is called a sampling design) and the variability in the target population of the characteristics upon which our estimate will be based.

There are two components to the sampling error: the sampling bias and the sampling variance. An estimator displays no sampling bias if, loosely speaking, the average of all its outputs, the estimates obtained by feeding it with all possible samples, matches the parameter's value. So, estimates taken individually may all be off from the parameter's value but on average be right on target; in such a case the estimator is said to be unbiased (and biased otherwise). The other component of sampling error is sampling variance which measures to what extent the estimates differ from one another.


A well-known analogy helps illustrate these concepts. A dart-thrower (the estimator) is told to throw a series of darts at a target; each strike corresponds to an estimate. We do not expect all hits to be on the bull's eye; each of these is a contributor to the (total) sampling error. There are essentially four possible scenarios for the hits as a whole, depending on the magnitude of the two components of the sampling error, bias and variance:




Bias and small variance



Unbiased and small variance



Bias and large variance



Unbiased and large variance

In practice, unfortunately, it is usually impossible to have an estimator which performs well on both components, i.e., an estimator with both low bias and variance. Usually, a low-variance estimator will turn out to be largely biased while a low-bias estimator will have large variance. Survey statisticians usually

rely on estimators which are known to have little to no bias; they thus prefer to reduce the component of sampling error due to bias to its strict minimum and do whatever they can afterwards about variance. This is what we suppose also; consequently, we will restrict the discussion on sampling errors to considerations about the sampling variance only.

In order to assess the sampling variance of an estimator, we would need to get our hands on the estimates arising from all possible samples. In practice, however, we have only *one* sample to work from, the one for which collection was carried out. Even though we cannot evaluate the sampling variance of an estimator, we can usually get an estimate of it based on the one sample at hand. The reader may be dubious as to how one can predict well the variability in results arising from all samples by using just one of those. While the general situation is somewhat intricate, variance estimation fundamentally rests upon the following observation. One can show that for simple estimators (like that for the mean), used in conjunction with a very simple selection mechanism of the sample, the sampling variance is a direct function of the population spread. So, in such a simple case the sampling variance is unknown to us *because* it is a (known) function of an unknown quantity: the population spread. If we knew the latter we would know the former. Can we get a good guess of that population spread on the ground of what we would know – a sample of values? Again, in this simplest case, the spread observed in the values within the sample provides a direct measure of the spread in the population. Indeed, in that specific setting one obtains a sample which is usually “well-balanced”: it contains about the same proportions of small, medium and large values as the population itself, making the observed spread a direct indication of the population’s spread. In more complex settings, no such simple relationship exists between population and sample spreads which could in turn be used to estimate sampling variance. There are nonetheless clever methods like the bootstrap, about which we will have more to say later, which succeed doing just what we have done in the simple case above: assessing sampling variance from just one observed sample.

In this chapter, we will explain why it is important to calculate the sampling variance and we will present different tools to do so for the NLSCY.

13.1 Terms Related to Sampling Error and Variance

There is sometimes confusion about what is meant by the terms “population variance”, “sampling variance”, “standard deviation”, and “standard error”. In this section we seek to clarify what each one is all about.

Unfortunately, the term “variance” as in “population variance” and “sampling variance” is used for two very different things. The variability observed in the values of a characteristic in the target population is often referred to (unfortunately) as the **population variance** (of the characteristic). For example, in the population of all 10-year-old boys in Canada, there is some variability in their measured height in centimetres.

But variance should strictly be used when there is an underlying random process at work (such as the random process in survey sampling by which samples are selected). Since the values of a characteristic in the population are fixed (and thus not the result of some random process), we propose to refer to their inherent variability rather as the **population spread** (of that characteristic) than as population variance. In the example above, we would say that there is some *spread* in the height measured in centimetres of all 10-year-old boys in Canada.

A mathematical definition of the population spread often used is:

$$\sigma^2 = \frac{\sum_{i=1}^N (y_i - \bar{Y})^2}{N - 1}$$

where y_i is the value of the characteristic Y for person i ;

\bar{Y} is the average of y_i in the population; and

N is the size of the population.

The population spread is simply then a parameter of the population, one of many others like the population's average value for Y , \bar{Y} , for instance. The reader will note that the symbol used to represent the population spread is $\sigma^2 = \sigma \times \sigma$ and not just σ , say. This is so we're reminded that the measure of spread is not expressed in the same units of measure as the variable itself, but rather in terms of its square. So, to obtain a measure of spread expressed in terms of the same units as the variable itself, it seems reasonable to take the square root of the population spread. And this is precisely what we do, actually; this yields what is known as the **standard deviation**.

We have already encountered the **sampling variance**. It is the variability that would be observed in the estimates of a same parameter if all possible samples were processed. Since the sample selection is a random process, it makes sense here to speak of "variance" as opposed to the situation above involving the population spread. Just like the population spread is not expressed in terms of the same unit as the variable whose variability it assesses but rather in terms of its square. So, if the estimate about personal income is expressed in terms of \$, then the sampling variance is expressed as $\2 . Again, it seems therefore natural to take the square root of it to restore comparability in terms of level with the estimate; this gives the **standard error**.

13.2 Coefficient of Variation

A measure of sampling variance provides us with a valuable indication as to the reliability of an estimate. As we saw, the standard error is a better tool still since it is expressed in terms of the same units of measure as the estimate itself. But standard error (just like variance) does suffer a major drawback: standard errors arising from different estimates are not comparable in situations where they actually should be. For example, is the estimate obtained for the province A less precise or better than the corresponding estimate for province B ?

To illustrate the drawback, suppose a sample is drawn to collect information to estimate the total of a characteristic Y . But with the same information we can also estimate the *mean* of Y by simply dividing the estimated total by the (assumed known) population size N . The point here is this: both estimates feed from the same sample information in the same way through the estimated total. Consequently, *if* standard errors for both the total and the mean *were* to be comparable, as means to evaluate *relative* precision among estimates, then we'd expect them to have the same value. But they won't have the same value because the standard error of the total will actually be N times larger than that of the mean since estimates of the mean and total taken from the same observed sample will all differ by the same factor: N .

A way to obtain a relative measure of sampling error, one which gives the same value in "comparable" sampling settings, is the coefficient of variation (CV). The coefficient of variation is defined as the standard error of the estimate $\hat{\theta}$ divided by the estimate itself, that is:

$$CV(\hat{\theta}) = \frac{s.e.(\hat{\theta})}{\hat{\theta}}$$

Now, contrary to the sampling variance associated with an estimate, the coefficient of variation allows the analyst to compare estimates of different magnitude or measured in different units on a

common scale for income-like⁸ variables. Going back to our earlier comparison of estimates of the mean and the totals, we can see that in such a case the CVs would be equal. Indeed, even though as we said the standard error of the total is N times greater than that of the mean, the corresponding denominator for the CV of the total is also N times bigger than that used to compute the CV of the mean. Consequently, the same N factor appears in both the numerator and denominator and thus cancels itself out.

Although CVs are useful for indicating the quality of estimates such as totals, there are some pitfalls that users should be aware of when using CVs to examine the quality of proportions. A few potential problems are outlined below.

Issue 1: Coefficients of variation for very small (or very large) proportions

Since the standard error of a proportion p is the same as the standard error of the proportion $(1-p)$, the CVs of p and $(1-p)$ may differ substantially because the denominators are p and $(1-p)$ respectively. One can imagine a scenario when p is very small giving a very large CV for p , but the CV of $(1-p)$ being excellent. Let's use the following example to illustrate. Suppose we have calculated the standard error of the estimates p and $(1-p)$ as 0.0475, for a value of p of 0.95. The CV for the value of p , 0.95 would be:

$$0.0475 / 0.95 = 0.05 \text{ or } 5\%, \text{ which is a very good CV.}$$

On the other hand, the CV for the proportion $(1-p)$ is:

$$0.0475 / 0.05 = 0.95 \text{ or } 95\%. \text{ This is a very poor CV.}$$

Issue 2: Applying rules about the assessment of sampling error in the case of proportions

Rules have long been circulating helping the user assess the validity of an estimate based on the magnitude of the sampling error as measured by the CV. Typically, these rules state that an estimate with a CV under 16.5% is of good quality – inferences based on these results are deemed trustworthy; a CV between 16.5% and 33% indicates acceptable quality – inferences based on these results require caution; a CV beyond 33% describes an estimate of poor quality – inferences based on these results are purely exploratory and should not otherwise be trusted. These rules can be quite useful to the data user in figuring out what to make of a survey's results but they have their limitations; consulting a survey statistician is certainly, though, the way to make the most of the results since other factors specific to the analysis other than the sampling error may determine the validity of the inferences made. In the case of proportions, these rules require greater caution on the part of the user, as the following examples show.

Example 1: An estimated proportion of 0.50 with a 99% confidence interval of 0.10 to 0.90 falls into the **marginal** category, using the previously published quality guidelines (the CV is 31%). But the confidence interval is so large, that the estimate is not really giving us much information.

Example 2: Suppose that we have 27,000 sampled individuals of which 44 have a characteristic we are interested in studying. Using the survey weights, we calculate that 0.16% of the population has this characteristic, with a corresponding CV of 34%. A CV of 34% is classified as **unacceptable** by applying the previously published quality guidelines. Although 0.16% is a very small proportion, if we construct its 99% confidence interval we get

⁸ By this we mean a positive quantity of interest which is continuous, as opposed to a dichotomous variable.

(0.02%, 0.30%). This means we have considerable confidence that the true rate is less than, say 0.5%. Of course, data users should be cautious as the true rate could be 0.05% or 0.25% or even 0.30%. Depending on the goal of the research, maybe a statement that the estimate is smaller than 0.5% is meaningful. Therefore, blindly applying the quality guidelines and throwing away an estimate because of its high CV may not be appropriate.

Consequently, if users wish to use CVs as a measure of sampling error when dealing with proportions, they are strongly encouraged to calculate the CVs for both the proportions p and $(1-p)$. CVs associated with proportions, particularly extreme proportions, can be misleading as the above examples illustrate. We do not actually recommend comparing different proportions in terms of CVs. We rather suggest constructing confidence intervals and rely on them to conduct proper inferences involving proportions.

Research is currently taking place to find better alternatives than the CVs for the extreme proportions (close to 0 and close to 1). However, for now, a solution which will meet all needs has yet to be found. Consequently, users must use caution if they wish to use CVs for proportions.

13.3 Importance of Reporting the Sampling Variance


As we saw in earlier sections, an estimate will almost assuredly differ from the true value were it to become known to us; this numerical difference between the estimate and the parameter is the error. We thus have to assume that with an estimate there's always an error that goes with it. **Consequently, some indication of the magnitude or extent of that error in the inference has to be provided to those consulting the estimates.** Indeed, stating survey estimates without corresponding measures of the error involved can be very misleading. One of the two components of that error is the sampling error, the other being the non-sampling error. The latter encompasses such things as response errors and processing errors. In practice the magnitude of the non-sampling errors is difficult, if not impossible to quantify. (When errors cannot be gauged in terms of their impact, as is the case with most non-sampling errors, then they must at least be clearly reported and described.) On the other hand, the sampling error *can* be quantified if we are able to mathematically keep track of the effect of randomness on the yield of estimates. We measure the sampling error associated to an estimate by estimating the sampling variance of the process that created it. Consequently, the sampling variance must be computed and provided to the users as a means to describe the quality of the estimates provided.

13.4 Sampling Variance Calculation

It would be difficult (not to say impossible) to derive an exact formula to calculate the sampling variance for the NLSCY due to the complex sample design, nonresponse adjustments, treatment of out-of-scope units, and the post-stratification. Actually, such a task could only be undertaken under such strong assumptions as to yield a framework too simplistic to be of any little use in practice. A very good way to approximate the sampling variance is to use the bootstrap method⁹.

To help grasp what the bootstrap is about, we need to introduce the concept of sampling distribution of an estimator. We've observed already that if all possible samples could be fed into the sample-information-extractor which is the estimator, then we would get all possible estimates. Suppose we did just that. We could then plot the frequency by which we observe each of these estimates; this is called the sampling distribution of the estimator. Here's an example of such a plot, a sampling distribution of an estimator.

⁹ The challenge here is to adapt successfully the bootstrap for survey sampling when it actually was designed for a non-survey setting in the beginning. Our understanding of the bootstrap in a survey setting, and how to adequately implement it, has kept evolving over the last 10 years or so, ever since actually the NLSCY first started using it.



We notice, for instance, that extremely small estimates are rare, so are very large ones, while the most frequent estimates being somewhere in the middle. In our example, estimates are clustered around the mean value of 5 (which would be the unknown value if the estimator is unbiased) and estimates greater than 5.275, for instance, are a rare event (they arise less than 0.5% of the time). The fact that not all estimates are identical but are rather spread over some range is the graphic description of the sampling variance. Indeed, the sampling variance is by definition a measure of the variability observed in the distribution of estimates in the plot. The sampling variance is thus a characteristic (one of many) of the sampling distribution of the estimator.

The bootstrap attempts to provide us with this sampling distribution, which again is determined by all possible estimates, by re-sampling a large number of times from just one sample, the one for which in practice we collect data through field work. Just like the sampling that has led to the observed sample is accompanied with survey weights, the re-sampling produces its own set of weights, the bootstrap weights.

For the NLSCY, a set of 1,000 bootstrap weights is available. To obtain an estimate of the sampling variance of an estimator using the bootstrap, one has to compute the 1,000 estimates that correspond to the 1,000 bootstrap weights. Indeed, one replaces the one set of survey weights used by each of the 1,000 sets of bootstrap weights. Just like one set of survey weights has yielded one estimate, 1,000 sets of bootstrap weights will yield 1,000 estimates. The computed variance of these 1,000 (bootstrap) estimates is precisely the estimate of the sampling variance of the estimator we get from the bootstrap. Algebraically put, the bootstrap variance estimate \hat{v} is computed as:

$$\hat{v} = \frac{1}{1000} \sum_{k=1}^{1000} (\theta_k - \theta^*)^2$$

where θ_k is the k-th bootstrap estimate and θ^* is the original sample-based estimate of θ .

(Note: in practice one can most of the times use in the formula the average of the θ_k 's as instead of θ^* , the sample-based estimate. This shortcut is used whenever one computes the bootstrap

variance using PROC SUMMARY from SAS on the set of bootstrap estimates using the VAR statistic.)

Two tools, both making use of the bootstrap weights, have been developed to help users calculate the sampling variance and the CV for their estimates. These tools are:

- An Excel spreadsheet, with a Visual Basic interface, that enables users to retrieve approximate sampling variances for proportions across a large number of domains, e.g., by age and by province.
- Macros to calculate the sampling variance, using the bootstrap weights.

The choice of tool to use depends on the type of analysis and the level of precision required.

In cycles prior to Cycle 5, a third variance approximation tool was available: CV look-up tables. Using a representative design effect (the design effect compares the variance of estimators from the NLSCY sample design to those from a simple random sample) users were able to obtain CVs for some domains, by age cohort or by province. These tables are no longer available since the Visual Basic interface and bootstrap macros are more flexible and more accurate.

13.4.1 Spreadsheet with Approximate Sampling Variances for Proportions

A set of spreadsheets is available to users to (approximately) calculate the sampling variance associated with estimates of proportions. Available in Excel format, the Visual Basic interface accesses results (calculated using replication methods) for thousands of domains. These domains include cross-tabulations of age, age groups, provinces, or regions. The sample sizes for each domain is also available.

Details on how the spreadsheets and interface were created, what they contain, and how to use them, can be found in separate documentation that accompanies these spreadsheets.

13.4.2 SAS and SPSS Macros to Calculate the Sampling Variance Using the Bootstrap Weights: Bootvar

SAS and SPSS macros have been developed to calculate the sampling variance using the bootstrap weights; they form what is known as Bootvar. The most current SAS version of Bootvar is 3.1 while in SPSS it is 3.0. Bootvar can be accessed through “research tools” at www.statcan.ca/english/rdc/whatdata.htm. Bootvar can compute, for any domains, variance estimates for such things as: totals, ratios, difference of ratios, linear and logistic regression coefficients. The sampling variance calculated using this method takes into account the sample design and also, the specificities of the variable of interest. Finally, as opposed to the spreadsheet, the user is not restricted to pre-defined domains.

This method has many advantages but requires more work from the researcher. The sampling variance calculation using these macros is more time consuming than using the spreadsheet. The user must first become familiar with the macros before using them. However, these macros have been developed in such a way that they are easy to use. The researcher must have access to the macros, to the data files and to the bootstrap weight files. Access to these tools is possible in a Statistics Canada Research Data Centre (RDC). Also, detailed documentation on how to use these SAS or SPSS macros is available in the RDCs.

Despite the time required to run these macros, Bootvar is strongly recommended over the VBA tool using Excel spreadsheets to obtain a sampling variance estimate of any estimate which must be published; it provides a more exact and suitable measure of the sampling variance.

Again, details on how to use these programs can be found in separate documentation that accompanies the programs and bootstrap weights.

13.4.3 Other Computer-based Tools

Other commercially available software can properly calculate the sampling variance from the bootstrap weights provided. SUDAAN (setting design = BRR) and WesVar are two such software available in the RDCs. STATA9 is another possibility.

To calculate the sampling variance for estimates not included in Bootvar, analysts may wish to write their own programs implementing the bootstrap method. However, this is not a trivial matter.

13.4.4 Taylor Linearization and Other Techniques

The bootstrap weight files contain variables indicating the primary sampling unit (PSU) and stratum from which the individual was selected. Some existing software packages (such as Stata, SUDAAN or SAS) have procedures that calculate sampling variance estimates using design information (stratum and PSU) and the survey weights. The technique is known by several names: Taylor Linearization or Binder or robust variance estimation. The problem with using these procedures with the NLSCY data is that they require at least two PSUs per stratum, and the NLSCY very often does not satisfy this requirement. Although collapsing of strata is possible, its effectiveness at this point is unclear, as a thorough comparison of sampling variances obtained this way to bootstrap sampling variances has not yet been done. Furthermore, several adjustments are needed to turn the design weights into the released weights and Taylor linearization can't account for the impact of these on the variance. Therefore, we recommend using one of the sampling variance tools described in this section (Visual Basic interface or the bootstrap weights) to obtain design-consistent estimates of sampling variance.

Lastly, software packages such as SAS or SPSS do compute a variance for estimates produced in their built-in procedures, e.g. PROC UNIVARIATE in SAS. However, many of these routines do not take into account the sample design, e.g. stratification, even using the WEIGHT statement, which means the variance calculated is not the sampling variance we're after (it is usually well underestimated this way). Therefore, these procedures are not recommended since they can lead to erroneous conclusions.

14.0 Direct Assessment

Background

Research on early childhood and youth development plays a significant role in the formulation of policy for young children and youth. Using various assessment tools in the National Longitudinal Survey of Children and Youth (NLSCY) will help to enhance the knowledge about developmental processes in early childhood and youth and provide relevant data on which to base policy directions for these stages.

Choices about the assessment tools to be included in the NLSCY were made on the basis of an extended literature review, development of a research framework on child development and learning, consultations with many experts in Canada and internationally, review of material on many different possible instruments and field testing of the most likely possibilities. The instruments selected for consideration were also reviewed using a number of criteria. The criteria included reliability and validity of the instrument, coverage of domains in the research framework, ability of the instrument to indicate normal development and developmental delays, the ease of administration by lay interviewers and the availability of the instrument in English or French (or ease of translation to French or English). The final decision was strongly influenced by key experts who had a history of providing advice to the NLSCY Team.

The NLSCY conducts direct assessments of children aged 4 and up. These assessments are described in this chapter.

14.1 The Peabody Picture Vocabulary Test - Revised

The PPVT-R was designed to measure receptive or hearing vocabulary and can be used for any age group, up to adult. The test was developed by Lloyd and Leota Dunn, at the University of Hawaii, and has been widely used in large-scale data collections as well as assessments. A French adaptation of the PPVT-R was developed by the test's authors and Claudia M. Thériault at St. Thomas University in Fredericton, New Brunswick. The French test is called the Échelle de vocabulaire en images de Peabody (EVIP).

For the NLSCY, the PPVT-R was used to measure school readiness for children in the 4- to 5-year-old age group. Verbal parental consent was required before the test was administered. If permission was granted, the interviewer then administered the test to the child in the home. The child looked at pictures on an easel and identified the picture that matched the word the interviewer read out.

A total raw score was calculated for each child who completed the PPVT-R by computing correct responses. A standardized score was also assigned to each child. Standard scores allow for comparisons of scores across age groups. Obviously, a 5-year-old would be expected to perform better on the PPVT-R than a 4-year-old and thus have a higher score. The standard score takes into account the child's age.

Standard scores for a test are usually based on the distribution of scores obtained from the entire population. In the absence of scores for the entire population, a representative sample distribution, called the norm sample, is more often used. Each cycle, the NLSCY yields a representative sample of children. Any of those samples would be a viable option to estimate the distribution of scores measured in the population. While each sample is selected probabilistically, albeit for different reference period, experts in the field of cognition might disagree as to whether differences between the estimated distributions from one sample to the other reflects a true population difference over time or simply resulting from sampling error.

For robustness, we had elected to use not one but all normative samples for each of the cycles of collection in the NLSCY, until the overall distribution had become relatively static. This point was reached and therefore the standardized scores of Cycle 6 were calculated using the same norms

that were used for Cycle 4 (v2) and Cycle 5¹⁰. The Cycle 6 norms are therefore based on the 28,214 records from Cycles 1 to 5 of the NLSCY main survey with PPVT-R raw scores. Some records with zero PPVT-R raw scores were excluded from the data used to create the norms: one record from Cycle 2, 49 records from Cycle 3, and six records from Cycle 4. These zero scores were probably incomplete tests so they are not reliable and would underestimate the true measure of ability (particularly in Cycle 3). To obtain the norms, each record was weighted by its cross-sectional weight divided by the average cross-sectional weight of records from the same cycle. The PPVT-R individuals in the norm sample were assigned standard scores so the mean of the standard scores was 100 and the standard deviation was 15 for all ages in months. Loess smoothing was applied to the data to ensure that the PPVT-R norms increase with age.

Reliability measures for the PPVT-R have been calculated based on the American norm sample (Dunn and Dunn, 1981).

14.1.1 Psychometric Properties of Scores

This section addresses the quality of the test itself as it applies to the survey population, as compared to the original population for which the test was developed. We find that the test still provides a reasonable assessment of the child's ability, and we outline the reasons below.

The Raw Score

One of the main advantages of the test in a survey context is that it is tailored to the child's age and performance so that not all of the questions need to be asked to determine the ability level. Based on the age of the respondent a starting question is selected, and the test proceeds with increasingly difficult questions. When the respondent appears to be answering at random – at least six out of the last eight questions are missed - the test stops and a score is derived based on the rank of the last question and the number of incorrect answers.

Questions are ranked by increasing order of difficulty and are designed to be equally spaced on the "difficulty scale". Originally the PPVT-R test was calibrated by using a representative sample of about 5,000 English speaking children. Similar efforts were undertaken to calibrate the French version. The difficulty items were calculated using the Rasch model. In the language of Item Response Theory (IRT), this is known as the one-parameter logistic model.

Since the calibration test was done some years ago, it is natural to expect some drift of the difficulty items, as the language itself evolves and some words become more or less common. To verify whether this is the case an IRT analysis of the items was done in Cycle 4, and derived scores based on the new difficulties of the items were created. For some of the items we did find some deviation from what was to be expected in the original test.

However, no systematic deviations were found in the differences for the measured outcome. Consistently no overestimating or underestimating of the child's ability was measured in any portion of the test, and the scores derived by using the IRT were consistent with the raw scores. The correlation coefficient between the two scores from the tests was 94% for the English version, and 96% for the French version, which is high by any standard. Therefore we are confident that the raw scores can be used as they are.

For a number of children (23) the test was not completed in the field as per the directives

10 Note that for cycles 1, 2 and 3, a different norm sample was used. While some slight variation will exist, these are well within sampling error that results from using different samples as the norm sample.

of the instrument and no score could be assigned by the application. These cases resulted in a score of zero despite a significant number of test questions being answered for some cases. We were able to derive approximate raw scores for these children using IRT. For the rest of the respondents – that is, 99% of them - we used the score that was produced by the application using the original rules of the PPVT-R test.

The Standardized Score

As described earlier, the standardized score is determined by using population distributions for each age. Strictly speaking, we can never know the population distribution, since applying the test to the whole population is not feasible. One way to deal with this is to use the sample that we have within an age group as representative of the population in that age group, and derive the necessary percentiles.

However, certain sample limitations exist that need to be addressed before the score can be standardized. By inspecting the percentiles for different ages, we would expect an increasing trend in the ability measurement with age. While there is an overall increasing trend, for a lot of ages the trend is reversed. This is due to the fact that the sample is not large enough for all age groups, and a lot of noise is introduced due to poor representation.

The better approach is to use the percentiles from the sample as a starting point, and smooth the progression with age until we are satisfied that we have a "natural" progression. We used the progression of the original percentiles from the PPVT-R handbook as an example of what degree of smoothing should be expected. Then we used the resulting points as the percentiles for standardization. We should note here that even though features of the norms were similar, the percentiles drifted upwards over the years, which, according to the experts, can be expected.

The test is usually applied to children whose effective age is 4 or 5 (note that the children's real age may include 3- and 6-year-olds if they are assessed early before their fourth birthday or after their sixth in the next calendar year).

Final Note

The PPVT-R scores used in the NLSCY are a valid measurement of ability. However, to minimize the potential for biased estimates when doing analysis, nonresponse should be handled, on a case by case basis. For more information about nonresponse, please see Chapters 10.0 and 13.0.

14.2 Number Knowledge Assessment

The purpose of the Number Knowledge assessment is to assess the development of children's understanding of numbers by examining their comprehension of the system of whole numbers. For the NLSCY, the assessment is administered to 4- and 5-year-old children.

The assessment was developed by Dr. Robbie Case from the University of Toronto, with colleagues, including Yukari Okamoto at the University of California at Santa Barbara. The assessment is constructed based on Dr. Case's theory of central conceptual structures for explaining the development of children's thought. Before his death in May 2002, Dr. Case was adapting the test for the NLSCY. Following Dr. Case's death, Yukari Okamoto assisted the NLSCY team in completing the adaptations of the assessment for the survey.

Theoretical Background

According to Dr. Case's theory, four developmental levels can be distinguished in children's understanding of numbers: pre-dimensional, uni-dimensional, bi-dimensional, and integrated bi-dimensional. Some degree of mastery of each level is required prior to continuing onto the next.

Typically the four levels are attained at the ages of 4, 6, 8, and 10. The pre-dimensional level assesses the ability to count by rote and to quantify small sets, using concrete objects. This knowledge is important for the uni-dimensional level where children deal with changes in quantity without objects than can be touched or seen. The uni-dimensional level assesses children's knowledge of the number sequence and ability to handle simple arithmetic problems. To solve the items, children must rely on a "mental counting line" in their heads. This "line" integrates their understanding of numbers and quantities. This assessment measures the essential prerequisites for successful school learning.

Assessment Description

In consultation with Dr. Case and Dr. Okamoto, the test was revised for the NLSCY. The assessment has been made continuous with three levels; some items were revised or dropped to make the test somewhat shorter. The original version of the test was discontinuous, i.e., the child had to pass sufficient items at any one level to go to the next level. Since it is accepted that we cannot expect a child to do well at a level without also doing well at the preceding level, it is sensible to stop administering the assessment after a certain number of missed items. The test was also programmed into the computer application, so that the stopping rule was automatically applied. The interviewer asked the child the question and then entered the answer. The application determines whether or not the child answered correctly.

The test is composed of 22 items. Some of them have two parts - a) and b). Children must pass both part a) and b) to earn a pass for these items. This convention was adopted because each two-part item gives children a choice between two alternatives and a child has a 50% chance of getting the right answer by guessing alone¹¹. Requiring children to pass two such items before they get a point increases confidence that children have the knowledge required by the item.

Children are not permitted to use a pencil and paper to answer the questions, which are given orally. Instead, the children must rely on a "mental counting line", which integrates the child's understanding of numbers and quantities. Children do have access to the various manipulative aids such as chips and a number card to help solve the problems.

The administration of the Number Knowledge assessment should take approximately 10 to 15 minutes.

Scoring

Three different types of scores have been made available for Cycle 6: the "Age Equivalent Score" (FKNCdS01), the "30-point-raw score" (FKNCfS01) and the "30-point-raw age-standardized score" (FKNCfS02).

The "Age Equivalent Score" is derived based on the child's responses to the items. The "Age Equivalent Score" assigns a point for each of the three levels passed and then the points are totalled (a maximum of one point for each level completed can be assigned). Passing a level means passing a certain number of items from that level – for instance, for the pre -dimensional level, three out of five items must be correct. A child failing to answer any questions at the first level will get the minimum (zero), while a child who answers all the questions of all three levels correctly receives the maximum (three).

Level 1 represents the proportion of correct responses for the pre-dimensional level. There are five items in this level. To reach the age equivalent of this level, the child must achieve a proportion of at least 0.6, i.e., get three out of five correct responses. Level 2 represents the proportion of correct responses for the uni-dimensional level. There are eight items in this level. To reach the age equivalent of this level, the child must achieve a proportion of at least 0.6, i.e., get five out of eight correct responses. Level 3 represents the proportion of correct responses for

11 For example, part a) may ask which of two piles of counting chips is bigger and part b) asks which pile is smaller.

the bi-dimensional level. There are nine items in this level. To reach the age equivalent of this level, the child must achieve a proportion of at least 0.6, i.e., get six out of nine correct responses.

Although the Number Knowledge test is made up of 22 items, a child who goes through the whole questionnaire is asked 30 questions, since some items have a) and b) parts. The 30-point raw score is simply the total number of correct answers among those 30 questions.

A 30-point raw age-standardized score was also assigned to each child. Standardized scores allow for comparisons of scores across age groups. Obviously, a 5-year-old would be expected to perform better on the Number Knowledge test than a 4-year-old and thus would have a higher score. The standardized score takes into account the child's age. The norms used for the standardization have been built using Number Knowledge 30-point raw scores from Cycle 4 and Cycle 5. To obtain the norms, each record was weighted by its cross-sectional weight divided by the average cross-sectional weight of records from the same cycle. The children in the norm sample were assigned standard scores so the mean of the standard scores was 100 and the standard deviation was 15 for all age groupings. This standardization was done for each age in months. Loess smoothing was applied to the data to ensure that the norms increase with age.

Evaluation of the Assessment

Analysis was conducted on the Number Knowledge data to validate this assessment. The analyses included: comparing age equivalent scores to the child's age, comparison with the Who Am I? and an analysis of the items and of nonresponse. As the scoring procedures were being developed the NLSCY team consulted with Dr. Okamoto to ensure that the procedures were consistent with Dr. Case's theories.

All the evidence validated the test and the test should provide data users with information about the child's acquisition of the necessary skills to succeed at math in school. However, this assessment is not free of nonresponse bias. Please see Chapters 10.0 and 13.0, for more information on nonresponse.

14.3 Who Am I?

The purpose of the Who Am I?¹² assessment is to evaluate the developmental level of young children from 3 to 7 years of age. For the NLSCY, the assessment is administered to 4- and 5-year-old children.

The assessment was developed by Dr. Molly de Lemos and her colleagues at the Australian Council for Educational Research (ACER). The NLSCY team worked closely with Dr. de Lemos to make some modifications to the assessment for the NLSCY (mainly dropping the drawing task) and to enhance the administration and scoring procedures for the NLSCY context.

Theoretical Background

The Who Am I? instrument assesses the developmental level of young children from the ages of 3 to 7 years old. The Who Am I? involves copying and writing tasks. The copying tasks in the assessment are designed to assess the child's ability to conceptualize and reconstruct a geometrical shape. The writing tasks assess the ability of the child to understand and use symbolic representations such as numbers, letters and words. The child's ability to complete the tasks depends on many factors including maturity, culture, experiences, and language skills.

The use of the ability to copy geometrical figures to assess the level of development in children has been long established. This type of assessment is included in measures of intelligence and

12 For more information about the Who Am I?, please see "Patterns of Young Children's Development: An International Comparison of Development as Assessed by *Who Am I?*" by Molly de Lemos (R-02-5E). This research paper was published by Human Resources and Social Development Canada.

development over a long period of time. Piaget's research on the development of spatial concepts in young children also provides evidence of the validity of copying tasks as a measure of developmental level.

Because the Who Am I? assesses nonverbal language, it can be used to assess children whose knowledge of English or French is limited. These children could be allowed to complete tasks in their mother tongue as well as English and French. Their scores in their mother tongue would provide information on their developmental stage; the score in English or French would give some idea of their development in that language. The NLSCY chose to only assess children in English or French for two reasons. First, it was felt that an assessment of the child's development in one of the official languages was an important indicator of the child's ability to function in the Canadian school system. Secondly, it would be operationally difficult to score questionnaires in the variety of languages spoken in Canada.

The tasks were developed based on research that indicates that copying skills are strongly associated with subsequent school achievement, are valid across different cultural groups and provide a reliable measure of development at the time of assessment. Also, children's attempts at early writing are linked to their growing understanding of the way spoken sounds are represented by print.

Assessment Description

The Who Am I? assessment is composed of three scales: a copying scale, a symbols scale and a drawing scale. The copying scale is composed of shapes (circle, cross, square, triangle and diamond) which the child attempts to reproduce. The symbols' scale is composed of a set of writing tasks (printing their name, printing some letters, numbers, words and a sentence) which the child attempts to complete. Children are only required to complete as much as they feel they can, but they are encouraged to at least attempt each task. For the drawing task, the child is asked to draw a picture of herself or himself. The drawing scale is not used in the NLSCY due to time constraints. Dr. Molly de Lemos was consulted before the drawing scale was dropped for the NLSCY.

The assessment consists of an appealing booklet in which the child completes the tasks as the assessor turns the pages and gives instructions. The booklet takes about 10 minutes to complete and is scored in Head Office. The child completes as much as he/she is able but is encouraged to produce at least a scribble for each task.

Scoring

For the NLSCY, the Who Am I? assessment is hand-scored by trained individuals at Statistics Canada. These individuals have been trained to recognize signs of each level in a child's responses. Scorers who cannot make a decision on a child's level because the work does not fit clearly into one level are asked to make a judgment about the child's level based on the score on other items. Scoring was done by a small number of people and was subject to quality control procedures. The head of the scoring team also met with a staff member from ACER to consult on scoring procedures.

All the items are rated on a scale from 1 to 4 by the scorers. If no attempt was made by the child, then an initial score of 0 is given. These items will be imputed later on in the process. Therefore, all items will eventually end up being given a score from 1 to 4. The Copying Scale Score (FWICdS02) is the sum of the scores attributed to the tasks related to reproducing a symbol. The Symbols Scale Score (FWICdS03) is the sum of the scores attributed to the tasks related to writing. Since there are five tasks for each of these scales, the Copying Scale Score and the Symbols Scale Score both range from 5 to 20.

In addition to the two scales retained in the NLSCY, there is also a combined total score, the Total Who Am I? scale score (FWICdS01), which is simply the total of the Copying Scale Score

and the Symbols Scale Score. It therefore ranges from 10 to 40 and gives a general overview of the child's developmental level.

Note again that items undergo imputation before being summed to form the scores.

For the first time, age-standardized scores will also be made available in Cycle 6: the Total Who Am I? Scale standardized score (FWICfS01), the Copying Scale standardized score (FWICfS02) and the Symbols Scale standardized score (FWICfS03). Standardized scores allow for comparisons of scores across age groups. Obviously, a 5-year-old would be expected to perform better on the Who Am I? test than a 4-year-old and thus have a higher score. The standardized scores take into account the child's age. The norms used for the standardization have been built using Who Am I? raw scores from Cycle 4 and Cycle 5. To obtain the norms, each record was weighted by its cross-sectional weight divided by the average cross-sectional weight of records from the same cycle. The children in the norm sample were assigned standard scores so the mean of the standard scores was 100 and the standard deviation was 15 for all age groupings. This standardization was done for each age in months. Loess smoothing was applied to the data to ensure that the norms increase with age.

Imputation

In summing scores on the Who Am I? tasks to obtain a total score for the copying and symbols scales, as well as a total score, it is necessary, according to Dr. de Lemos, to allocate a score in cases in which responses have been recorded as 0 (no attempt).

In most cases, it is assumed that no attempt indicates that the child is unable to do the task. From a developmental point of view, this is equivalent to a scribble. For the construction of norms, no attempt responses were considered to be equivalent to a scribble, and were allocated a score of 1.

It was, however, noted that, in some cases, children who were capable of more advanced responses on previous items did not attempt some of the more difficult items, particularly the diamond and the sentence. In such cases, allocating a score of 1 would lead to an underestimate of the child's developmental level. For this reason, a procedure was used for dealing with cases in which the child makes no attempt. This involved assigning a score based on the score to other items. For example, if a child had a score of 4 on the square and did not attempt the diamond then a score of 3 would be applied to the diamond.

Dr. de Lemos felt that imputation was necessary to make the NLSCY data more consistent with data collected with the Who Am I? in other studies. In most cases, the Who Am I? is administered by the child's teacher or an ACER researcher trained in child development. The NLSCY uses lay interviewers who only have a short time, in the interview setting, to develop rapport with the child. This made it harder for the interviewers to convince the children to attempt the more difficult items. The imputation rules attempt to adjust the scores to better reflect the child's developmental level.

Evaluation of the assessment

Analysis was conducted on the Who Am I? data to determine whether this assessment was valid. The analyses included: comparing age equivalent scores to the child's age, comparison with the Number Knowledge, comparison with Who Am I? in other studies, an analysis of the items and of nonresponse. As the scoring procedures were being developed the NLSCY team consulted with Dr. de Lemos.

All of the evidence indicates that the test was valid and should provide data users with information about the child's developmental level. This assessment is not free of nonresponse bias. Please see Chapters 10.0 and 11.0, for more information on nonresponse.

14.4 Mathematics Computation Exercise

The Mathematics Computation Exercise administered to the child is a shortened version of the Mathematics Computation Test of the standardized Canadian Achievement Tests, Second Edition (CAT/2). The CAT/2 is a series of tests designed to measure achievement in basic academic skills.

The CAT/2 mathematical operations test measures the student's ability to do addition, subtraction, multiplication and division operations on whole numbers, decimals, fractions, negatives and exponents. Problem solving involving percentages and the order of operations are also measured. The short version of the test developed for the purposes of the NLSCY consists of 20 questions at each level, since Cycle 5.

For each level, the test administered at Cycle 6 is the same as the one used at Cycle 5. No modifications were made to the number of questions or to the questions themselves. The Cycle 6 math tests were all administered between January and June 2005. Therefore, the adjustment made in Cycle 5 to take into account the earlier assessment doesn't have to be repeated for Cycle 6.

Scoring

Each child who took the mathematics test was given a raw (gross) score (FMACS01), a scaled score referred to as the classical scaled score (FMACS02) and an IRT scaled score (FMACS03). The raw (gross) score is obtained simply by adding the number of correct answers. The classically derived scale score and the IRT scaled score are described as follows.

In Cycle 2, an IRT approach was used successfully to derive scores for the reading comprehension tests. Item Response Theory (IRT) is a measurement system commonly used in psychometric and educational testing. The IRT test scoring involves calculating either the most likely or the expected value of the ability of the examinee. The probability of a correct response to a question is assumed to be a certain logistic function of the examinee's ability. This probability is an S-shaped curve over the range of abilities. Its shape depends on the difficulty of the question, and sometimes also the discriminating power of the question (in the two-parameter IRT model) and the chance of a hypothetical no-ability examinee guessing correctly (in the three-parameter model, for multiple-choice questions).

Unlike the approach of the classical theory, IRT makes it possible to scale the scores without preset population standards. Using common test items linking grades, standards are estimated from the entire population of children taking the test for this cycle. Scores are derived, ranking each child within a level, and then the scores are vertically scaled to reflect the progression of scores throughout all the levels. In order to ensure comparability from year to year, each sample from each cycle must represent equivalent populations.

The three-parameter logistical model was chosen for the math tests. The three-parameter model takes into consideration both the difficulty and the discrimination of the item and also considers the pseudo-guessing component. In this way, IRT takes into consideration the pattern of responses. Two children with the same raw (gross) score will not have the same scaled score unless they answered exactly the same way. For example, a child who only answered the five easiest questions correctly would have a lower scaled score than one who only answered the five hardest questions correctly.

This IRT score differs from the other scaled score reported for the math test as it provides a greater precision in the estimates of the primary latent trait (which, in factor analysis, is the leading factor, which would be estimated by the linear weighting of items that explains as much of the variance of the items as possible.). Unlike the other reported scale score, this score is not referenced to an external population of expected performance, but is instead measured against

the expected performance of the current population as estimated by all test subjects. Past rank test analysis performed using both methods of scoring showed no significant difference between the two measurements.

The classical scaled score is derived from standards (norms) established by the Canadian Test Centre (CTC) in 1992. The CTC developed these standards from a sample of Canadian children from all 10 provinces (although the test has been developed in English only, and so the sample represents only the English schools), which is referred to as the normative sample. The children from the normative sample received the complete test. The scaled scores are units of a single scale with equidistant intervals that cover all of the grade levels. The scale was developed using a Thurstone procedure derived from the classical testing theory.

The fact that a short test was used for children in the NLSCY sample meant that it was not possible to directly associate the CTC scaled scores with the raw (gross) scores obtained in the survey. For this reason, the CTC normative sample was used to calculate the percentile rank for each raw (gross) score on our shortened version of the test. For example, using level 6, we find in the short test a percentile rank of 0.94% for a raw (gross) score of 1. On the complete test, the percentile ranks of 0.55% and 0.99% correspond to raw (gross) scores of 3 and 4 and to scaled scores of 315 and 319 respectively. After linear interpolation, we obtain a scaled score of 318 for the gross score of 1 on the short version of the test.

The table below shows the relation between the raw (gross) scores and the scaled scores by grade for the NLSCY mathematics test.

Relation Between Raw Scores and Scaled Scores (Classical) by Grade for the Cycle 4 Mathematics Test

Raw Score	Classical Scaled Score						
	Level 4	Level 5	Level 6	Level 7	Level 8	Level 9	Level 10
0	267	294	311	330	361	397	406
1	285	306	318	338	376	423	419
2	301	324	332	359	401	449	430
3	314	339	347	381	425	477	443
4	327	355	365	405	443	504	475
5	339	370	383	426	464	530	495
6	350	382	397	444	480	554	518
7	361	392	409	461	494	574	536
8	371	403	421	477	506	589	565
9	380	414	433	492	517	605	581
10	388	425	445	506	529	623	597
11	396	434	456	518	540	641	619
12	405	443	468	529	557	659	636
13	416	453	480	541	570	678	662
14	425	464	495	550	583	696	681
15	434	478	510	559	597	717	703
16	445	489	527	574	614	739	724
17	458	503	544	594	637	760	751
18	475	522	564	611	664	781	791
19	497	540	584	636	684	803	830
20	524	568	622	674	729	825	871

14.5 Problem-Solving Exercise (16- and 17-Year-Olds)

The Problem-Solving Exercise was introduced in Cycle 6 for youth aged 16 and 17. This new assessment is a combination of the cognitive measures used in prior cycles for this age group. In an attempt to continue the measurement of development in children, it was felt that a more comprehensive measure of ability was required to see how children readied themselves to take on the challenges ahead. This point in transition is very important as certain educational decisions are starting to affect the career paths of children. Specific cognitive abilities, such as reading comprehension, problem-solving and decision-making are known to have a pivotal role in the choices and the opportunities presented to youth at this juncture.

Strategy and Revision

The Problem-Solving Exercise (Booklet 32) is a 20-item assessment. The measure is a combination of Booklets 30 and 31 that were used in the previous cycle. It was decided for Cycle 6 that the two booklets would be combined and the respondents would all be given the same assessment. In addition, two questions were taken from levels 7 and 9 of the math tests.

The questions found in this measure were taken from a pre-existing instrument that had already been developed and tested. These items were tested for the Youth in Transition Survey (YITS). Because of the copyright restrictions of the questions and the sensitivity of having the NLSCY administer similar questions to potentially overlapping populations, the NLSCY was given permission to use only items that were tested but excluded from the main YITS questionnaire. The targeted population used was different from that being assessed for the NLSCY, but it was felt that the cognitive construct was still appropriate for this cohort.

Methodology for Scoring

For the exercise, nine of the 20 questions are hand-scored by trained individuals at Statistics Canada. These individuals have been trained to score the items according to the scoring guide provided to them. Scoring was done by a small number of people and was subject to quality control procedures. The head of the scoring team also met with a staff member of the subject matter team to consult on scoring procedures.

When the scorers mark the complex items, a value of '1' or '2' or '9' is given. A score of '1' is given for full credit, a score of '2' is given for no credit and a score of '9' is given for missing values. In those circumstances where items have been attempted and then crossed out, a score of '2' (no credit) is given, as opposed to a '9' (missing value). This includes erased or crossed-out work, unless it is clear that the erased or crossed-out work is correct, in which case it is given a value of '1'. If the respondent has written something, but the scorer cannot identify what is written, e.g., because it has been scribbled out too successfully, then the item is given a score of '2'.

The remaining items in the booklet are data-captured by trained staff at Statistics Canada and analyzed by the NLSCY methodologists. Taking the scores from the scored items and the data captured responses, analysis is conducted and an overall score is derived.

To obtain the Problem Solving Exercise score (FMAYdS01), the three-parameter model from Item Response Theory was used. For free-response items, the pseudo-chance parameter (to model guessing in multiple-choice items) was fixed at a value of zero. The Maximum Likelihood Estimation (MLE) of the three item parameters (discrimination, difficulty and pseudo-chance) followed by the *Expected A Priori* (EAP) estimation of the score was performed in an iterative process until the score converged sufficiently. The statistical software SAS was used to perform these computations.

As a final step, a lower bound was placed on the score, which raised the scores of the bottom 35 respondents. Removing or lowering this bound decreased the correlation of the score with all of

the math test scores from the first three cycles of the NLSCY. These 35 respondents' overly low scores appear to be an artifact of guessing or the low-stakes nature of the test.

Treatment of attempted items

Before the iterative estimation process began, raw (gross) scores for each item were calculated. Incorrect answers were scored as zero and correct answers were scored as 1, which is standard for the IRT. Partially correct answers were scored as marks received divided by the maximum possible mark, e.g. 1/2 for half-marks.

Treatment of items with no response

Items without a response ("unanswered items") can be either "omitted" items or "not-reached" items. Omitted items are those that the examinee probably saw but did not answer. In the scoring of the Problem Solving Exercise, an unanswered item earlier in the test than the last item attempted was considered to be an omitted item. The first and second items after the last item attempted were also considered to be omitted items. In this case, the respondent probably saw the questions, decided that they were too difficult, and stopped taking the test.

Usually, examinees omit items because they do not know the correct answer. Therefore, an omitted free-response item was given a zero mark. An omitted multiple-choice item was given a mark of one divided by the number of choices.

Not-reached items are those that the examinee probably did not see. These items do not provide any information about the ability of the respondent. In the Cognitive Measure scoring, all items up to and including the item after the last attempted item were considered to be reached items. Items not reached by a respondent were ignored in the estimation of that respondent's ability.

Note: The Problem Solving Exercise is not timed. The respondent can take as much time as needed to complete the test.

Scores

(FMAYdS01)

The scores presented on the data file have a mean of 0 and a standard deviation of 1. This scale is standard in IRT.

(FMAYfS02)

This score is sometimes called the "raw score" and represents the number of items that have been answered correctly

14.6 Literacy Assessment (18- and 19-year-olds)

This direct assessment that measures the literacy abilities of youth aged 18 and 19 was added to the NLSCY in Cycle 6. This assessment is made up of 36 items taken from an adaptive, abridged version of the International Adult Literacy and Skills Survey (IALSS) and the Adult Literacy and Life Skills Survey (ALL). The main purpose of this assessment is to establish how well these youth use printed information to function in society. These items were used because they are measures of cognitive ability that have already been established. The Literacy assessment covers two types of literacy: prose literacy and document literacy.

- ◆ Prose Literacy: is the knowledge and skills needed to understand and use information from texts such as editorials, new stories, poems and fiction.
- ◆ Document Literacy: is the knowledge and skills required to locate and use information contained in various formats such as tables, forms, graphs and diagrams.

It was important to include this measure of literacy in the NLSCY given the changing meaning of this concept. Definitions of reading and literacy have changed over time in parallel with changes

in our society, economy, and culture. The growing acceptance of the importance of lifelong learning has expanded the views and demands of reading and literacy. Literacy can be viewed as a developing set of skills, knowledge, and strategies that individuals build on throughout their lives in various contexts and through interaction with their peers and with the larger communities in which they participate.

Methodology for Scoring

This assessment is hand-scored by trained individuals at Statistics Canada. These individuals have been trained to score the items according to the scoring guide provided to them. Scoring was done by a small number of people and was subject to quality-control procedures. The head of the scoring team also met with a staff member of the subject matter team to consult on scoring procedures.

For the literacy assessment, there are three possible marks for each item. A value of '1' is given for full credit, a value of '7' is given for no credit and a value of '0' is given for a missing value. These values are recorded on the score sheet at the end of each booklet. Once completed, the score sheets are data captured and sent to the NLSCY team for analysis.

The literacy score (FLIYfS01) is simply the number of correct ('full credit') answers among the 36 items.

Although the questions of the NLSCY Literacy Assessment were selected from among the questions of 2003 International Adult Literacy and Skills Survey, the NLSCY Literacy Assessment score and the IALSS document literacy score are not comparable. Differences in the way the assessment was conducted, in the marking process and in the scoring methodology are such that the two scores should not be compared.

Note: The literacy assessment is administered in the youth's home by the interviewer and it is not timed. The respondent can take as much time as needed to complete the test.

14.7 Numeracy Assessment (20- and 21-year-olds)

This assessment was included in Cycle 6 to measure the numeracy skills of the oldest respondents. Similar to the literacy assessment, the 32 items included in the measure were taken from an adapted, abridged version of IALSS and ALL. As with the literacy assessment, the main purpose of this measure is to determine how well these youth use printed information to function in society.

Numeracy refers to the knowledge and skills required to effectively manage mathematical demands in diverse situations. Some researchers have identified a link between literacy and numeracy and together they are key determinants of workplace success.

As with the prose and document tasks, quantitative tasks require individuals to match information in a question or a directive with information stated in one or more texts where a text could be either continuous or non-continuous. In addition, quantitative tasks may require respondents to deal with plausible distracters when extracting information for an arithmetic operation. Individuals are also required to process some type of information. While the type of information varies for the prose and document tasks, requested information is always an amount in quantitative tasks.

Methodology for Scoring

Similar to the literacy assessment, this measure is hand-scored by trained individuals at Statistics Canada. These individuals have been trained to score the items according to the scoring guide provided to them. Scoring was done by a small number of people and was subject to quality-control procedures. The head of the scoring team also met with a staff member of the subject matter team to consult on scoring procedures.

For the numeracy assessment, there are three possible marks for each item. A value of '1' is given for full credit, a value of '7' is given for no credit and a value of '0' is given for a missing value. These values are recorded on the score sheet at the end of each booklet. Once completed, the score sheets are data captured and sent to the NLSCY team for analysis.

The numeracy score (FNUYfS01) is simply the number of correct ('full credit') answers among the 32 items.

Although the questions of the NLSCY Numeracy Assessment were selected from among the questions of the 2003 IALSS, the NLSCY Numeracy Assessment score and the IALSS numeracy score are not comparable. Differences in the way the assessment was conducted, in the marking process and in the scoring methodology are such that the two scores should not be compared.

Note: The numeracy assessment is administered in the youth's home by the interviewer and it is not timed. The respondent can take as much time as needed to complete the test.

15.0 Survey of Northern Children, Cycle 6

15.1 Introduction

The Survey of Northern Children (SNC) was conducted by Statistics Canada between February and April 2005, as a complementary survey to the National Longitudinal Survey of Children and Youth (NLSCY). The survey was undertaken with the co-operation and support of Human Resources and Social Development Canada.

The SNC is a census type cross-sectional survey of children living in the Yukon and Nunavut who were born in 1999 and who were enrolled in Senior Kindergarten in September 2004. The survey was conducted at the same time as Cycle 6 of the NLSCY, using the same interviewers as well as a similar survey instrument.

The development plan for the SNC was based on the assumption that the Ministries of Education for the Yukon, Northwest Territories (NWT) and Nunavut would provide Statistics Canada with enrolment lists, including the phone number of the parent or guardian. Unfortunately, the Northwest Territories was unable to provide sample frame information before the start of the collection period.

15.2 Background

National Longitudinal Survey of Children and Youth, Northern Component, Cycles 1, 2 and 3

The collection of information about Canadian children living in the territories began with Cycle 1 of the NLSCY in 1994 and 1995. Cycle 2 data were collected between 1996 and 1997, and the Cycle 3 collection took place between the fall of 1998 and the spring of 1999.

Because both the NLSCY and the National Population Health Survey (NPHS) were interested in gathering data for residents of the Yukon and Northwest Territories, the two surveys were combined and shortened in Cycles 1, 2 and 3 in order to keep the interview to a reasonable length. Interviews were conducted by interviewers who had been hired and trained by the Yukon and NWT Bureaus of Statistics. Households were contacted by telephone or visited in person, depending on their location, and interviewers used paper and pencil questionnaires.

As in all cycles of the northern surveys, information on the NLSCY portion was provided by a parent. A number of sections which are on the main NLSCY questionnaires, such as temperament, family functioning, and child care, were excluded from the northern survey while other, more relevant questions were added, including whether or not respondents were First Nations persons, and whether an adoption was an Aboriginal custom adoption.

National Longitudinal Survey of Children and Youth, Northern Component, Cycle 4

In 2001-2002, data for the territories were collected by Statistics Canada interviewers for the first time, along with the collection of the NLSCY in the provinces. The sample for the SNC was based on the responding children from the previous cycle.

Data collection was done from the regional offices, and interviewers used a computer-assisted telephone interviewing (CATI) application. This meant that data could be processed jointly and more quickly with the NLSCY.

Cycle 4 Summary

The tracking of respondents in the territories can be very difficult due to the typically transient nature of the population. In longitudinal terms, this results in serious coverage issues since children who no longer reside in the territories are no longer considered part of the sample. Many respondents were lost over the four cycles due to this type of attrition and it became evident following the completion of Cycle 4 that data quality, particularly the longitudinal aspect, was a serious issue in the territories. Data for Cycle 4 could only be made available for the Yukon and Northwest Territories; low response rates prevented the release of data for Nunavut.

Cycle 5 Summary

Difficulties with collection in the Yukon, NWT and Nunavut, accompanied by poor response rates, lack of knowledge about the sample design and problems encountered with data processing all served to jeopardize the accuracy of both the longitudinal and cross-sectional estimates.

In an effort to maintain the momentum provided by the ongoing cycles of the NLSCY, while at the same time creating a database of quality information about children, data collection in the North underwent a thorough review in which a number of alternate collection options were considered. It was decided that a new sampling frame that would provide current, reliable information was required.

In the summer of 2002, discussions began with each of the territories to obtain contact information about 5-year-old children from school board registration records. The SNC, Cycle 5 Yukon experience has shown that files from the Ministry of Education work extremely well as a frame. Unfortunately, the Northwest Territories and Nunavut were unable to provide sample frame information before the start of the collection period, and so data are not available for those territories.

At that time, an equally reliable and current frame had not been found for other age groups in the territories, although it was hoped that this issue would be addressed by the future Aboriginal Children's Survey.

Cycle 6

Difficulties with collection in the North continued to be a problem, particularly in Nunavut. Although the survey started out as a telephone interview survey, the response rate was so low that the survey was converted to personal interviewing and interviewers were flown into the communities to track down respondents and conduct the interviews in person.

Telephone response rates were determined to be low because many northern respondents do not have a phone and therefore, the phone number supplied to the school board was for the community centre or another family member with a phone. However, with personal interviewing a good response rate was obtained, but at a greater collection expense.

Next Cycle

In the year 2006, the Survey of Northern Children will be combined with the Aboriginal Children's Survey. Combining these two surveys will strengthen the understanding of how our young children in the North are developing and will provide valuable information to guide future policy development. It will also support our current commitment under the Early Childhood Development Agreement to report on a common set of indicators of young children's well-being.

15.3 Establishing the Survey of Northern Children

Objectives

There is growing recognition that the first five years of life are critical to a child's development. The early years shape long-term outcomes related to scholastic achievement, to employment success, to health, to quality of life and to the ability to adapt.

With this in mind, this survey covers a comprehensive range of topics including the health of the children, information on their physical development, learning and behaviour and data on their social environment (family, friends, schools and communities). As well, the survey has a question on Aboriginal identity that will offer the opportunity to compare how questions were answered for Aboriginal children vs. non-Aboriginal children.

Content

In general, the SNC, Cycle 6 application was similar to that of the NLSCY for the provinces with a few exceptions.

- 1) Only questions that applied to 5-year-olds were asked of the person most knowledgeable (PMK) about the child or the PMK's spouse;
- 2) Since, data collection was done by telephone using a CATI application (with the exception of 191 respondents in Nunavut and 47 respondents in the Yukon who were interviewed in a CAPI environment), direct measures used in the NLSCY were not part of the interview for the territories. Therefore, the following NLSCY tests were not included:
 - Peabody Picture Vocabulary Test - Revised
 - Number Knowledge Test
 - Who Am I?
- 3) Finally, in order to reduce respondent burden as well as remove questions that were not relevant, the following modules were removed from the application:
 - Neighbourhood Safety
 - Work after Birth
 - Temperament
 - Sleep
 - Child Care
 - Custody
 - Family Custody History

For further details on concepts and definitions, data collection methods and data processing steps used for the SNC (which were identical to those used for the NLSCY), please refer to the appropriate chapters in this document.

Survey Methodology

The target population of the SNC, Cycle 6 consisted of children who were born in 1999 and who were enrolled in Senior Kindergarten in September 2004. Since this was a census type survey, there was no need for a sample design, sample selection or sample allocation.

The Ministry of Education in the Yukon provided a sample frame of 326 children. Of this number, six children were considered to be out of scope and were not included in the weighting process. Of the 320 children remaining in the sample file, 286 respondents (the PMK or the PMK's spouse) responded to the Survey of Northern Children, Cycle 6.

The Ministry of Education in Nunavut provided a sample file of 683 children. Of this number 38 children were considered to be out of scope. Of the 645 children remaining in the sample file, 485 respondents (the PMK or the PMK's spouse) responded to the Survey of Northern Children, Cycle 6.

Weighting

The Survey of Northern Children, Cycle 6 is a census type survey, not a probability sample, which means that the weighting is quite simple. In addition, we were required to produce only cross-sectional weights since the survey is no longer longitudinal.

Cross-sectional weights were calculated for all children who responded to the SNC in Cycle 6 and again, since this was a census type survey, the initial cross-sectional weight assigned to each child was 1.

Nonresponse Adjustment

In order to calculate the cross-sectional weight a simple nonresponse adjustment was made to adjust the weights of the respondents to account for the selected children who did not respond.

The only available information for nonrespondent children is that which is provided on the enrollment lists from the school boards, thus, it is hard to make inferences about these children. The only derived variable that proved significant is whether the child lived in Whitehorse or elsewhere. This creates two different groups. The nonresponse adjustment (NRA) factor in each group is the sum of the respondents and nonrespondents divided by the number of respondents, i.e., it is greater than or equal to 1.

Post-stratification

The target population for this survey is 5-year-old children, i.e., children who were born in 1999 and who were enrolled in Senior Kindergarten in September 2004. This is different than targeting all 5-year-old children (including those not enrolled in schools) in Nunavut and the Yukon. For this reason, the cross-sectional weights will not be post-stratified to conform to known population totals since those totals would include all 5-year-old children in the Yukon and Nunavut. This survey was a census type survey, which means that each child's record is counted once and has equal weight in the survey results.

Details of the method used to calculate these weights are presented in Chapter 12.0.

Data Quality

Response Rates

In the SNC, Cycle 6, a respondent is defined as a child who has at least one of the following components completed: the adult component or the child component.

Of the children who were part of the SNC Cycle 6, the Yukon sample produced a response rate of 89.4% and the Nunavut sample produced a response rate of 75.1%.

Table 1 – Cross-sectional Response Rates

Territory	Number of Children	Number of Respondents	Cross-sectional Response Rate
Yukon	320	286	89.4%
Nunavut	645	485	75.1%
Total	965	771	79.8%

Partial Response Rates

According to the definition given above, only one component needed to be completed in order for a household to be considered as a responding household. Almost all respondents completed both components.

Table 2 – Component Response Rate

Territory	Number of Respondents	Both Components Completed	Only Child Component Completed	Only Adult Component Completed
Yukon	286	276	10	0
Nunavut	485	443	42	0
Total	771	719	52	0

Survey Errors

The estimates derived from this survey are based on a census of children; therefore there is no sampling error.

Errors which are not related to sampling may occur at almost every phase of a survey operation. Interviewers may misunderstand instructions, respondents may make errors in answering questions, the answers may be incorrectly entered on the questionnaire and errors may be introduced in the processing and tabulation of the data. These are all examples of non-sampling errors.

There were many reasons why some households did not respond to the survey. In some cases, the interviewers were unable to trace the household, or to make contact with a selected household during the collection period. In other cases, the household refused to participate in the survey.

Guidelines for Tabulation, Analysis and Release

Estimates generated from this survey are subject to meeting the guidelines for analysis and release outlined in Chapter 17.0 of this document.

16.0 Analytic Issues

This chapter provides users with an overview of the various analytic issues that should be considered when analysing data from the National Longitudinal Survey of Children and Youth (NLSCY). Some of the points mentioned in this chapter have already been explored in greater detail in previous chapters. The purpose of this chapter is to highlight some of the key data analysis issues, which are:

- Statistics Canada recommends that the survey weights be used at analysis;
- attention must be paid when combining several cycles of data since repeated measures across time for an individual are likely to be correlated;
- standardized or normalized weights can lead to incorrect variance estimates;
- analysis should be accompanied by an estimate of sampling error;
- suggestions for dealing with missing data.

For a detailed description of the NLSCY sample, see Chapter 5.0 on Survey Methodology – Sample; for a description of how the survey weights are calculated, see Chapter 11.0 on Weighting and Treatment of Nonresponse; for a description of how to estimate sampling error, see Chapter 13.0 – Variance Estimation; for more on data quality, including a detailed description of the various sources of non-sampling errors in a survey, e.g., nonresponse, response, under coverage and processing errors, see Chapter 12.0 – Data Quality, Response Rates and Coverage.

16.1 How a Complex Sample Design Affects Analysis

Data analysis involves summarizing the data and interpreting their meaning in a way that provides clear answers to questions that initiated the survey. Sometimes the analyst simply wishes to describe the sample, but more often he or she wants to use the sample to describe some population.

When making inferences about a population that was surveyed, Statistics Canada recommends that the survey weights be used (either cross-sectional, or longitudinal, depending on the analysis). The reason for this is that for some characteristic of interest the distribution in the sample is unlikely to reflect the distribution of the characteristic in the population due to the complex sample design. Only by applying the survey weights can the population's distribution be preserved.

For example, stratification and clustering (both present in the NLSCY sample design) lead to unequal probabilities of selection: the probability that a child in the population is sampled by the NLSCY depends on the age of the child, the child's province of residence, etc. (In the sample there is a disproportionately number of children from small provinces.) Unequal nonresponse rates within the population can also lead to unequal representation of children in the sample. And clustering in the sample leads to the non-independence of units: children belonging to the same household are not independent.

Suppose that the analyst wants the distribution of children across Canada, i.e., by province, for the original cohort. The population of inference is: children aged 0-11 as of December 31st, 1994 who were living in one of the 10 provinces at the time of Cycle 1 collection (1994/1995). Two different sets of longitudinal weights could be used: the 'funnel' weights (for children who have responded to every cycle) or the non-funnel weights (for children who responded at Cycles 1 and 6, but not necessarily all in-between). The table below illustrates the difference between weighted and unweighted estimates of the number and proportion of children in Canada, using the funnel weights.

Table: The Distribution of Children in the Population, for the Original Cohort, Weighted versus Unweighted Estimates using Cycle 6 Funnel Weights (fwtcwd1l)

Province	Number of children (unweighted)	Percentage of children (unweighted)	Number of children (weighted)	Percentage of children (weighted)
Newfoundland and Labrador	689	6.00	88,986	1.91
Prince Edward Island	339	2.95	23,148	0.50
Nova Scotia	839	7.31	144,088	3.09
New Brunswick	710	6.18	115,131	2.47
Quebec	2,108	18.36	1,090,582	23.41
Ontario	2,834	24.68	1,773,616	38.08
Manitoba	905	7.88	182,869	3.93
Saskatchewan	958	8.34	173,611	3.73
Alberta	1,109	9.66	489,913	10.52
British Columbia	992	8.64	576,125	12.37
Total	11,483	100.00	4,658,069	100.00

Without the weights, the analyst would incorrectly conclude that 22.44% of children reside in the Atlantic provinces when, in fact, the true number is only 7.97%. The unweighted proportions reflect the fact that the sample has a disproportionate number of children from the smaller provinces (to ensure adequate sample size in small provinces). Without the weights, the analyst would also incorrectly conclude that there are only 11,483 children in the population when, in fact, there are over 4.6 million.

16.2 Unit of Analysis

In the NLSCY, the unit of analysis is always the child. While some household data were collected, no estimates can be produced at the household level; all estimates must be at the child level. For example, the number of children living in single-parent households can be estimated but **not** the number of single-parent households.

16.3 Longitudinal versus Cross-sectional Analysis

With the NLSCY, users have the choice of longitudinal or cross-sectional analysis.

Longitudinal Analysis

The population represented by the longitudinal weights is the population at the time of the child's initial selection. For children belonging to the original cohort (aged 10 to 21 at Cycle 6) two longitudinal weights are available, labelled FWTCW01L and FWTCWd1L. The first set of weights is for children who responded (or their PMK responded) to Cycles 1 and 6, but not necessarily all cycles in between. The second set of longitudinal weights, called "funnel" weights, apply only to those children who responded (or their PMK responded) to every Cycle 1 through 6.

For example, if an analyst were interested in the characteristics at Cycle 6 of children introduced in Cycle 1, but not interested in their data from Cycles 2 to 5, then the first set of longitudinal weights would be appropriate. However, if the analyst wanted to look at the data from all cycles, then it would be better to use the second set of longitudinal weights. (This analysis would be an example of repeated measures, explained in the next section). It should be pointed out that the two sets of longitudinal weights have been tested for some key variables and that the estimates by either set are similar.

Cross-sectional Analysis

Cross-sectional analysis refers to analysis where the population of interest is some population occurring after the initial longitudinal population. For Cycle 6, cross-sectional weights are only calculated for children aged 0 to 5 (for a description of the sample composition of these children, see Chapter 5.0 Survey Methodology – Sample). It is not recommended that the original cohort be used to make inferences about the cross-sectional population of all 10- to 21-year-olds living in Canada at the time of Cycle 6 collection since this sample has never been topped up for immigrants who arrived after 1994. (For more details on the impact of immigration, see Chapter 12.0, Data Quality, Response Rates and Coverage.)

16.4 Simple Weighted Estimates (Totals, Averages and Proportions)

This section explains how to use survey weights to generate estimates for simple descriptive statistics. Totals, averages and proportions are typically estimated for a wide range of characteristics collected from the sample units. These characteristics, often referred to as variables, may be categorical or qualitative, e.g. sex, or they may be quantitative, e.g. birthweight. Depending on the type of data, different statistics and different statistical procedures for the production of estimates are appropriate.

Estimates can be produced for the whole survey population or for specific subgroups or domains of the population, e.g. provinces. The following estimators can be applied to any probability sample design – simple or complex.

16.4.1 Estimating a Population Total

The estimate of the total number of units in the survey population is calculated by adding together the final adjusted weights of the responding units:

$$\hat{N} = \sum_{i \in S_r} w_i$$

where i is the i^{th} responding unit in the sample and w_i is its final adjusted weight and S_r is the set of responding units.

For quantitative data, the estimate of a total value is the product of the final weight, w_i , and the value, y_i , for each responding unit, summed over all responding units:

$$\hat{Y} = \sum_{i \in S_r} w_i y_i$$

16.4.2 Estimating a Population Average

For quantitative data, the estimate of an average value in the population is obtained by adding together the product of the sample value and the weight for each responding unit, and dividing this amount by the sum of the weights. In other words, the estimate of the average in the population is the estimate of the total value for quantitative data divided by the estimate of the total number of units in the population.

$$\begin{aligned}\hat{Y} &= \frac{\sum_{i \in S_r} w_i y_i}{\sum_{i \in S_r} w_i} \\ &= \frac{\hat{Y}}{\hat{N}}\end{aligned}$$

16.4.3 Estimating a Population Proportion

For qualitative data, the estimate of the proportion of units in the survey population having a given characteristic is obtained by adding together the weights for the units having that characteristic, and dividing this total by the sum of the weights for all respondents. In other words, the estimate of the proportion in the population is the estimate of the total number of units possessing the given characteristic divided by the estimate of the total number of units in the population.

$$\begin{aligned}\hat{P} &= \frac{\sum_{i \in S_r \cap C} w_i}{\sum_{i \in S_r} w_i} \\ &= \frac{\hat{N}_C}{\hat{N}}\end{aligned}$$

where C is the collection of units having the given characteristic.

16.4.4 Estimating for Domains of the Population

Estimates may be required for domains such as age group or sex.

The size of the population for a domain of interest for either qualitative or quantitative data is estimated as follows:

$$\hat{N}_{domain} = \sum_{i \in S_r \cap domain} w_i$$

The estimate of a domain total for quantitative data is:

$$\hat{Y}_{domain} = \sum_{i \in S_r \cap domain} w_i y_i$$

The estimate of a domain average for quantitative data is:

$$\begin{aligned}\hat{Y}_{domain} &= \frac{\sum_{i \in S_r \cap domain} w_i y_i}{\sum_{i \in S_r \cap domain} w_i} \\ &= \frac{\hat{Y}_{domain}}{\hat{N}_{domain}}\end{aligned}$$

The estimate of a domain proportion for qualitative or quantitative data is:

$$\begin{aligned}\hat{P}_{domain} &= \frac{\sum_{i \in S_r \cap domain \cap C} w_i}{\sum_{i \in S_r \cap domain} w_i} \\ &= \frac{\hat{N}_{domain \cap C}}{\hat{N}_{domain}}\end{aligned}$$

16.5 Normalized Weights

In order to ensure that survey estimates of the characteristics of the finite survey population are approximately unbiased with respect to the survey design, each observation has a weight. However, using the survey weight with certain procedures of software not specialized in the analysis of survey data, e.g. SAS and SPSS, can lead to erroneous conclusions. This is because the software package interprets the sum of the weights to be the number of observed units in the sample, and consequently overestimates the statistical power of the test.

When performing statistical tests with model-based software, the analyst should at least re-scale the original survey weights to ensure that the sum of the normalized weights is equal to the sample size. (Most SAS and SPSS procedures have options to normalize weights.) The number of observed units is then at least correct. However, since the model-based software still treats the units as if they were independently and identically distributed, the sample variance estimates produced are likely to be incorrect.

The use of normalized weights can be seen to be an incomplete implementation of the design-based approach. Normalized weights should only be used when the statistical analysis cannot be conducted properly using a design-based software, or if there is insufficient design information to correctly calculate sampling variances, e.g., there are no bootstrap weights, or if the analyst is simply running preliminary analyses.

How to normalize?

Normalized weights are calculated by dividing the final survey weight for each unit to be analyzed by the (unweighted) mean of the survey weights of all units being analyzed:

$$w_i^{norm} = \frac{w_i}{\bar{w}}$$

For example, suppose that there are six children in the sample with final survey weights of 1, 3, 4 and 6. The normalized weights are presented in the table below.

Table: Example of normalized weights

Observation number	Final Survey Weight	Normalized Weight
1	1	0.25
2	3	0.75
3	4	1.00
4	4	1.00
5	6	1.50
6	6	1.50
Total	24	6

Normalization should only be performed on respondents. For example, if the characteristic of interest is missing for two of the units in the sample, then the sum of the normalized weights should equal 4, not 6 (see the table below).

Table: Normalized weights in the presence of nonrespondent units

Observation number	Response Status for the variable of interest	Survey Weight	Incorrect Normalized Weight (calculated using nonrespondents)	Correct Normalized Weight (respondents only)
1	Respondent	1	0.25	0.33
2	Respondent	3	0.75	1.00
3	Respondent	4	1.00	1.33
4	Respondent	4	1.00	1.33
5	Nonrespondent	6	Excluded from the analysis	Excluded from the analysis
6	Nonrespondent	6		
Total		24	3	4

Normalization should also be done separately for each domain of analysis, to ensure that the sum of the normalized weights respects the sample size by domain. Suppose that the analysis is by province and that two respondents belong to province A and four belong to province B. The sum of the normalized weights for province A should equal 2 and the sum of the normalized weights for province B should equal 4 (see the table below).

Table: Normalized weights, by domain of analysis

Observation number	Response Status for the variable of interest	Province	Survey Weight	Normalized Weight (by Province)
1	Respondent	A	1	0.50
2	Respondent	A	3	1.50
Total			4	2
3	Respondent	B	4	0.80
4	Respondent	B	4	0.80
5	Respondent	B	6	1.20
6	Respondent	B	6	1.20
Total			20	4

The problem with normalization

In most surveys with a complex design, the effective number of units in the sample is smaller than the actual number due to the correlation of sampled units (which is often the result of clustering). In such cases, normalization leads to:

- an over-estimation of the effective number of units,
- an under-estimation of the variability,
- too many significant results (the true size of the test is larger than expected).

Some users of normalized weights consequently use a more conservative level of test (usually 1% instead of 5%) before declaring a result significant. But such a rule of thumb is flawed: sometimes it is too strict, sometimes not strict enough.

The example below illustrates how the estimated p-value generated using normalized weights can be incorrect, using a:

- test of independence with PROC FREQ in SAS,
- logistic regression with PROC LOGISTIC in SAS.

This example uses synthetic data from a Statistics Canada health survey.

Example 1: Is a married person's date of birth a predictor of divorce?

In this example, the analyst wants to know if a person's marital status and when the person was born during the year are related. Is it necessary to be born in the first three months of the year in order to be married? Are people born in the last three months of the year more likely to be unmarried?

Results after normalization:

SAS: a value of $\chi^2 = 24.33$ ($p=0.0038$)

Conclusion: The analyst would conclude that marital status and birth date are linked.

Results with design-based software:

SUDAAN: a value of $\chi^2 = 14.95$ ($p=0.0955$)

Conclusion: The analyst would correctly conclude that marital status and birth date are independent.

Conclusion

With model-based software, normalization is an attempt to try to correct a number of procedures, but it constitutes an incomplete implementation of the design-based approach, since it takes account of the survey weights, but not other aspects of the design (such as units not being independently and identically distributed).

Normalized weights often lead to an under-estimation of the sampling variance of the estimates and consequently declarations of significant differences where none exist. A rule of thumb is often used to compensate, but this adjustment is sometimes too conservative, sometimes not conservative enough.

When calculating normalized weights, the domain of analysis and number of respondents should be accounted for.

With design-based software (such as SUDAAN or Bootvar), normalization is not required.

16.6 Repeated Measures

Repeated measures are multiple observations of the same variables from the same sampled unit. Repeated measures arise when a sample is followed over time – such as in a longitudinal survey – and the same information is collected at multiple time points, e.g. height and weight. In this case dependence among repeated observations from the same individual should be accounted for in the estimate of sampling variance.

When variance estimation is performed using the bootstrap weights (as is recommended for the NLSCY), there is a simple way of factoring in the correlation structure of multiple observations from the same sampled child: prepare the analysis file such that each record is one event or observation, where the survey weight and bootstrap weights associated with that record are the survey weight and bootstrap weights for the child who experienced the event or produced the observation.

For example, suppose that a researcher is interested in repeated measures x and y , say height and weight, over the first three cycles of the NLSCY. Let x_{ij} represent the height of child ' i ' at cycle ' j ' and y_{ij} represent the weight of child ' i ' at cycle ' j '. Let w_1 be the child's survey weight, e.g. the non-funnel longitudinal weights, and let $bs1_1$, $bs2_1$, etc. be the bootstrap weights for the child. Suppose that for the first child, we have data from Cycles 1, 2 and 3; for the second child we only have data from Cycle 1; for the third child we have data from Cycles 1 and 2; and for the fourth child, we only have data from Cycles 2 and 3. Then the input file would be constructed as below, and the estimated sampling variance using the bootstrap weights would be calculated as usual.

Record	Child	Survey Weight, w	Variable x	Variable y	Bootstrap weight, replicate 1	Bootstrap weight, replicate 2
1	1	w_1	x_{11}	y_{11}	$bs1_1$	$bs2_1$
2	1	w_1	x_{12}	y_{12}	$bs1_1$	$bs2_1$
3	1	w_1	x_{13}	y_{13}	$bs1_1$	$bs2_1$
4	2	w_2	x_{21}	y_{21}	$bs1_2$	$bs2_2$
5	3	w_3	x_{31}	y_{31}	$bs1_3$	$bs2_3$
6	3	w_3	x_{32}	y_{32}	$bs1_3$	$bs2_3$
7	4	w_4	x_{42}	y_{42}	$bs1_4$	$bs2_4$
8	4	w_4	x_{43}	y_{43}	$bs1_4$	$bs2_4$

Note that the sum of the survey weights would be much greater than the total number of children in the longitudinal population. The following section discusses issues that the analyst should be kept in mind when pooling data.

16.7 Pooling Data

Analysts who pool data across cycles should be aware that unless it is done carefully, the resultant analysis could be incorrect. The issues are as follows:

- Pooling can result in a child appearing more than once in the data, i.e., repeated measures.
- To avoid repeated measures, the analyst would have to combine independent samples. This can lead to the analyst having to pick one measure from across several cycles, or deriving a new combined measure.
- It may be difficult to define the reference population that is represented by the pooled sample and therefore about which inferences are being made.
- The analyst may have to calculate new weights for the pooled sample.

These issues are explained below.

If dependent samples are pooled, resulting in some children appearing more than once in the pooled sample, then the correlation structure needs to be factored into the estimates of sampling variance in order for test statistics to be correct, i.e., the analyst is dealing with repeated measures.

Any pooling across cycles of the NLSCY's original cohort will result in a dependent pooled sample and thus repeated measures, for example, combining 0- to 11-year-olds in Cycle 1 with 2- to 13-year-olds in Cycle 2 (these are the same children).

Note that just because a child has a cross-sectional weight greater than zero does not mean that the child is not also followed longitudinally. One easy way to identify if a child appears more than once in the pooled sample is to check if a child's identifier (variable PERSRUK) appears more than once.

If independent samples are pooled, then the child will only appear once on the file. For example, a new independent sample of 1-year-olds is selected at each cycle of the NLSCY, so samples of

1-year-olds could be easily pooled. The analyst could then simply use the cross-sectional weight and bootstrap weights for each baby (from the relevant cycle), and the reference population for the analysis would be all 1-year-olds who were born during the reference years covered by the pooled cycles.

This method treats each different year of birth as a stratum and allows users to easily use existing survey weights (no adjustments are necessary). And there would be only one measure per child, from the cycle when the child was a 1-year-old.

Often, though, there are many possible ways to combine data across several cycles. For example, suppose that the analyst is interested in 0- to 4-year-olds in Cycles 1 and 2. The analyst could simply group together all 0- to 4-year-olds from Cycle 1 and all 0- to 4-year-olds in Cycle 2, in which case some children would appear more than once since the 0- to 2-year-olds in Cycle 1 are 2- to 4-year-olds in Cycle 2. The analyst would be dealing with repeated measures.

If the analyst only wanted a child to appear once in the file, then he or she could select all 0- to 4-year-olds in Cycle 1 and combine them with the independent sample of 0- to 1-year-olds in Cycle 2. The analyst could then use the cross-sectional weights for 0- to 4-year-olds from Cycle 1 and the cross-sectional weights for 0- to 1-year-olds in Cycle 2. The reference population for analysis would be all children who were 0-1 in Cycles 1 and 2, and 2-4 years old in Cycle 1.

Alternatively, the analyst could derive a new measure that combines all the information across the cycles, and then pool the samples independently so that each child only appears once in the final file, with the new derived measure.

To illustrate these different options, take the example presented earlier (0- to 4-year-olds in Cycles 1 and 2) and suppose that the analyst is interested in the general health of the child (excellent, very good, good, fair, poor). This is variable AHLCQ01 in Cycle 1 and variable BHLCQ01 in Cycle 2. For the children in the pooled sample who are 0-2 years old in Cycle 1 and 2-4 years old in Cycle 2, the analyst has several options: these children could appear twice in the file (repeated measures), or the analyst could use the Cycle 1 data only, or the Cycle 2 data only, or the analyst could derive one measure that combines the information from both cycles.

Combining data across cycles may result in the analyst having to recalculate new weights for the new pooled sample. For example, suppose that the researcher is interested in 0- to 5-year-olds in Cycles 2 and 3. The analyst may decide to pool the 0- to 5-year-olds in Cycle 2 with the 0- to 5-year-olds in Cycle 3 who were not present at Cycle 2. For this combined sample, the analyst could use the cross-sectional weights for 0- to 5-year-olds from Cycle 2. However, since the analyst only chose a subset of the Cycle 3 children, the Cycle 3 weights would have to be adjusted. For example, they could be rescaled so that for the subsample of 0- to 5-year-olds selected the weights sum to the known population totals of 0- to 5-year-olds at Cycle 3.

One other issue to remember when analyzing data from pooled samples is that for a particular age, the number of children in the sample may vary substantially from one cycle to another. For example, Cycle 3 has an unusually large sample of 5-year-olds.

16.8 Nonresponse

Like any other survey, the NLSCY is subject to nonresponse. There are two main types of nonresponse: total and partial.

Total nonresponse is the complete lack of data for a selected and eligible child due to factors such as refusal to take part in the survey or inability to trace the child. Total nonresponse is taken into account and corrected by the sampling weights. See Chapter 11.0 for details.

Partial nonresponse is when specific questions are not answered. Note that not all questions have to be answered for a child or youth to be considered a survey respondent at Cycle 6. For example, 90% of respondent children complete the PPVT (for more examples of partial nonresponse, see Chapter 11.0) Some missing questions are imputed, but most are not.

Nonresponse is a concern for analysts because if it is not properly corrected for, it can lead to biased estimates. Conclusions based on biased estimates can be erroneous. When analysing NLSCY data it is common to be faced with partial nonresponse to some variables of interest. There are a few ways to deal with this situation:

1. Only keep records for which complete information is available for your main analysis, while in a separate side analysis profiles your nonrespondents. For instance, you can note that your respondent group is more likely to live in certain provinces or have certain income levels, and that these differ from your nonrespondents. Being upfront about describing the differences between the two groups alerts readers of your results that nonresponse was an issue with the data and helps them properly interpret your analytical findings.
2. Report partial nonresponse as a category
3. This approach means you report the nonresponse category as a valid category in tables or in models. This is especially well-suited to categorical data, and most of the NLSCY data are in fact categorical.
4. Reweight the records with a response to account for the partial nonrespondents
5. Ignore the partial nonresponse records, but the weights of the respondents are increased to account for the nonrespondents. This is an especially interesting strategy when an entire component of the survey is missing (for instance, the self-complete component). If however you only have a few missing data here and there, this may not be the best strategy because it throws out records with any partial nonresponse.
6. The weight adjustment applied to the respondents is simply the inverse of the response probability, which is often taken to be the weighted response rate.
7. Note that the weights for NLSCY are post-stratified to known counts by age, sex and province. When reweighting to adjust for partial nonresponse, these control totals are not respected anymore. You may choose to re-poststratify after adjusting for nonresponse or instead of reweighting you can use the imputation approach (see next point).
8. Impute partial nonresponse (replace missing values with replacement values)
 - a. The advantage over reweighting is that all records are kept, which means no data are discarded. Only the missing values in each record are filled in.
 - b. It is important to quote the imputation rate with your analytical results and give information about the imputation strategy used.

For an example of how to assess and report partial nonresponse, see Appendix II: Partial Nonresponse Analysis. For those looking for more information on nonresponse, some references are given below. The list is by no means exhaustive.

Nonresponse overview

- Survey Methods and Practices (Statistics Canada, 2003)
 - Different chapters discuss nonresponse related issues
 - This book can be ordered from Statistics Canada's website (www.statcan.ca)
- Sampling: Design and Analysis (Lohr, S., Duxbury Press, 1999)
 - Chapter 11.0 is devoted entirely to nonresponse

Nonresponse treated in more detail

- Incomplete Data in Sample Surveys (Madow, W.G., et al., New York: Academic Press, 1983)
- Nonresponse in Household Surveys (Groves, R. and Cooper, M., New York: Wiley, 1998)
- Statistical Analysis with Missing Data (Little, R.J.A. and Rubin, D.B., New York: Wiley, 1987)

More information on how to deal with partial nonresponse in NLSCY is also available through a workshop offered by Statistics Canada. The slides of this workshop are available in the Research Data Centres. As well, the workshop can be presented by a methodologist from Statistics Canada upon request.

16.9 Other Sources of Non-sampling Errors

Besides nonresponse, the analyst should keep in mind the effect that other non-sampling errors can have on the analysis – errors that could potentially bias the results – such as those due to: undercoverage of the population, processing errors, response errors, etc. Nonsampling errors are described in detail in Chapter 12.0 – Data Quality, Response Rates and Coverage.

16.10 Computing the Variance with Certain Software Applications

SAS and SPSS, software applications commonly used by analysts, are able to compute point estimates correctly using sampling weights. However, with the exception of some SAS procedures, these applications could not take into account the NLSCY's sample design (including the complex sample design and weight adjustments for nonresponse and post-stratification) when estimating the sampling variance of a point estimate. As a result, many software applications would underestimate the NLSCY's sampling variance, sometimes substantially. For this reason, the analyst is strongly encouraged to use the Bootstrap weights for variance estimation.

Some software applications can use these Bootstrap weights: SUDAAN, WesVar and STATA9 take into account the sample design in calculating the variance, using the Balanced Repeated Replication (BRR) method. The creation of BRR weights differ in theory from the creation of Bootstrap weights, but the variance estimator is the same. As a result, the NLSCY Bootstrap weights can be used with these applications.

There is a stand alone version of SUDAAN, as well as an integrated version with SAS. The latter gives the flexibility to use the SUDAAN procedures within SAS. With its nine procedures, SUDAAN can produce estimates of means, ratios and totals, independence tests, linear, log-linear and logistic regressions and survival tests. SUDAAN can read SAS and SPSS files, as well as other common types of files.

WesVar uses a "point and click" approach, which makes it easy to learn. The types of analysis are more limited than SUDAAN, but WesVar still allows the variance estimation of means, ratios and totals, independence tests, as well as linear and logistic regressions. WesVar can read SAS, SPSS and other common type of files, but the application requires an additional step to create a WesVar file before proceeding with the analysis. Bootstrap weights can be used with other applications which offer the required programming environment and the desired analytical tools. SAS and SPSS macros have been developed by the NLSCY to use the Bootstrap weights to produce variance estimates based on the sample design.

The reader will find useful information on using bootstrap weights with SUDAAN and WesVar at <http://www.statcan.ca/english/freepub/12-002-XIE/2004002/pdf/phillips.pdf> For more details on estimating sampling variance, including details on another tool called Bootvar, refer to Chapter 13.0.

16.11 Coefficients of Variation for Proportions

Coefficients of variation (CV) have been widely used for a long time to measure the quality of estimates such as totals, proportions or others. However, when the CV is used to assess the quality of an estimated proportion, the analyst must proceed with great care. The CV is not always an appropriate measure of quality for estimated proportions. More details about this are available in Chapter 13.0 on Variance Estimation and Chapter 17.0 on Guidelines for Tabulations, Analysis and Release.

16.12 Standard Deviation Versus Standard Error

There is sometimes confusion about the terms standard deviation and standard error. For clear definitions of these terms and how they apply to the NLSCY, see Chapter 13.0 on Variance Estimation.

16.13 Understanding the Difference Between "Not stated", "Don't know", "Refusal" and "Valid skip"

All questions in the NLSCY do not apply to all children. When working with NLSCY data, a question that was not intended for a particular child will have the response "Valid skip". For a question that was intended for an individual and no answer was provided, "Not stated", "Don't know", or "Refusal" appears in the data file. When analysing particular populations, the analyst should take care to ensure that the questions of interest are applicable. In the case of examining nonresponse, the "Valid skip" cases should **not** be treated as nonrespondents - it is not that the questions were not answered; it is that they do not apply. Occasionally, "Valid skip" can take a specific meaning such as "0" or "Not in school". The analyst should review the questionnaire to know the details.

17.0 Guidelines for Tabulation, Analysis and Release

This chapter of the documentation outlines the guidelines to be adhered to by users tabulating, analyzing, publishing or otherwise releasing any data derived from the survey microdata files. With the aid of these guidelines, users of microdata should be able to produce the same figures as those produced by Statistics Canada and, at the same time, will be able to develop currently unpublished figures in a manner consistent with these established guidelines.

17.1 Rounding Guidelines

In order that estimates for publication or other release derived from the National Longitudinal Survey of Children and Youth (NLSCY) microdata files correspond to those produced by Statistics Canada, users are urged to adhere to the following guidelines regarding the rounding of such estimates:

- a) Estimates in the main body of a statistical table are to be rounded to the nearest hundred units using the normal rounding technique. In normal rounding, if the first or only digit to be dropped is 0 to 4, the last digit to be retained is not changed. If the first or only digit to be dropped is 5 to 9, the last digit to be retained is raised by one. For example, in normal rounding to the nearest 100, if the last two digits are between 00 and 49, they are changed to 00 and the preceding digit (the hundreds digit) is left unchanged. If the last digits are between 50 and 99 they are changed to 00 and the preceding digit is incremented by 1. For example, an estimated total of 21,352 would be rounded to 21,400.
- b) Marginal sub-totals and totals in statistical tables are to be derived from corresponding unrounded components and then are to be rounded themselves to the nearest 100 units using normal rounding.
- c) Averages, proportions, rates and percentages are to be computed from unrounded components, i.e., numerators and/or denominators, and then are to be rounded themselves to one decimal using normal rounding.
- d) Sums and differences of aggregates (or ratios) are to be derived from their corresponding unrounded components and then are to be rounded themselves to the nearest 100 units (or the nearest one decimal) using normal rounding.
- e) In instances where, due to technical or other limitations, a rounding technique other than normal rounding is used resulting in estimates to be published or otherwise released which differ from corresponding estimates published by Statistics Canada, users are urged to note the reason for such differences in the publication or release document(s).
- f) Under no circumstances are unrounded estimates to be published or otherwise released by users. Unrounded estimates imply greater precision than actually exists.

17.2 Sample Weighting Guidelines for Tabulation

In survey estimation, each sample unit represents not only itself, but several other units in the survey population. For the NLSCY, the survey weight assigned to each child reflects the number of children represented by a particular respondent child.

When producing simple estimates, including the production of ordinary statistical tables, users should apply the proper survey weight. If proper weights are not used, the estimates derived from the microdata file cannot be considered to be representative of the survey population, and will not correspond to those produced by Statistics Canada.

17.3 Guidelines for Statistical Modelling

Sample Design

As mentioned earlier, the NLSCY is based upon a complex sample design, with stratification, multiple stages of selection, and unequal probabilities of selection of respondents. Using data from such complex surveys presents challenges to analysts because the survey design and the selection probabilities affect the estimation and variance calculation procedures that should be used. In order for survey estimates and analyses to be free from bias, the appropriate survey weights should be used whenever possible. For more details about the survey weights, refer to Chapter 11.0.

Variance Estimates

While many analysis procedures found in statistical packages allow weights to be used, the meaning or definition of the weight in these procedures may differ from that which is appropriate in a sample survey framework, with the result that while in many cases the estimates produced by the packages are correct, the variance estimates that are calculated are poor. Users should estimate design-consistent variances using the Bootstrap weights and tools described in Chapter 13.0.

17.4 Release Guidelines

Data users must not release or publish any estimate that would allow the identification of a specific respondent or reveal any individual's responses. For this reason, estimates (for example, the cells in a cross-tabulation) should have at least five contributing respondents.

Apart from the above requirement, all estimates can be considered releasable, but before releasing and/or publishing any estimate from the NLSCY, users should first determine the quality of the estimate. This means that the standard error associated with the estimate must be calculated (see Chapter 13.0). Users should also consider how non-sampling errors discussed in Chapter 11.0 may affect the estimate.

Also, the number of children who contribute to the calculation of the estimate should be determined. If this number is small, the standard error associated with the weighted estimate is expected to be large, and the estimate is likely unreliable for most purposes. When considering proportions, one can certainly infer from the survey that a certain characteristic is rare, but the true rate cannot usually be determined from the survey data. For example, one can use the NLSCY to estimate that 1 out of 1,500 children have some specific health problem, but the true rate may be twice this estimate or half this estimate. In such instances, releasing a less exact estimate, i.e., the rate is estimated to be less than 0.5%, would be preferable as this is the level of precision that the survey can legitimately claim.

Appendix I

Guidelines for Researchers and Analysts Using the National Longitudinal Survey of Children and Youth

Some analysts and researchers using the National Longitudinal Survey of Children and Youth (NLSCY) database have expressed a need for guidelines to help them plan their analyses and report their findings. The purpose of this document is to respond to those requests.

This document is made up of two main sections. The first section concerns the research proposal. It gives the reader recommendations on different methodological aspects to consider when submitting a research proposal using the NLSCY as a source of data. The second section concerns research papers and reports. It deals with recommendations on what to consider when writing a paper using the NLSCY data. Many elements included in the section on preparing a research proposal are also found in the section on writing a paper. These two components can be used together, or as separate documents.

Before you submit a research proposal for review:

Methodological considerations

Before undertaking any analysis using the NLSCY data, researchers and analysts should first familiarize themselves with the complexity of the NLSCY and the resulting implications for analysis. The purpose of this document is to facilitate their work by clearly identifying the key methodological issues to be considered when using NLSCY data.

This document identifies several important methodological elements to be considered when submitting a research proposal. A companion document specifies the elements to consider when submitting a paper for review. Authors are encouraged to use these documents to ensure that they have addressed the relevant elements before submitting their research proposal or their paper.

The NLSCY data can be used in many ways. The main objective of the NLSCY is to allow inferences to be made about a population, using a probabilistic sample. This document has been written with this objective in mind. When NLSCY data are used with objectives other than making inferences about the population (for example a case study), some of the elements described in this document might not apply. However, for such cases, caveats provided by the author will help to put the analytical framework into perspective for the reader.

Elements of the Analytical Framework

There are six main elements to be considered in preparing a research proposal or paper using the NLSCY. These include:

- 1) Data sources,
- 2) Factors affecting the analysis,
- 3) Variables,
- 4) Type of analysis,
- 5) Variance estimation, and
- 6) Methods of analysis.

1) Data sources

All sources of data to be used in the analysis can be specified as follows:

- a Specify the main source of data to be used in the analysis.
 - NLSCY
 - Other (specify)

- b Indicate what other sources of data, if any, will be used in the analysis and whether these data will be included as raw data or in tabular form.
- c If the analysis is to be limited to a subgroup or domain, provide a description of the subgroup or domain, e.g., age groups, provinces and variables with certain characteristics.
- d Specify the cycle or cycles of the NLSCY to be used.

2) Factors affecting the analysis

The research proposal should include a description of factors that may restrict or affect the analysis:

- a Description of the target population
 - Provide a clear definition of the target population of the NLSCY.
 - If the target population differs from the NLSCY definition, include a statement about the potential impact on your analysis.
 - If comparative sources will be used, include a statement about how their target populations differ from the NLSCY population.
- b Treatment of nonresponse
 - If some variables used in the analysis have nonresponse, include a statement about the level of nonresponse, if known, and its potential impact on the analysis.
 - Specify how partial nonresponse will be handled:
 - Imputation
 - Re-weighting
 - Reported as a value
 - Ignored, analysis to be done with the respondents only.
 - Analysis of characteristics of nonrespondents versus respondents has to be done to identify possible biases.
- c Data limitations
 - Provide the sample sizes, overall and for all sub-domains, where this information is known. Sample sizes will be needed that are sufficiently large both to respect confidentiality and to give reliable estimates.
 - Indicate if any other limitations are foreseen with the use of the NLSCY in your project.

3) Variables

- Provide a preliminary list of the variables in the NLSCY file to be used in the analysis.
- Indicate both predictor and outcome variables to be considered, to the extent that this is known.

Note that extensive information about variables can be learned before accessing the master files by studying questions in the questionnaires (on the Statistics Canada web site) or examining variable lists in the data files (via the Data Liberation Initiative at university libraries).

4) Type of analysis

- a Indicate the kind of study planned, whether longitudinal, cross-sectional, or both. Note that if both kinds are included in the analysis, the target population may differ from one type to the other.

- b Specify the kind of survey weights to be used, whether longitudinal, cross-sectional or both. Note that if estimates of both cross-sectional and longitudinal populations are to be analyzed, make sure to use the appropriate weights for each analysis.
- c If survey weights were not to be used, include an explanation of why not. Note that it is unlikely that the use of survey weights is irrelevant to the analysis.

5) Variance estimation

Various methods are available for estimating precision when making inferences, including the measurement of the variances and/or coefficients of variation (CV). The research proposal should include some indication of the approach to be used, if possible. Options include:

- Approximations using the CV look-up tables (available for the first 4 cycles)
- Use of the NLSCY Excel Interface with CVs for many domains of interest
- Use of the bootstrap weights with the Bootvar program, SUDAAN, or some other program that incorporates the bootstrap weights
- No estimation of variance or coefficient of variation. Note that this would imply that no statistical inferences are being made.
- Use of other software (specify software: _____). Note that very few software programs are capable of handling the complex survey design of the NLSCY when estimating the variance.

6) Methods of analysis

- a Present a description of planned analytical methods.
- b Describe the statistical techniques to be used to determine whether the estimates are statistically significant.
- c Plan to include confidence intervals based on appropriate variance calculation in the analysis.

Before you submit a paper for review:

Methodological considerations

Before undertaking any analysis using the NLSCY data, researchers and analysts should first familiarize themselves with the complexity of the NLSCY and the resulting implications for analysis. The purpose of this document is to facilitate their work by clearly identifying the key methodological issues to be considered when using NLSCY data.

This document identifies several important methodological elements to be considered when submitting a paper for review. A companion document specifies the elements to consider when submitting a research proposal. Authors are encouraged to use these documents to ensure that they have addressed the relevant elements before submitting their research proposal or their paper.

The NLSCY data can be used in many ways. The main objective of the NLSCY is to allow inferences to be made about a population, using a probabilistic sample. This document has been written with this objective in mind. When NLSCY data are used with objectives other than making inferences about the population (for example a case study), some of the elements described in this document might not apply. However, for such cases, caveats provided by the author will help to put the analytical framework into perspective for the reader.

Elements of the Analytical Framework

There are six main elements to be considered in preparing a research proposal or paper using the NLSCY. These include:

- 1) Data sources,
- 2) Factors affecting the analysis,
- 3) Variables,
- 4) Type of analysis,
- 5) Variance estimation, and
- 6) Methods of analysis.

1) Data sources

All sources of data to be used in the analysis can be specified as follows:

- a Specify the main source of data to be used in the analysis.
 - NLSCY
 - Other (specify)
- b Indicate what other sources of data, if any, were used in the analysis and whether these data were included as raw data or in tabular form.
- c If the analysis was limited to a subgroup or domain, provide a description of the subgroup or domain, e.g., age groups, provinces and variables with certain characteristics.
- d Specify the cycle or cycles of the NLSCY were used.

2) Factors affecting the analysis

The paper should include a description of factors that restricted or affected the analysis:

- a Description of the target population
 - Provide a clear definition of the target population of the NLSCY.
 - If the target population differed from the NLSCY definition, include a statement about the potential impact on the analysis.
 - If comparative sources were used, include a statement about how their target populations differed from the NLSCY population.
- b Treatment of nonresponse (if any)
 - If some variables used in the analysis have nonresponse, include a statement about the level of nonresponse and the impact on the analysis.
 - Specify how partial nonresponse was handled:
 - Imputation
 - Re-weighting
 - Reported as a value
 - Ignored, analysis done with the respondents only.
 - Analysis of nonrespondents versus respondents should be done to identify possible biases.
- c Data limitations
 - Provide the sample sizes, overall and for all sub-domains.
 - Ensure that the sample sizes used in the report are sufficient both to respect confidentiality and to give reliable estimates.
 - Indicate if there are any other limitations with the use of the NLSCY in your project, e.g., with the variables used.

3) Variables

- Describe the variables in the NLSCY file that were used in the analysis.

4) Type of analysis

- a. Indicate the kind of study, whether longitudinal, cross-sectional, or both. Note that if both kinds were included in the analysis, the target population may differ from one type to the other.
- b. Specify the kind of survey weights used, whether longitudinal, cross-sectional or both. If estimates for both cross-sectional and longitudinal populations were reported, ensure that the appropriate weights were used for each analysis.
- c. If survey weights were not used, include an explanation of why not. It is unlikely that the use of survey weights is irrelevant to the analysis.

5) Variance estimation

Describe the method of estimating precision when making inferences, including the measurement of the variances and/or coefficients of variation (CV) used:

- Approximations using the CV look-up tables (available for the Cycles 1 to 4)
- Use of the NLSCY Excel Interface with CVs for many domains of interest
- Use of the bootstrap weights with the Bootvar program, SUDAAN, or some other program that incorporates the bootstrap weights
- No estimation of variance or coefficient of variation was done. Note that this would imply that no statistical inferences can be made in the paper.
- Use of other software (specify software: _____). Note that very few software programs are capable of handling the complex survey design of the NLSCY when estimating the variance.

6) Methods of analysis

- a. Present a description of all analytical methods used.
- b. Describe the statistical techniques used to determine whether the estimates were statistically significant.
- c. Include confidence intervals based on appropriate variance calculation.

Summary

A reviewer of your paper, who has access to the same data as you do, should be able to reproduce perfectly your results and reach the same conclusions, given the methodology you have used is sound and well explained in your paper.

Appendix II

Partial Nonresponse Analysis

The following Partial Nonresponse Analysis, which is based on the Cycle 5 findings of the **National Longitudinal Survey of Children and Youth**, was conducted to accompany the release entitled “National Longitudinal Survey of Children and Youth: Home environment, income and child behaviour” which appeared in *The Daily* on February 21, 2005.

Like all surveys, the NLSCY must deal with nonresponse. There are two main categories of nonresponse: total nonresponse and partial nonresponse. Total nonresponse is the complete absence of data, or too little data to be considered a response, for a sampled unit. Design weights provided with the data files have been adjusted to take into account the total nonresponse.

Partial nonresponse is the absence of information for certain questions only, with the person selected having nonetheless adequately answered enough questions to be considered a respondent. The purpose of this document is to assess partial nonresponse for the variables in the NLSCY that were used in the report noted above.

The attached table presents the percentage of respondents aged 2 to 5 in Cycle 1 in 1994/1995 and 10 to 13 in Cycle 5 in 2002/2003 for whom data are available for each of the predictor and outcome variables under study in the report.

In 1994/1995, all responses were supplied by the reporting parents. Overall, response rates for these variables were very high, ranging from 96% to 98% for the predictor variables (family functioning, maternal depression, punitive parenting), and from 87% to 94% for the outcome variables (child aggressive behaviour, child anxiety, child prosocial behaviour).

In 2002/2003, responses for two predictor variables were supplied by the parents, and overall response rates were again very high, 96% for family functioning and 95% for maternal depression. The remaining responses were supplied by the children themselves, using self-complete questionnaires. Response rates were lower than for parent-reported information, ranging from 74% to 78% for predictor variables (punitive parenting, nurturing parenting, parental monitoring) and from 76% to 81% for the behaviour outcome variables (aggressive behaviour, anxiety, prosocial behaviour, self-esteem).

In an effort to identify possible sources of nonresponse bias in the data, response rates were compared for females and males, for low-income and higher-income households, and for the five regions of Canada. The results of these analyses follow:

- No sex differences in response rates were found for any of the variables.
- One difference in response rates was found between low-income and higher-income respondents. The response rate was lower for low-income than higher-income respondents for the nurturing parenting variable (68% compared with 75%). Though not large, this difference was statistically significant ($p < 0.05$).
- Regional differences in response rate were found for 1994/1995 child anxiety, 1994/1995 punitive parenting, and 2002/2003 maternal depression. Respondents in the Prairie Region had a significantly lower response rate for 1994/1995 child anxiety than those in the Atlantic Region (92% compared with 97%). No other differences were statistically significant for this variable. Respondents in the Prairie Region had a significantly lower response rate for 1994/1995 punitive parenting practices than those in the Atlantic Region and those in Quebec (94% compared with 99% and 98% respectively). No other differences were statistically significant for this variable. Finally, respondents in Ontario had a significantly lower response

rate for 2002/2003 maternal depression than those in the Atlantic Region (93% compared with 98%). No other differences were statistically significant for this variable.

No specific adjustments were made in the analyses for these variations in nonresponse rate, and findings should be interpreted with these limitations in mind.

Response rates for scales 1994/1995 (aged 2 to 5) and 2002/2003 (aged 10 to 13) by sex, income status, and region of residence

Outcome measure	Overall response rate	Sex		Income status		Region of residence in 2002				
		Female	Male	Low	Higher	Atlantic	Québec	Ontario	Prairie	British Columbia
Aggressive behaviour score 1994/1995	92%	93%	92%	91%	93%	95%	93%	92%	92%	92%
Aggressive behaviour score 2002/2003	77%	78%	76%	74%	78%	78%	78%	77%	75%	77%
Anxiety score 1994/1995 ¹	94%	94%	94%	93%	94%	97%	95%	93%	92%	94%
Anxiety score 2002/2003	76%	77%	76%	75%	77%	77%	78%	77%	74%	76%
Prosocial behaviour score 1994/1995	87%	86%	88%	86%	88%	91%	88%	87%	88%	86%
Prosocial behaviour score 2002/2003	76%	76%	76%	75%	76%	76%	78%	77%	73%	75%
Self esteem score 2002/2003	81%	81%	82%	79%	82%	83%	82%	82%	78%	84%
Family functioning score 1994/1995	98%	98%	98%	97%	98%	99%	99%	98%	97%	98%
Family functioning score 2002/2003	96%	96%	96%	94%	96%	98%	96%	95%	97%	96%
Maternal depression score 1994/1995	98%	98%	98%	99%	98%	98%	97%	99%	97%	98%
Maternal depression score 2002/2003 ²	95%	95%	95%	92%	96%	98%	96%	93%	97%	97%
Punitive parenting score 1994/1995 ³	96%	96%	96%	95%	97%	99%	98%	96%	94%	95%
Punitive parenting score 2002/2003	78%	79%	78%	74%	79%	81%	77%	80%	76%	80%
Nurturing parenting score 2002/2003 ⁴	74%	74%	73%	68%	75%	74%	76%	73%	71%	76%
Parental monitoring score 2002/2003	76%	77%	75%	70%	77%	77%	77%	76%	72%	77%

Notes

Bold print indicates statistically significant differences at $p < .05$

1. Overall response rate was lower for the Prairie Region than for the Atlantic Region
2. Overall response rate was lower for Ontario than for the Atlantic Region
3. Overall response rate was lower for the Prairie Region than for the Atlantic Region or for Quebec
4. Overall response rate was lower for low-income than for higher-income respondents

Appendix III

Concordance of Processing Variable Names to Dissemination Variable Names for the Self-complete Questionnaires, National Longitudinal Survey of Children and Youth, Cycle 6

Variable	Size	Type	Book20	Book21	Book22	Book23
PERSRUK	14	C	PERSRUK	PERSRUK	PERSRUK	PERSRUK
FMMCQ01	3	N	EFFAGEYR	EFFAGEYR	EFFAGEYR	EFFAGEYR
FMMCbQ1A	3	N	AGE	AGE	AGE	AGE
FMMCQ02	1	C	GENDER	GENDER	GENDER	GENDER
FFFCQ01	1	N	A1	A1	A1	A1
FFFCQ02	1	N	A2	A2	A2	A2
FFFCQ03	1	N	A3	A3	A3	A3
FFFCQ04	1	N	A4	A4	A4	A4
FFFCcQ4A	1	N			A5	A5
FFFCQ05	2	N	A5	A5	A6	A6
FFFCQ06	2	N	A6_2			
FFFCbQ13	2	N		A6_2	A7_2	A7_2
FFFCbQ14	2	N		A7_2	A8_2	A8_2
FFFCbQ15	1	N		A8	A9	A9
FFFCQ07	1	N	A8	A10	A13	A13
FFFCQ08A	1	N	A9A	A11A	A14A	A14A
FFFCQ08B	1	N	A9B	A11B	A14B	A14B
FFFCQ08C	1	N	A9C	A11C	A14C	A14C
FFFCQ08D	1	N	A9D	A11D	A14D	A14D
FFFCQ08E	1	N	A9E	A11E	A14E	A14E
FFFCQ08F	1	N	A9F	A11F	A14F	A14F
FFFCQ08G	1	N	A9G	A11G	A14G	A14G
FFFCQ08H	1	N	A9H	A11H	A14H	A14H
FFFCc08I	1	N	A9I	A11I	A14I	A14I
FFFCQ08J	1	N	A9J	A11J		
FFFCQ08K	1	N	A9K	A11K	A14J	A14J
FFFCQ08L	1	N	A9L	A11L	A14K	A14K
FFFCQ08M	1	N	A9M	A11M	A14L	A14L
FFFCQ08N	1	N	A9N	A11N	A14M	A14M
FFFCQ09	1	N	A10	A12		
FFFCQ12	2	N	A11	A13		
FFFCd12A	1	N			A15	A15
FDRCdQ05	1	N	A7A	A9A	A10A	A10A
FDRCdQ09	1	N	A7B	A9B	A10B	A10B
FFFCd16C	1	N		A9C	A10C	A10C
FFFCc16D	1	N	A7C	A9D	A10D	A10D
FFFCc16E	1	N	A7D	A9E	A10E	A10E
FFFCc19A	1	N			A11A	A11A_1
FFFCc19B	1	N			A11B	A11B_2
FFFCc19C	1	N			A11C	A11C_3
FFFCc19D	1	N			A11D	A11D_4
FFFCc20A	1	N			A12A	A12A
FFFCc20B	1	N			A12B	A12B

FFFCc20C	1	N			A12C	A12C
FSCCQ01	1	N	B1	B1	B1	
FSCCbQ18	1	N		B2	B2	
FSCCb19A	1	N		B3A	B3A	
FSCCb19B	1	N		B3B	B3B	
FSCCb19C	1	N		B3C	B3C	
FSCCb19D	1	N		B3D	B3D	
FSCCc19F	1	N		B3E	B3E	
FSCCb19E	1	N		B3F	B3F	
FSCCb20A	1	N		B4A	B4A	
FSCCb20B	1	N		B4B	B4B	
FSCCb20C	1	N		B4C	B4C	
FSCCb20E	1	N		B4D	B4D	
FSCCc20H	1	N		B4E	B4E	
FSCCc20I	1	N		B4F	B4F	
FSCCc20J	1	N		B4G	B4G	
FSCCb20G	1	N		B4H	B4H	
FSCCQ02	1	N	B2	B5	B5	
FSCCcQ3a	1	N	B3A	B6A	B6A	
FSCCcQ03	1	N	B3B	B6B	B6B	
FSCCcQ3b	1	N	B3C	B6C	B6C	
FSCCcQ3D	1	N	B3D	B6D	B6D	
FSCCcQ3C	1	N		B6E	B6E	
FSCCcQ3E	1	N		B6F	B6F	
FSCCcQ3F	1	N		B6G	B6G	
FSCCd3G	1	N			B6H	
FSCCc21A	1	N	B4A	B7A	B7A	
FSCCc21B	1	N		B7D	B7D	
FSCCc21C	1	N	B4B	B7B	B7B	
FSCCc21D	1	N	B4C	B7C	B7C	
FSCCc21E	1	N		B7E	B7E	
FSCCc21F	1	N		B7F	B7F	
FSCCcQ26	1	N			B8	
FSCCcQ27	1	N			B9	
FSCCcQ10	1	N	B5F	B8	B10	B7
FSCCbQ22	1	N		B9A	B12A	
FSCCc22A	1	N		B9B	B12B	
FSCCcQ28	1	N			B13	
FSCCcQ29	1	N			B14	
FSCCQ12	1	N	B6A	B10A	B15A	
FSCCQ11	2	N	B6B	B10B	B15B	
FSCCcQ16	2	N	B6C	B10C	B15C	
FSCCQ17	2	N	B6D	B10D	B15D	
FSCCcQ30	2	N		B11	B16	
FSCCQ13	2	N	B7A	B12A	B17A	
FSCCQ14	1	N	B7B	B12B	B17B	
FSCCQ15	1	N	B7C	B12C	B17C	
FSCCcQ24	2	N		B13	B18	
FSCCc31A	1	N			B11A	
FSCCc31B	1	N			B11B	

FSCCc31C	1	N				B11C
FSCCc31D	1	N				B11D
FSCCc31E	1	N				B11E
FSCCQ05	1	N	B5A			
FSCCQ06	1	N	B5B			
FSCCcQ07	1	N	B5C			
FSCCQ08	1	N	B5D			
FSCCQ09	1	N	B5E			
FAMCQ01A	1	N	C1A	C1A	C1A	B1A
FAMCQ01B	1	N	C1B	C1B	C1B	B1B
FAMCQ01C	1	N	C1C	C1C	C1C	B1C
FAMCQ01D	1	N	C1D	C1D	C1D	B1D
FAMCcQ1E	1	N	C1E	C1E	C1E	B1E
FAMCe25A	1	N	C2A	C2A	C2A	B2A
FAMCe25B	1	N	C2B	C2B	C2B	B2B
FAMCe25C	1	N	C2C	C2C	C2C	B2C
FAMCe25D	1	N	C2D	C2D	C2D	B2D
FAMCe25E	1	N	C2E	C2E	C2E	B2E
FAMCe25F	1	N	C2F	C2F	C2F	B2F
FAMCe25G	1	N	C2G	C2G	C2G	B2G
FAMCe25H	1	N	C2H	C2H	C2H	B2H
FAMCe25I	1	N	C2I	C2I	C2I	B2I
FAMCe25J	1	N	C2J	C2J	C2J	B2J
FAMCe25K	1	N	C2K	C2K	C2K	B2K
FAMCe25L	1	N	C2L	C2L	C2L	B2L
FAMCe25M	1	N	C2M	C2M	C2M	B2M
FAMCe25N	1	N	C2N	C2N	C2N	B2N
FAMCe25O	1	N	C2O	C2O	C2O	B2O
FAMCbQ02	1	N		C3	C3	B3
FAMCcQ03	1	N		C4	C4	B4
FAMCcQ6A	1	N		C5A	C7A	
FAMCcQ6B	1	N		C5B	C7B	
FAMCdQ6C	1	N				B8A
FAMCcQ7A	1	N		C6A	C8A	
FAMCcQ7B	1	N		C6B	C8B	
FAMCdQ7C	1	N				B8B
FAMCcQ8A	1	N		C7A	C9A	
FAMCcQ8B	1	N		C7B	C9B	
FAMCdQ8C	1	N				B8C
FAMCdQ4A	1	N			C5A	
FAMCcQ4B	1	N			C5B	
FAMCcQ4C	1	N			C5C	
FAMCcQ4D	1	N			C5D	
FAMCeQ4E	1	N			C5E	
FAMCe26A	1	N				B5A
FAMCe26B	1	N				B5B
FAMCe26C	1	N				B5C
FAMCe26D	1	N				B5D
FAMCe26E	1	N				B5E
FAMCe26F	1	N				B5F

FAMCcQ5A	1	N			C6A	B6A
FAMCcQ5B	1	N			C6B	B6B
FAMCcQ5C	1	N			C6C	B6C
FAMCcQ5D	1	N			C6D	B6D
FAMCdQ09	1	N				B9
FAMCdQ10	1	N				B10
FFBCQ01A	1	N	D1A	D1A	D1A	
FFBCc01B	1	N	D1B	D1B	D1B	
FFBCQ01C	1	N	D1C	D1C	D1C	
FFBCQ01D	1	N	D1D	D1D	D1D	
FFBCQ01E	1	N	D1E	D1E	D1E	
FFBCQ01F	1	N	D1F	D1F	D1F	
FFBCQ01G	1	N	D1G	D1G	D1G	
FFBCQ01H	1	N	D1H	D1H	D1H	
FFBCQ01I	1	N	D1I	D1I	D1I	
FFBCQ01J	1	N	D1J	D1J	D1J	
FFBCQ01K	1	N	D1K	D1K	D1K	
FFBCQ01L	1	N	D1L	D1L	D1L	
FFBCQ01M	1	N	D1M	D1M	D1M	
FFBCQ01P	1	N	D1N	D1N	D1N	
FFBCQ01Q	1	N	D1O	D1O	D1O	
FFBCQ01R	1	N	D1P	D1P	D1P	
FFBCQ01S	1	N	D1Q	D1Q	D1Q	
FFBCQ01T	1	N	D1R	D1R	D1R	
FFBCQ01U	1	N	D1S	D1S	D1S	
FFBCc01V	1	N	D1T	D1T	D1T	
FFBCQ01W	1	N	D1U	D1U	D1U	
FFBCd01X	1	N	D1V	D1V	D1V	
FFBCQ01Z	1	N	D1W	D1W	D1W	
FFBCQ1AA	1	N	D1X	D1X	D1X	
FFBCQ1BB	1	N	D1Y	D1Y	D1Y	
FFBCQ1CC	1	N	D1Z	D1Z	D1Z	
FFBCQ1DD	1	N	D1AA	D1AA	D1AA	
FFBCQ1FF	1	N	D1BB	D1BB	D1BB	
FFBCQ1GG	1	N	D1CC	D1CC	D1CC	
FFBCQ1HH	1	N	D1EE	D1EE	D1EE	
FFBCQ1JJ	1	N	D1DD	D1DD	D1DD	
FFBCQ1LL	1	N	D1FF	D1FF	D1FF	
FFBCQ1MM	1	N	D1GG	D1GG	D1GG	
FFBCQ1NN	1	N	D1HH	D1HH	D1HH	
FFBCQ1OO	1	N	D1II	D1II	D1II	
FFBCQ1PP	1	N	D1JJ	D1JJ	D1JJ	
FFBCQ1QQ	1	N	D1KK	D1KK	D1KK	
FFBCQ1RR	1	N	D1LL	D1LL	D1LL	
FFBCQ1SS	1	N	D1MM	D1MM	D1MM	
FFBCQ1TT	1	N	D1NN	D1NN	D1NN	
FFBCc1UU	1	N	D1OO	D1OO	D1OO	
FFBCd10A	1	N				C1A
FFBCd10B	1	N				C1B
FFBCd10C	1	N				C1C

FFBCd10D	1	N				C1D
FFBCd10E	1	N				C1E
FFBCd10F	1	N				C1F
FFBCd10G	1	N				C1G
FFBCd10H	1	N				C1H
FFBCd10I	1	N				C1I
FFBCd10J	1	N				C1J
FFBCd10K	1	N				C1K
FFBCd10L	1	N				C1L
FFBCcQ04	1	N		D2	D2	C2
FFBCcQ4A	1	N		D3	D3	C3
FFBCcQ05	1	N		D4	D4	C4
FFBCcQ07	1	N		D5	D5	C5
FFBCcQ08	1	N		D6	D6	C6
FFBCbQ2A	1	N	D2A			
FFBCbQ2B	1	N	D2B	D7A	D7A	C7A
FFBCbQ2C	1	N	D2C			
FFBCbQ2D	1	N	D2D			
FFBCbQ2E	1	N	D2E	D7B	D7B	C7B
FFBCbQ2F	1	N	D2F	D7C	D7C	C7C
FFBCcQ2H	1	N				C7D
FFBCeQ2O	1	N		D7D	D7D	C7E
FFBCbQ2P	1	N		D7E	D7E	C7F
FFBCdQ2Z	1	N				C7G
FFBCe2ZZ	1	N		D7F	D7F	C7H
FFBCbQ2Y	1	N		D7G	D7G	C7I
FFBCe2BB	1	N		D7H	D7H	C7J
FFBCcQ3A	1	N		D8	D8	C8
FFBCQ03	1	N	D3			
FATCbQ1A	1	N	H1A	E1A		
FATCc1AA	1	N			E1A	
FATCbQ1B	1	N	H1B	E1B		
FATCc1BB	1	N			E1B	
FATCbQ1C	1	N	H1C	E1C		
FATCc1CC	1	N			E1C	
FATCbQ1D	1	N	H1D	E1D		
FATCc1DD	1	N			E1D	
FATCbQ1E	1	N	H1E	E1E		
FATCc1EE	1	N			E1E	
FATCbQ1I	1	N	H1F	E1F		
FATCc1II	1	N			E1F	
FATCbQ1F	1	N	H1G			
FATCeQ20	2	N	H2	E2	E2	
FATCbQ04	1	N		E3	E3	
FATCdQ07	2	N	H3	E7	E7	
FATCQ03	2	N	H4	E4		
FATCeQ12	1	N	H5A	E8A	E8A	
FATCe12B	1	N	H5B	E8B	E8B	
FATCe12C	1	N	H5C	E8C	E8C	
FATCeQ13	1	N	H6A	E9A	E9A	

FATCe13B	1	N	H6B	E9B	E9B	
FATCe13C	1	N	H6C	E9C	E9C	
FATCeQ21	2	N	H7	E10	E10	
FATCeQ22	1	N	H8	E11	E11	
FATCdQ10	2	N		E12	E12	
FATCbQ11	2	N		E13		
FATCdQ5A	1	N		E5A	E5A	
FATCbQ5B	1	N		E5B	E5B	
FATCbQ5C	1	N		E5C	E5C	
FATCbQ5D	1	N		E5D	E5D	
FATCbQ5E	1	N		E5E	E5E	
FATCbQ5F	1	N		E5F	E5F	
FATCbQ5G	1	N		E5G	E5G	
FATCdQ05	1	N		E6	E6	
FATCc14A	1	N			E4A	
FATCc14B	1	N			E4B	
FATCc14C	1	N			E4C	
FATCc14D	1	N			E4D	
FDRCdQ01	2	N	G1	F1	F1	D1
FDRCQ04	2	N		F2	F2	D2
FDRCQ03	2	N	G2_2	F3_2	F3_2	
FDRCcQ06	1	N	G3			
FDRCdQ6A	2	N		F4	F4	D3
FDRCdQ07	2	N	G4	F5	F5	
FDRCdQ9A	1	N		F6	F6	
FDRCbQ9B	2	N		F7	F7	
FDRCdQ9C	2	N		F8	F8	D4
FDRCdQ15	2	N		F9	F9	D5
FDRCd14A	2	N		F11A_3	F11A_3	
FDRCc14C	2	N		F11B_3	F11B_3	
FDRCc14B	2	N		F11C_3	F11C_3	
FDRCd14F	2	N		F11D_3	F11D_3	
FDRCc14D	2	N		F11E_3	F11E_3	
FDRCQ10	1	N	G5			
FDRCQ12	2	N	G6			
FDRCc18A	2	N		F10A	F10A	D6A
FDRCd18B	2	N		F10B	F10B	D6B
FDRCd18C	2	N		F10C	F10C	D6C
FDRCc18D	2	N		F10D	F10D	D6D
FDRCdQ19	1	N				D7
FDRCdQ20	1	N				D8
FPMCdQ04	2	N	E1	G1	G1	G1
FPMCcQ5A	1	N	E2A	G2A	G2A	G2A
FPMCcQ5B	1	N	E2B	G2B	G2B	G2B
FPMCcQ5C	1	N	E2C	G2C	G2C	G2C
FPMCcQ06	1	N	E3	G3	G3	G3
FPMCdQ6A	2	N				G4A
FPMCdQ6B	2	N				G4B
FPMCdQ6C	1	N				G5A
FPMCdQ6D	1	N				G5B

FPMCdQ6E	1	N				G5C
FPMCdQ6F	1	N				G5D
FPMCdQ6G	1	N				G5E
FPMCdQ6H	1	N				G5F
FPMCdQ6I	1	N				G5G
FPMCdQ6J	1	N				G5H
FPMCdQ6K	1	N				G5I
FPMCdQ6L	1	N				G5J
FPMCdQ07	2	N	E4	G4	G4	G6
FPMCcQ8A	1	N	E5A	G5A	G5A	G7A
FPMCcQ8B	1	N	E5B	G5B	G5B	G7B
FPMCcQ8C	1	N	E5C	G5C	G5C	G7C
FPMCcQ09	1	N	E6	G6	G6	G8
FPMCdQ9A	2	N				G9A
FPMCdQ9B	2	N				G9B
FPMCdQ9C	1	N				G10A
FPMCdQ9D	1	N				G10B
FPMCdQ9E	1	N				G10C
FPMCdQ9F	1	N				G10D
FPMCdQ9G	1	N				G10E
FPMCdQ9H	1	N				G10F
FPMCdQ9I	1	N				G10G
FPMCdQ9J	1	N				G10H
FPMCdQ9K	1	N				G10I
FPMCdQ9L	1	N				G10J
FPMCcQ10	1	N	E7	G7	G7	G12
FPMCd11A	2	N	E8	G8	G8	G13
FPMCc11B	2	N	E9	G9	G9	G14
FPMCcQ1A	1	N	E10A	G10A	G10A	
FPMCcQ1B	1	N	E10B	G10B	G10B	
FPMCcQ1C	1	N	E10C	G10C	G10C	
FPMCcQ1D	1	N	E10D	G10D	G10D	
FPMCcQ1E	1	N	E10E	G10E	G10E	
FPMCcQ1F	1	N	E10F	G10F	G10F	G11A
FPMCcQ1G	1	N	E10G	G10G	G10G	
FPMCcQ1H	1	N	E10H	G10H	G10H	
FPMCcQ1I	1	N	E10I	G10I	G10I	
FPMCcQ1J	1	N	E10J	G10J	G10J	
FPMCcQ1R	1	N	E10K	G10K	G10K	
FPMCcQ1K	1	N	E10L	G10L	G10L	
FPMCcQ1L	1	N	E10M	G10M	G10M	
FPMCcQ1M	1	N	E10N	G10N	G10N	
FPMCcQ1N	1	N	E10O	G10O	G10O	
FPMCcQ1O	1	N	E10P	G10P	G10P	
FPMCcQ1P	1	N	E10Q	G10Q	G10Q	
FPMCcQ1Q	1	N	E10R	G10R	G10R	
FPMCdQ1S	1	N	E10S	G10S	G10S	
FPMCdQ1T	1	N	E10T	G10T	G10T	G11B
FPMCdQ9U	1	N				G11C
FPMCdQ9V	1	N				G11D

FPMCbQ3A	1	N	E11A	G11A		
FPMCbQ3C	1	N	E11B	G11B		
FPMCbQ3E	1	N	E11C	G11C		
FPMCd12A	1	N				G15
FPMCd12B	1	N				G16
FPMCd12C	1	N				G17A
FPMCd12D	1	N				G17B
FPMCd12E	1	N				G17C
FPMCd12F	1	N				G17D
FPMCd12G	1	N				G17E
FPMCd12H	1	N				G17F
FPMCd12I	1	N				G17G
FPMCd12J	1	N				G17H
FPMCd12K	1	N				G17I
FPMCd12L	1	N				G17J
FHTCcQ03	1	N		H1	H1	
FHTCbQ3A	1	N		H4A	H4A	
FHTCbQ3B	1	N		H4B	H4B	
FHTCbQ3C	1	N		H4C	H4C	
FHTCbQ3G	1	N		H4D	H4D	
FHTCbQ04	1	N		H5		
FHTCbQ05	1	N		H6		
FHTCdQ06	1	N		H7	H5	
FHTCeQ5A	1	N		H8	H6	E8
FHTCeQ5B	1	N				E8_1A
FHTCeQ5C	1	N				E8_1B
FHTCeQ5D	1	N				E8_1C
FHTCeQ5E	1	N				E8_1D
FHTCdQ5F	1	N				E8_1E
FHTCdQ5G	1	N				E8_2A
FHTCdQ5H	1	N				E8_2B
FHTCdQ5I	1	N				E8_2C
FHTCdQ5J	1	N				E8_2D
FHTCeQ5K	1	N				E8_3A
FHTCeQ5L	1	N				E8_3B
FHTCeQ5M	1	N				E8_3C
FHTCeQ5N	1	N				E8_3D
FHTCeQ5O	1	N				E8_3E
FPUCQ01	1	N	F1	H9	H7	E3
FPUCQ02	1	N	F2	H10	H8	E4
FPUCQ03	1	N	F3	H11		
FPUCdQ04	1	N	F4	H12	H10	E6
FPUCQ05	1	N	F5	H13	H11	E7
FDACcQ1A	2	N		H14_2	H12_2	F1_2
FDACcQ1B	1	N		H15	H13	F2
FDACcQ1C	1	N			H14	
FDACcQ02	2	N		H16	H15	F3
FDACcQ2A	1	N			H16	F4
FDACcQ05	1	N			H17	F5
FDACdQ23	1	N				F8_1

FDACeQ26	1	N				F9
FDACeQ25	2	N				F11
FDACcQ06	2	N			H18	F6
FDACcQ07	2	N			H19_1	F7
FDACcQ08	1	N			H20	
FDACcQ09	1	N			H21	
FDACd22A	1	N				F10A
FDACd22B	1	N				F10B
FDACd22C	1	N				F10C
FDACd22D	1	N				F10D
FDACd22E	1	N				F10E
FDACd22F	1	N				F10F
FDACd22G	1	N				F10G
FDACd22H	1	N				F10H
FDACbQ4C	1	N		H17A		
FDACbQ4D	1	N		H17B		
FDACbQ4E	1	N		H17C		
FDACbQ4F	1	N		H17D		
FWKCbQ02	1	N		I1		
FWKCcQ4A	1	N			I1A	
FWKCcQ4B	1	N			I1B	
FWKCcQ4C	1	N			I1C	
FWKCcQ4D	1	N			I1D	
FWKCe06A	1	N			I2A	
FWKCe06B	1	N			I2B	
FWKCe06C	1	N			I2C	
FWKCe06D	1	N			I2D	
FWKCe06E	1	N			I2E	
FWKCe06F	1	N			I2F	
FWKCcQ9A	2	N			I3A_3	
FWKCcQ9B	2	N			I3B_3	
FWKCeQ10	1	N		I2		
FWKCeQ11	1	N			I4	
FWKCd12A	1	N			I5A	
FWKCd12B	1	N			I5B	
FWKCd12C	1	N			I5C	
FWKCd12D	1	N			I5D	
FWKCd13A	1	N			I6A	
FWKCd13B	1	N			I6B	
FWKCd13C	1	N			I6C	
FWKCd13D	1	N			I6D	
FWKCd13E	1	N			I6E	
FWKCd13F	1	N			I6F	
FWKCd13G	1	N			I6G	
FWKCd13H	1	N			I6H	
FWKCd13I	1	N			I6I	
FPUCcQ3A	3	N			H9	E5
FHTCbQ01	4.2	N		H2	H2	E1
FHTCbQ02	7.3	N		H3	H3	E2
FEQYeS06	2	N	RA	RA	RA	RA

FEQYeS07	2	N	ER	ER	ER	ER
FEQYeS08	2	N	SM	SM	SM	SM
FEQYeS09	2	N	AD	AD	AD	AD
FEQYeS10	2	N	GM	GM	GM	GM
FEQYeS04	2	N	EQ4	EQ4	EQ4	EQ4
FEQYeS05	2	N		EQ5	EQ5	EQ5
FBMleS01	7.3	N		selfbmi	selfbmi	selfbmi
FBMleS02	2	N		selfbmi2	selfbmi2	selfbmi2
FBMleS03	1	N		selfbmi3	selfbmi3	selfbmi3
FFFcS01	2	N	frfs01	frfs01	frfs01	frfs01
FAMcS02	2	N	abms01	abms01	abms01	abms01
FFBcS01	2	N	behs01	behs01	behs01	
FFBCdS02	2	N	behs02	behs02	behs02	
FFBcS03	2	N	behs03	behs03	behs03	
FFBCdS04	2	N	behs04	behs04	behs04	
FFBcS05	2	N	behs05	behs05	behs05	
FFBcS07	2	N	behs07	behs07	behs07	
FPMCcS1	2	N	pars01	pars01	pars01	
FPMCbS2B	2	N	pars02	pars02	pars02	
FPMCcS3	2	N	pars03	pars03	pars03	
FPMCdS4	2	N				pars04
FPMCdS5	2	N				pars05
FHTCbS1b	2	N				deps01
FWTCw01L	12.4	N				
Fwtc wd1l	12.4	N				

Appendix IV

Concordance of Variables Names from Cycle 6 and Variable Names from Booklet 24 (ages 18 to 19, Self-complete Questionnaires) of Cycle 5

Module	C6 Variable Name	C6 Collection Name	C5 Variable Name	C5 Booklet 24 Self-complete
ABM	FAMYfQ01	ABMQ1	EAMCQ01A	A1-A
ABM	FAMYfQ02	ABMQ2	EAMCQ01B	A1-B
ABM	FAMYfQ03	ABMQ3	EAMCQ01C	A1-C
ABM	FAMYfQ04	ABMQ4	EAMCQ01D	A1-D
ABM	FAMYfQ05	ABMQ5	EAMCQ01E	A1-E
ABM	FAMYfQ6A	ABMQ6A	EAMCe26A	A10-A
ABM	FAMYfQ6B	ABMQ6B	EAMCe26B	A10-B
ABM	FAMYfQ6C	ABMQ6C	EAMCe26D	A10-C
ABM	FAMYfQ6E	ABMQ6E	EAMCe26E	A10-D
ABM	FAMYfQ6G	ABMQ6G	EAMCe26F	A10-E
RLY	FRLYfQ01	rlyq1	EDACcQ2A	C4
RLY	FRLYfQ02	rlyq2	EDACcQ1B	C2
RLY	FRLYfQ03	rlyq3	EDACcQ1C	C2
RLY	FRLYfQ07	rlyq7	EDACcQ05	C5
RLY	FRLYfQ08	rlyq8	EDACcQ06	C6
RLY	FRLYfQ09	rlyq9	EDACcQ07	C7
RLY	FRLYfQ10	rlyq10	EDACdQ23	C8
RLY	FRLYfQ11	rlyq11	EAMCeQ21	A6
RLY	FRLYfQ12	rlyq12	EAMCeQ22	A7
RLY	FRLYfQ13	rlyq13	EAMCeQ23	A8
RLY	FRLYfQ14	rlyq14	EAMCeQ24	A9
RLY	FRLYf15A	rlyq15a	EDACd22A	C9
RLY	FRLYf15B	rlyq15b	EDACd22B	C9
RLY	FRLYf15C	rlyq15c	EDACd22C	C9
RLY	FRLYf15F	rlyq15f	EDACd22D	C9
RLY	FRLYf15G	rlyq15g	EDACd22E	C9
RLY	FRLYf15H	rlyq15h	EDACd22F	C9
RLY	FRLYf15I	rlyq15i	EDACd22H	C9
RLY	FRLYfQ16	rlyq16	EDACeQ26	C10
RLY	FRLYf17A	rlyq17a	EDACeQ25	C11
RLY	FRLYf17B	rlyq17b	EDACeQ25	C11
RLY	FRLYf17C	rlyq17c	EDACeQ25	C11
RLY	FRLYf17D	rlyq17d	EDACeQ25	C11
RLY	FRLYf17E	rlyq17e	EDACeQ25	C11
RLY	FRLYf17F	rlyq17f	EDACeQ25	C11
RLY	FRLYf17G	rlyq17g	EDACeQ25	C11
RLY	FRLYf17H	rlyq17h	EDACeQ25	C11
RLY	FRLYf17I	rlyq17i	EDACeQ25	C11
RLY	FRLYf17J	rlyq17j	EDACeQ25	C11
HLTY	FHTYfQ06	hlyq6	EDRCdQ01	E1
HLTY	FHTYfQ07	hlyq7	EDRCQ04	E2
HLTY	FHTYfQ08	hlyq8	EDRCdQ6A	E3
HLTY	FHTYfQ10	hlyq10	EDRCdQ15	E5
HLTY	FHTYfQ11	hlyq11	EDRCdQ15	E5
HLTY	FHTYfQ12	hlyq12	EDRCc18A	E6

HLTY	FHTYfQ13	hltyq13	EDRCc18A	E6
HLTY	FHTYf14A	hltyq14a	EFBCd10A	B1-A
HLTY	FHTYf14B	hltyq14b	EFBCd10B	B1-B
HLTY	FHTYf14C	hltyq14c	EFBCd10C	B1-C
HLTY	FHTYf14D	hltyq14d	EFBCd10D	B1-D
HLTY	FHTYf14E	hltyq14e	EFBCd10E	B1-E
HLTY	FHTYf14F	hltyq14f	EFBCd10F	B1-F
HLTY	FHTYf14G	hltyq14g	EFBCd10G	B1-G
HLTY	FHTYf14H	hltyq14h	EFBCd10H	B1-H
HLTY	FHTYf14J	hltyq14j	EFBCd10I	B1-I
HLTY	FHTYf14K	hltyq14k	EFBCd10J	B1-J
HLTY	FHTYf14M	hltyq14m	EFBCd10K	B1-K
HLTY	FHTYf14N	hltyq14n	EFBCd10L	B1-L
FBH	FFBYfQ01	fbhq1	EFBCcQ4A	B2-5
FBH	FFBYfQ02	fbhq2	EFBCcQ4A	B2
FBH	FFBYfQ03	fbhq3	EFBCcQ05	B3
FBH	FFBYfQ04	fbhq4	EFBCcQ07	B4
FBH	FFBYfQ05	fbhq5	EFBCcQ08	B5
FBH	FFBYfQ06	fbhq6	EFBCcQ2H	B6-B
FBH	FFBYfQ07	fbhq7	EFBCdQ2Z	B6-E
FBH	FFBYfQ08	fbhq8	EDRCdQ19	E7
FBH	FFBYfQ09	fbhq9	EDRCdQ20	E8
FBH	FFBYfQ10	fbhq10	EFBCcQ3A	B7
SUPY	FSPYfQ14	supyq4	EPMCcQ06	D3
SUPY	FSPYfQ15	supyq5	EPMCcQ09	D8