Microdata User Guide

National Longitudinal Survey of Children and Youth

Cycle 8

September 2008 to July 2009

Statistics Canada Statistique Canada

Canada

Table of Contents

1.0 Introduction

The National Longitudinal Survey of Children and Youth (NLSCY), Cycle 8 was conducted from September 2008 to July 2009 by Statistics Canada in partnership with Human Resources and Skills 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 dataset 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.gc.ca

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

Human Resources and Skills Development Canada

Manon Monette
Manager
Knowledge and Data Management Directorate
Human Resources and Skills Development Canada
Telephone: 819-994-9565
E-mail: manon.monette@hrsdc.rhdcc.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 conducted by Statistics Canada and sponsored by Human Resources and Skills 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; their physical development, learning and behaviour; and 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 by policy-making organizations.

Survey population
In Cycle 8, a representative sample of Canadian children aged 0 to 7 years from each of the provinces was surveyed for longitudinal and cross-sectional purposes. The cohort of children and youth aged 14 to 25 years was surveyed for longitudinal purposes.

Target population
The NLSCY's objectives are to produce longitudinal and cross-sectional estimates. Therefore, several populations are targeted in the Cycle 8 sample. Please see Chapter 5.0 for more detailed information about the sample.

- Cross-sectionally, the Cycle 8 sample represents all children who were 0 to 7 years old on December 31, 2008.

- Longitudinally, we have four cohorts, representing 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 any province during collection for Cycle 1 in 1994/1995. These children are now 14 to 25 years old at Cycle 8.

    2) The second cohort represents all children who were 0 to 1 year old as of December 31st, 2002, and who were living in any province during collection for Cycle 5 in 2002/2003. These children are now 6 to 7 years old at Cycle 8.

    3) The third cohort represents all children who were 0 to 1 year old as of December 31st, 2004, and who were living in any province during collection for Cycle 6 in 2004/2005. These children are now 4 to 5 years old at Cycle 8.

    4) The fourth cohort represents all children who were 0 to 1 year old as of December 31st, 2006, and who were living in any province during collection for Cycle 7 in 2006/2007. These children are now 2 to 3 years old at Cycle 8.

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
7	September 2006	July 2007
8	September 2008	July 2009

Data release strategy
The Cycle 8 data are being released in four files: *Longitudinal Cohort – Child* (14 to 17 years), *Longitudinal Cohort – Youth* (16 to 25 years), *Early Child Development Cohort* (0 to 7 years) and the Self-complete File (14 to 17 years).

- *Longitudinal Cohort – Child*: These children from the original longitudinal cohort were 0 to 3 years old in the Cycle 1 of the NLSCY. In Cycle 8, they are 14 to 17 years old.

- *Longitudinal Cohort – Youth*: These children from the original longitudinal cohort were 2 to 11 years old in Cycle 1. In Cycle 8, they are 16 to 25 years old. This file contains data collected for youth and their households, as well as variables, such as Birth weight, that were brought forward from previous cycles and do not change over time. See Chapter 8.0 for more information about the content of the various questionnaires.

- *Early Childhood Development Cohort*: This file contains data collected from the person most knowledgeable (PMK) about the child, for children aged 0 to 7 years in Cycle 8.

- *Self-complete File:* This file contains data collected from the children (aged 14 to 17) by paper questionnaire. See Chapter 8.0 for more information about the content of various questionnaires.

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, such as the Peabody Picture Vocabulary Test – Revised (PPVT–R) and math tests.

4.0 Main changes to Cycle 8 since Cycle 7

This chapter outlines the main changes to the survey since Cycle 7. A more detailed explanation will be found in the individual chapters as referenced below.

4.1 Content changes

At 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 "h" 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 8:

Self-Complete Booklet
The Self-Complete Booklet for 12- to 13-year-olds, Booklet 21, has been retired. The youngest respondents for the longitudinal cohort in Cycle 8 are the 14- to 15-year-olds.

Political Engagement (ages 20 – 25)
In addition to asking youth if they voted in a previous election, youth were asked if they believed they had a duty to vote and how often they talked about political events with others (HPOLhQ1 to HPOLhQ7).

Abilities (ages 16 – 25)
The Abilities questions were added in Cycle 8 but had previously appeared in Cycle 5 on the self-complete paper questionnaires (HABYhQ1 to HABYhQ6).

Family Formation and Fertility (ages 18 – 25)
A new section of questions was added to gather information about the factors influencing young adults' transition to marriage and parenthood (HFFYhQ1 to HFFYhQ16).

4.2 Methodology changes

At Cycle 7, ECD children ranged in age from 0 to 9, while at Cycle 8, the ECD children range in age from 0 to 7.

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); Section 5.3 provides details on how sampling was performed at each cycle, along with response rates; and Section 5.4 describes how the various NSLCY 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 8

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 any province. 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 as a result of 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, such as 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 has been selected at every cycle and followed for at least three cycles (prior to Cycle 7, ECD children were only followed until age 4 to 5, but at Cycle 8, the ECD children range in age from 0 to 7 years old). 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 0- and 1-year-olds in Cycle 3 and some 5-year-olds in Cycles 3 and 4 who were drawn from Birth Registry data.

At Cycle 8, the NLSCY sample consists of children aged 0 to 7 years old (ECD children) and 14 to 25 years (original cohort). A child's effective age at Cycle 8 is as of December 31, 2008. Thus, 0-year-olds are born in 2008, and 1-year-olds are born in 2007.

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

Figure 1
Age of children at each cycle, original cohort versus ECD cohorts

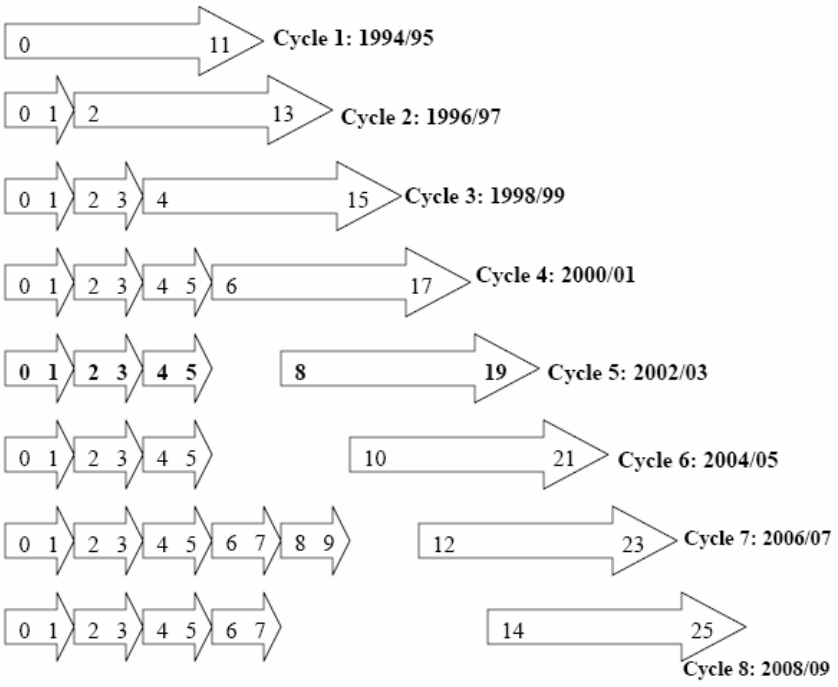

Notes: Ages of children in years are shown in arrows.
Longer arrows represent the original cohort and shorter arrows represent the early childhood development (ECD) cohorts.
For details on how sampling was performed at each cycle, see Section 5.3.
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

5.2 The Labour Force Survey

The LFS is a monthly survey that collects labour market data from a national sample of more than 52,000 dwellings (as of the 2004 redesign). At 10-year intervals, following alternate population censuses, 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.

The sample for the original cohort is a mixture of pre-1994 LFS design and the 1994 redesign. A part of ECD children surveyed at Cycle 8 come from the 1994 LFS design: the ECD cohort sampled in Cycle 5 use the 1994 design; the Cycle 6 cohort is a mixture of the 1994 and 2004 design; the Cycles 7 and 8 cohorts uses the 2004 design. The following sections provide details on the 1994 and 2004 LFS redesigns.

5.2.1 Target population

The LFS sample (1994) is representative of the civilian, non-institutionalized population aged 15 or over in Canada's 10 provinces. Specifically excluded from the survey's coverage are residents of the Yukon, Nunavut and the Northwest Territories, people living on Indian reserves, full-time members of the Canadian

Armed Forces and inmates of institutions. These groups collectively represent an exclusion of approximately 2% of the population aged 15 or over.

5.2.2 Stratification

The LFS sample design is based on a stratified, multistage design using probability sampling at all stages of the design. The design principles of the LFS are the same for each province. The stratification for the 1994 redesign is as follows:

Primary strata
Provinces are divided into economic regions (ERs) and Employment Insurance economic regions (EIERs). ERs are geographic areas with a more or less homogeneous economic structure formed on the basis of federal–provincial agreements. They are relatively stable over time. 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 Skills Development Canada.

The intersections of the two types of regions form the primary strata for the LFS. Then, substratification takes place within these primary strata (Section 5.2.3). Census metropolitan areas (CMAs) constitute a third set of regions; they, too, are stratified in the current LFS design, as each CMA is also an EIER.

Types of areas
The primary strata (intersections of ERs and EIERs) 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 from the largest CMAs to the smallest villages categorized by the 1991 Census as urban (1,000 people or more). 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 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 list is to ensure better representation of apartment dwellers in the sample and to minimize the effect of growth in clusters, which is due to construction of new apartment buildings. In the major cities, the apartment strata are further divided into low-income and regular strata.

Where it is possible and/or necessary, the urban area frame is further divided into regular strata, high-income strata, and strata with low population density. Most urban areas fall into the regular urban strata; in fact, these account for the majority of Canada's population. High-income strata are found in major urban

areas, whereas strata with low-population densities consist of small towns that are geographically scattered.

In rural areas, the population density can vary greatly from relatively high to low, resulting in the formation of strata that reflect these variations. The different stratification strategies for rural areas were based not only on the 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 generally include a set of one or more city blocks, called block faces. The selection of a sample of clusters (always 6 or a multiple of 6 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 6 or 24 clusters are sampled within each town.

For urban apartment strata, rather than defining clusters, the apartment building is used as 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 within remote strata, the number of settlements sampled is 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 list of all private dwellings in the cluster is prepared. From the list, a sample of dwellings is 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 6 or 8 dwellings, depending on the size of the city. In the urban apartment frame, each cluster yields 5 dwellings, whereas in the rural areas and urban EAs, each cluster yields 10 dwellings. In all clusters, dwellings are sampled systematically. This represents the final stage of sampling.

5.2.5 Changes introduced in the 2004 redesign

The 2004 redesign was developed in the context of a restricted budget framework. Unlike the redesign in 1994, the questionnaire and collection application were not modified. The new sample design was gradually introduced beginning in November 2004.

To reduce survey operating costs, two major changes to the methodology were introduced. Before November 2004, the first of six interviews was conducted in person. To reduce collection costs, the first interview is now done by telephone for certain dwellings in urban areas. The second change aims at reducing the cost of listing (listing consists of compiling a list of residential addresses within a selected PSU).

To reduce the cost of listing and improve the survey frame coverage, the Address Register (AR) was used. The AR is a database containing the addresses of dwellings in urban centres. In 2005, the AR included approximately 13 million addresses. The majority of these addresses were reported to be valid residential dwellings during the 2001 Census. Other addresses found on the AR are obtained through updates from administrative files.

In addition to these two major changes, other improvements were made to the sample design. In the past, the Statistics Canada geographical database did not cover the entire territory of all 10 provinces. This database is required to establish the boundaries of the PSUs. The quality of the geographical database has greatly improved since the 1994 redesign. For the first time in 2004, the LFS was able to define the boundaries of PSUs for the entire territory of the 10 provinces.

Another improvement is that in order to better control the sample distribution, and in turn, collection costs, a new strategy was implemented for regions with high collection costs. The LFS also introduced methods to target the immigrant population in large centres and the Aboriginal population in the four Western provinces. To decrease the maintenance costs associated with the survey frame, the new sample design no longer contains a survey frame of apartments. Lastly, the sample of small rural areas is now selected using a two-stage design rather than the previous three-stage design.

5.2.6 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 non-response 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.

5.2.7 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 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 aged 15 or over.

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 8, the NLSCY sample consists of:

- a new ECD cohort of 0- to 1-year-old children selected at Cycle 8
- a top-up sample of new 2- to 5-year-olds selected at Cycle 8.

Plus the following returning samples:

- returning 2- to 3-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 7
- returning 4- to 5-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 6
- returning 4- to 7-year old children from the top-up sample of 2- to 5-year-olds selected at Cycle 7
- returning 6- to 7-year-old children from the top-up sample of 2- to 5-year-old children selected at Cycle 6
- returning 6- to 7-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 5
- returning 14- to 25-year-old children from the original cohort of 0- to 11-year-olds selected at Cycle 1.

All children were sampled from the LFS. Note that at Cycle 8, there are no children aged 8 to 13 years old.

In addition to explaining how sampling is 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 cross-sectionally out-of-scope, e.g., they were deceased, had left the country or had too many cycles of non-response.

For details on how sampling was performed at each cycle, please refer to a cycle's User's Guide.

5.3.1 Original cohort, Cycles 1 to 8

The following describes the composition of the original cohort at each cycle (see Figure 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 with varying precision by province—provincial estimates for children from

age 0 to 11, by two-year age groupings: 0 to 1, 2 to 3, 4 to 5, 6 to 7, 8 to 9 and 10 to 11. A maximum of four children per household were selected. Households were sampled from the following sources:

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

At the end of Cycle 1, there were 22,831 respondent children in the NLSCY. The child-level response rate was 86.5% (see Table 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 child-level response rate for collection was 91.5% (see Table 1). The cumulative, longitudinal response rate for children in the original cohort was 79.1% (see Table 2).

Cycle 3
At Cycle 3, 185 children were excluded from the sample because at the end of Cycle 2 they were either cross-sectionally out-of-scope (71) or hard refusals (114). Children who are cross-sectionally out-of-scope include those who died, whose age was not in-scope, who had permanently left the country, or who had moved to an Indian reserve. Thus, of the 16,903 children sampled for the original cohort, collection was performed on 16,718. The child-level response rate for collection was 89.2% (see Table 1). The cumulative, longitudinal response rate for children in the original cohort was 76.0% (see Table 2).

Cycle 4
At Cycle 4, to make collection more efficient, it was decided that households with two or more consecutive cycles of non-response would be dropped from collection (along with households with one cycle of non-response followed by the status “Temporarily moved”). Consequently, 1,086 children were dropped from collection at Cycle 4 because at the end of Cycle 3 they were either cross-sectionally out-of-scope (106) or had had two or more cycles of non-response (980). A total of 15,632 children were surveyed. The child-level response rate for collection was 84.5% (see Table 1). The cumulative, longitudinal response rate for children in the original cohort was 67.8% (see Table 2).

Cycle 5
At Cycle 5, it was decided that 18- and 19-year-olds would be dropped from collection only after three consecutive cycles of non-response (versus two for younger children). The reason for this is that at age 18, the youth becomes the sole respondent, whereas before age 18 the primary respondent is the person most knowledgeable (PMK), who is typically the mother.

At Cycle 5, 469 children were dropped from collection because at the end of Cycle 4 they were either cross-sectionally out-of-scope (32) or had too many consecutive cycles of non-response (437). A total of 15,163 children were surveyed at Cycle 5. The child-level response rate for collection was 81.3% (see Table 1). The cumulative, longitudinal response rate for children in the original cohort was 63.1% (see Table 2).

Cycle 6
At Cycle 6, 1,506 children were dropped from collection because they had too many consecutive cycles of non-response by the end of Cycle 5. A total of 13,657 children were surveyed at Cycle 6. The child-level response rate for collection was 82.4% (see Table 1). The cumulative, longitudinal response rate for children in the original cohort was 57.6% (see Table 2).

Cycle 7
At Cycle 7, 613 children were dropped from collection because at the end of Cycle 6 they were either cross-sectionally out-of-scope (11) or had too many consecutive cycles of non-response (602). A total of 13,709 children were surveyed at Cycle 7. The child-level response rate for collection was 80.5% (see Table 1). The cumulative, longitudinal response rate for children in the original cohort was 56.6% (see Table 2).

Note that at Cycle 7 a new rule was applied for returning children who were 18 or older: the PMK's history of non-response was ignored when deciding if the child should be sent to collection or not. This means that some children who had been dropped from collection in previous cycles were sent for collection in Cycle 7.

Cycle 8
At Cycle 8, 56 children were dropped from collection because at the end of Cycle 7 they were either cross-sectionally out-of-scope (17) or had too many consecutive cycles of non-response (39). A total of 15,056 children were surveyed at Cycle 8. The child-level response rate for collection was 68.0% (see Table 1). The cumulative, longitudinal response rate for children in the original cohort was 52.7% (see Table 2).

Note that the rule that was applied at Cycle 7 for returning children who were 18 or older, was also applied at Cycle 8. For the 14- to 17-year-old children, the ones that last responded at Cycle 5, 6 or 7 were surveyed at Cycle 8.

5.3.2 Early childhood development cohorts present at Cycle 8

The ECD children present at Cycle 8 were first sampled in Cycles 5, 6, 7 and 8. When the first ECD cohort of 0- to 1-year-olds was selected at Cycle 2, the rule was a maximum of one child per household, except for twins, in which case both were sampled.[1] At Cycle 5, the rule changed to one child per household without exception.[2] There were no twins in the returning ECD cohorts at Cycle 8.

Prior to Cycle 7, for the ECD samples, only respondents from the previous cycle were surveyed at subsequent cycles. At Cycle 7, this rule was modified so that non-respondents from previous cycles were surveyed, so long as they did not have 2 or more consecutive cycles of non-response. At Cycle 8, this latter condition was dropped so that all non-respondents from previous cycles were surveyed, regardless of the number of consecutive cycles of non-response.

1. 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. No Cycle 2 ECD children are present in the Cycle 8 sample.
2. For the Cycle 7 sample, it was decided that one of the returning twins was dropped for returning ECD cohorts. The original cohort continues to have a maximum of two children per household.

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 non-respondents. Consequently, there were 3,252 respondents at Cycle 5, returning as 2- to 3-year-olds, who were surveyed at Cycle 6. The response rate was 88.6% (see Table 1). The cumulative longitudinal response rate was 65.3% (see Table 2).

At the end of Cycle 6, 4 children were cross-sectionally out-of-scope and 33 were non-respondents. Consequently, there were 3,215 who responded to a previous cycle, returning as 4- to 5-year-olds, who were surveyed at Cycle 7. The response rate was 85.9% (see Table 1). The cumulative longitudinal response rate was 62.4% (see Table 2).

Finally, there were 3,214 respondents from a previous cycle, returning as 6- to 7-year-olds, who were surveyed at Cycle 8. The response rate was 80.5% (see Table 1). The cumulative longitudinal response rate was 58.8% (see Table 2).

Cycle 6 ECD cohort
At Cycle 6, a sample of 0- 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).

At the end of Cycle 6, 21 children were cross-sectionally out-of-scope and 143 were non-respondents. Consequently, there were 5,631 children aged 2 to 7 years old who were surveyed at Cycle 7. The response rate was 83.0% (see Table 1). The cumulative longitudinal response rate was 79.7% (see Table 2).

Finally, there were 5,039 children aged 4 to 7 years old who were surveyed at Cycle 8. The response rate was 76.9% (see Table 1). The cumulative longitudinal response rate was 75.1% (see Table 2).

Cycle 7 ECD cohort
At Cycle 7, a sample of 0- to 5-year-olds was selected from the LFS. The total sample size was 5,843 children and households. At the end of Cycle 7 collection, there were 4,691 responding children. The response rate was 80.8% (see Table 1).

At the end of Cycle 7, 35 children were cross-sectionally out-of-scope and 11 were non-respondents. Consequently, there were 5,797 children aged 2 to 7 years old who were surveyed at Cycle 8. The response rate was 79.1% (see Table 1). The cumulative longitudinal response rate was 78.2% (see Table 2).

Cycle 8 ECD cohort
At Cycle 8, a sample of 0- to 5-year-olds was selected from the LFS. The total sample size was 6,685 children and households. At the end of Cycle 8 collection, there were 5,065 responding children. The response rate was 76.0% (see Table 1).

Table 1 Child-level response at collection

Sampling type	Sampling cohort	Survey cycle	Age group in years	Sample size		Sample reduction		Out-of-scope Dropped from previous cycle		Non-respondents Dropped from previous cycle		In-scope		Respondents		In-scope rate (%)	Response rate (%)
				hhs	child	hhs	child	hhs	child	hhs	child	hhs	child	hhs	child	hhs	child
1994 Original cohort	1	1	0-11	43,751	..	..	..	..	..	..	..	15,502	26,409	13,439	22,831	35.4	86.5
		2	2-13	11,188	16,903	25,588	..	5,345	..	1,677	..	11,140	16,816	10,216	15,391	99.6	91.5
		3	4-15	11,032	16,718	0	0	38	71	73	114	10,937	16,563	9,801	14,777	99.1	89.2
		4	6-17	10,449	15,632	0	0	65	106	618	980	10,418	15,588	8,834	13,176	99.7	84.5
		5	8-19	10,355	15,163	0	0	24	32	286	437	10,320	15,113	8,582	12,280	99.7	81.3
		6	10-21	9,881	13,657	0	0	0	0	878	1,506	9,816	13,572	8,201	11,178	99.3	82.4
		7	12-23	10,522	13,709	0	0	7	11	406	602	10,454	13,616	8,561	10,966	99.4	80.5
		8	14-25	12,021	15,056	0	0	13	17	28	39	11,981	15,007	8,468	10,208	99.7	68.0
1996 ECD & NB Buy-in	2	2	0-1	5,592	..	..	..	..	..	..	..	4,929	5,087	4,496	4,634	88.1	91.1
		3	2-3	3,992	4,046	558	..	598	..	444	..	3,950	4,004	3,592	3,640	98.9	90.9
		4	4-5	3,577	3,610	520	540	25	25	34	35	3,552	3,585	3,023	3,052	99.3	85.1
1998 ECD & 5 top-up	3	3	0-1 & 5	16,812	..	..	..	..	..	..	..	15,929	16,263	13,256	13,546	94.7	83.3
		4	2-3	7,941	8,118	6,935	..	516	..	1,420	..	7,896	8,070	6,956	7,111	99.4	88.1
		5	4-5	6,960	7,115	0	22	41	41	940	940	6,919	7,073	6,208	6,340	99.4	89.6
		7	8-9	6,016	6,016	..	..	..	..	..	..	5,988	5,988	5,321	5,321	99.5	88.9
2000 ECD & 5 top-up	4	4	0-1 & 5	9,439	..	..	..	..	..	..	..	9,116	9,192	6,908	6,961	96.6	75.7
		5	2-3	3,788	3,841	4,405	..	125	..	1,121	..	3,776	3,829	3,281	3,324	99.7	86.8
		6	4-5	3,280	3,323	0	10	11	11	497	497	3,270	3,313	2,931	2,964	99.7	89.5
		7	6-7	3,231	3,231	0	43	3	3	46	46	3,217	3,217	2,882	2,882	99.5	89.6
2002 ECD	5	5	0-1	4,492	4,492	..	..	..	..	..	..	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
		7	4-5	3,215	3,215	0	0	4	4	33	33	3,189	3,189	2,740	2,740	99.2	85.9
		8	6-7	3,214	3,214	0	0	26	26	3	3	3,205	3,205	2,580	2,580	99.7	80.5
2004 ECD & 2-5 top-up	6	6	0-5	5,795	5,795	..	..	..	..	..	..	5,763	5,763	4,684	4,684	99.4	81.3
		7	2-7	5,631	5,631	0	0	21	21	143	143	5,600	5,600	4,650	4,650	99.4	83.0
		8	4-7	5,039	5,039	0	0	31	31	688	688	5,006	5,006	3,852	3,852	99.3	76.9
2006 ECD & 2-5 top-up	7	7	0-5	5,843	5,843	..	..	..	..	..	..	5,808	5,808	4,691	4,691	99.3	80.8
		8	2-7	5,797	5,797	0	0	35	35	11	11	5,769	5,769	4,561	4,561	99.5	79.1
2008 ECD & 2-5 top-up	8	8	0-5	6,685	6,685	..	..	..	..	..	..	6,666	6,666	5,065	5,065	99.7	76.0

Notes:
.. Not available
hhs stands for "household", child stands for "respondent children".
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Table 2 Child-level longitudinal response

Longitudinal											
Sampling type	Sampling cohort	Survey cycle	Age group in years	Sample size		In-scope		Respondents		In-scope rate (%)	Response rate (%)
				hhs	child	hhs	child	hhs	child	hhs	child
1994 Original cohort	1	1	0-11	18,163	..	12,818	19,487	11,141	16,903	70.6	86.7
		2	2-13	18,210	..	12,863	19,481	10,220	15,403	70.6	79.1
		3	4-15	18,165	..	12,818	19,481	9,810	14,796	70.6	76.0
		4	6-17	18,265	..	12,912	19,435	8,839	13,168	70.7	67.8
		5	8-19	18,481	..	13,134	19,481	8,592	12,300	71.1	63.1
		6	10-21	18,885	..	13,532	19,474	8,222	11,210	71.7	57.6
		7	12-23	19,402	..	14,409	19,474	8,597	11,016	72.4	56.6
		8	14-25	20,043	..	14,690	19,474	8,510	10,268	73.3	52.7
1996 ECD cohort	2	2	0-1	5,271	..	4,673	4,733	4,100	4,154	88.7	87.8
		3	2-3	5,271	..	4,671	4,731	3,595	3,643	88.6	77.0
		4	4-5	4,620	..	3,846	3,880	2,899	2,928	83.2	75.5
1998 ECD cohort	3	3	0-1	9,877	..	9,361	9,559	7,949	8,126	94.8	85.0
		4	2-3	9,877	..	9,141	9,334	6,794	6,946	92.5	74.4
		5	4-5	9,877	..	9,141	9,334	6,060	6,189	92.5	66.3
		7	8-9	9.877	..	9,141	9,141	5,325	5,325	92.5	58.3
2000 ECD cohort	4	4	0-1	5,034	..	4,909	4,985	3,788	3,841	97.5	77.1
		5	2-3	5,034	..	4,907	4,983	3,279	3,322	97.5	66.7
		6	4-5	5,034	..	4,911	4,987	2,932	2,965	97.6	59.5
		7	6-7	5,034	..	4,911	4,911	2,885	2,885	97.6	58.7
2002 ECD cohort	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	2,867	2,867	97.8	65.3
		7	4-5	4,492	4,492	4,392	4,392	2,741	2,741	97.8	62.4
		8	6-7	4,492	4,492	4,392	4,392	2,582	2,582	97.8	58.8
2004 ECD cohort & 2-5 Top-up	6	6	0-1	4,356	4,356	4,343	4,343	3,521	3,521	99.7	81.1
		7	2-3	4,356	4,356	4,343	4,343	3,463	3,463	99.7	79.7
		8	4-5	4,356	4,356	4,343	4,343	3,263	3,263	99.7	75.1
2006 ECD cohort & 2-5 Top-up	7	7	0-1	4,997	4,997	4,975	4,975	4,015	4,015	99.6	80.7
		8	2-3	4,997	4,997	4,973	4,973	3,888	3,888	99.5	78.2
2008 ECD & 2-5 Top-up	8	8	0-1	5,482	5,482	5,464	5,464	4,107	4,107	99.7	75.2

Notes:
This table contains data for longitudinal children only, i.e., children who are followed through time.
.. Not available
hhs stands for "household", child stands for "respondent children".
Source: Statistics Canada, National Longitudinal Survey of Children and Youth

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 are 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 non-response, to match to demographic counts by age, 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. This is because with time populations change as a result of 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 31, 1994, and who lived in any province during collection of Cycle 1, in 1994/1995. A cross-sectional population at Cycle 8 could be children aged 0 to 11 as of December 31, 1994, who lived in any province in 2008, 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, as 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 to make inferences about the cross-sectional populations after Cycle 4. However, the ECD cohorts are designed for both cross-sectional and longitudinal purposes. 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 15.0.)

It should also be noted that all final weights are adjusted for non-response 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 original and various 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'.

The original cohort, selected at Cycle 1

- Longitudinal population comprises children aged 0 to 11 as of December 31, 1994, who were living in any province 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 8, these children were aged 14 to 25 (as of December 31, 2008). Sample reductions were made at Cycle 2. The children dropped between Cycle 1 and Cycle 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 8

- The longitudinal population of the ECD cohort selected at Cycle 2 comprises children aged 0 to 1 as of December 31, 1996, who were living in any province 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 8 sample.

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

- The longitudinal population of the ECD cohort selected at Cycle 3 comprises children aged 0 to 1 as of December 31, 1998, who were living in any province at the time of Cycle 3 collection (1998/1999).

  At Cycle 3, a longitudinal sample of children aged 0 to 1 was selected from the LFS and Birth Registry data. This cohort was followed for four cycles: at Cycle 4, they were 2 to 3 years old, at Cycle 5, they were 4 to 5 years old, and at Cycle 7, they were 8 to 9 years old. They are not present in the Cycle 8 sample.

  Weights were produced for this longitudinal population at Cycles 3, 4, 5 and 7.

- The longitudinal population of the ECD cohort selected at Cycle 4 comprises children aged 0 to 1 as of December 31, 2000, who were living in any province 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. This cohort was followed for four cycles. They are not present in the Cycle 8 sample.

  Weights were produced for this longitudinal population at Cycles 4, 5, 6 and 7.

- The longitudinal population of the ECD cohort selected at Cycle 5 comprises children aged 0 to 1 as of December 31, 2002, who were living in any province 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 8, these children were 6 to 7 years old.

Weights were produced for this longitudinal population at Cycles 5, 6, 7 and 8.

- The longitudinal population of the ECD cohort selected at Cycle 6 comprises children aged 0 to 1 as of December 31, 2004, who were living in any province 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. The returning 4- to 5- year-olds were surveyed at Cycle 8.

  Weights were produced for this longitudinal population at Cycles 6, 7 and 8

- The longitudinal population of the ECD cohort selected at Cycle 7 comprises children aged 0 to 1 as of December 31, 2006, who were living in any province at the time of Cycle 7 collection (2006/2007).

  At Cycle 7, a longitudinal sample of children aged 0 to 1 was selected from the LFS. The returning 2- to 3- year-olds were surveyed at Cycle 8.

  Weights were produced for this longitudinal population at Cycles 7 and 8.

5.4.2 Cohorts and their cross-sectional populations

The original cohort (at Cycle 8, aged 14 to 25 as of December 31, 2008)

As 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. Nevertheless, the original cohort can be used to make inferences about the longitudinal population defined in 5.4.1 and the following cross-sectional populations:

- Cycle 2 cross-sectional population: children aged 2 to 13 as of December 31, 1996, who were living in any province 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 31, 1998, who were living in any province 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 31, 2000, who were living in any province 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 31, 1996, who were living in any province 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 31, 1998, who were living in any province 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 31, 2000, who were living in any province 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 31, 2002, who were living in any province 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 31, 2004, who were living in any province at the time of Cycle 6 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 at Cycle 6 (from the LFS).

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

- Cycle 7 cross-sectional population: children aged 0 to 9 as of December 31, 2006, who were living in any province at the time of Cycle 7 collection (2006/2007).

  - the ECD cohort of 0- to 1-year-old children selected at Cycle 7
  - returning 2- to 3-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 6
  - returning 4- to 5-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 5
  - a new top-up of 2- to 5-year-old children selected at Cycle 7 (from the LFS)
  - returning 4- to 7-year-old children from the top-up sample of 2- to 5-year-old children selected at Cycle 6
  - returning 6- to 7-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 4
  - returning 8- to 9-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 3.

  Cross-sectional weights were produced for this population at Cycle 7.

- Cycle 8 cross-sectional population: children aged 0 to 7 as of December 31, 2008, who were living in any province at the time of Cycle 8 collection (2008/2009).

  - the ECD cohort of 0- to 1-year-old children selected at Cycle 8
  - returning 2- to 3-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 7
  - returning 4- to 5-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 6
  - a new top-up of 2- to 5-year-old children selected at Cycle 8 (from the LFS)
  - returning 4- to 7-year-old children from the top-up sample of 2- to 5-year-old children selected at Cycle 7
  - returning 6- to 7-year-old children from the top-up sample of 2- to 5-year-old children selected at Cycle 6
  - returning 6- to 7-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 5.

  Cross-sectional weights were produced for this population at Cycle 8.

6.0 Data collection

Data for Cycle 8 of the National Longitudinal Survey of Children and Youth (NLSCY) were collected from the fall of 2008 to the summer of 2009.

The collection was divided into three waves as described below. The date of the interview is recorded in variable HMMHgQ06.

Collection period	Age group
September to January, Wave 1	Ages 0 to 6 and 18- to 25-year-olds who do not have selected siblings aged 14 to 17
January to April, Wave 2	Ages 7, 14 to 17 and 18 to 25 who were not collected in the first collection period
March to July, 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, helping with data quality and ensuring that respondents answer only the questions appropriate to their situations. 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 'survey instruments'—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 their relationships to one another, gather tracing information, and record basic demographic characteristics such as sex, date of birth, marital status and relationships between household members. Sex and date of birth are updated for new respondents only.

The person most knowledgeable (PMK) about the child was also identified in this component. The PMKs provide the information for all selected children in the household and then give information about themselves and their spouses or partners. 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 stepfamily, 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 was selected per household.

The PMK was selected once the information about the relationships between household members had been collected.

6.2 Child component

A Child component was created for each selected child from 0 to 17 years of age, except for youth aged 16 or 17 years old who are living independently. The PMK 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 was 17 years old or younger, except for youth aged 16 or 17 years old who are living independently. 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 8

The number of children and youth sampled in Cycle 8 is shown by age and province in the following tables.

Number of sampled children and response rate, by age at Cycle 8

Age as of December 31, 2008 (Years)	Sampled	In-scope	Respondents	Cycle 8 response rate (%)
	Number			
0	2,400	2,389	1,808	75.7
1	3,082	3,074	2,298	74.8
2	2,064	2,060	1,643	79.8
3	3,516	3,495	2,729	78.1
4	2,291	2,278	1,761	77.3
5	3,113	3,094	2,369	76.6
6	1,992	1,984	1,636	82.5
7	2,279	2,272	1,814	79.8
8	1	0	0	n/a
11	1	0	0	n/a
14	1,528	1,527	1,227	80.4
15	1,606	1,602	1,274	79.5
16	1,136	1,134	890	78.5
17	1,102	1,101	880	79.9
18	1,304	1,302	850	65.3
19	1,219	1,213	785	64.7
20	1,178	1,172	689	58.8
21	1,183	1,182	677	57.3
22	1,232	1,221	757	62.0
23	1,186	1,185	713	60.2
24	1,250	1,246	755	60.6
25	1,132	1,122	711	63.4
Total	**35,795**	**35,653**	**26,266**	**73.7**

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Number of sampled children and response rate, by province of residence at Cycle 8

Province	Sampled	In-scope	Respondents	Cycle 8 response rate (%)
	Number			
Newfoundland and Labrador	2,104	2,098	1,704	81.2
Prince Edward Island	1,513	1,511	1,167	77.2
Nova Scotia	2,347	2,339	1,744	74.6
New Brunswick	2,318	2,312	1,778	76.9
Quebec	5,658	5,642	4,344	77.0
Ontario	8,579	8,559	6,059	70.8
Manitoba	2,710	2,697	1,940	71.9
Saskatchewan	2,868	2,857	2,219	77.7
Alberta	4,244	4,227	2,988	70.7
British Columbia	3,375	3,360	2,272	67.6
Outside the 10 provinces	79	51	51	100.0
Total	**35,795**	**35,653**	**26,266**	**73.7**

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

6.6 Direct assessments

A variety of direct assessments are administered to the selected respondents. 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- and 5-year-olds	Computer-assisted interview
Who Am I?	4- and 5-year-olds	Paper questionnaire
Number Knowledge	4- and 5-year-olds	Computer-assisted interview
Mathematics Computation Exercise	7- to 15-year-olds in grades 2 to 10	Paper questionnaire
Problem-solving Exercise	16- and 17-year-olds	Paper questionnaire
Literacy	18- and 19-year-olds	Paper questionnaire
Numeracy	20- and 21-year-olds	Paper questionnaire

6.7 Self-complete questionnaires – Ages 14 to 17

Respondents from 14 to 17 years of age completed a paper questionnaire on various aspects of their lives. They were given the questionnaire during the interview and asked to complete it themselves. To ensure confidentiality, each respondent placed the completed questionnaire in an envelope, sealed the envelope and gave it to the interviewer.

The Self-complete questionnaires consisted of two booklets, one for each age group. The table below shows the topics covered by each age group in the booklet. The questions for each subject may be different for each age group. The booklets are reproduced in Book 2 of the *National Longitudinal Survey of Children & Youth, Cycle 8 Survey Instruments 2008/2009*.

Topic	Self-complete section by age and booklet	
	14 and 15 years Booklet #22	16 and 17 years Booklet #23
Friends and family	A	A
School	B	...
About me	C	B
Feelings and behaviours	D	C
Activities	E	...
Smoking, drinking and drugs	F	D
My parent(s)	G	G
Puberty	H	...
Dating / My relationships	H	F
Health	H	E
Work	I	...
Thank you	J	H

… not applicable
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

6.8 Collection personnel (training, supervision and control)

The NLSCY is conducted by Statistics Canada interviewers. A number of them have worked on one or more cycles of the NLSCY.

In each cycle, senior interviewers are responsible for ensuring that NLSCY interviewers are familiar with the survey's concepts and procedures. The senior interviewers ensure that prompt follow-up action is taken in the case of refusal and other kinds of non-response. If necessary, the non-response cases are transferred to a senior interviewer and reassigned. The senior interviewers, in turn, report to the program managers in Statistics Canada's regional offices.

For the NLSCY, a combination of classroom training and self-study materials is used to ensure that interviewers and supervisors have a proper understanding of the survey concepts. In the self-study portion, which precedes the classroom training, the program managers, senior interviewers and interviewers read the interviewer's manual for the survey. The classroom training is initially given by Head Office to the senior interviewers who subsequently train all the interviewers in their respective areas.

6.8.1 Interviewing in non-official languages

The NLSCY questionnaires are only available in English and French. If a respondent wishes to be interviewed in another language, the case will be given to an interviewer who speaks the respondent's language, if possible. In Cycle 8, fewer than 80 cases were not completed because of a language barrier.

7.0 Data processing

The main outputs of the National Longitudinal Survey of Children and Youth (NLSCY), Cycle 8 are clean data files that researchers can use. This chapter presents a brief summary of some of the processing steps involved in producing these files.

The processing of the NLSCY, Cycle 8 data was done using the divisional 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 are:

- Cleanup
- 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 each household, except for the 14- to 17-year-olds' Self-complete questionnaires and the direct assessment, was collected in a face-to-face or telephone interview using a computer-assisted interviewing (CAI) application. As such, it was possible to build edits and checks into the questionnaires for the various household CAI components to ensure that high-quality information was collected. Below are specific examples of the types of edits used in the NLSCY's CAI application:

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

Range edits
These were built into the CAI system to deal with questions asking for numeric values. If values entered are outside the range, the system generates a pop-up window that states the error and instructs 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 is significantly high or low, a pop-up message will 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 PMKs are asked whether they use daycare or babysitting. Based on the response given, the flow of the questions could be different. If child care is used, the CAI system continues with a series of questions about the specific child care method(s) used for the child. If not, the CAI system automatically skips this series of questions.

General consistency edits
Some consistency edits were included as part of the CAI system to allow interviewers to return to previous questions to correct for inconsistencies. Instructions are displayed to interviewers for handling or correcting problems such as incomplete or incorrect data. The system generates a pop-up window that states the error and instructs 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:

- Self-complete questionnaires for 14- to 17-year-olds
- Who Am I? for 4- to 5-year-olds
- Mathematics Test for 7 year-olds and 14- to 15-year-olds who are in grade 2 and higher
- Problem Solving Exercise for 16- and 17-year-olds
- Literacy assessment for 18- and 19-year-olds
- Numeracy assessment for 20- and 21-year-olds.

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

In cases where more than one response was checked off by the respondent, a value of "Not stated" was entered. This is a change from Cycle 7 when the operators were instructed to accept the first response. Errors remaining within the questionnaires were edited at a later stage.

7.3 Cleanup

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 are split between response and non-response based on predetermined criteria.

The responding and non-responding questionnaires were reviewed and analyzed. Based on the analysis, specifications were created to determine which records would be dropped because of non-response. Essentially, if a record was missing key information or had more than half the questions unanswered, it was 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 inadequate information, the household was dropped from the responding sample and treated as a non-response.

Each file contains records for children who are now deceased or who have moved out of the country, but who will be kept on the files for weighting purposes. There are 114 ECD records, 143 longitudinal records and 189 youth records that have a longitudinal weight (HWTCW01L) and all

other variables have been set to 'Not stated'. For those children in this group who are 16 and 17 they will appear on both the longitudinal and youth files (53 cases).

7.4 Age and gender edits

In this step, the age variables are verified for new sample and new household members in returning households. This is a change from Cycle 6 in which verification was done on all household members. Starting in Cycle 7, the application did not allow updates to the returning sample and household members. Also, the age is verified to be consistent with the age cohort. The respondent's sex is also verified to be consistent with the previous cycle.

7.5 Relationship edits

The relationship edit step validates the relationships among the members of the household and creates the family-derived variables. Standard edits are made to the relationship information entered for all members of a given household; some inconsistencies are corrected automatically by an application using a set of rules, whereas others are flagged for manual review and recoding. The relationship edits produce a related set of derived variables.

7.6 Pre-edits

For all records where values were missing (blank) from the collection, the value of "9," "99," "999," etc. 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 destrung 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 that are skipped based on flows are converted from "Not stated" to "Valid skip" codes ("6," "96," "996," etc.).

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

There were 24 16- and 17-year-olds living independently in Cycle 8 (HDMCfD03 = 82). All variables in the longitudinal file were set to "Valid skip" for these respondents, except for some variables in the Demographic section.

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, PMKs who had worked in the previous 12 months were asked a series of open-ended questions about their 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?

Questions about career aspirations were asked in the Youth Questionnaire for 18- to 21-year-olds:

- What kind of career would you be interested in having or what kind of work would you be interested in doing 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 that describe the nature of the respondent's work.

How they are coded
These open-ended questions were coded using various standard classifications. For Cycle 8 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)[3]. Grouped versions of these codes are available on the data file (HLFPgD7A and HLFPfD8A for the PMK, and HLFSgD7A and HLFSfD8A for the spouse or 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. 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 Activities section, for children aged 3 to 7 years old who reported using a computer at home (HACCeQ4E), there was a consistency edit that compared these children and those who reported having a computer in their home (HACCeQ4D). If the child reported using a computer at home, then the response to "Is there a computer in your home?" (HACCeQ4D) was set to "Yes" in the edit.

A change was made to the consistency edits for Motor and Social Development questions in Cycle 7. In previous cycles a correlation was made between Milestone responses and Motor and Social Development responses. For example, if a child was not able to sit up (MILQ1) then an assumption was made that the child could not sit up for 10 minutes without support (HMSCQ16). In Cycle 8 no correlation was made. This affected several questions in the Motor and Social Development block.

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 non-response' or 'partial response'.

Imputation
For a few variables on the NLSCY file, rather than using a special non-response code, imputation was used. 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.

3. Information about classification and concordance to previous classifications can be found at: www.statcan.gc.ca/concepts/index-eng.htm.

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

All imputation flag variables on the NLSCY data file have an "I" as the fifth or sixth character of the variable name. For example, the imputation flag variable for the income of the PMK is named HINPI1AA.

7.11 Creation of derived variables and longitudinal edits

Deriving variables from more than one data item
A number of variables have been derived by combining questions on the questionnaire, or by combining other derived variables to facilitate data analysis. For example, in the Labour Force section, current working status is derived from questions about respondents' current work situations. In the Adult Education section, there are questions about whether respondents are 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 (HLFPD51, HLFSD51).

Longitudinal derived variables
In previous cycles a derived variable was created to indicate changes between data reported in the current and previous cycles for family structure, PMK and spouse changes. Starting in Cycle 7, this derived variable does not exist; rather the PMK and spouse identifiers have been added to the data files for Cycle 4 through Cycle 8.

Inconsistencies between cycles for the child's height and the child's weight were flagged. The variables that identify these inconsistencies contain a "Z" as the fifth character of the variable name.

At this step, data were brought forward from previous cycles for variables that were asked of the respondent only once, such as Birth weight, Breast-feeding information and Sociodemographic information. Data were also brought forward for derived variables that were from previous cycles and that were derived from data that did not change over time.

Derived variable name
All derived variables on the NLSCY data file have a "D" as the fifth or sixth character of the variable name. For example, the name of the derived variable for primary care arrangements is HCRCgD2A.

7.12 Standard coding structures

Standards have been developed for the coding structure of NLSCY variables to explain certain situations consistently across all variables. The following describes various situations and the codes used to describe them.

Refusals
During a CAI interview, the respondent may choose to refuse to answer 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 to Head Office.

On the NLSCY data file, an item that is refused is indicated by a code "8". For a variable that is one digit long, the code is "8"; it is "98" for a two-digit variable, "998" for a three-digit variable, 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", "97", "997", etc.

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

In other 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 (HCRCg1GA) was applicable only for children for whom this type of care is used (HCRCgQ1G = 1). Otherwise, a code "996" was used for this question.

Occasionally, an entire section of the questionnaire or the entire questionnaire was not applicable. For example, the Motor and Social Development section was applicable only to children 0 to 3 years old. For all children 4 years old and older, the motor and social development variables have been set to a "Valid skip".

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 was set to "Not stated". The "Not stated" code indicates that the answer to the question is unknown. "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 "Refusal," it was assumed that the line of questioning was sensitive and that the respondent would probably 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 adequately informed to answer further questions. As part of the NLSCY processing system, it was decided that all of these subsequent questions would be assigned a "Not stated" code. This code means that the question was not asked of the respondent. In some cases, it is not even known whether the question was applicable to the respondent.

- In certain cases, a questionnaire was not started or was ended prematurely. For example, an interruption occurred, or the respondent decided that he/she wished to terminate the interview. If there was enough information collected to establish the household as a responding household, 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 when it was known that a certain section or a certain questionnaire was not applicable, and these questions were set to "Valid skip".
- Some paper questionnaires were returned partially complete. If enough information was collected about the respondent, all remaining incomplete items on the questionnaire were set to "Not stated." The one exception was when it was known that a certain section or a certain questionnaire was not applicable, and these questions were set to "Valid skip".

- "Not stated" codes were used occasionally as a result of consistency edits. When the relationship between groups of variables was checked for consistency, if an error existed, 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", "99", "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 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 Cycle 1
"B" indicates Cycle 2
"C" indicates Cycle 3
"D" indicates Cycle 4
"E" indicates Cycle 5
"F" indicates Cycle 6
"G" indicates Cycle 7
"H" indicates Cycle 8

The **second and third** characters, which together comprise an acronym, 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 section.

The **fourth** character refers to the collection unit or the unit to which the variable refers. There are seven possibilities[4].

"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
"M" if the variable refers to the mother
"O" if the variable refers to other household members 15 years of age or older

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, respectively.

4. It should be noted that whereas variables do exist for various units of analysis, 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 is 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. However, it is not possible to produce estimates of the number of low-income households or depressed PMKs.

For example,

"h" the lower-case letter refers to the NLSCY cycle in which the variable first appeared on the file or the cycle (for example, "h" refers to Cycle 8) in which changes to a previously asked question were made.

"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.

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 8, the variable was treated as new and contains an "h". For example, in Cycle 8 the variable HCRChQ6B was asked of children aged 6 and 7 only, resulting in a change of coverage. For children 5 and under a new question was asked for the age in months (HCRChQ6A).

7.14 Examples of variable names

The following examples illustrate the naming convention used for variables included in the NLSCY data file.

Variable name	Refers to
HLFSQ02	Q2 in the Labour Force section for the spouse/partner
H	Cycle 8 variable
LF	Labour Force section
S	Spouse/partner
Q	An item asked directly on the questionnaire
02	The identification number of the item

Variable name	Refers to
HPRCS03	A Positive interaction score on the Parenting scale for a 2- to 15-year-old child
H	Cycle 8 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. For additional information on questionnaire sections, please see Chapter 8.0.

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, 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	Sociodemographic	Child Questionnaire Adult Questionnaire (section for the PMK and spouse/partner)
HL	Health	Adult Questionnaire (section for the PMK and spouse/partner) Child Questionnaire Youth Questionnaire
CH	Adult chronic conditions	Adult Questionnaire (Health section for the PMK and spouse/partner)
RS	Restriction of activities	Adult Questionnaire (Health section for the PMK and spouse/partner)
DP	Depression scale	Parent Questionnaire (this scale was administered to the PMK)
ED	Education	Child Questionnaire Adult Questionnaire (section for the PMK and spouse/partner) Youth Questionnaire
LF	Labour force	Adult Questionnaire (section for the PMK and spouse/partner) Youth Questionnaire
IN	Income	Adult Questionnaire (Household income and personal income of the PMK) Youth Questionnaire
FN	Family functioning	Adult Questionnaire (section for the PMK)
MD	Medical/biological	Child Questionnaire
TM	Temperament	Child Questionnaire
LT	Literacy	Child Questionnaire
AC	Activities	Child Questionnaire Youth Questionnaire
BE	Behaviour	Child Questionnaire
MS	Motor and social development	Child Questionnaire
WB	Work after birth	Child Questionnaire

Acronym	Variable	Collected or derived from
RL	Social relationship	Child Questionnaire Youth Questionnaire
PR	Parenting style	Child Questionnaire
CR	Child care	Child Questionnaire
SL	Sleep	Child Questionnaire
PB	Positive behaviour	Child Questionnaire
CM	Communication scale	Child Questionnaire
ML	Development milestones	Child Questionnaire
CS	Custody	Child Questionnaire
AG	Ages and stages	Child Questionnaire
SF	Neighbourhood safety	Adult Questionnaire (section for the PMK or spouse/partner)
SP	Social support	Adult Questionnaire (section for the PMK or spouse/partner) Youth Questionnaire
AS	Aspirations	Child Questionnaire Youth Questionnaire
EQ	Emotional Quotient	Self-complete questionnaires Youth Questionnaire
AM	About me	Self-complete questionnaires Youth Questionnaire
FB	Feelings and behaviour	Self-complete questionnaires Youth Questionnaire
MO	Moving out of parental home	Youth Questionnaire
FF	Family formation and fertility	Youth Questionnaire
AB	Abilities	Youth Questionnaire
PO	Political engagement	Youth Questionnaire
FF	Friends and family	Self-complete questionnaires
SC	School	Self-complete questionnaires
PM	My parents and me	Self-complete questionnaires
PU	Puberty	Self-complete questionnaires
DR	Smoking, drinking and drugs	Self-complete questionnaires
AT	Activities	Self-complete questionnaires
HT	Health	Self-complete questionnaires
WK	Work and sources of money	Self-complete questionnaires
DA	Dating	Self-complete questionnaires
PP	PPVT–R Test	Peabody Picture Vocabulary Test
WI	Who am I?	Direct Measure (Who Am I? Questionnaire)

KN	Number knowledge	Direct Measure (Number knowledge Questionnaire)
MA	Math Computation Test	Math Tests and Problem Solving Exercise
LT	Literacy assessment	Literacy assessment Questionnaire
NU	Numeracy assessment	Numeracy assessment Questionnaire
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 8:

- *Early Childhood Development Cohort* (0 to 7 years) (Adult, Child and Household questionnaires)
- *Longitudinal Cohort – Child* (14 to 17 years) (Adult, Child and Household questionnaires)
- *Longitudinal Cohort – Youth* (16 to 25 years) (Youth and Household questionnaires; also static variables brought forward from previous cycles)
- *Self-complete Questionnaires* (14 to 17 years old)

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 and youth. 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, which consists of researchers in the area of child development and the social sciences
- federal departments
- representatives responsible for child development programs in the provinces and territories.

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, the child's parent(s), and the characteristics of the family, the neighbourhood, and 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 NLSCY includes several scales that will be discussed briefly in this chapter. Information about the scales included on the NLSCY and the validation of the scale scores are presented in greater detail in 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 The household roster is completed at the beginning of the interview and collects basic demographic information for all household members and their relationships to everyone 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 selected children in the household. No Adult component is generated for youth aged 18 years and older or for youth aged 16 or 17 years old who are living independently.

Child Questions about the selected child are asked of the PMK. The Child component is completed for each selected child aged 0 to 7 years and 14 to 17 years. The only sections of the Child Questionnaire asked about youth aged 16 and 17 are the Expectations and Aspirations section, the Custody section, and the Sociodemographic section.

Youth In this section, respondents aged 16 to 25 answer questions about themselves in a computer-assisted interview (CAI).

Self-completes Respondents aged 14 to 17 answer questions about themselves in a paper questionnaire.

Direct Assessments
Several Direct Assessments are completed by the children and youth; these are described in Chapter 14.0.

8.2 Demographic variables

The demographic variables are collected in the Household component. Basic demographic information, such as age, sex 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.

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, in terms of both 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, which is 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, such as 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 is carried out using the North American Industry Classification System (NAICS) 2002 and the National Occupational Classification for Statistics 2001 (NOC-S).

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

HLFPfD5A / HLFSfD5A:	NAICS 2002 code for PMK's/spouse's current job
HLFPfD6A / HLFSfD6A:	NOC-S 2001 code for PMK's/spouse's main job
HLFPgD7A / HLFSgD7A:	Standard industry code for current job (NAICS 2002) – grouped
HLFPfD8A / HLFSfD8A:	Standard occupation code for current job (NOC-S 2001) – grouped

Income
In the Income section of the survey, amount of income and its sources are collected for each household. There are also a few questions about the perceptions of the PMK or the PMK's spouse regarding how well he/she thinks they are 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 (HINHgD4A) has been created to compare the household income with the pre-tax low-income cut-off (LICO).[5] The LICO is used to distinguish low-income family units from other family units. A family unit is considered to have low income when its income is below the cut-off for its family size and its community. The variable HINHD03A gives the value of the LICO by geographic area.

Adult Health
This section asks PMKs and their spouses about their general health, chronic conditions and restriction of activities, and includes questions on smoking and drinking. Questions on smoking were included because research indicates that parental smoking may predict smoking among children. Questions on alcohol consumption were included 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, such as allergies, asthma and high blood pressure. The derived variables (HCHPgD01 and HCHSgD01) indicate that the respondents answered "Yes," and that they have at least one of the long-term conditions.

Restriction of activities
PMKs and their spouses are asked a series of questions about whether or not their activities are restricted at home, work, school, etc. Derived variables (HRSPdD01 and HRSSdD01) were created stating whether or not the PMK or spouse reported an activity restriction.

Maternal History
These questions focus on pregnancy history. They are asked only of those being interviewed for the first time. The questions on pregnancy and birth were provided by Dr. J.-F. Saucier, Ste. Justine Hospital, Montréal, and were later modified by the project team.

Depression scale
A Depression scale (HDPPS01) is administered to the PMK as part of the Adult Questionnaire.

Family Functioning
The objective of this 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 to 15 years old.

Neighbourhood Safety
This section gathers information about the respondents' satisfaction with their neighbourhood as a place to raise children, including perception of the extent of danger, problems, and social cohesion or 'neighbourliness.' Two scales are created in this section: Neighbourhood Safety score (HSFHhS5), indicating the degree of perceived neighbourhood safety, and Neighbours score (HSFHhS6), indicating the degree of neighbour cohesiveness.

5. For more information, please see *Low income cut-offs for 2006 and low income measures 2005*, Statistics Canada Catalogue no. 75F0002MIE2007004.

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.

Sociodemographics
The objective of this section is to gather information on immigration, ethnic background, the language profile and religious affiliation of the PMK and spouse. The information allows for analysis of various components of the Canadian population and permits identification of visible minorities.

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 (HMMCQ01). This is done to ensure that the children stay in the age groups to which they were assigned, regardless of whether collection takes place before or after their birthday. For Cycle 8, the effective age is calculated as 2008 minus the year of birth. For example, children born in 2004 would have an effective age of 4 years (2008 minus 2004). Note that the actual age of children at the time of the interview is sometimes different from their 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 on the age of the child; more information is collected for the older children with 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, and number of school changes and residential moves.

For children in Grade 1 or higher, additional questions are asked about other aspects of the educational experience such as 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 to 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.

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 to 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 (HHLCcD2A). 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 nine attributes: vision, hearing, speech, mobility (ability to get around), dexterity (use of hands and fingers), cognition (memory and thinking), emotion (feelings), 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
This section is completed for children in the 0- to 5-year-old 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 and longer-term developmental difficulties than average.

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. HMDCD06 gives the gestational age in days and HMDCD07 indicates whether 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 (HMDCD08) to indicate whether 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 co-ordination
3) Fine motor: hand and finger co-ordination
4) Problem-solving: doing different activities with objects, drawing
5) Personal–Social: solitary and social play, dressing and feeding self.

The gross motor portion of the ASQs is not included, as this concept is covered in other portions of the survey. 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.

Milestones
The questions in this section are included 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

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 7, with changes to reflect age-appropriate activities.

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. Information is collected about how the children spend their time, what their personal interests are and to which degree they interact with peers.

Several questions are included for children 4 to 7 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 PMKs and their spouses/partners.

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.

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 have been 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 (HMSCS01, HMSCS02 and HMSCdS03) 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 based on 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 children's family arrangements, including whether or not their parents are married, separated or divorced and the age of the children when parents separated/divorced.

Child Care
This section provides basic information about the methods of care currently provided for the child and 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.

Questions about parental interactions with child care provider (HCRCgQ20, HCRCg22A, HCRCg22B) were adapted from the General Social Survey (Statistics Canada), as were the questions on the preferred form of child care (HCRCg27, HCRCg28A - HCRCg28K). Additional questions on the preferred form of child care (HCRCgQ29, HCRCg30A - HCRCg30I) were developed by the project team. Questions on the criteria for selecting child care (HCRCg25A - HCRCg25K, HCRCg26A - HCRCg26K) and the cost of child care (HCRCg40A, HCRCg40B) were adapted from the Early Childhood Longitudinal Study (US National Center for Education Statistics). The availability of financial assistance/subsidy questions (HCRCgQ36, HCRCgQ38, HCRCgQ39) were derived from the survey "Awareness and Attitudes Regarding Early Learning and Child Care" (EKOS).

Beginning in Cycle 7, respondents who replied that they did not use child care while working or studying were asked if they used regular child care for other reasons. As this was a significant change to how the data were collected in the child care section, a decision was made by the NLSCY project team to create two sets of derived variables in the child care section; one series was created for those who reported using child care while working and studying and a second series of derived variables was created which included all respondents using child care. For example, the derived variable, Total Number of Care Arrangements Used, HCRCgD03, includes respondents using child care to allow PMK and spouse to work or study. It was also created as HCRCgD08, but the latter derived variable includes all respondents who use child care. Two sets of derived variables were also created for each of the following: Number of Hours per Week Spent in all Care Arrangements; Primary Type of Child Care; Number of Hours per Week Spent in Primary Care Arrangement; and Length of Time the Primary Care arrangement has been used in Months.

Expectations and Aspirations (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. 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.

Sociodemographics
The questions in this section gather sociodemographic information on ethnicity, country of origin, language(s) spoken, and Aboriginal identity.

8.5 Youth Questionnaire (ages 16 to 25)

Moving Out of the Parental Home (ages 18 to 25)
There are numerous transitions that a youth goes through from adolescence to adulthood. Undoubtedly, one of these major transitions is when youth leave their parental homes for the first time to live independently in their own residences.

Some of these youth live away from home permanently or temporarily so they can attend school or work. 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 vary depending on information collected previously and during the Household component.

The questions were designed on the advice of Dr. Dianne Looker, Chair, Department of Sociology, Acadia University, who 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 25)
This section collects information on the youths' education experiences. The first few questions establish the current educational status. Based on this information, respondents are streamed to the questions applicable to their situations. 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)

4) currently in postsecondary (for youth who are attending a postsecondary 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 that are involved in a youth choosing to continue at or leave school.

As there are many respondents of different ages and similar education statuses, the youth are streamed to the correct questions based on previous and current education statuses rather than age.

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.

Youth Labour Force (ages 18 to 25)
The questions in this section are similar to those asked of youth aged 16 and 17. However, more questions were taken from the Adult Labour Force section, because the respondents were older. Some of the youth may be working as their main activity and the questions need to reflect this possibility.

These questions collect information 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).

Variables derived from the labour force
Several labour-force derived variables have been created for the youth aged 18 to 25. They include HLYYeD5A, HLYYeD6A, HLYYgD7A and HLYYeD8A.

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 to gain insight into the degree to which young people plan for their 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 HASYfD03 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.

Youth Income (ages 18 to 25)
These questions are similar to those asked on the Adult Questionnaire. The youth are asked about their personal income if single and about their household income if married or living common law.

At this age, many youth are moving out of the parental home for the first time. This section includes a series of questions concerning payment of housing/shelter expenses and financial debt and savings. Determining whether they are paying for shelter, paying off debt and have savings or investments is important data to collect in order to assess how youth adjust to financial responsibilities.

Youth Health (ages 16 to 25)
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 related to the sleep patterns of the youth. Sleep is an important indicator of the attitudes of youth towards their bodies and how they take care of themselves. The amount of sleep reported can be used to help understand whether youth are successfully balancing the demands of work, school, volunteering, sports, etc.

Youth Health (ages 18 to 25)
This section includes questions that ask about height, weight, and risky behaviours such as smoking, drinking alcohol and drug use.

The Depression scale was included in this section (HHTYfS01). As a number of transitions take place during early adulthood, it can be an extremely stressful time for youth, and strong emotions may be evoked. Research suggests that many people suffer from depression. Because good mental health is as important as good physical health, it is important to gather information on both.

Feelings and Behaviours (ages 18 to 19 and 22 to 23)
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 questions that ask about engaging in risky behaviour such as stealing, fighting, drinking and driving, and gang membership.

Relationships (ages 18 to 25)
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 25 depending on their age and marital status.

Youth Activities (ages 16 to 25)
Young adulthood can be a time of high involvement in a variety of activities that are not related to school. It is important to measure these activities to understand how this involvement contributes to good outcomes. Of particular interest to researchers is the degree to which youth engage in life-long learning and establish their own autonomy.

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. Youth receive different questions depending on their age.

For youth aged 16 to 18, questions are included about youths' access to a vehicle and whether or not they have 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.

Political Engagement (ages 20 to 25)
A new series of questions on Political Engagement were added in Cycle 8. In addition to asking youth if they voted in a previous election, youth were asked if they believed they had a duty to vote and how often they talked about political events with others. The content was suggested by André Blais of the University of Montreal and adapted by the NLSCY project team.

About Me (ages 18 to 19 and 22 to 23)
The questions at the beginning of this section are intended to establish the level of self-esteem the youth experiences. The General Self–image score is derived from the responses (HAMYfS01).

Furthermore, additional questions are asked about any painful events youth may have experienced within the past two years. Included are events such as a painful break-up with a boyfriend or 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 in this section are similar to those in the Self-complete questionnaires given to the younger children.

Abilities (ages 16 to 25)
The Abilities questions were added in Cycle 8 but had previously appeared in Cycle 5 on the self-complete paper questionnaires. These questions help to determine youth's self-rated abilities in using a computer, writing, reading, oral communication, problem-solving and mathematics.

Emotional Quotient (ages 20 to 21 and 24 to 25)
The Emotional Quotient scale was developed by Dr. Reuven BarOn and Dr. James D.A. Parker. 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 version is similar to the 15-item scale asked of the 14- to 17-year-olds on the Self-complete questionnaires.

Youth Social Support (ages 18 to 25)
In this section the questions vary depending on the age of the youth. These questions comprise the Social Support scale from the Adult component and collect information on the youth's social support network. These questions establish the perceived amount of support that youth receive from family and friends. This section also contains questions about their relationship with their mother and father.

Family Formation and Fertility (18 to 25)
As the original NLSCY cohort moves from adolescence to adulthood, questions about marriage and common-law relationships and plans for starting a family take on relevance for this age group. In Cycle 8, a new section of questions was added to gather information about the factors influencing young adults' transition to marriage and parenthood. If single, youth are asked if they expect to marry or live in a common-law relationship. Questions are also asked about factors that may influence youth to have a child or more children. Youth who indicate that it is unlikely they will have children are asked the reasons why.

8.6 Self-complete questionnaires (ages 14 to 17)

The objective of these questionnaires is to collect information directly from the youth on a variety of aspects of their lives to supplement information obtained from the parent. The questionnaires also collect information 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 report 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 in a face-to-face interview.

Friends and Family (ages 14 to 17)
The objective of this section is to determine how well youth feel they get 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 (HFFcS01) 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 Buhmeter's Network of Relationships Inventory.

School (ages 14 to 15)
This section asks about attitude of youth towards school, how well they are doing at school, the importance of good grades, feelings of safety and acceptance at school, and the perception of whether the teacher is fair and provides 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 14 to 17)
These questions are used to determine the youth's overall self-esteem. A score is calculated (HAMcS02) based on the answers to these questions.

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 youth aged 20 to 21 and 24 to 25. This measure was selected because it assesses the respondent's social, personal and emotional abilities rather than their behaviours.

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

Feelings and Behaviours (ages 14 to 17)

Behaviour checklist (ages 14 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 created.

Risky behaviours (ages 14 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 14 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 25. A score (HHTCbS1b) is calculated based on these questions.

My Parent(s) (ages 14 to 17)
This section aims to capture the youth's relationship with parents or guardians from several different angles. Questions are geared to uncover the amounts of understanding, fairness and affection received from each parent or guardian as well as conflict resolution practices and parental supervision. Youths' impressions of the relationship and conflict resolution skills of their parents/guardians are also addressed.

My Parents and Me (ages 14 to 15)
Three scales are created using these questions:
1) Parental Nurturance (HPMCcS1)
2) Parental Rejection (HPMCbS2b)
3) Parental Monitoring (HPMCcS3).

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 (HPMCdS4)
2) Conflict Resolution scale – Father (HPMCdS5).

Smoking, Drinking and Drugs (ages 14 to 17)
This section asks questions to determine whether 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 driver who has been drinking 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 14 to 15)
The objective is to determine the extent of the youth's 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. The activities are also covered on the CAI parent-report Child Questionnaire for younger children.

Literacy Activities (ages 14 and 15)
These questions ask 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 the 16- to 25-year-olds in the CAI questionnaire.

Health (ages 14 to 17)
Youth are asked to report on their height and weight, symptoms of stress, 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.

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

Questions on puberty are asked as it is an important marker of physical development. These questions ask youth about key physiological indicators and their perceptions of their own puberty. They were provided by Dr. Richard Tremblay from the University of Montreal and based on the Pubertal Development Scale (PDS) by Petersen and Crockett.

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 (ages 14 to 15)
Youth are asked about work during the school year and about work last summer. They are asked questions about their job(s), such as hours worked and whether work reduces the amount of time they spend studying.

Dating (ages 14 to 17)

This section asks youth about their experiences with a boyfriend/girlfriend and their sexual activity. 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, it was determined that it was important to assess three parenting behaviours using a scale on the Child Questionnaire. The scale measures 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 order individuals. For example, in the case of children with higher than average Positive interaction scores on the Parenting scales, the person most knowledgeable (PMK) reported having more positive encounters with the child, e.g., laughed with them more and praised them more. The score for a particular factor is usually based on a series of items, as a single item usually cannot measure the construct or factor with adequate precision.

During the development of the NLSCY, when considering what scales should be used to measure a particular concept, scales were selected that had been used in other studies. In this way, the psychometric properties of the measures produced by each scale were associated with reputable 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. Second, scale scores were calculated based on this factor structure. Third, 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 that 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 used 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 found in the literature was expected. In practice, however, some scales showed a different number of factors because, in some cases, factors combined, whereas 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 and multiplied by the child's normalized survey weight. An individual's statistical weight is normalized by dividing his/her weight (_WTCW01C)[6] 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 0 (zero). A score of 0 indicates that the child has no problems for all factors in the Behaviour scale except for the prosocial factor, where a score of 0 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 totalling the values of the items that made up that factor (including imputed values). If too many of the values of any items included in the factor were unreported, the score was set to “Missing”. A value was missing if the parent refused to answer or did not know the answer to the item.

6. 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 HPRCS01 in Cycle 8 begins with the letter "H" on the microdata file and is referred to here as _PRCS01.

9.2.2 Data transformation using optimal scaling

Factor analysis requires that the data have the property of interval or ratio data, whereby the distance between each answer category of the question is 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 for the scales used in the NLSCY.

Therefore, 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 et al. (Young 1981), which is a variant of Fisher's optimal scaling technique. The method is presented as a means of transforming nominal or ordinal data to data that are expressed at the interval or ratio level so that statistical techniques, which are appropriately applied only to interval and ratio data, may be used.

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, values were rescaled 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 Questionnaire was the ineffective parenting factor. 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?
_PRCgQ13	How often do you get angry when you punish your child?
_PRCgQ14	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 the 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 0 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 0 represents the absence of a problem. However, exceptions to this 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 value for this item were reversed, i.e., “All the time” was scored as 0 and “More than half the time” as 1, and on the other end of the scale, “Never” was scored as 4. In the documentation for each scale, any item that was reversed for the scoring algorithm because of a negative loading is indicated.

9.3.4 Non-response codes

When the score was being calculated for each factor, it is possible that one or more of the items making up the score had a non-response code (“Don't know,” “Refusal” or “Not stated”). If any of the items had a non-response 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 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 that 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 score for the Cronbach's Alpha for each factor 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 scale

Objectives and overview
The Depression scale was administered to the PMK as part of the Parent Questionnaire. Survey questions for this scale (_DPPQ12A to _DPPQ12L) are a shorter version of the Depression 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 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 scale was reduced to 12 questions by Dr. M. Boyle of the Chedoke-McMaster Hospital, McMaster University.

This 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. Introducing policies in this area could make a difference.

Overview of the results for the PMK Depression Scale for 0- to 15-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
PMK depression (_DPPS01)	_DPPQ12A, _DPPQ12B, _DPPQ12C, _DPPQ12D, _DPPQ12E, _DPPQ12F*, _DPPQ12G, _DPPQ12H*, _DPPQ12I, _DPPQ12J*, _DPPQ12K, _DPPQ12L	3,893	0 to 1	0.818
		4,153	2 to 3	0.803
		3,888	4 to 5	0.821
		3,332	6 to 7	0.850
		2,430	14 to 15	0.853

**Indicates that the values have been reversed.*

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.5.2 Family Functioning scale

Objectives and overview
Questions related to family functioning, i.e., _FNHhQ1A to _FNHhQ1L, 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 _FNHhQ1M, 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. 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 the PMK 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 0, 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 (_FNHhQ1A, _FNHhQ1C, _FNHhQ1E, _FNHhQ1G, _FNHhQ1I, and _FNHhQ1K). The total score (_FNHhS01) may therefore vary between 0 and 36, a high score indicating family dysfunction.

Overview of the results for the Family Functioning Scale for 0- to 15-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Family functioning (_FNHhS01)	_FNHhQ1A*, _FNHhQ1B, _FNHhQ1C*, _FNHhQ1D, _FNHhQ1E*, _FNHhQ1F, _FNHhQ1G*, _FNHhQ1H, _FNHhQ1I*, _FNHhQ1J, _FNHhQ1K*, _FNHhQ1L	3,853	0 to 1	0.914
		4,127	2 to 3	0.910
		3,879	4 to 5	0.916
		3,314	6 to 7	0.907
		2,410	14 to 15	0.917

**Indicates that the values have been reversed.*
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.5.3 Neighbourhood Safety scale

Objectives and overview
The objective of the scale is to gather information on the respondents' satisfaction with their 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 _SFHhQ01, _SFHhQ02 and _SFHhQ5A to _SFHhQ6E 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 _SFHhQ03 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 eight cycles of the NLSCY. In Cycle 1, three scales were created: neighbourhood safety (_SFHhQ5A to _SFHhQ5C), neighbours (_SFHhQ6A to _SFHhQ6E) 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 _SFHhQ5A to _SFHhQ5C and without questions ASFHQ07A to ASFHQ07F. Also, the questions that made up the Neighbourhood Problems scale in Cycle 1 (ASFHQ07A to ASFHQ07F) were dropped after Cycle 4.

Overview of the results for the Neighbourhood Safety Scales for 0- to 15-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Neighbourhood safety (_SFHhS5)	_SFHhQ5A*, _SFHhQ5B*, _SFHhQ5C*	3,885	0 to 1	0.732
		4,151	2 to 3	0.712
		3,865	4 to 5	0.701
		3,325	6 to 7	0.735
		2,423	14 to 15	0.743
Neighbours (_SFHhS6)	_SFHhQ6A*, _SFHhQ6B*, _SFHhQ6C*, _SFHhQ6D*, _SFHhQ6E*	3,225	0 to 1	0.905
		3,540	2 to 3	0.905
		3,367	4 to 5	0.913
		2,978	6 to 7	0.913
		2,153	14 to 15	0.899

**Indicates that the values have been reversed.*
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.5.4 Social Support scale

Objectives and overview
This section is asked of the PMK with children or 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. Because of 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), 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, because of 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 eight 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 because of 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 subscale (items _SPHh01H and _SPHh01I) 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, which is 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."

Overview of the results for the Social Support Scale for 0- to 15-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Social support (_SPHhS01)	_SPHh01A*, _SPHh01B, _SPHh01C, _SPHh01D*, _SPHh01E*, _SPHh01F, _SPHh01H, _SPHh01I*	3,913	0 to 1	0.895
		4,164	2 to 3	0.898
		3,893	4 to 5	0.906
		3,337	6 to 7	0.909
		2,426	14 to 15	0.906

**Indicates that the values have been reversed.*
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.5.5 Behaviour scales

Objectives
The objective of the Behaviour scales is to assess aspects of the behaviour of children 2 to 11 years of age. In Cycle 8, there were no 8- to 11-year-olds sampled, consequently these questions were asked of children 2 to 7 years of age.

Separation anxiety (aged 2 to 3 years)
Includes items _BECQ6CC, _BEC6DD1, _BEC8LL1, _BEC8PP1 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 4 to 7 years)
Includes items _BECQ6AA, _BECQ6FF, _BECd6JJ and _BECQ6NN from the Ontario Child Health Study (OCHS).

Hyperactivity (aged 2 to 7)
Includes items _BECQ6B, _BECQ6I, _BECQ6P and _BECQ6W from the OCHS and _BECQ6QQ and _BEDQ8HH 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 7 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 4 to 7 years)
Includes items _BECQ6J, _BECQ6R, _BECQ6Z, _BECQ6LL and _BECQ6TT from Lagerspetz, Bjorngvist and Peltonen of Finland.

Physical aggression (aged 2 to 7 years)
Includes items _BECQ6X from the Montreal Longitudinal Survey and _BECQ6G, _BECQ6AA and _BECQ6NN from the OCHS.

Inattention (Aged 2 to 7 years)
Includes items _BECQ6P from the OCHS and _BECQ6QQ from the Montreal Longitudinal Survey.

Prosocial behaviour (Aged 6 to 7 years)
Includes items _BECQ6A, _BECQ6H, _BECQ6M, _BECQ6GG and _BECQ6OO from the OCHS and _BECQ6D, _BECQ6U, _BECQ6BB, _BECQ6SS and _BECQ6UU 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- to 5-year-olds were excluded from this scale and were asked the questions in the Positive behaviour section instead.

Overview of the results for the Behaviour Scales for 2- to 3-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Hyperactivity – inattention (_BECeS01)	_BECQ6B, _BECQ6I, _BECQ6P, _BECQ6W, _BECQ6QQ, _BEDQ8HH	4,292	2 to 3	0.727
Emotional disorder – anxiety (_BECdS03)	_BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM, _BECQ6RR	4,316	2 to 3	0.638
Physical aggression – opposition (_BECS04)	_BECQ6G, _BECQ6X, _BECQ6NN, _BECQ6R1, _BECQ8E1, _BECQ8T1, _BECQ8Z1, _BECQ8J1	4,277	2 to 3	0.723
Separation anxiety (_BECS05)	_BECQ6CC, _BEC6DD1, _BEC8LL1, _BEC8PP1, _BEC8TT1	4,317	2 to 3	0.579

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Overview of the results for the Behaviour Scales for 4- to 7-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Hyperactivity – Inattention (_BECdS06)	_BECQ6B, _BECQ6I, _BECQ6P, _BECQ6S, _BECQ6W, _BECQ6QQ, _BEDQ8HH	4,059 3,412	4 to 5 6 to 7	0.809 0.817
Emotional disorder – Anxiety (_BECdS08)	_BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM, _BECQ6RR, _BECQ6CC	4,072 3,422	4 to 5 6 to 7	0.691 0.737
Conduct disorder – Physical aggression (_BECdS09)	_BECQ6G, _BECQ6X, _BECQ6AA, _BECQ6FF, _BECd6JJ, _BECQ6NN	4,067 3,425	4 to 5 6 to 7	0.774 0.735
Indirect aggression (_BECS10)	_BECQ6J, _BECQ6R, _BECQ6Z, _BECQ6LL, _BECQ6TT	3,973 3,217	4 to 5 6 to 7	0.632 0.696
Prosocial behaviour (_BECdS07)	_BECQ6A, _BECQ6D, _BECQ6H, _BECQ6M, _BECQ6U, _BECQ6BB, _BECQ6GG, _BECQ6OO, _BECQ6SS, _BECQ6UU	3,138	6 to 7	0.831

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.5.6 Motor and Social Development scale

Motor and Social Development section
The Motor and Social Development section of the Child 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). 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 or less, received a standardized score based on Cycle 1 data and a standardized score based on the Cycle 8 data.

Standardized scores based on Cycle 8 norms
Each child aged 4 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 were found and were used to produce a normalized score with a mean of 0 and a standard deviation of 1. This score was adjusted such that the mean MSD score was 100 and the standard deviation was 15. Therefore, children who are 4 months old have an average MSD score of 100, children who are aged 5 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 in their age group. Therefore the average of the MSD scores on the data file by age in months may not be exactly 100. Using the standardized score (_MSCdS03) makes it possible to compare scores of children across the 4- to 47-month-old age group, without having to control for age.

This score was not calculated for children aged 0 to 3 months old, as there were not enough respondent children (in this age group) 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 old. 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 8 data.

Overall there are no major differences between the scores found using Cycle 8 norms and the scores found using Cycle 1 norms. The score calculated using Cycle 1 norms should be used to compare scores over cycles; it is available for all cycles of data.

The Motor and Social Development scale questions have remained unchanged throughout the eight 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.

9.5.7 Parenting scales

Objectives and overview
The objective of this set of questions is to measure certain parenting practices. Specifically, two scales were used. The first was divided into three sub-scales designed to measure the positive interaction, ineffectiveness and consistency of the parenting of the child. The second scale was designed to measure rational parental practices.

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

Questions _PRCQ01 to _PRCQ18 and _PRCQ21 to _PRCQ24 on positive interaction, ineffectiveness and 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.

Overview of the results for the Parenting Scales for 0- to 1-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Positive interaction (_PRCS01)	_PRCQ01, _PRCQ02, _PRCQ03, _PRCQ06, _PRCQ07	3,945	0 to 1	0.622
Ineffective parenting (_PRCS02)	_PRCQ04, _PRCQ05	3,954	0 to 1	0.343

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Overview of the results for the Parenting Scales for 2- to 7-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Positive interaction (_PRCS03)	_PRCQ01, _PRCQ02, _PRCQ03, _PRCQ06, _PRCQ07	4,267	2 to 3	0.679
		4,012	4 to 5	0.677
		3,396	6 to 7	0.682
Ineffective parenting (_PRCgS04)	_PRCQ04, _PRCQ08*, _PRCQ09, _PRCgQ13, _PRCgQ14, _PRCQ15, _PRCQ18	4,225	2 to 3	0.607
		3,965	4 to 5	0.630
		3,350	6 to 7	0.672
Consistent parenting (_PRCgS05)	_PRCe10, _PRCe11, _PRCgQ12*, _PRCgQ16*, _PRCgQ17*	4,240	2 to 3	0.723
		3,989	4 to 5	0.650
		3,378	6 to 7	0.520
Rational parenting (_PRCS06)	_PRCQ21, _PRCQ22*, _PRCQ23, _PRCQ24*	4,243	2 to 3	0.562
		3,989	4 to 5	0.528
		3,383	6 to 7	0.568

**Indicates that the values have been reversed.*

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Overview of the results for the Parenting Scale for 14- to 15-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Conflict resolution (_PRCbS09)	_PRCb30A*, _PRCb30B, _PRCb30C, _PRCb30D, _PRCb30E, _PRCb30F, _PRCb30G, _PRCb30H*	2,434	14 to 15	0.746

**Indicates that the values have been reversed.*

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.5.8 Ages and Stages scale

The Ages and Stages component was generated for all children 3 to 47 months with an effective age of 3 years or less. 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

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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-complete 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.

Overview of the results for the youth-reported Depression Scale for 16- to 25-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Depression (_HTCbS1B)	_FBCd10A, _FBCd10B, _FBCd10C, _FBCd10D, _FBCd10E, _FBCd10F*, _FBCd10G, _FBCd10H*, _FBCd10I, _FBCd10J*, _FBCd10K, _FBCd10L	1,358	16 to 17	0.833
Depression (_HTYfS01)	_ HTYf14A, _HTYf14B, _HTYf14C, _HTYf14D, _HTYf14E, _HTYf14F*, _HTYf14G, _HTYf14H* _HTYf14J, _HTYf14K*, _HTYf14M, _HTYf14N	1,622	18 to 19	0.799
		1,354	20 to 21	0.805
		1,464	22 to 23	0.816
		1,460	24 to 25	0.832

**Indicates that the values have been reversed.*

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.6.2 Neighbourhood Structure scale

Objectives and overview

The objective of this scale is to gather information on the respondents' satisfaction with their 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.

Overview of the results for the Neighbourhood Structure Scale for 16- to 17-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Neighbourhood structure (_ACYdS01)	_ACYd13A*, _ACYd13B*, _ACYd13C*, _ACYd13D*, _ACYd13F*, _ACYd13G*	1,495	16 to 17	0.809

**Indicates that the values have been reversed*

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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 aged 14 to 17 in the Self-complete paper questionnaire and of youth aged 18 to 19 and 22 to 23 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.

Overview of the results for the General Self-image Scale for 14- to 23-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
General self-image (_AMcS02)	_AMCQ01A, _AMCQ01B, _AMCQ01C, _AMCQ01D	2,091	14 to 15	0.830
		1,373	16 to 17	0.817
General self-image (_AMYfS01)	_AMYfQ01, _AMYfQ02, _AMYfQ03, _AMYfQ04	1,624	18 to 19	0.764
General self-image (_AMYfS01)	_AMYfQ01, _AMYfQ02, _AMYfQ03, _AMYfQ04	1,465	22 to 23	0.778

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.6.4 Emotional Quotient scale

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. Second, they address the respondent's social, personal, and emotional 'abilities,' as opposed to their behaviours.

The EQ-i:YV measure was asked of youth aged 14 to 17, and EQ-i was asked of youth 20 to 21 and 24 to 25. Youth 14 to 17 were asked in the Self-complete 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 0. 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 EQ continuum, whereas the highest scores represent the positive end. For example, 33 on total EQ would mean that the individual is much more emotionally intelligent that 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. Although each standard deviation is set at 15, the deviations are set at 10 points around the mean values to differentiate 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 (underdeveloped emotional/social capacity)
70 to 79	Very low (extremely underdeveloped emotional/social capacity)
Under 70	Markedly low (atypically impaired emotional/social capacity)

Overview of the results for the EQ-i:YV scales for 14- to 17-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Intrapersonal (_EQYeS06)	_AMCe25A, _AMCe25F, _AMCe25K	2,089 1,372	14 to 15 16 to 17	0.847 0.837
Interpersonal (_EQYeS07)	_AMCe25B, _AMCe25G, _AMCe25L	2,082 1,370	14 to 15 16 to 17	0.643 0.663
Stress management (_EQYeS08)	_AMCe25C*, _AMCe25H*, _AMCe25M*	2,078 1,366	14 to 15 16 to 17	0.670 0.708
Adaptability (_EQYeS09)	_AMCe25D, _AMCe25I, _AMCe25N	2,085 1,371	14 to 15 16 to 17	0.773 0.782
General mood (_EQYeS10)	_AMCe25E, _AMCe25J, _AMCe25O	2,082 1,370	14 to 15 16 to 17	0.654 0.646
Emotional Quotient (EQ4) (_EQYeS04)	_AMCe25A, _AMCe25F, _AMCe25K, _AMCe25B, _AMCe25G, _AMCe25L, _AMCe25C*, _AMCe25H*, _AMCe25M*, _AMCe25D, _AMCe25I, _AMCe25N	2,052 1,349	14 to 15 16 to 17	0.697 0.716
Emotional Quotient (EQ5) (_EQYeS05)	_AMCe25A, _AMCe25F, _AMCe25K, _AMCe25B, _AMCe25G, _AMCe25L, _AMCe25C*, _AMCe25H*, _AMCe25M*, _AMCe25D, _AMCe25I, _AMCe25N, _AMCe25E, _AMCe25J, _AMCe25O	2,046 1,343	14 to 15 16 to 17	0.774 0.776

**Indicates that the values have been reversed.*
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Overview of the results for the EQ-i: scales for 20- to 21-year-olds and 24- to 25-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Interpersonal (_EQYfS11)	_EQYfQ01, _EQYfQ06, _EQYfQ11, _EQYfQ16	1,358 1,465	20 to 21 24 to 25	0.691 0.712
Intrapersonal (_EQYfS12)	_EQYfQ02*, _EQYfQ07*, _EQYfQ12*, _EQYfQ17*	1,341 1,442	20 to 21 24 to 25	0.590 0.589
Stress management (_EQYfS13)	_EQYfQ03*, _EQYfQ08*, _EQYfQ13*, _EQYfQ18*	1,358 1,460	20 to 21 24 to 25	0.700 0.702
Adaptability (_EQYfS14)	_EQYfQ04, _EQYfQ09, _EQYfQ14, _EQYfQ19	1,352 1,459	20 to 21 24 to 25	0.639 0.703
General mood (_EQYfS15)	_EQYfQ05, _EQYfQ10, _EQYfQ15, _EQYfQ20	1,355 1,461	20 to 21 24 to 25	0.650 0.707
Emotional Quotient (EQ4) (_EQYfS16)	_EQYfQ02*, _EQYfQ07*, _EQYfQ12*, _EQYfQ17*, _EQYfQ01, _EQYfQ06, _EQYfQ11, _EQYfQ16, _EQYfQ03*, _EQYfQ08*, _EQYfQ13*, _EQYfQ18*, _EQYfQ04, _EQYfQ09, _EQYfQ14, _EQYfQ19	1,334 1,435	20 to 21 24 to 25	0.792 0.787
Emotional Quotient (EQ5) (_EQYfS17)	_EQYfQ02*, _EQYfQ07*, _EQYfQ12*, _EQYfQ17*, _EQYfQ01, _EQYfQ06, _EQYfQ11, _EQYfQ16, _EQYfQ03*, _EQYfQ08*, _EQYfQ13*, _EQYfQ18*, _EQYfQ04, _EQYfQ09, _EQYfQ14, _EQYfQ19, _EQYfQ05, _EQYfQ10, _EQYfQ15, _EQYfQ20	1,333 1,434	20 to 21 24 to 25	0.838 0.843

**Indicates that the values have been reversed.*
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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 whereas these and the 15 subscales appear in the EQ-i. The subscales are bulleted below under each of the composite scales, as follows:

1) **Intrapersonal competencies – Self-awareness and self-expression**
   These competencies include the following subcomponents that govern our ability to be aware of ourselves, to understand our strengths and weaknesses, and to express our thoughts and feelings non-destructively.
   - 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 non-destructively.
- 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 subcomponents 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 subcomponents 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 subcomponents 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 changes 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 subcomponents 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, Reuven. 2004. "The Bar-On Emotional Quotient Inventory (EQ-i): Rationale, description, and summary of psychometric properties." *Measurement of emotional intelligence*: *Common ground and controversy*. Glenn Geher (ed.). Nova Science Publishers. Hauppauge, New York. pp. 111-142.

2) BarOn, Reuven, and Parker, James D.A. 2000. *Emotional Quotient Inventory: Youth Version* (*EQ-i:YV*): *Technical manual*. Multi-Health Systems. Toronto.

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. Because of 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.

Overview of the results for the Social Support Scale for 18- to 19-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (Standardized)
Social support (_SPYeS01)	_SPYeQ1A*, _SPYeQ1B, _SPYeQ1C, _SPYeQ1D*, _SPYeQ1E*, _SPYeQ1F, _SPYeQ1G, _SPYeQ1H*	1,627	18 to 19	0.840

**Indicates that the values have been reversed.*

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.6.6 Friends scale

The Friends scale is intended to measure how well youth feel they get along with their 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.

Overview of the results for the Friends Scale for 14- to 17-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Friends (_FFcS01)	_FFCQ01, _FFCQ02,	2,059	14 to 15	0.845
	_FFCQ03, _FFCQ04	1,365	16 to 17	0.865

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.6.7 My Parents and Me scales

The Parenting scales are intended to complement the Parenting section in the parent-reported Child Questionnaire by gathering information from children regarding their perception of their relationships with parents. For the Self-complete 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.

Overview of the results for the Parent and Me Scales for 14- to 15-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Parental nurturance (_PMCcS1)	_PMCcQ1A, _PMCcQ1D, _PMCcQ1K, _PMCcQ1H, _PMCcQ1I, _PMCcQ1M, _PMCcQ1Q	2,025	14 to 15	0.911
Parental rejection (_PMCbS2B)	_PMCcQ1C, _PMCcQ1G, _PMCcQ1J, _PMCcQ1L, _PMCcQ1O, _PMCcQ1P, _PMCcQ1R	1,997	14 to 15	0.767
Parental monitoring (_PMCcS3)	_PMCcQ1B, _PMCcQ1F, _PMCcQ1N, _PMCcQ1E, _PMCdQ1T	2,049	14 to 15	0.415

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.6.8 Conflict Resolution scale

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 their mothers and the other score refers to the relationship between the youth and their fathers. A high score indicates an elevated number of disagreements between the youth and their parents.

Overview of the results for the Conflict Resolution Scales for 16- to 17-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Conflict resolution – Mother (_PMCdS4)	_PMCdQ6C*, _PMCdQ6D, _PMCdQ6E, _PMCdQ6F, _PMCdQ6G, _PMCdQ6H, _PMCdQ6I, _PMCdQ6J*, _PMCdQ6K, _PMCdQ6L	1,326	16 to 17	0.840
Conflict resolution – Father (_PMCdS5)	_PMCdQ9C*, _PMCdQ9D, _PMCdQ9E, _PMCdQ9F, _PMCdQ9G, _PMCdQ9H, _PMCdQ9I, _PMCdQ9J*, _PMCdQ9K, _PMCdQ9L	1,282	16 to 17	0.837

**Indicates that the values have been reversed.*
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.6.9 Behaviour scales

This section replicates the behaviour checklist included in the Child Questionnaire completed by the PMK in Cycle 8. All youth aged 14 to 15 answer these questions in the Self-complete questionnaire. It is intended to provide indicators of the following behaviours: conduct disorder, hyperactivity, inattention, physical aggression, indirect aggression, emotional disorder, anxiety, prosocial behaviours and behaviours related to property offences.

Overview of the results for the Behaviour Scales for 14- to 15-year-olds

Score	Items included	Sample size	Universe (age in years)	Cronbach's Alpha (standardized)
Indirect aggression (_FBcS01)	_FBCQ01J, _FBCQ01R, _FBCQ01Z, _FBCQ1LL, _FBCQ1TT	2,070	14 to 15	0.702
Emotional disorder – Anxiety (_FBCdS02)	_FBCQ01F, _FBCQ01K, _FBCQ01Q, _FBCc01V, _FBCQ1CC, _FBCQ1MM, _FBCQ1RR	2,069	14 to 15	0.810
Conduct disorder – Physical aggression (_FBcS03)	_FBCQ01G, _FBCd01X, _FBCQ1AA, _FBCQ1FF, _FBCQ1JJ, _FBCQ1NN	2,068	14 to 15	0.763
Hyperactivity – Inattention (_FBCdS04)	_FBCc01B, _FBCQ01I, _FBCQ01P, _FBCQ01S, _FBCQ01W, _FBCQ1HH, _FBCQ1QQ	2,051	14 to 15	0.783
Prosocial behaviour (_FBcS05)	_FBCQ01A, _FBCQ01D, _FBCQ01H, _FBCQ01M, _FBCQ01U, _FBCQ1BB, _FBCQ1GG, _FBCQ1OO, _FBCQ1SS, _FBCc1UU	2,047	14 to 15	0.871
Property offences (_FBcS07)	_FBCQ01C, _FBCQ01E, _FBCQ01L, _FBCQ01T, _FBCQ1DD, _FBCQ1PP	2,060	14 to 15	0.701

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.7 Summary of the Cycle 8 scales

9.7.1 Parent-reported scales

Score variable	Scale name	Universe
_DPPS01	Depression	PMK of children 0 to 15 years
_FNHhS01	Family Functioning	PMK of children 0 to 15 years
_SFHhS5	Neighbourhood Safety	PMK of children 0 to 15 years
_SFHhS6	Neighbours	PMK of children 0 to 15 years
_SPHhS01	Social Support	PMK of children 0 to 15 years
_BECeS01	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 7 years
_BECdS07	Prosocial Behaviour	PMK of children 6 to 7 years
_BECdS08	Emotional Disorder – Anxiety	PMK of children 4 to 7 years
_BECdS09	Conduct Disorder – Physical Aggression	PMK of children 4 to 7 years
_BECS10	Indirect Aggression	PMK of children 4 to 7 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 8 norms	PMK of children 4 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 7 years
_PRCgS04	Ineffective Parenting	PMK of children 2 to 7 years
_PRCgS05	Consistent Parenting	PMK of children 2 to 7 years
_PRCS06	Rational Parenting	PMK of children 2 to 7 years
_PRCbS09	Conflict Resolution	PMK of children 14 to 15 years

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.7.2 Self-complete scales (reported by child or youth)

Score variable	Scale name	Universe
_FFcS01	Friends	Children/Youth 14 to 17 years
_AMcS02	General Self-image	Children/Youth 14 to 17 years
_FBcS01	Indirect Aggression	Children/Youth 14 to 15 years
_FBCdS02	Emotional Disorder – Anxiety	Children/Youth 14 to 15 years
_FBcS03	Physical Aggression – Conduct Disorder	Children/Youth 14 to 15 years
_FBCdS04	Hyperactivity – Inattention	Children/Youth 14 to 15 years
_FBcS05	Prosocial Behaviour	Children/Youth 14 to 15 years
_FBcS07	Property Offences	Children/Youth 14 to 15 years
_PMCcS1	Parental Nurturance	Children/Youth 14 to 15 years
_PMCbS2b	Parental Rejection	Children/Youth 14 to 15 years
_PMCcS3	Parental Monitoring	Children/Youth 14 to 15 years
_EQYeS04	Emotional Quotient (4 factors)	Children/Youth 14 to 17 years
_EQYeS05	Emotional Quotient (5 factors)	Children/Youth 14 to 17 years
_EQYeS06	Intrapersonal Skills	Children/Youth 14 to 17 years
_EQYeS07	Interpersonal Skills	Children/Youth 14 to 17 years
_EQYeS08	Stress Management Skills	Children/Youth 14 to 17 years
_EQYeS09	Adaptability Skills	Children/Youth 14 to 17 years
_EQYeS10	General Mood	Children/Youth 14 to 17 years
_PMCdS4	Conflict Resolution – Mother	Youth 16 to 17 years
_PMCdS5	Conflict Resolution – Father	Youth 16 to 17 years
_HTCbS1B	Depression	Youth 16 to 17 years

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

9.7.3 Youth-reported scales (self-reported)

Score variable	Scale name	Universe
_ACYdS01	Neighbourhood Structure	Youth 16 to 17 years
_HTYfS01	Depression	Youth 18 to 25 years
_SPYeS01	Social Support	Youth 18 to 19 years
_AMYfS01	General Self-image	Youth 18 to 23 years
_EQYfS11	Interpersonal	Youth 20 to 21 and 24 to 25 years
_EQYfS12	Intrapersonal	Youth 20 to 21 and 24 to 25 years
_EQYfS13	Stress Management	Youth 20 to 21 and 24 to 25 years
_EQYfS14	Adaptability	Youth 20 to 21 and 24 to 25 years
_EQYfS15	General mood	Youth 20 to 21 and 24 to 25 years
_EQYfS16	Emotional Quotient (EQ4)	Youth 20 to 21 and 24 to 25 years
_EQYfS17	Emotional Quotient (EQ5)	Youth 20 to 21 and 24 to 25 years

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

10.0 Imputation

The definition of a respondent is a child or youth, who has completed at least one of the Child, Youth or Adult components. Among these respondents, there exist many cases of partial non-response. 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 Income 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 fifth or sixth character of the variable name. For example, the name of the imputation flag for PMK income (HINPeD04) is HINPeI03.

10.1 Household income imputation

Several income questions were asked during the NLSCY household interview. Information on income, broken down into three sources, was collected 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. The same three income sources were asked with respect to all other members of the household at least 15 years of age; these amounts were not reported for each such member individually but for all other members, as a single unit. Information on income, broken down into four sources, was also asked. Those four income sources are Child Tax Benefit/National Child Benefit, social assistance, child and spousal support, and all other sources. The total household income represents the sum of these 13 sources of income. .

Income is a sensitive topic. As a result, some respondents refused to provide answers to the detailed income questions. Among those who refused, some respondents did provide estimates of their total household income or an estimate of their income using ranges. For those who provided answers to the detailed income questions, 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 non-response 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 older and those that only have youth aged 16 and 17 who live independently. Of the 19,619 eligible households, at least one income variable was imputed for a total of 4,488 households. The 19,619 eligible households actually correspond to 20,305 children or youth who are split in two files depending on their effective age, as follows: 1) the longitudinal file for children or youth part of the original cohort selected in Cycle 1 who are aged from 14 to 17 years in Cycle 8 and 2) the early childhood development (ECD) file for children or youth who are part of the birth cohorts selected in Cycles 5, 6, 7 and 8 and who are aged from 0 to 7 years in Cycle 8. The longitudinal file contains 4,247 children or youth of which 801 were imputed, whereas the ECD file contains 16,058 children of which 3,822 were imputed. Overall, 4,623 of the 20,305 children or youth were imputed. At Cycle 8, it was decided to make the income section non-proxy; this means that if someone other than the PMK was answering the questionnaire then the whole income section was skipped. This decision created 254 non-responses to the income section. We were able to find 160 of them in the previous cycle, but for 94 of them we had no income information to help with imputation.

The most critical piece of information in the Adult Income section is the total household income. Our imputation strategy for Cycle 8 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, as follows: 1) households that provided an estimated income, 2) households that responded in Cycle 7, 3) households that reported their income in ranges and 4) remaining households (including households missing only income sources for non-PMK and non-spouse members of the household aged 15 and over).

1) 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 13 sources of income were imputed simultaneously 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 13 sources of income may vary slightly from the provided estimated income.

2) Imputation of households that responded in Cycle 7
To preserve longitudinal coherence through time, the imputation of the total income of households that responded in Cycle 7 was performed by nearest neighbour trend imputation, excluding from the donor pool households with extreme income trends from Cycle 7 to Cycle 8. When a recipient household reported its total income in ranges, we ensured that the imputed total income respected the specified range. Similar to 1) above, the donor household's distribution of income sources was used to impute the 13 sources of income for the recipient household (donor ratio imputation).

3) 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 used directly to impute the missing income sources for the recipient. The missing sources of income were imputed all at once, using the same donor household and making sure that the total household income respected the specified range.

4) Imputation of the remaining households
Little information on the total income was available for the remaining households. As in 3) above, 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 used directly to impute the missing income sources for the recipient. However, the imputation was performed in three steps this time, as follows: i) imputation of the PMK's income sources, ii) imputation of the spouse's income sources and iii) imputation of the "other household member" and household level income sources.

Please note that in Cycle 8 a portion of the households dealt with at this step were those that, when prompted for an estimate of the total household income from all sources, provided an estimate within $1,000 of the sum of their reported PMK, spouse and household income (10 sources in total). It was felt that the respondent in these households likely either did not know what these other household members earned and therefore did not venture a guess, or did not consider them to be contributing to the household income. Therefore, their estimated income was disregarded. As mentioned above, the missing sources of income were imputed simultaneously by a classic nearest neighbour approach. In all, 51 of the 903 Step 4 households were of this type.

Breakdown of imputed households, Cycle 8, by imputation step

Imputation step	Households	
	Number	%
1) Households that provided an estimated income	1,253	27.9
2) Households that responded in Cycle 7	1,734	38.6
3) Households that reported their income in ranges	598	13.3
4) Remaining households	903	20.1
Total	**4,488**	**100**

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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

Total household income flag (HINHgI03):

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 (HINPeI03 and HINSeI03):

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 (HINPI1AA, HINPI1AB, HINPI1AC, HINSI1AA, HINSI1AB, HINSI1AC, HINHI1AD, HINHI1AE, HINHI1AF, and HINHI1AG):

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, Cycle 8

Variable	Imputation flag	Overall imputation rate	Imputation rate for longitudinal file (ages 14 to 17)	Imputation rate for early childhood development file (ages 0 to 7)
		%		
HINPc1AA (PMK income from wages and salaries)	HINPI1AA	14.4	12.8	14.8
HINPc1AB (PMK income from self-employment)	HINPI1AB	11.5	10.2	11.9
HINPc1AC (PMK income from Employment Insurance benefits)	HINPI1AC	12.9	10.4	13.5
HINPeD04 (Total personal income for PMK)	HINPeI03	17.6	14.5	18.4
HINSc1AA (Spouse income from wages and salaries)	HINSI1AA	18.3[1]	15.6[1]	18.9[1]
HINSc1AB (Spouse income from self-employment)	HINSI1AB	14.3[1]	13.2[1]	14.5[1]
HINSc1AC (Spouse income from Employment Insurance benefits)	HINSI1AC	13.0[1]	11.0[1]	13.5[1]
HINSeD04 (Total personal income for spouse)	HINSeI03	20.2[1]	17.0[1]	20.9[1]
HINOg1AA (Income from wages and salaries for all other 15+ household members)	HINOI1AA	20.3[2]	15.3[2]	29.2[2]
HINOg1AB (Income from self-employment for all other 15+ household members)	HINOI1AB	12.3[2]	10.2[2]	15.8[2]
HINOg1AC (Income from Employment Insurance benefits for all other 15+ household members)	HINOI1AC	12.1[2]	9.7[2]	16.3[2]
HINHeD3P (Total personal income for all other 15+ household members)	HINHeI3P	21.1[2]	15.9[2]	30.4[2]
HINHe1AD (Household income from child tax benefits)	HINHI1AD	18.1	15.0	18.9
HINHe1AE (Household income from social assistance)	HINHI1AE	11.4	9.5	11.9
HINHe1AF (Household income from child and spousal support)	HINHI1AF	11.7	9.9	12.2
HINHe1AG (Household income from other sources)	HINHI1AG	12.0	10.3	12.4
HINHgQ03 (Total household income)	HINHgI03	22.8	18.9	23.8

1. Households where there was no spouse were not included in the calculation of the imputation rate.
2. Households where the only household members at least 15 years of age are the PMK and spouse were not included in the calculation of the imputation rate.

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

10.2 Youth income imputation, 16- to 17-year-olds

Information on income, broken down in five sources, was asked as part of the Youth component for those who were 16 to 17 years old in Cycle 8. The youths were asked their income from odd jobs, employers, scholarships or bursaries, parents and then all 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 who refused, some respondents did provide estimates of their income using ranges. Income imputation was carried out to fill in the holes resulting from partial non-response.

Imputation of income was considered for all respondents who were 16 to 17 years old and had completed the Youth component. Of the 1,746 respondent youth aged 16 to 17 years in Cycle 8, 1,614 completed the Youth component. The remaining 132 youth completed only 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 218 had at least one source of income imputed.

Imputation was carried out for each of the five sources of income. Imputation was done, for most cases, 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 is taken into consideration in the donor selection so that the imputed total income respects the specified range. When only one of the five income sources requires imputation and there is a total income provided in ranges, a plausible value randomly chosen from a uniform distribution of possible values is imputed.

The imputation flags provide information on how the imputation is done. The descriptions 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

Breakdown of imputed youth aged 16 to 17 by imputation method, Cycle 8

Imputation method[1]	Number of youth[2]	%
Donor imputation by income range	74	33.9
Plausible value imputation	103	47.3
Donor imputation	41	18.8
Total	**218**	**100.0**

1. Imputation flag HINYeI02.
2. Respondents who had at least one of their sources of income imputed.

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Imputation rates for the income variables for youth aged 16 to 17, Cycle 8

Variable	Imputation flag	Imputation rate (%)
HINYeQ1A (Income from odd jobs)	HINYeI1A	4.4
HINYdQ1B (Income from employers)	HINYeI1B	5.0
HINYeQ1E (Income from scholarships or bursaries)	HINYeI1E	1.6
HINYeQ1C (Income from parents)	HINYeI1C	8.3
HINYdQ1D (Income from other sources)	HINYeI1D	3.9
HINYeD01 (Total youth income)	HINYeI02	13.5

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

10.3 Youth income imputation, 18 years and older

Information on income was asked as part of the Youth component for those who were 18 years old and older. The youths were asked their total income amount as well as their sources of income, but not the particular amount attributed to each source. Youths could choose from the following list of income sources:

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 detailed income questions. Among those who refused, some respondents did provide estimates of their income using ranges. For those who provided answers to the detailed income questions, 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 in the holes resulting from partial non-response and to rectify, when possible, these incoherencies.

Imputation of income was considered for all respondents who were 18 years old and older. Of the 5,937 respondent youths aged 18 years and older in Cycle 8, 694 had their total personal income imputed and 20 youths had their list of income sources imputed (1 of whom did not require income imputation). Also, as in Cycle 7, the youth household income variable (HIYYgQ2B), which was asked of all youths with a spouse, was released. Of the 1,103 youths with a spouse, 191 had their household income 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 is taken into consideration in the donor selection so that the imputed total income respects the specified range.

The imputation flags provide information on how the imputation is done. The descriptions of the flag values follow.

Youth Income Sources (HIYYfQ1A to HIYYfQ1O)
Imputation flag = 0 No imputation
Imputation flag = 1 Donor imputation

Youth Personal Income (HIYYfD01)
Imputation flag = 0 No imputation
Imputation flag = 1 Donor imputation by income range
Imputation flag = 2 Donor imputation

Youth Household Income (HIYYgQ2B)
Imputation flag = 0 No imputation
Imputation flag = 1 Donor imputation

Breakdown of imputed youth aged 18 and older by imputation method, Cycle 8[1]

Imputation method[2]	Number of youth[3]	%
Donor imputation by income range	528	76.1
Donor imputation	166	23.9
Total	**694**	**100.0**

1. This table covers personal income, the only variable imputed using ranges provided by the respondent. The imputation rates for the sources and household income are presented in the next table.
2. Imputation flag HIYYfI2A.
3. Respondents who had their total personal income imputed.

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Imputation rates for the income variables for youth aged 18 and older, Cycle 8

Variable	Imputation flag	Imputation rate (%)
HIYYfQ1A to HIYYfQ1O (Income sources)	HIYYfI1	0.3
HIYYfD01 (Total personal income)	HIYYfI2A	11.7
HIYYgQ2B (Total household income)	HIYYgI2B	17.3

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

10.4 Motor and Social Development (MSD) scale imputation

The imputation of the Motor and Social Development (MSD) scale in Cycle 8 uses the same methodology as the one used in Cycles 5, 6 and 7. Prior to Cycle 5, to obtain the raw MSD score for a child (variable HMSCS01), 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. In Cycle 5, we decided we could obtain a reasonably accurate score making use of 13 or 14 valid responses and imputing the missing items. Since then, we have proceeded in the following way.

Specifically, if a child had 13 or 14 valid responses, a donor record was chosen at random from among the children having complete responses and the same response pattern to the common questions. When one item was imputed, 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 choosing a donor among the children having complete responses 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 that was to be imputed. For example, an 8-month-old child missing HMSCQ21 had potential donors aged 7 to 9 months whose PMKs were asked the same 15 questions (HMSCQ12 to HMSCQ26) and had the same pattern of "Yes" and "No" responses for HMSCQ12 to HMSCQ20 and HMSCQ22 to HMSCQ26.

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, 439 additional MSD scores were obtained by having at least one response imputed; 386 had exactly one response imputed and 53 had exactly two responses imputed. This represents 5.2% of all eligible children.

The imputation flag variables HMSCIS1A and HMSCIS1B identify which MSD questions were imputed. A value of 0 for both of these flags means that no imputation was done for the MSD questions.

11.0 Weighting and treatment of non-response

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 8 is a combination of samples selected in Cycles 1, 5, 6, 7 and 8. 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 and does not change over time; a cross-sectional population may refer to some other time period. Differences between the longitudinal and cross-sectional populations are due to births, deaths, immigration and emigration.

The NLSCY produces three sets of weights at each cycle, two longitudinal (funnel and non-funnel) and one cross-sectional. Funnel weights are assigned to longitudinal children who have responded at every cycle, while non-funnel weights are assigned to longitudinal children who responded at the most recent cycle, but not necessarily at all previous cycles.

Survey weights are calculated by taking the child's design weight and making adjustments for survey non-response and post-stratification to ensure that the final survey weights sum to known counts of children by age, sex and province. The design weight is the inverse of the probability of selection, that is, the probability that a child in the population is selected by the NLSCY sample.

This chapter explains how the survey weights are calculated and what populations they refer to.

11.1 Weights available at Cycle 8

The following describes the reference populations for the various weights produced at Cycle 8. For information on survey weights and their reference populations for previous cycles, see Chapter 5.0.

11.1.1 Longitudinal weights: funnel weights (variable HWTCWd1L) and non-funnel weights (variable HWTCW01L)

Funnel and non-funnel longitudinal weights were assigned to respondents at Cycle 8 who belonged to:

- 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 5;
- the ECD cohort of 0- to 1-year-old children selected at Cycle 6;
- the ECD cohort of 0- to 1-year-old children selected at Cycle 7.

For the ECD cohorts, the funnel weights, _wtcwd1l, were first produced at Cycle 7. Prior to Cycle 7, children were surveyed only if they responded to all previous cycles. Therefore, the longitudinal weights, _wtcw01l, received by ECD children prior to Cycle 7 could be interpreted as funnel weights. However, starting in Cycle 7, ECD children were eligible to be surveyed whether or not they responded at previous cycles. That is why, funnel and non-funnel weights, _wtcwd1l and _wtcw01l, respectively, could be computed for ECD children.

At a child's cycle of introduction, only cross-sectional weights are assigned. It is at subsequent cycles that returning children receive longitudinal weights. Thus, at Cycle 1, children in the original cohort received a cross-sectional weight, and only at subsequent cycles did they receive longitudinal weights. Similarly, in the case of ECD children who are first surveyed at age 0 to 1, at their cycle of introduction they receive a cross-sectional weight, and at subsequent cycles they receive both cross-sectional and longitudinal weights.

Also, for children in the original cohort, only those who responded at Cycle 1 receive longitudinal weights. For the ECD cohorts, longitudinal weights are only assigned to those who were first sampled at age 0 to 1, but they need not have responded at that cycle of introduction. Any ECD child who was sampled for the first time at age 2 to 5 (i.e., top-up samples at Cycles 6, 7 and 8) will only receive a cross-sectional weight since these top-ups were for cross-sectional purposes.

11.1.2 Longitudinal populations

The following describes the longitudinal populations of the original cohort and the ECD cohorts. Children belonging to a particular cohort can be identified using the variable MEMCYCLE which indicates the cycle when the children first entered the survey.

The original cohort

The longitudinal population for the original cohort is defined as children aged 0 to 11 years old as of December 31, 1994, who were living in any province at the time of Cycle 1 collection (1994/1995).

These individuals can be identified on the data files by the condition MEMCYCLE = 01.

The ECD cohort introduced in Cycle 5

The longitudinal population of the ECD cohort selected at Cycle 5 is defined as children aged 0 to 1 year old as of December 31, 2002, who were living in any province at the time of Cycle 5 collection (2002/2003).

These individuals can be identified on the data files by the condition MEMCYCLE = 05.

The ECD cohort introduced in Cycle 6

The longitudinal population of the ECD cohort selected at Cycle 6 is defined as children aged 0 to 1 year old as of December 31, 2004, who were living in any province at the time of Cycle 6 collection (2004/2005).

These individuals can be identified on the data files by the condition MEMCYCLE = 06.

The ECD cohort introduced in Cycle 7

The longitudinal population of the ECD cohort selected at Cycle 7 is defined as children aged 0 to 1 year old as of December 31, 2006, who were living in any province at the time of Cycle 7 collection (2006/2007).

These individuals can be identified on the data files by the condition MEMCYCLE = 07.

11.1.3 Cross-sectional weights (variable HWTCW01C)

At Cycle 8, cross-sectional weights were assigned to ECD respondent children who could be used to represent the Cycle 8 cross-sectional population of 0- to 7-year-olds. Since the original cohort has never been topped-up for immigrants, it should not be used for cross-sectional analyses.

At Cycle 8, the cross-sectional sample of ECD children aged 0 to 7 years old consists of the following:

- the ECD cohort of 0- to 1-year-old children selected at Cycle 8,
- returning 2- to 3-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 7,
- returning 4- to 5-year-old children from the ECD cohort of 0- to 1-year-old children selected at Cycle 6,
- a new top-up of 2- to 5-year-old children (selected from the LFS),
- returning 4- to 7-year-old children from the top-up sample of 2- to 5-year-old children selected at Cycle 7,

The cross-sectional population at Cycle 8 is defined as children aged 0 to 7 years old as of December 31, 2008, who were living in any province at the time of Cycle 8 collection (2008/2009).

11.2 Weighting method

Some details of the weighting method are described below. NLSCY users who are interested in knowing more details 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 child's design weight, where the design weight is equal to the inverse of the child's probability of selection. In this chapter, we denote the NLSCY design weight as $w_{\text{NLSCY design}}$.

All children still in the NLSCY at Cycle 8 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 the NLSCY, we further select households from the LFS and then children within the 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.

A child's final NLSCY survey weight is obtained by applying non-response and post-stratification adjustments to the NLSCY design weight.

11.2.2 First adjustment: Non-response adjustment

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

More precisely, the goal of the non-response adjustment is to inflate the NLSCY design weights of the respondents so that their non-response adjusted weights add up to the sum of the NLSCY design weights for everyone in the original sample. In other words, the non-response adjustment aims to give the NLSCY design weights of the non-respondents to the respondents in an intelligent way.

Children in the original cohort drawn at Cycle 1 have experienced eight cycles of non-response by Cycle 8. The Cycle 5 ECD cohort has been affected by non-response over four cycles (Cycles 5, 6, 7 and 8), the Cycle 6 ECD cohort has suffered three cycles of non-response, and the Cycle 7 ECD cohort have lost respondents over two cycles. The Cycle 8 ECD sample has only been subjected to one cycle of non-response. The non-response adjustment aims to adjust for all of this non-response.

To decide how to assign the weight of the non-respondents 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 8 is determined through logistic regression models, using a similar approach that was applied at Cycle 7. The model produces predicted response probabilities (values between 0 and 1) for everyone. These probabilities are divided into several RHGs using cluster analysis. The number of RHGs is determined so that the overall homogeneity of the clusters is maximized and certain constraints (size of the adjustment factor, minimum number of respondents in each RHG) are satisfied. This yields reasonable, reliable adjustment factors.

For the Cycle 8 weighting, the non-response adjustment was performed in two phases, using two different logistic models. Non-respondents were divided in two groups: the non-contact and the others (mostly refusals). With the first adjustment, the weights of the cases that we were not able to contact were given to cases for which a contact was made (respondents and other non-respondents) while with the second adjustment, the weights of the other non-respondents were given to the respondents.

Ten logistic regression models were built to model the non-response. Models were derived for the funnel and non-funnel longitudinal response of the original cohort; the funnel and non-funnel longitudinal response of the ECD children from Cycles 5, 6, and 7 and the cross-sectional response of the ECD children from Cycles 5, 6, 7 and 8. Separate models were constructed since the original cohort has suffered many more

cycles of non-response than the ECD cohorts, and there is every reason to believe that non-response behaviour varies as the number of times the individual has been interviewed increases. For each of the aforementioned non-response adjustments, two different models were constructed for the two groups of non-respondents.

11.2.3 Second adjustment: Poststratification

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 purpose of this adjustment is to ensure that the sum of the weights match known population totals. The post-stratification totals depend on the population of reference.

All final survey weights are post-stratified. For Cycle 8 cross-sectional weighting, the reference year to calculate a child's effective age is 2008. The post-stratification counts refer to January 2009, so that we have a reliable count of children of a given age as of December 31, 2008. In a similar vein, the funnel and non-funnel longitudinal weights are post-stratified to January counts following the reference year of interest (listed in 11.1.2).

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 Cycles 6, 7 and 8 differs from the method at previous cycles

The non-response adjustment methodology employed at Cycles 6, 7 and 8 differs from that used at previous cycles. At Cycle 6, we changed from segmentation modelling to logistic regression modelling. The result is fewer, more robust and discriminating RHGs. Also at Cycle 6, we designed a non-response model for the longitudinal weights that is truly longitudinal. At Cycle 8, cluster analysis was introduced to create the RHGs and the non-response adjustments were divided in two phases to account for the two types of non-respondents.

For more information about Cycle 6 weighting, see the Cycle 6 User's Guide.

11.3 Applying the weighting method

11.3.1 Non-funnel longitudinal weighting

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 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: non-response adjustment
Two sets of non-response 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"), where available, with the original cohort model having an additional co-operation score as an independent

variable (a person's co-operation is calculated as a percentage of questions answered at a previous cycle of the NLSCY). The non-response weight adjustment is calculated for each RHG, using the formula presented earlier.

For the original and ECD cohorts, the number of RHGs used was between 6 and 13. Although the ECD cohorts were modeled together, their non-response adjustments were calculated separately, for the simple reason that they each represent their own distinct reference population.

Second adjustment: post-stratification adjustment
For the original cohort, the reference population is the population of all children aged 0 to 11 years old as of December 31, 1994. For the ECD cohorts, the reference populations are listed in 11.1.2. Each group was post-stratified to the relevant age-sex-province population counts, for the January following the reference year.

11.3.2 Funnel longitudinal weighting

Definition of a funnel respondent
A funnel respondent is a longitudinal respondent at Cycle 8 (defined above) who was also a respondent at all previous cycles. The following indicates all the cycles where a particular cohort must have responded to.

Cohort	Cycles must have responded to
Original Cohort	1, 2, 3, 4, 5, 6, 7, 8
ECD Cohort from Cycle 5	5, 6, 7, 8
ECD Cohort from Cycle 6	6, 7, 8
ECD Cohort from Cycle 7	7, 8

Notice that the all-cycle respondents (i.e. the funnel respondents) are a subset of the Cycle 8 respondents.

First adjustment: non-response adjustment
Once again the non-response adjustment is based on the creation of RHGs. The non-response classes for the original cohort are created using LFS variables and a co-operation score. The non-response classes for the ECD cohorts are created using LFS variables. Using the model output, RHGs are generated. Note that these RHGs are different from the RHGs created for longitudinal weighting, as Cycle 8 respondents who did not respond in all of the previous cycles have a different non-response mechanism than respondents to all cycles. The adjustment factor is computed for each RHG.

Second adjustment: post-stratification adjustment
For the original cohort, the reference population is the population of all children aged 0 to 11 years as of December 31, 1994. For the ECD cohorts, the reference populations are listed in 11.1.2. Each group was post-stratified to the relevant age-sex-province population counts, for the January following the reference year.

11.3.3 Cross-sectional weighting

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: non-response adjustment
Children in the NLSCY, aged 0 to 7 years in 2008, were given a cross-sectional weight. Non-response models were derived using LFS variables. The estimated probabilities of response from the models were then combined to form the RHGs. Within each RHG, a non-response adjustment was then computed.

Second adjustment: post-stratification
The reference population is children aged 0 to 7 years as of December 31, 2008. The adjustment is computed for each age-sex-province combination, for January 2009 counts.

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 http://www.statcan.gc.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 Chapter 13.0.

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 non-response 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 non-response and non-response bias

In surveys, non-response results from the inability to obtain a set of measurements for a given unit in the sample. Non-response can be classified into two types, total (unit) non-response and partial (item) non-response. Unit non-response arises when none of the survey measurements for a given unit are available. Such a unit is labelled a non-respondent. Item non-response 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 non-response and Section 12.4 discusses item non-response.

Non-response is a situation that can lead to bias in the survey estimates. Biased estimates can result if non-respondents have significantly different characteristics from respondents. Both the amount of non-response and the degree to which the non-respondents would have reported different answers than the respondents affect the amount of bias in the estimate. We are unable to accurately measure what the non-respondents would have reported, but we can measure the level of non-response. 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 non-response 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 non-response. However, this practice is based on certain assumptions, and it does not guarantee that there is no bias because of non-response.

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 non-respondents 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 non-respondents may differ from the respondents in important unobserved or unknown ways.

Non-response cumulates over time. As we have fewer and fewer participants, the estimated sampling error increases, and 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 subsample 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, assume the initial sample contained 20 girls with autism in some province, yet none of these 20 responded at Cycle 8. 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 given in terms of the statistical unit (for the NLSCY, this is the child) and 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 are not part of the Cycle 8 sample sent to collection. For example, units that were hard refusals in a previous cycle are removed from the Cycle 8 sample, and their absence affects data quality. Such units are considered in the estimation phase rates, but not in the collection phase rates. Some child-level collection phase rates are given in Chapters 5.0 and 6.0. 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 non-response from influential units is more damaging to the survey estimate than non-response from less influential units. Influential units have, for example, large design weights in social surveys or a large influence on estimates in business surveys. 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 component with a set of key questions answered is considered **completed**. Note that substantial item non-response (Section 12.4) can be present within components classified as completed.

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 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 and older is considered a **respondent** if the Youth component of the survey is completed. There is no Adult component administered for youth 18 and 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. Note that 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 **non-respondent** is an in-scope child or youth who does not meet our response criteria. Non-response 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 **estimation phase response rate** is defined as the number of respondent children or youth over the number of in-scope children or youth in the initial sample.

The **collection phase response rate** is defined as the number of respondent children or youth over the number of in-scope children or youth among the units that were sent to the field for collection.

12.3.2 Cross-sectional response rates

The cross-sectional Cycle 8 sample is the aggregation of children introduced in four different cohorts, namely Cycles 5, 6, 7 and 8.

Unweighted cross-sectional response rates, Cycle 8, by sample cohort

Effective age (Years)	Initial sample		Sent to collection		Respondents	Collection phase response rate (%)	Estimation phase response rate (%)
	Sample size	In-scope	Sample size	In-scope			
0 and 1	5,479	5,463	5,479	5,463	4,106	75.2	75.2
2 and 3 returners	4,991	4,946	4,964	4,940	3,886	78.7	78.6
2 and 3 top-up	615	615	615	615	486	79.0	79.0
4 and 5 returners	4,857	4,788	4,816	4,785	3,658	76.4	76.4
4 and 5 top-up	588	587	588	587	472	80.4	80.4
6 and 7 returners	5,568	5,406	4,271	4,256	3,450	81.1	63.8
Total	**22,098**	**21,805**	**20,733**	**20,646**	**16,058**	**77.8**	**73.6**

Column definitions:

Effective age – This is based on year of birth. Those with effective age = 0 were born in 2008; those with effective age = 1 were born in 2007; etc.

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

Initial sample/In-scope – This is the count of cross-sectionally in-scope children at Cycle 8. This count forms the denominator of the estimation phase response rate. See Section 12.3.1 for a complete definition.

Sent to collection/Sample size – This is the count of children that were sent to collection at Cycle 8.

Sent to collection/In-scope – This is the count of cross-sectionally in-scope children among the children that were sent to collection at Cycle 8. This count forms the denominator of the collection phase response rate.

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

Collection phase response rate – Ratio of the number of respondents to the number of in-scope children sent to collection.

Estimation phase response rate – Ratio of the number of respondents to the number of in-scope children in the initial sample.

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

The table above illustrates the difference between the collection phase and estimation phase response rates. For example, for the returners with effective age of 6 or 7, the collection phase response rate is 81.1%. However, when one takes into account that several hundred in-scope children were dropped from the Cycle 8 sample - mainly because of non-response at Cycle 5 – the estimation phase response rate is 63.8%, substantially lower.

At Cycle 7, although we started bringing back non-respondents from previous cycles, we did not bring back those who were non-respondents at their initial cycle. At Cycle 5, there were 1,142 initial non-respondents. These were all cross-sectionally in-scope at Cycle 7, but they were not sent to the field, resulting in the large difference between the initial sample and those sent to collection, which contributed to the low estimation response rate.

By contrast, for the 4-5 returners at Cycle 8, those who did not initially respond at Cycle 6 were returned to the field. There were 852 initial non-respondents at Cycle 6 who were sent to the field at Cycle 8.

The Cycle 7 collection strategy resulted in low estimation response rates even though collection response rates were high. At Cycle 8, the two are essentially the same, with the exception of the 6-7 year old returners because initial Cycle 5 non-respondents continued to not be sent to the field.

The next table shows the estimated phase rates by province.

Unweighted cross-sectional estimation phase response rates, Cycle 8, by province

Province	In-scope sample	Respondents	Estimation phase response rate (%)
Newfoundland and Labrador	1,435	1,192	83.1
Prince Edward Island	1,171	893	76.3
Nova Scotia	1,380	1,063	77.0
New Brunswick	1,517	1,154	76.1
Quebec	2,957	2,280	77.1
Ontario	5,064	3,511	69.3
Manitoba	1,787	1,274	71.3
Saskatchewan	1,824	1,405	77.0
Alberta	2,516	1,811	72.0
British Columbia	2,154	1,475	68.5
Canada	**21,805**	**16,058**	**73.6**

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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 8, known only for respondents.

12.3.3 Longitudinal attrition rates for Cycle 1 cohort

In a longitudinal survey, the longitudinal response rate shows the 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 non-response at Cycle 1 for children introduced at Cycle 1. Specifically, the actual number of

eligible children in non-responding 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 respondents remaining in the survey relative to the number of **respondents** at Cycle 1.

Unweighted longitudinal attrition rates for Cycle 1 children, by effective age

Effective age at Cycle 8 (Years)	Number of respondents in Cycle 1	Collected data for Cycle 8		Possess longitudinal weight at Cycle 8[1]		Possess funnel weight at Cycle 8[2]	
		Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1
14 to 17[3]	6,967	4,271	61.3	4,414	63.4	3,624	52.0
18 to 25	9,936	5,937	59.8	6,073	61.1	3,848	38.7
Total	16,903	10,208	60.4	10,487	62.0	7,472	44.2

Notes: The denominator for all the percentages shown in this table is the number of responding children in Cycle 1 who were followed in Cycle 2. Note that because the sample size has been reduced from Cycle 1 to Cycle 2, not all Cycle 1 responding children are considered longitudinal.

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. As a result, it is possible to have children who have not responded in a particular cycle but have responded in the current cycle.

1. Those receiving a longitudinal weight at Cycle 8 can be classified into two categories: those with reported data and those without reported data. Children who have died or moved outside the 10 provinces belong in the second category, those without reported data. These children still have a longitudinal weight because 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 8" column and are a subset of those given a longitudinal weight.
2. This column shows the number of children who have never been considered non-respondents. That is, they have received a longitudinal weight on each cycle's master file. The majority of these children reported data in each cycle. However, a much smaller number, those who have died or moved outside the 10 provinces, retain a weight without reported data. See Chapter 11.0 for more details about the funnel weight, variable HWTCWd1L. Note that the small discrepancy between the total number of units that possess a funnel weight and the total number of units that are "respondents in all cycles" in the following table is due to the change in the concept of "age" that occurred at Cycle 4. A few units that have a funnel weight had been considered longitudinally out-of-scope for Cycle 4 only, and hence were not assigned a longitudinal weight at that Cycle.
3. For those in the first row of the table, the primary respondent is the PMK. Youth 18 and older respond for themselves.

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

The following table shows attrition throughout the different cycles. "Respondents" are those with a longitudinal weight including those without reported data. The proportion of the Cycle 1 respondents that had a longitudinal weight dropped by almost 5% – from 66.9% to 62.0% – between Cycle 7 and Cycle 8. While this drop was only 1% between Cycle 6 and Cycle 7, the drop between Cycle 7 and Cycle 8 is similar to the drop observed at each of the cycles previous to Cycle 7.

Unweighted longitudinal attrition rates for Cycle 1 children, Cycles 2 to 8, by province

Province in Cycle 1	Respondents in Cycle 1	Respondents in Cycle 2		Respondents in Cycle 3		Respondents in Cycle 4		Respondents in Cycle 5		Respondents in Cycle 6		Respondents in Cycle 7		Respondents in Cycle 8		Respondents in all cycles	
	Number	Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1	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	679	71.5	636	66.9	491	51.7
Prince Edward Island	467	443	94.9	434	92.9	392	83.9	364	77.9	339	72.6	340	72.8	305	65.3	227	48.6
Nova Scotia	1,191	1,068	89.7	1,085	91.1	988	83.0	903	75.8	839	70.4	809	67.9	748	62.8	536	45.0
New Brunswick	1,070	958	89.5	958	89.5	836	78.1	792	74.0	710	66.4	685	64.0	678	63.4	449	42.0
Quebec	3,182	2,944	92.5	2,844	89.4	2,522	79.3	2,361	74.2	2,108	66.2	2,147	67.5	2,093	65.8	1,453	45.7
Ontario	4,342	3,899	89.8	3,760	86.6	3,318	76.4	3,104	71.5	2,834	65.3	2,804	64.6	2,598	59.8	1,809	41.7
Manitoba	1,232	1,161	94.2	1,112	90.3	1,019	82.7	1,004	81.5	905	73.5	868	70.5	743	60.3	568	46.1
Saskatchewan	1,413	1,305	92.4	1,257	89.0	1,073	75.9	1,002	70.9	958	67.8	971	68.7	906	64.1	640	45.3
Alberta	1,599	1,465	91.6	1,420	88.8	1,242	77.7	1,162	72.7	1,109	69.4	1,083	67.7	968	60.5	686	42.9
British Columbia	1,457	1,333	91.5	1,282	88.0	1,143	78.4	1,076	73.9	992	68.1	923	63.3	812	55.7	596	40.9
Canada	16,903	15,468	91.5	14,997	88.7	13,310	78.7	12,523	74.1	11,483	67.9	11,309	66.9	10,487	62.0	7,455	44.1

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

12.3.4 Longitudinal response rates for children selected in Cycles 5 and 6

The estimation phase response rates given in the "2 and 3 returners", "4 and 5 returners", and "6 and 7 returners" rows in the table in section 12.3.2 above give a good indication of the longitudinal response rates for children selected in Cycle 7, 6 and 5 respectively.

12.4 Partial non-response

The previous section on total non-response dealt with the issues of representation of the sample for cross-sectional or longitudinal estimation. Whereas 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 non-response are measured in this survey and they are usually not corrected through an adjustment of survey weights.

A person may provide enough information to qualify as a respondent, but some of the questions (variables) may still be not answered, resulting in partial non-response. 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 non-response 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 non-response. Another cause of partial non-response is when the telephone portion is collected but the other components are missing.

Usually, the nature of partial non-response 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 non-response.

Item non-response 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 non-response is left uncorrected except where specifically noted by imputation flags. Item non-response is detailed in the code book with categories such as "Refusal" or "Not stated." The "Don't know" category is regarded as a non-response 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, whereas 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-complete 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 non-response 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 non-response, 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 non-response affects the data they are attempting to analyse. As in the case of total non-response, partial non-response 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 non-response, for instance, reweighting and imputation. Users are strongly encouraged to make themselves aware of the extent of the non-response in the analysis they are doing and, if appropriate, take corrective measures to compensate for the non-response. At minimum, they should detail the impact of component or item non-response in their findings. This is also discussed later in Chapter 15.0.

The following sections will explore the issue of component non-response 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. This is not the total non-response for the component. For example, for the PPVT-R the component response rate is 83.9%. The total and partial non-response rates together mean that 64.5% of the 5,375 in-scope children for which a PPVT is desired completed the test.

Further, as this guide has been written before the release of the master file, small differences may exist between what is stated here and what is 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 8, but 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.

Child component response rate

Number eligible	Number answered	Component response rate (%)
20,305	20,212	99.5

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Among survey respondents, it is exceedingly rare that the entire Child component is missing. 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 concerns children, not adults. In households with two selected children, there is a single PMK. We desire PMK information from 19,619 individuals. In terms of children, we desire 20,305 children to have PMK information.

Person most knowledgeable component response rate

Number eligible	Number answered	Component response rate (%)
20,305	19,817	97.6

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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

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 concerns children, not adults. In households with two selected children, there is one spouse answering. We desire spouse information from 16,594 individuals. In terms of children, we desire 17,141 children to have spouse information. Children living with a single parent are not eligible for the spouse component.

Spouse component response rate

Number eligible	Number answered	Component response rate (%)
17,141	16,560	96.6

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Among survey respondents, it is rare that the entire spouse component is missing. 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 and 17. Youth 18 and older must complete the youth component to be considered a response.

Youth component response rate

Number eligible	Number answered	Component response rate (%)
1,746	1,590	91.1

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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 Peabody Picture Vocabulary Test – Revised

The PPVT–R was administered to children aged 4 or 5 years. The Peabody Picture Vocabulary Test – Revised (PPVT-R) was designed to measure receptive or hearing vocabulary.

Peabody Picture Vocabulary Test – Revised response rate

Number eligible	Number answered	Component response rate (%)
4,130	3,467	83.9

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

12.4.6 Number Knowledge

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

Number Knowledge response rate

Number eligible	Answered enough to get a score	Component response rate (%)
4,130	3,485	84.4

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

12.4.7 Who Am I?

The Who Am I? assessment was administered to children aged 4 or 5 years. The purpose of the Who am I? assessment is to evaluate the developmental level of young children.

Who Am I? response rate

Number eligible	Answered enough to get a score	Component response rate (%)
4,130	3,461	83.8

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

12.4.8 Mathematics tests

The NLSCY mathematics tests are made of 20 computational questions answered in the home by respondents aged 7 to 15. The level of test (ranging from 2 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.

Mathematics tests response rate

Number eligible	Answered enough to get a score	Component response rate (%)
4,315	3,470	80.4

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

12.4.9 Problem solving

The Problem solving assessment for 16- and 17-year-olds in the NLSCY consisted of 18 questions to assess strengths in reading comprehension, problem-solving and decision-making. It also tests some mathematical skills.

Problem solving response rate

Number eligible	Answered enough to get a score	Component response rate (%)
1,770	1,413	79.8

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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.

Literacy assessment response rate

Number eligible	Answered enough to get a score	Component response rate (%)
1,635	1,240	75.8

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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.

Numeracy assessment response rate

Number eligible	Answered enough to get a score	Component response rate (%)
1,366	1,011	74.0

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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.

Self-complete component response rate

Effective age (Years)	Number eligible	Number answered	Component response rate (%)
14 to 15	2,501	2,239	89.5
16 to 17	1,746	1,502	86.0
Total	**4,247**	**3,741**	**88.1**

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

12.5 Cycle non-response

Certain longitudinal respondents do not participate in every cycle. This is cycle non-response. 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, 5, 7 and 8, but not in Cycles 2, 3, 4 and Cycle 6.

If data from every cycle are crucial, the analyst's consideration can be limited to children without cycle non-response and use the funnel longitudinal weights for this group, variable HWTCWd1L.

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 5 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 this distinction clear. 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 inches (see chart in Section 12.14.4). 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 because of 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, 0 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.

Proxy rates, Person most knowledgeable and Spouse components

Component	Eligible children	Children with proxy responses	Proxy rate (%)
	Number		
PMK	20,305	257	1.3
Spouse	17,141	11,221	65.5

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

It is rare that the spouse responds to the PMK component, but it is common for the PMK to 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

Because of 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 a highest level of education lower than was reported in the previous cycle, the highest level of education is set to the previous reported value.

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 but 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:

$$\text{BMI} = \frac{\text{weight in kilograms}}{(\text{height in metres}) \times (\text{height in metres})}$$

For children under 2 years of age, the use of BMI for analysis is not recommended, and as such it is not derived for these children. For children of at least 2 years of age, the following height and weight variables are used to derive BMI from NLSCY data:

- PMK reported for 2- to 7-year-olds: HHLCQ03B (height) and HHLCQ04A (weight) yield the respondent's BMI score, HHLCeS01.
- Self-complete for 14- to 17-year-olds: HHTCbQ01 (height) and HHTCbQ02 (weight) yield the respondent's BMI score, HHLCeS01.
- Self-completed by the youth in the youth component for 16- to 25-year-olds: HHTYeD01 (height) and HHTYeD02 (weight) yield the respondent's BMI score, HHTYeS03. Analyses of the BMI for women in this group are not recommended: The NLSCY doesn't identify the ones that are pregnant or lactating, which can have an important impact on the BMI.

By calculating a BMI score, this score can then be compared with others to see into which percentile it falls. Differing cut-offs 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 cut-offs for children, youth, and adults. Similarly, Tim Cole et al. have proposed international cut-offs 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 one of the following: 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 cut-offs are based on the person's age broken down into one-month intervals. Consequently, in processing the NLSCY data, the Age in months variable (HMMCdQ1B) was used to derive the cut-offs. The percentile ranges proposed by the CDC can potentially be used for 0- to 19-year-olds. The release name for this variable is HHLCeD03 for those aged 17 and younger and HHTYeD05 for those aged 18 and older.

More information on the CDC BMI cut-offs for children and youth can be obtained at the website http://www.cdc.gov/growthcharts.

12.14.3 Body mass index: International cut-offs

A set of international BMI cut-offs 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 cut-offs classify BMI scores into three categories: normal, overweight and obese. Similar to the CDC cut-offs, these categories are age- and sex-specific. In contrast to the CDC cut-offs, the international cut-offs 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 and Singapore). The release name for this variable is HHLCeD02 for those aged 17 or younger and HHTYeD04 for those aged 18 and older. After 17 years of age, the World Health Organization cut-offs for adults are used.

More information on the BMI cut-offs proposed by Tim Cole et al. can be obtained at the website http://www.bmj.com.

12.14.4 Body mass index: Data quality

Several issues affect the quality of the BMI scores. First, there is a higher rate of non-response for the BMI variables, as both the Height and Weight variables must contain valid responses to calculate a score. In particular, there is considerable amount of missing data on young children's heights. The matter has been investigated, and measures will be taken to improve the response rates as well as the accuracy of these measurements in the future.

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 they may 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 weight or both were 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. Although the results below apply for Cycle 6, similar behaviors have been noticed for other variables and at other cycles, including at Cycle 8.

Cycle 6: Unweighted counts of some reported heights in inches

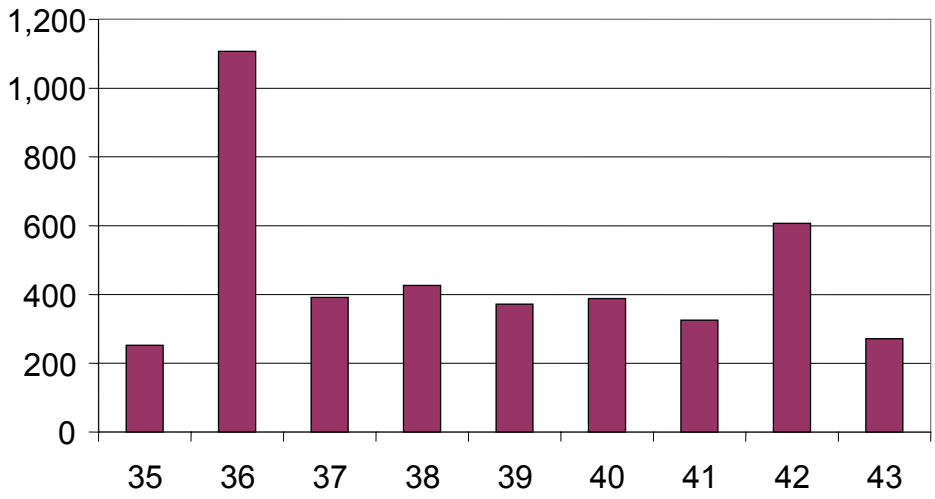

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

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 (PMK) 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, however, the PMK can change. For a given child, perhaps the mother was the PMK at Cycle 1, the father at Cycle 2 and Cycle 3, and the mother again at Cycle 4, Cycle 5, Cycle 6, Cycle 7 and Cycle 8. Observed changes in the child's characteristics may result from different people answering these questions. For example, although the child's environment is not truly different, the person answering the questions has changed and, naturally, has a different perspective. The variable HDMCD06 shows the relationship of the PMK to the child. The variables PPERSC4, PPERSC5, PPERSC6, PPERSC7 and PPERSRUK provide the ID of the PMK in Cycles 4, 5, 6, 7 and 8, respectively.

12.17 Coverage of Canadian children by the NLSCY sample

The survey population is the population represented by the sample, and about which inferences can be made when survey weights are used at analysis. For example, the survey population for the original cohort is children aged 0 to 11 as of December 31, 1994, who were living in any province at the time of Cycle 1 collection (1994/1995).

Recall that all children receive longitudinal weights (except for top-up samples) but only ECD children receive cross-sectional weights. The various survey populations covered by the NLSCY longitudinal and cross-sectional weights are described in Chapter 5.0.

NLSCY data users should be aware of certain exclusions and the uneven coverage of some Canadian children, such as immigrants, since these could affect their analyses.

12.17.1 LFS exclusions

Some children are excluded from the LFS and therefore from NLSCY samples that were drawn from the LFS. The LFS only covers the civilian, non-institutionalized population in Canada's 10 provinces. It excludes the Yukon, Nunavut and Northwest Territories, people living on Indian reserves, full-time members of the Canadian Armed Forces and inmates of institutions.

12.17.2 Coverage of recent immigrants

Since the original cohort has never been topped up for immigrants who arrived after 1994/1995, it should not be used for cross-sectional analyses. 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 recent immigrants was so great that the original cohort should not be used to make inferences about the cross-sectional populations after Cycle 4.

In the case of the ECD cohorts, some topping-up has been performed in some provinces, for some ages, and therefore some new immigrants are present, but the sample design does not specifically ensure an adequate sample of new immigrants. Therefore, caution should be exercised when using the ECD sample to make inferences about immigrants.

Specifically, before Cycle 7, the ECD cohorts were only followed for three cycles (until age 4 to 5) so, even in the absence of top-ups, it was felt that they could be used for cross-sectional purposes. At Cycles 6, 7 and 8, the NSLCY topped-up 2 to 5 year-olds in some provinces (not Ontario or Quebec), and consequently some recent immigrants were picked up with these top-ups. However, the size of the top-up samples was typically small.

The Cycle 8 ECD children therefore have some new immigrants as a result of the Cycle 8 top-up of 2 to 5 year-olds, the returning Cycle 7 top-up children who are 4 to 7 years-old, and the returning Cycle 6 top-up children who are 6 to 7 years-old.

Thus, for the ECD children present at Cycle 8, while some recent immigrant children are present, the sample design does not ensure that the sample is representative of all immigrant children aged 0 to 7.

12.17.3 Coverage by birth month

Because of the way that the NLSCY samples babies from the LFS, babies born at the end of the calendar year typically have a lower probability of selection than those born at the beginning of the year. This unequal distribution in the sample by birth month became pronounced at Cycles 6 and 7 and weight adjustments were performed: at Cycle 6, a uniform adjustment was added to the survey weights for 0 to 1 year-olds, at Cycle 7, the birth-month weight adjustment for 0 to 1 year-olds was refined. At Cycle 8, the methodology from Cycle 7 was used.

12.17.4 Coverage by birth order in the original cohort

At Cycle 2, the original cohort's sample was cut for budget purposes. 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. 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 households with at least one 0- or 1-year-old child selected from the LFS after the 1994 redesign, 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 the original cohort, 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. There is a corresponding slightly decreased proportion of children aged 2 to 11 (age at Cycle 1) with two or more younger siblings.

12.18 Data validation

For Cycles 7 and 8, an additional validation step was used to the survey process in order to check consistency of estimates over several cycles. The tool, referred to as the pre-release validation tool, automatically detected large discrepancies in estimates for consecutive cycles for all variables common to those cycles. This proved to be very helpful in identifying and correcting errors before the dissemination of the data and should improve the overall data quality of the release files.

12.19 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 that 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 results in different numerical guesses, i.e., in 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 the ability to measure mathematically the magnitude of the sampling error. By definition, a census has no sampling error because the only possible sample is the whole population (and therefore, the only estimate we get corresponds exactly to the parameter's value).

Even if sampling error can 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 is; 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 (specified through **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 matches the parameter's value. The outputs are the estimates obtained by including all possible samples. So, estimates taken individually may all be off from the parameter's value but on average be 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 the hits 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:

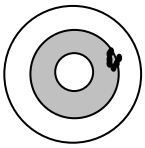

Bias and small variance

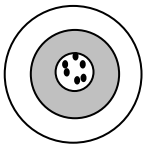

Unbiased and small variance

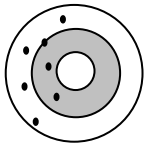

Bias and large variance

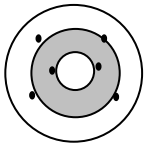

Unbiased and large variance

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 be largely biased whereas a low-bias estimator will have large variance. Survey statisticians usually rely on

estimators that are known to have little to no bias; they prefer to reduce the component of sampling error resulting from bias to a strict minimum and do whatever they can afterwards about variance. We also prefer this and will restrict the discussion on sampling errors to considerations about the sampling variance only.

To assess the sampling variance of an estimator, we need to get estimates that arise 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.

Using variance estimation, we can accurately predict the variability in results that arise from all samples by using just one sample. Whereas the general situation is somewhat intricate, variance estimation rests fundamentally 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 sampling variance, we would know the population spread. Can we get a good guess of that population spread on the basis of the 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 that 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 in assessing sampling variance from one observed sample, as we have done above.

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 term means.

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 be used strictly when there is an underlying random process at work (such as the random process in survey sampling by which samples are selected). Because 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 as the '**population spread**' (of that characteristic) rather 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.

The following mathematical definition of the population spread is often used:

$$\sigma^2 = \frac{\sum_{i=1}^{N} (y_i - \overline{Y})^2}{N-1}$$

where y_i is the value of the characteristic *Y* for person i

$\overline{Y}$ is the average of y_i in the population

N is the size of the population.

The population spread is then, simply a parameter of the population, one of many others like the population's average value for *Y*, $\overline{Y}$, 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. 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. As 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 as 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, the spread associated with other values is expressed in terms of its square. For example, if the estimate about personal income is expressed in terms of $, then the sampling variance is expressed as $\$^2$. Again, it is logical to take the square root of the value to restore comparability in terms of level with the estimate; the square root of the value is called the **standard error**.

13.2 Coefficient of variation

A measure of sampling variance provides us with a valuable indication of the reliability of an estimate. As we saw, the standard error is a better tool than sampling variance, as it is expressed in terms of the same units of measure as the estimate itself. Standard error (like variance) has a major drawback, however, because the standard errors arising from different estimates are not comparable in similar situations. 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*. Using the same information, we can also estimate the **mean** of *Y* by dividing the estimated total by the (assumed known) population size *N*. Nevertheless, both estimates are derived 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 comparable tools to evaluate relative precision among estimates, we would expect them to have the same value. They have different values, however, because the standard error of the total will actually be *N* times larger than that of the mean; estimates of the mean and total taken from the same observed sample 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:

$$CV\left(\hat{\theta}\right) = \frac{s.e.\left(\hat{\theta}\right)}{\hat{\theta}}$$

Now, in contrast 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-like[7] 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 $(1-p)$ is 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, which yields a p value of 0.95. The CV for the value of p, 0.95 would be:

0.0475 / 0.95 = 0.05 or 5%, which is a very good CV.

On the other hand, the CV for the proportion $(1-p)$ is

0.0475 / 0.05 = 0.95 or 95%, which is a very poor CV.

Issue 2: Applying rules about the assessment of sampling error in the case of proportions

Rules have been established to help 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 trustworthy); a CV between 16.5% and 33% indicates **acceptable quality** (inferences based on these results require caution) and a CV greater than 33% describes an estimate of **poor quality** (inferences based on these results are purely exploratory and untrustworthy). These rules can be useful to the data user in determining what to make of a survey's results, but they have their limitations. Consulting a survey statistician is certainly the best way to interpret the results, because factors specific to the analysis other than the sampling error can influence the validity of the inferences. 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%). The confidence interval is so large that the estimate is giving us little information.

7. By this we mean a positive quantity of interest that is continuous, as opposed to a dichotomous variable.

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 (the CV is 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 (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%, 0.25% or even 0.30%. Depending on the goal of the research, a statement that the estimate is smaller than 0.5% may be meaningful. Therefore, blindly applying the quality guidelines and rejecting an estimate because of its high CV is inappropriate.

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. Rather, we suggest constructing confidence intervals and relying on them to conduct proper inferences on 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 that 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, and the other is the non-sampling error. The latter encompasses such things as response errors and processing errors. On one hand, the magnitude of the non-sampling errors is difficult, if not impossible, to quantify in practice. (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 with 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 because of the complex sample design, non-response adjustments, treatment of out-of-scope units and 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 use in practice.

A very good way to approximate the sampling variance is to use the 'bootstrap' method.[8]

The idea behind the Bootstap method is to select random sub-samples from the full sample in such a way that each of the sub-samples (or replicates) follows the same design as the full sample. The final weights for units in each replicate are recalculated following the same weighting steps used for the full sample (see Chapter 11.0). These Bootstrap weights are used to calculate a population estimate for each replicate. The variance among the replicate estimates for a given characteristic is an estimate of the sampling variance of the full sample population estimate.

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 had all possible samples. 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.

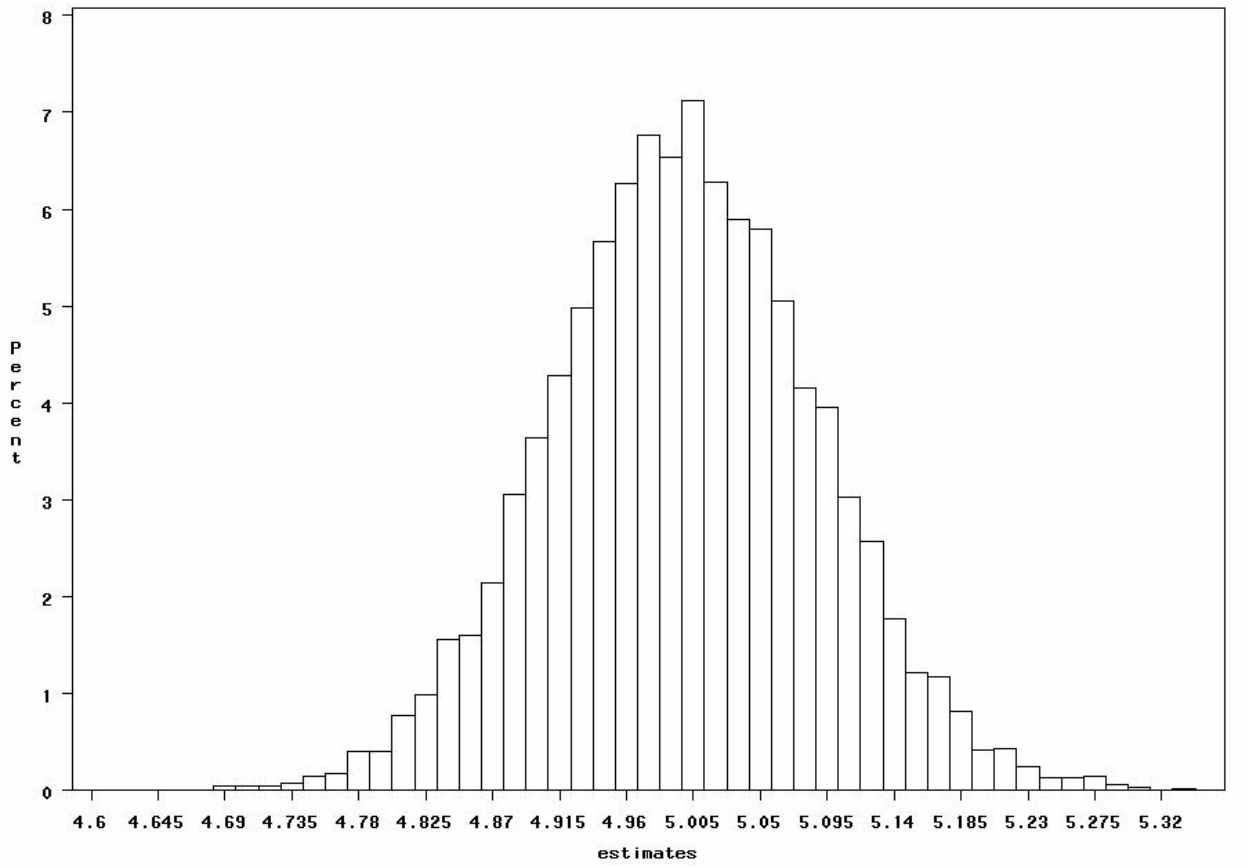

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

We notice, for instance, that extremely small estimates are rare, as are very large ones, whereas the most frequent estimates are 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 rare events (they arise less than

8. The challenge here is to adapt successfully the bootstrap for survey sampling although it was initially designed for non-survey setting. Our understanding of the bootstrap in a survey setting, and how to adequately implement it, has been evolving over the last 10 years or so, ever since the NLSCY first started using it.

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 true sampling distribution of a given estimator is based on all possible samples. The bootstrap representation of that sampling distribution is based on resampling from the one observed sample, i.e., the one for which we collect data through field work. When that correspondence exists, one can estimate a given parameter of the sampling distribution (like its variance) by computing it for the bootstrap distribution. Just as with the original sampling which resulted in the survey weights, the resampling leads to a set of 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:

$$\hat{v} = \frac{1}{1000}\sum_{k=1}^{1000}\left(\theta_k - \theta^*\right)^2$$

where θ_k is the k-th bootstrap estimate and θ^* is the original sample-based estimate of θ.

(Note: in practice, one can use the average of the θ_k's instead of θ^*, the sample-based estimate, in the formula most of the time. 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:

- 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 (not available for Cycle 8).
- 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. Because the Visual Basic interface and bootstrap macros are more flexible and more accurate, these tables are no longer provided.

13.4.1 Spreadsheet with approximate sampling variances for proportions

A set of spreadsheets is available to users to calculate the approximate 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 3.2, whereas in SPSS it is 3.0. Bootvar can be accessed through the link *Research tools* at http://www.statcan.gc.ca/rdc-cdr/bootvar_sas-eng.htm. Bootvar can compute, for any domains, variance estimates for such things as totals, ratios, difference of ratios, and linear and logistic regression coefficients. The sampling variance calculated using this method takes into account the sample design and the specificities of the variable of interest. Finally, by using Bootvar rather than the spreadsheet, the user is not restricted to predefined domains.[9]

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 become familiar with the macros before using them. However, macros were developed in such a way that they are easy to use. The researcher must have access to the macros, data files and 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, which uses Excel spreadsheets, to obtain a sampling variance estimate of any estimate that must be published; it provides a more exact and suitable measure of the sampling variance.

Details on how to use these programs can be found in the 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. For example, SUDAAN (setting design = BRR), WesVar and STATA9.

9. It should be noted that with each cycle we have increasing attrition of the original NLSCY longitudinal sample. Attrition can result from various circumstances, for example the inability to trace the children and their parents, older children or their parents refusing to participate in the survey, or children moving outside of the country. Consequently, some post strata, which were originally calibrated to census counts based on children's ages by sex and province for the year 1994 (e.g. 5-year-old females in Prince Edward Island in 1994), may now contain missing/zero values in certain bootstrap replicates. This will have some impact on the variance estimates when we examine the data at other levels of aggregation for Cycle 8 data of the original longitudinal sample (e.g. by age only; by sex only; by province only; nationally). As a result, when looking at counts at these higher levels of aggregation, though the NLSCY estimates will match the Census counts they may be accompanied by non-zero variance estimates.

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 undertaking.

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, such as 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 strata is possible, at this point its effectiveness is unclear, because 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. Many of these routines, however, do not take into account the sample design, e.g., stratification, even using the WEIGHT statement, which means the calculated variance is not the sampling variance we're after (it is usually well underestimated this way). Therefore, these procedures are not recommended because they can lead to erroneous conclusions.

But with SAS 9.2, the procedures PROC SURVEYFREQ, PROC SURVEYLOGISTIC, PROC SURVEYMEANS and PROC SURVEYREG can properly calculate the sampling variance from the Bootstrap weights provided. This could be done with the REPWEIGHTS statement when specifying VARMETHOD=BRR in the PROC statement.

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 to 21. These assessments are described in this chapter.

14.1 The Peabody Picture Vocabulary Test – Revised

The Peabody Picture Vocabulary Test – Revised (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 based on the number of 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 have a higher score. The standard score, however, 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. Whereas each sample is selected probabilistically, albeit for a different reference period, experts in the field of cognition measures might disagree as to whether differences between the estimated distributions from one sample to the other reflect a true population difference over time or simply result 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 the standardized scores of Cycle 8 were calculated using the same norms that were

used for Cycle 4 (v2), Cycle 5, Cycle 6, and Cycle 7.[10] Therefore, the Cycle 8 norms are based on the 28,214 records from Cycles 1 to 5 of the NLSCY main survey with PPVT–R raw scores. Some records with 0 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 were excluded. These 0 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 a test given in a survey context is that it is tailored to the child's age and performance so that only selected 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 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.

Because the calibration test was created 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 coefficients between the two scores from the tests were 94% for the English version and 96% for the French version, which were high by any standard. Therefore, we were confident that the raw scores could be used as they are.

10. Note that for Cycle 1, Cycle 2 and Cycle 3, a different norm sample was used. Whereas some slight variation existed, it was well within the sampling error that resulted from using different samples as the norm sample.

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, as 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. Whereas the trend is increasing overall, for many age groups the trend is decreasing because the sample is not large enough and poor representation introduces a lot of noise.

A 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 data may include those for 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, non-response should be handled on a case-by-case basis. For more information about non-response, please see Chapter 12.0.

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: predimensional, unidimensional, bidimensional, and integrated bidimensional. 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 predimensional level assesses the ability to count by rote and to quantify small sets, using concrete objects. This knowledge is important for the unidimensional level, where children deal with changes in quantity without objects that can be touched or seen. The unidimensional 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 slightly 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. As it is accepted that we cannot expect a child to do well at a level if the child has not done 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 entered the answer. The application determined 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 alone.[11] Requiring children to pass both parts 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 8, the Age-equivalent score (HKNCdS01), the 30-point-raw score (HKNCfS03) and the 30-point-raw age-standardized score (HKNCfS02).

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 predimensional level, three out of five items must be correct. A child failing to answer any questions at the first level will get the minimum (0), whereas 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 predimensional level. There are five items in this level. To reach the Age-equivalent score 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 unidimensional level. There are eight items in this level. To reach the Age-equivalent score 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 the bidimensional level. There are nine items in this level. To reach the Age-equivalent score 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, as some items have a) and b) parts. The 30-point raw score is simply the total number of correct answers among those 30 questions.

11. For example, part a) may ask which of two piles of counting chips is bigger and part b) asks which pile is smaller.

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 Cycles 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 the Age-equivalent score to the child's age, comparison with the Who Am I? and an analysis of the items and of non-response. 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 non-response bias. Please see Chapter 12.0, for more information on non-response.

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 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 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

12. For more information about the Who Am I? assessment, 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 Skills Development Canada. (www.hrsdc.gc.ca).

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 her or himself. The Drawing scale is not used in the NLSCY because of time constraints. Dr. Molly de Lemos was consulted before the Drawing scale was dropped from the NLSCY.

The assessment consists of an appealing booklet in which the children complete 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 children complete as much as they are able but are 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.

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 (HWICdS02) is the sum of the scores attributed to the tasks related to reproducing a symbol. The Symbols scale score (HWICdS03) is the sum of the scores attributed to the tasks related to writing. As 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 (HWICdS01), 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.

The Age-standardized scores were first made available in Cycle 6. The Cycle 8 variables are: the total Who Am I? scale standardized score (HWICfS04), the Copying scale standardized score (HWICfS05) and the Symbols scale standardized score (HWICfS06). 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, the Number Knowledge, and Who Am I? 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 non-response bias. Please see Chapter 12.0, for more information on non-response.

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. Since Cycle 5, the short version of the test developed for the purposes of the NLSCY consists of 20 questions at each level.

For each level, the test administered at Cycle 8 is the same as the one used at Cycle 7. No modifications were made to the number of questions or to the questions themselves.

Scoring
Each child who took the Mathematics Test was given a raw (gross) score (HMACS01), and a scaled score referred to as the classical scaled score (HMACfS02) and a variable (HMACgD03) that provides information on the total number of questions answered by the respondent and excludes questions with missing values.

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 (the test was developed in English only, so the sample represents only children in English schools). This sample 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% corresponds to a raw (gross) score of 1. On the complete test, the percentile rank of 0.55% corresponds to a raw (gross) score of 3 and a scaled score of 315, and the percentile rank of 0.99% corresponds to raw (gross) scores of 4 and a scaled score of 319. 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 level for the NLSCY Mathematics Test.

Relation between raw scores and classical scaled scores for the Mathematics Test, by level

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

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

In Cycle 5, the mathematics assessments for children from levels 2 to 10 (under the age of 16) were administered much earlier than in previous cycles. The impact for this change was mentioned in the Cycle 5 user guide as follows:

"The raw (gross) scores measured during collection are affected by the varying collection reference time for the children and youth who took the test. Consequently, the CTC equivalent score will reflect that same effect. This is an unfair assessment of individuals as their ability is directly related to the learned curriculum. An estimated learning growth model, based on all the children tested over time, was used to compensate for these differences and a corrected score (EMACfS02) has been produced."

Users must note that with the exception of Cycle 5, a similar methodology was applied to all other remaining cycle scores (i.e. AMACS02-DMACS02, FMACfS02, GMACfS02 and HMACfS02).

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 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) was revised from a 20-item assessment in Cycle 6 to an 18-item assessment in Cycle 7. This measure is a combination of Booklets 30 and 31 that were used in cycle 5. Questions 1 and 2 were removed from the Cycle 7 assessment and the remaining questions were renumbered. The two items (questions 1 and 2) were taken from the levels 7 and 9 Mathematics tests which do not allow calculator use. However respondents are allowed to use calculators in the completion of the Problem Solving Exercise. In Cycle 6, the two mathematic computation items could not be used to help link the Problem solving assessment with the mathematics test, therefore a decision was made to retire these items in Cycle 7. The test administered at Cycle 8 is the same as the one used at Cycle 7.

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 18 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 0 through 3 or 9 is given. A score of 0 is given for no credit, 1 is given for partial credit and 2 is given for full credit. Except for one question that has 2 for a superior partial credit and 3 for full credit. A score of 9 is given for missing values. In circumstances where items have been attempted and crossed out, a score of 0 (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 or partially correct, in which case it is given a value of 1, 2 or 3. 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 0. The complex items were taken from PISA (OECD's Programme for International Student Assessment) which allows for part marks in their scoring guide. Thus, this scoring approach was adopted for Cycle 7 and Cycle 8. Please note that in Cycle 6 complex questions marked by scorers were given a value of 1, 2 or 9. A score of 1 was given for a full credit, a score of 2 was given for no credit and a score of 9 was given for missing values. No marks were awarded for partial credits.

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.

Scores

(HMAYgS04)
This score is sometimes called the 'raw score' and represents the sum of the scores attributed to the 18 questions, without including the missing values.

(HMAYgD02)
In Cycle 7, a new variable was introduced. This variable provides information on the total number of questions answered by the respondent, and excludes questions with missing values.

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** 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, 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 (HLTYfS01) is simply the number of correct (‘full credit’) answers among the 36 items. An additional variable is derived (HLTYgD02) which provides information on the total number of questions answered by the respondent, and excludes questions with missing values.

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, the marking process and the scoring methodology are such that it is inadvisable to compare the two scores.

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. Whereas the type of information varies for the prose and document tasks, information is always requested as part of 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 (HNUYfS01) is simply the number of correct (‘full credit’) answers among the 32 items. An additional variable is derived (HNUYgD02) which provides information on the total number of questions answered by the respondent, and excludes questions with missing values.

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 it is inadvisable to compare the two scores.

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 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 these key data analysis issues:

- Statistics Canada recommends that the survey weights be used at analysis, whenever possible.
- Attention must be paid when combining several cycles of data as 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 should be made 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 non-response; for a description of how to estimate sampling error, see Chapter 13.0 – Variance estimation; and for more on data quality, including a detailed description of the various sources of non-sampling errors in a survey, e.g., non-response, response, undercoverage and processing errors, see Chapter 12.0 – Data quality, response rates and coverage.

A 2-day NLSCY data analysis workshop was held at the Toronto Research Data Centre in the fall of 2008 and covered many of the topics addressed in this chapter. The slides of the workshop are available at: http://www.utoronto.ca/rdc/events.html#NLSCY2008.

For assistance with a particular analytical question, please contact the Data Analysis Research Centre at Statistics Canada, by sending your question to: georgia.roberts@statcan.gc.ca.

15.1 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). Because of the complex sample design, the distribution of a characteristic of interest in the sample is probably different from its distribution in the population. Only by applying the survey weights can the population's distribution be preserved.

Stratification and clustering (both present in the NLSCY sample design) lead to unequal probabilities of selection. For example, 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 disproportionate number of children from small provinces.) Unequal non-response rates within the population can also lead to unequal representation of children in the sample. Finally, clustering in the sample leads to the statistical non-independence of units: the characteristics of children belonging to the same household are correlated.

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 to 11 as of December 31, 1994, who were living in any province 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 8, 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.

Distribution of children in the population, original cohort, weighted versus unweighted estimates using Cycle 8 funnel weights[1]

Province	Unweighted		Weighted	
	Frequency	Percent	Frequency	Percent
Newfoundland and Labrador	491	6.57	88,986	1.91
Prince Edward Island	227	3.04	23,148	0.50
Nova Scotia	536	7.17	144,088	3.09
New Brunswick	450	6.02	115,131	2.47
Quebec	1,457	19.50	1,090,582	23.41
Ontario	1,814	24.28	1,773,616	38.08
Manitoba	569	7.62	182,869	3.93
Saskatchewan	643	8.61	173,611	3.73
Alberta	688	9.21	489,913	10.52
British Columbia	597	7.99	576,125	12.37
Total	**7,472**	**100.00**	**4,658,069**	**100.00**

1. HWTCWd1L. See section 15.3 below.
Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Without the weights, the analyst would incorrectly conclude that 22.8% 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).

15.2 Unit of analysis

In the NLSCY, the unit of analysis is always the child. While some household data are 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.

15.3 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 14 to 25 at Cycle 8), two longitudinal weights are available, labelled HWTCW01L and HWTCWd1L. The first set of longitudinal weights is for children who responded (or their PMK responded) to Cycles 1 and 8, 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 8.

For example, if an analyst were interested in the characteristics at Cycle 8 of children introduced in Cycle 1, but not interested in their data from Cycles 2 to 7, 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 at or after the initial longitudinal population. For Cycle 8, cross-sectional weights are calculated only for children aged 0 to 7 (for a description of the sample composition of these children, see Chapter 5.0 or Chapter 11.0). It is not recommended that the original cohort be used to make inferences about the cross-sectional population of all 14- to 25-year-olds living in Canada at the time of Cycle 8 collection, as this sample has never been topped-up to include immigrants who arrived after 1994. (For more details on the impact of immigration, see Chapter 12.0.)

15.4 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., birth weight). 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, whether simple or complex.

15.4.1 Estimating a population total

The estimate of the total number of units in the survey population is calculated by adding together the final 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, w_i is its final survey weight (i.e., the design weight adjusted for non-response and post-stratification as described in Chapter 11.0) 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:

$$\hat{Y} = \sum_{i \in S_r} w_i y_i$$

15.4.2 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.

$$\hat{\bar{Y}} = \frac{\sum_{i \in S_r} w_i y_i}{\sum_{i \in S_r} w_i} = \frac{\hat{Y}}{\hat{N}}$$

15.4.3 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:

$$\hat{P} = \frac{\sum_{i \in S_r \cap C} w_i}{\sum_{i \in S_r} w_i} = \frac{\hat{N}_C}{\hat{N}}$$

where *C* is the collection of units having the given characteristic.

15.4.4 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:

$$\hat{\bar{Y}}_{domain} = \frac{\sum_{i \in S_r \cap domain} w_i y_i}{\sum_{i \in S_r \cap domain} w_i} = \frac{\hat{Y}_{domain}}{\hat{N}_{domain}}$$

The estimate of a domain proportion for qualitative or quantitative data is:

$$\hat{P}_{domain} = \frac{\sum_{i \in S_r \cap domain \cap C} w_i}{\sum_{i \in S_r \cap domain} w_i} = \frac{\hat{N}_{domain \cap C}}{\hat{N}_{domain}}$$

15.5 Normalized weights

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. 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, it is recommended that the analyst rescale 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.) This corrects the number of observed units. However, because 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. It is recommended that normalized weights 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:

$$w_i^{norm} = \frac{w_i}{\overline{w}}$$

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.

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).

Example of normalized weights in the presence of non-respondent units

Observation number	Response status for the variable of interest	Survey weight	Incorrect normalized weight (calculated using non-respondents)	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	Non-respondent	6	Excluded from the analysis	Excluded from the analysis
6	Non-respondent	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).

Example of 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
Subtotal			**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
Subtotal			**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, because of the correlation of sampled units (which is often the result of clustering). In such cases, normalization leads to:

- an overestimation of the effective number of units
- an underestimation of the variability
- too many significant results.

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
- a 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 whether 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, because 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 underestimate the sampling variance of the estimates and result in declarations of significant differences where none exist. A rule of thumb is often used to compensate, but this adjustment is sometimes too conservative and 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.

15.6 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 so 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.

Example of construction of input file

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 keep in mind when pooling data.

15.7 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 needs 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 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 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 will require 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 0 does not mean that the child is not also followed longitudinally. An easy way to identify whether a child appears more than once in the pooled sample is to check whether 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). 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, because 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- to 1-year-olds in Cycles 1 and 2 and 2- to 4-year-olds in Cycle 1.

Alternatively, the analyst could derive a new measure that combines all the information across the cycles and 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 to 2 years old in Cycle 1 and 2 to 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, as the analyst chose only 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 selected 0- to 5-year-olds 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.

15.8 Non-response

Like any other survey, the NLSCY is subject to non-response. There are two main types of non-response, total and partial.

> **Total non-response** is the complete lack of data for a selected and eligible child as the result of factors such as refusal to take part in the survey or inability to trace the child. Total non-response is taken into account and corrected by the sampling weights. See Chapter 11.0 for details.
>
> **Partial non-response** 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 8. For examples of partial non-response, see Chapter 11.0. Some missing questions are imputed, but most are not.

Non-response 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 non-response for some variables of interest. There are a few ways to deal with this situation:

1. Keep only records for which complete information is available for your main analysis, and keep your non-respondent profiles for a separate side analysis. 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 non-respondents. Being upfront about describing the differences between the two groups alerts readers of your results that non-response was an issue with the data and helps them properly interpret your analytical findings.
2. Report partial non-response as a category. This approach means you report the non-response 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.
3. Reweight the records with a response to account for the partial non-respondents.
4. Ignore the partial non-response records, but increase the weights of the respondents to account for the non-respondents. This is an especially interesting strategy when an entire component of the survey is missing (for instance, the Self-complete questionnaire). If, however, you only have a few missing data here and there, this may not be the best strategy because it rejects records with any partial non-response.
5. Apply a weight adjustment to the respondents, which is simply the inverse of the response probability, which is often taken to be the weighted response rate.
6. Note that the weights for NLSCY are post-stratified to known counts by age, sex and province. When reweighting to adjust for partial non-response, these control totals are no longer respected. You may choose to repost-stratify after adjusting for non-response, or instead of reweighing you may use the imputation approach (see next point).
7. Impute partial non-response (replace missing values with replacement values). The advantage of imputing partial non-response over reweighting is that all records are kept, which means no data are discarded. Only the missing values in each record are filled in. 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 non-response, see Appendix II. For those looking for more information on non-response, some references are given below. The list is by no means exhaustive.

Non-response overview

- Survey Methods and Practices (Statistics Canada 2003)

  - Different chapters discuss non-response related issues.
  - This book can be ordered from Statistics Canada's website (http://www.statcan.gc.ca).
- Sampling: Design and Analysis (Lohr, S., Duxbury Press, 1999)
  - Chapter 11.0 is devoted entirely to non-response.

Non-response 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, Second Edition (Little, R.J.A. and Rubin, D.B., New York: Wiley, 2002)

15.9 Other sources of non-sampling errors

Besides non-response, 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 resulting from undercoverage of the population, processing errors, response errors, etc. Non-sampling errors are described in detail in Chapter 12.0.

15.10 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 non-response 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 integrated version 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 and 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 that 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.gc.ca/pub/12-002-x/12-002-x2004002-eng.pdf. For more details on estimating sampling variance, including details on another tool called Bootvar, refer to Chapter 13.0.

15.11 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.

15.12 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

15.13 Understanding the difference between “Not stated” , “Don’t know”, “Refusal” and “Valid skip”

Not all questions in the NLSCY 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. When examining non-response, the “Valid skip” cases should **not** be treated as non-respondents — 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.

16.0 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.

16.1 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) Round estimates in the main body of a statistical table 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 1. 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) Derive marginal subtotals and totals in statistical tables from corresponding unrounded components and subtotals and totals to the nearest 100 units using normal rounding.

c) Compute averages, proportions, rates and percentages from unrounded components, i.e., numerators and/or denominators; then **round averages, proportions, rates and percentages to 1 decimal** using normal rounding.

d) Derive sums and differences of aggregates (or ratios) from their corresponding unrounded components; then round sums and differences of aggregates (or ratios) to the nearest 100 units (or the nearest 1 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 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 documents.

f) Under no circumstances are unrounded estimates to be published or otherwise released by users. Unrounded estimates imply greater precision than actually exists.

16.2 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.

16.3 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. For more details about the sample design, please refer to Chapter 5.0. 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 whereas 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.

16.4 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 (i.e. the reliability of the estimate). This means that the standard error associated with the estimate must be calculated (Chapter 13.0). Users should also consider how non-sampling errors discussed in Chapter 11.0 may affect the estimate.

A standard measure of the quality of estimates used in the release of NLSCY data is the coefficient of variation (described in section 13.2). Here is a table outlining the quality level guidelines of the estimates using the coefficient of variation:

Quality Level Guidelines

Quality Level of Estimate	Guidelines
1) Acceptable	Estimates have low coefficients of variation in the range of 0.0% to 16.5%. No warning is required.
2) Marginal	Estimates have high coefficients of variation in the range of 16.6% to 33.3%. These estimates are flagged with a superscript 'E'. They are also accompanied by a warning to caution users about the high levels of error, associated with the estimates.
3) Unacceptable	Estimates have very high coefficients of variation in excess of 33.3%. These estimates will be suppressed with an 'F' because they are of unacceptable quality. Conclusions based on these estimates will be unreliable, and most likely invalid.

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 will probably be large, and the estimate will probably be 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.

16.5 Modelling NLSCY data and bootstrap weights

The NLSCY is a survey with a complex sample design (please see section 5.0 of the NLSCY User Guide for an overview). As a result, some of the sampling methods used result in a dependency of some kind in the selection of units (like stratification and/or clustering i.e., the formation of groups through which units are selected). The complexity of a survey design, like the one used for the NLSCY, has a significant bearing on how we compute the variance associated to an estimate.

The use of standardized weights alone (i.e. normalized weights) will not estimate the total variance correctly; they cannot account for any dependencies in the selection of units.

In order to account for the impact of the sample design on the variance, the bootstrapping method, a replicate-based variance estimation process, is what we advocate to users.

Bootstrap weights are provided on a separate file and can be merged with the data file using the PERSRUK variable. All that remains to be done is for users to redo the analysis they just performed using the release weight but this time using (in turn) each of the bootstrap weights. Users can make use of the Bootvar programs available in SAS and SPSS for this purpose[13]. There are also some software packages that carry out this final step on the user's behalf, e.g. SUDAAN (using SUDAAN's "BRR" option).

Using the bootstrapping method increases the accuracy of the variance estimate by reflecting the various components of a complex design in the overall variance estimate. This is especially important if there are significant results close to the rejection threshold. Not taking into account the survey design will result in inaccurate variance estimates. This will in turn impact the test statistics and could lead to incorrectly establishing statistical significance where none exists.

13. Bootvar can be accessed through the link *Research tools* at www.statcan.gc.ca/rdc-cdr/bootvar_sas-eng.htm or at www.statcan.gc.ca/rdc-cdr/bootvar_spss-eng.htm.

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.

I. 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, 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:

1) data sources
2) factors affecting the analysis
3) variables
4) type of analysis
5) variance estimation
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 non-response

- If some variables used in the analysis have non-response, include a statement about the level of non-response, if known, and its potential impact on the analysis.
- Specify how partial non-response will be handled:
    - imputation
    - reweighting
    - reported as a value
    - ignored and analysis to be done with the respondents only.
- Analyze characteristics of non-respondents versus respondents to identify possible biases.

c Data limitations

- Provide the sample sizes, overall and for all subdomains, where this information is known. Sample sizes will be needed that are sufficiently large both to respect confidentiality and to give reliable estimates.
- Indicate whether 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 website) 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 the following:

- 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: _____)

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.

II. 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, 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:

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 that 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 non-response (if any)
- If some variables used in the analysis have non-response, include a statement about the level of non-response and the impact on the analysis.
- Specify how partial non-response was handled:
  - imputation
  - reweighting
  - reported as a value
  - ignored, analysis done with the respondents only.
- Analysis of non-respondents versus respondents should be done to identify possible biases.

c Data limitations
- Provide the sample sizes, overall and for all subdomains.
- Ensure that the sample sizes used in the report are sufficient both to respect confidentiality and to give reliable estimates.
- Indicate whether 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 following measurement of the variances and/or coefficients of variation (CV):

- 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: _____)

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.

III. Examples of proper citation of NLSCY products:

How to cite an NLSCY master file (remote data access users):

Statistics Canada. 2006-2007. National Longitudinal Survey of Children and Youth: NLSCY2006_C7_LONG_Master. Statistics Canada.

How to cite an NLSCY master file (accessed through a Research Data Centre):

Statistics Canada. 2006-2007. National Longitudinal Survey of Children and Youth: NLSCY2006_C7_ECD_Master. Statistics Canada. Using University of Alberta Research Data Centre. Released Month dd, 2008. http://www.statcan.gc.ca/cgi-bin/imdb/p2SV.pl?Function=getSurvey&SDDS=4450&lang=en&db=imdb&dbg=f&adm=8&dis=2 (November 9, 2010).

How to cite the NLSCY User Guide:

Statistics Canada. n.d. *Microdata User Guide, National Longitudinal Survey of Children and Youth, Cycle 7, September 2006 to July 2007*. http://www.statcan.gc.ca/imdb-bmdi/document/4450_D4_T9_V7-eng.pdf (November 9, 2010).

For further information:

Statistics Canada. 2006. *How to Cite Statistics Canada Products*. Statistics Canada Catalogue no. 12-591-XWE. Ottawa. Version updated March 31. Ottawa. http://www.statcan.gc.ca/english/freepub/12-591-XIE/12-591-XIE2006001.htm (November 9, 2010).

Appendix II Partial non-response analysis

As noted in Section 12.4, analysts using NLSCY data should be aware of how partial non-response affects the data they are attempting to analyse. Partial non-response may be higher for respondents with a particular characteristic. This may lead to bias, which can cast analytical results into question. There are techniques available to deal with partial non-response, such as reweighting and imputation. At a minimum, users should study the extent of component or item non-response in their results and include these findings in their report. An example of such an analysis is presented below.

The following partial non-response 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 non-response. There are two main categories of non-response, total non-response and partial non-response. Total non-response 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 non-response.

Partial non-response 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 non-response 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 those aged 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 non-response 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 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 non-response 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

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.

Appendix III Concordance of processing variable names to dissemination variable names for the Self-complete questionnaires, National Longitudinal Survey of Children and Youth, Cycle 8

Notes:
1. Blank cells = not applicable
2. In "Type" column, C = character, N = numeric

Variable	Size	Type	Book 22	Book 23
PERSRUK	14	C	PERSRUK	PERSRUK
HMMCQ01	3	N	EFFAGEYR	EFFAGEYR
HMMCbQ1A	3	N	AGE	AGE
HMMCQ02	1	C	GENDER	GENDER
HFFCQ01	1	N	A1	A1
HFFCQ02	1	N	A2	A2
HFFCQ03	1	N	A3	A3
HFFCQ04	1	N	A4	A4
HFFCcQ4A	1	N	A5	A5
HFFCQ05	2	N	A6	A6
HFFCbQ13	2	N	A7B	A7B
HFFCbQ14	2	N	A8B	A8B
HFFCbQ15	1	N	A9	A9
HFFCQ07	1	N	A13	A13
HFFCQ08A	1	N	A14A	A14A
HFFCQ08B	1	N	A14B	A14B
HFFCQ08C	1	N	A14C	A14C
HFFCQ08D	1	N	A14D	A14D
HFFCQ08E	1	N	A14E	A14E
HFFCQ08F	1	N	A14F	A14F
HFFCQ08G	1	N	A14G	A14G
HFFCQ08H	1	N	A14H	A14H
HFFCc08I	1	N	A14I	A14I
HFFCQ08K	1	N	A14J	A14J
HFFCQ08L	1	N	A14K	A14K
HFFCQ08M	1	N	A14L	A14L
HFFCQ08N	1	N	A14M	A14M
HFFCd12A	1	N	A15	A15
HDRCdQ05	1	N	A10A	A10A
HDRCdQ09	1	N	A10B	A10B
HFFCd16C	1	N	A10C	A10C
HFFCc16D	1	N	A10D	A10D
HFFCc16E	1	N	A10E	A10E
HFFCc19A	1	N	A11A	A11A
HFFCc19B	1	N	A11B	A11B
HFFCc19C	1	N	A11C	A11C

Variable	Size	Type	Book 22	Book 23
HFFCc19D	1	N	A11D	A11D
HFFCc20A	1	N	A12A	A12A
HFFCc20B	1	N	A12B	A12B
HFFCc20C	1	N	A12C	A12C
HSCCQ01	1	N	B1	
HSCCbQ18	1	N	B2	
HSCCb19A	1	N	B3A	
HSCCb19B	1	N	B3B	
HSCCb19C	1	N	B3C	
HSCCb19D	1	N	B3D	
HSCCc19F	1	N	B3E	
HSCCb19E	1	N	B3F	
HSCCb20A	1	N	B4A	
HSCCb20B	1	N	B4B	
HSCCb20C	1	N	B4C	
HSCCb20E	1	N	B4D	
HSCCc20H	1	N	B4E	
HSCCc20I	1	N	B4F	
HSCCc20J	1	N	B4G	
HSCCb20G	1	N	B4H	
HSCCQ02	1	N	B5	
HSCCcQ3a	1	N	B6A	
HSCCcQ03	1	N	B6B	
HSCCcQ3b	1	N	B6C	
HSCCcQ3D	1	N	B6D	
HSCCcQ3C	1	N	B6E	
HSCCcQ3E	1	N	B6F	
HSCCcQ3F	1	N	B6G	
HSCCd3G	1	N	B6H	
HSCCc21A	1	N	B7A	
HSCCc21B	1	N	B7D	
HSCCc21C	1	N	B7B	
HSCCc21D	1	N	B7C	
HSCCc21E	1	N	B7E	
HSCCc21F	1	N	B7F	
HSCCcQ26	1	N	B8	
HSCCcQ27	1	N	B9	
HSCCcQ10	1	N	B10	B7
HSCCbQ22	1	N	B12A	
HSCCc22A	1	N	B12B	
HSCCcQ28	1	N	B13	
HSCCcQ29	1	N	B14	
HSCCQ12	1	N	B15A	
HSCCQ11	2	N	B15B	

Variable	Size	Type	Book 22	Book 23
HSCCcQ16	2	N	B15C	
HSCCQ17	2	N	B15D	
HSCCcQ30	2	N	B16	
HSCCQ13	2	N	B17A	
HSCCQ14	1	N	B17B	
HSCCQ15	1	N	B17C	
HSCCcQ24	2	N	B18	
HSCCc31A	1	N	B11A	
HSCCc31B	1	N	B11B	
HSCCc31C	1	N	B11C	
HSCCc31D	1	N	B11D	
HSCCc31E	1	N	B11E	
HAMCQ01A	1	N	C1A	B1A
HAMCQ01B	1	N	C1B	B1B
HAMCQ01C	1	N	C1C	B1C
HAMCQ01D	1	N	C1D	B1D
HAMCcQ1E	1	N	C1E	B1E
HAMCe25A	1	N	C2A	B2A
HAMCe25B	1	N	C2B	B2B
HAMCe25C	1	N	C2C	B2C
HAMCe25D	1	N	C2D	B2D
HAMCe25E	1	N	C2E	B2E
HAMCe25F	1	N	C2F	B2F
HAMCe25G	1	N	C2G	B2G
HAMCe25H	1	N	C2H	B2H
HAMCe25I	1	N	C2I	B2I
HAMCe25J	1	N	C2J	B2J
HAMCe25K	1	N	C2K	B2K
HAMCe25L	1	N	C2L	B2L
HAMCe25M	1	N	C2M	B2M
HAMCe25N	1	N	C2N	B2N
HAMCe25O	1	N	C2O	B2O
HAMCbQ02	1	N	C3	B3
HAMCcQ03	1	N	C4	B4
HAMCcQ6A	1	N	C7A	
HAMCcQ6B	1	N	C7B	
HAMCdQ6C	1	N		B8A
HAMCcQ7A	1	N	C8A	
HAMCcQ7B	1	N	C8B	
HAMCdQ7C	1	N		B8B
HAMCcQ8A	1	N	C9A	
HAMCcQ8B	1	N	C9B	
HAMCdQ8C	1	N		B8C
HAMCdQ4A	1	N	C5A	

Variable	Size	Type	Book 22	Book 23
HAMCcQ4B	1	N	C5B	
HAMCcQ4C	1	N	C5C	
HAMCcQ4D	1	N	C5D	
HAMCeQ4E	1	N	C5E	
HAMCe26A	1	N		B5A
HAMCe26B	1	N		B5B
HAMCe26C	1	N		B5C
HAMCe26D	1	N		B5D
HAMCe26E	1	N		B5E
HAMCe26F	1	N		B5F
HAMCcQ5A	1	N	C6A	B6A
HAMCcQ5B	1	N	C6B	B6B
HAMCcQ5C	1	N	C6C	B6C
HAMCcQ5D	1	N	C6D	B6D
HAMCdQ09	1	N		B9
HAMCdQ10	1	N		B10
HFBCQ01A	1	N	D1A	
HFBCc01B	1	N	D1B	
HFBCQ01C	1	N	D1C	
HFBCQ01D	1	N	D1D	
HFBCQ01E	1	N	D1E	
HFBCQ01F	1	N	D1F	
HFBCQ01G	1	N	D1G	
HFBCQ01H	1	N	D1H	
HFBCQ01I	1	N	D1I	
HFBCQ01J	1	N	D1J	
HFBCQ01K	1	N	D1K	
HFBCQ01L	1	N	D1L	
HFBCQ01M	1	N	D1M	
HFBCQ01P	1	N	D1N	
HFBCQ01Q	1	N	D1O	
HFBCQ01R	1	N	D1P	
HFBCQ01S	1	N	D1Q	
HFBCQ01T	1	N	D1R	
HFBCQ01U	1	N	D1S	
HFBCc01V	1	N	D1T	
HFBCQ01W	1	N	D1U	
HFBCd01X	1	N	D1V	
HFBCQ01Z	1	N	D1W	
HFBCQ1AA	1	N	D1X	
HFBCQ1BB	1	N	D1Y	
HFBCQ1CC	1	N	D1Z	
HFBCQ1DD	1	N	D1AA	
HFBCQ1FF	1	N	D1BB	

Variable	Size	Type	Book 22	Book 23
HFBCQ1GG	1	N	D1CC	
HFBCQ1HH	1	N	D1EE	
HFBCQ1JJ	1	N	D1DD	
HFBCQ1LL	1	N	D1FF	
HFBCQ1MM	1	N	D1GG	
HFBCQ1NN	1	N	D1HH	
HFBCQ1OO	1	N	D1II	
HFBCQ1PP	1	N	D1JJ	
HFBCQ1QQ	1	N	D1KK	
HFBCQ1RR	1	N	D1LL	
HFBCQ1SS	1	N	D1MM	
HFBCQ1TT	1	N	D1NN	
HFBCc1UU	1	N	D1OO	
HFBCd10A	1	N		C1A
HFBCd10B	1	N		C1B
HFBCd10C	1	N		C1C
HFBCd10D	1	N		C1D
HFBCd10E	1	N		C1E
HFBCd10F	1	N		C1F
HFBCd10G	1	N		C1G
HFBCd10H	1	N		C1H
HFBCd10I	1	N		C1I
HFBCd10J	1	N		C1J
HFBCd10K	1	N		C1K
HFBCd10L	1	N		C1L
HFBCcQ04	1	N	D2	C2
HFBCcQ4A	1	N	D3	C3
HFBCcQ05	1	N	D4	C4
HFBCcQ07	1	N	D5	C5
HFBCcQ08	1	N	D6	C6
HFBCbQ2B	1	N	D7A	C7A
HFBCbQ2E	1	N	D7B	C7B
HFBCbQ2F	1	N	D7C	C7C
HFBCcQ2H	1	N		C7D
HFBCeQ2O	1	N	D7D	C7E
HFBCbQ2P	1	N	D7E	C7F
HFBCdQ2Z	1	N		C7G
HFBCe2ZZ	1	N	D7F	C7H
HFBCbQ2Y	1	N	D7G	C7I
HFBCe2BB	1	N	D7H	C7J
HFBCcQ3A	1	N	D8	C8
HATCc1AA	1	N	E1A	
HATCc1BB	1	N	E1B	
HATCc1CC	1	N	E1C	

Variable	Size	Type	Book 22	Book 23
HATCc1DD	1	N	E1D	
HATCc1EE	1	N	E1E	
HATCc1II	1	N	E1F	
HATCeQ20	2	N	E2	
HATCbQ04	1	N	E3	
HATCdQ07	2	N	E7	
HATCeQ12	1	N	E8A	
HATCe12B	1	N	E8B	
HATCe12C	1	N	E8C	
HATCeQ13	1	N	E9A	
HATCe13B	1	N	E9B	
HATCe13C	1	N	E9C	
HATCeQ21	2	N	E10	
HATCeQ22	1	N	E11	
HATCdQ10	2	N	E12	
HATCdQ5A	1	N	E5A	
HATCbQ5B	1	N	E5B	
HATCbQ5C	1	N	E5C	
HATCbQ5D	1	N	E5D	
HATCbQ5E	1	N	E5E	
HATCbQ5F	1	N	E5F	
HATCbQ5G	1	N	E5G	
HATCdQ05	1	N	E6	
HATCc14A	1	N	E4A	
HATCc14B	1	N	E4B	
HATCc14C	1	N	E4C	
HATCc14D	1	N	E4D	
HDRCdQ01	2	N	F1	D1
HDRCQ04	2	N	F2	D2
HDRCQ03	2	N	F3B	
HDRCdQ6A	2	N	F4	D3
HDRCdQ07	2	N	F5	
HDRCdQ9A	1	N	F6	
HDRCbQ9B	2	N	F7	
HDRCdQ9C	2	N	F8	D4
HDRCdQ15	2	N	F9	D5
HDRCd14A	2	N	F11AB	
HDRCc14C	2	N	F11BB	
HDRCc14B	2	N	F11CB	
HDRCd14F	2	N	F11DB	
HDRCc14D	2	N	F11EB	
HDRCc18A	2	N	F10A	D6A
HDRCd18B	2	N	F10B	D6B
HDRCd18C	2	N	F10C	D6C

Variable	Size	Type	Book 22	Book 23
HDRCc18D	2	N	F10D	D6D
HDRCdQ19	1	N		D7
HDRCdQ20	1	N		D8
HPMCdQ04	2	N	G1	G1
HPMCcQ5A	1	N	G2A	G2A
HPMCcQ5B	1	N	G2B	G2B
HPMCcQ5C	1	N	G2C	G2C
HPMCcQ06	1	N	G3	G3
HPMCdQ6A	2	N		G4A
HPMCdQ6B	2	N		G4B
HPMCdQ6C	1	N		G5A
HPMCdQ6D	1	N		G5B
HPMCdQ6E	1	N		G5C
HPMCdQ6F	1	N		G5D
HPMCdQ6G	1	N		G5E
HPMCdQ6H	1	N		G5F
HPMCdQ6I	1	N		G5G
HPMCdQ6J	1	N		G5H
HPMCdQ6K	1	N		G5I
HPMCdQ6L	1	N		G5J
HPMCdQ07	2	N	G4	G6
HPMCcQ8A	1	N	G5A	G7A
HPMCcQ8B	1	N	G5B	G7B
HPMCcQ8C	1	N	G5C	G7C
HPMCcQ09	1	N	G6	G8
HPMCdQ9A	2	N		G9A
HPMCdQ9B	2	N		G9B
HPMCdQ9C	1	N		G10A
HPMCdQ9D	1	N		G10B
HPMCdQ9E	1	N		G10C
HPMCdQ9F	1	N		G10D
HPMCdQ9G	1	N		G10E
HPMCdQ9H	1	N		G10F
HPMCdQ9I	1	N		G10G
HPMCdQ9J	1	N		G10H
HPMCdQ9K	1	N		G10I
HPMCdQ9L	1	N		G10J
HPMCcQ10	1	N	G7	G12
HPMCd11A	2	N	G8	G13
HPMCc11B	2	N	G9	G14
HPMCcQ1A	1	N	G10A	
HPMCcQ1B	1	N	G10B	
HPMCcQ1C	1	N	G10C	
HPMCcQ1D	1	N	G10D	

Variable	Size	Type	Book 22	Book 23
HPMCcQ1E	1	N	G10E	
HPMCcQ1F	1	N	G10F	G11A
HPMCcQ1G	1	N	G10G	
HPMCcQ1H	1	N	G10H	
HPMCcQ1I	1	N	G10I	
HPMCcQ1J	1	N	G10J	
HPMCcQ1R	1	N	G10K	
HPMCcQ1K	1	N	G10L	
HPMCcQ1L	1	N	G10M	
HPMCcQ1M	1	N	G10N	
HPMCcQ1N	1	N	G10O	
HPMCcQ1O	1	N	G10P	
HPMCcQ1P	1	N	G10Q	
HPMCcQ1Q	1	N	G10R	
HPMCdQ1S	1	N	G10S	
HPMCdQ1T	1	N	G10T	G11B
HPMCdQ9U	1	N		G11C
HPMCdQ9V	1	N		G11D
HPMCd12A	1	N		G15
HPMCd12B	1	N		G16
HPMCd12C	1	N		G17A
HPMCd12D	1	N		G17B
HPMCd12E	1	N		G17C
HPMCd12F	1	N		G17D
HPMCd12G	1	N		G17E
HPMCd12H	1	N		G17F
HPMCd12I	1	N		G17G
HPMCd12J	1	N		G17H
HPMCd12K	1	N		G17I
HPMCd12L	1	N		G17J
HHTCcQ03	1	N	H1	
HHTCbQ3A	1	N	H4A	
HHTCbQ3B	1	N	H4B	
HHTCbQ3C	1	N	H4C	
HHTCbQ3G	1	N	H4D	
HHTCdQ06	1	N	H5	
HHTCeQ5A	1	N	H6	E8
HHTCeQ5B	1	N		E8AA
HHTCeQ5C	1	N		E8AB
HHTCeQ5D	1	N		E8AC
HHTCeQ5E	1	N		E8AD
HHTCdQ5F	1	N		E8AE
HHTCdQ5G	1	N		E8BA
HHTCdQ5H	1	N		E8BB

Variable	Size	Type	Book 22	Book 23
HHTCdQ5I	1	N		E8BC
HHTCdQ5J	1	N		E8BD
HHTCeQ5K	1	N		E8CA
HHTCeQ5L	1	N		E8CB
HHTCeQ5M	1	N		E8CC
HHTCeQ5N	1	N		E8CD
HHTCeQ5O	1	N		E8CE
HPUCQ01	1	N	H7	E3
HPUCQ02	1	N	H8	E4
HPUCcQ3A	3	N	H9	E5
HPUCdQ04	1	N	H10	E6
HPUCQ05	1	N	H11	E7
HDACcQ1A	2	N	H12B	F1B
HDACcQ1B	1	N	H13	F2
HDACcQ1C	1	N	H14	F3
HDACcQ02	2	N	H15	F4
HDACcQ2A	1	N	H16	F5
HDACcQ05	1	N	H17	F6
HDACdQ23	1	N		F9
HDACeQ26	1	N		F11
HDACeQ25	2	N		F12
HDACcQ06	2	N	H18	F7
HDACcQ07	2	N	H19A	F8
HDACcQ08	1	N	H20	
HDACcQ09	1	N	H21	
HDACd22A	1	N		F10A
HDACd22B	1	N		F10B
HDACd22C	1	N		F10C
HDACd22D	1	N		F10D
HDACd22E	1	N		F10E
HDACd22F	1	N		F10F
HDACd22G	1	N		F10G
HDACd22H	1	N		F10H
HWKCcQ4A	1	N	I1A	
HWKCcQ4B	1	N	I1B	
HWKCcQ4C	1	N	I1C	
HWKCcQ4D	1	N	I1D	
HWKCe06A	1	N	I2A	
HWKCe06B	1	N	I2B	
HWKCe06C	1	N	I2C	
HWKCe06D	1	N	I2D	
HWKCe06E	1	N	I2E	
HWKCe06F	1	N	I2F	
HWKCcQ9A	2	N	I3AA	

Variable	Size	Type	Book 22	Book 23
HWKCcQ9B	2	N	I3BB	
HWKCeQ11	1	N	I4	
HWKCd12A	1	N	I5A	
HWKCd12B	1	N	I5B	
HWKCd12C	1	N	I5C	
HWKCd12D	1	N	I5D	
HWKCd13A	1	N	I6A	
HWKCd13B	1	N	I6B	
HWKCd13C	1	N	I6C	
HWKCd13D	1	N	I6D	
HWKCd13E	1	N	I6E	
HWKCd13F	1	N	I6F	
HWKCd13G	1	N	I6G	
HWKCd13H	1	N	I6H	
HWKCd13I	1	N	I6I	
HHTCbQ01	4.2	N	HEIGHT	HEIGHT
HHTCbQ02	7.3	N	WEIGHT	WEIGHT
HBMIeS01	7.3	N	SELFBMI	SELFBMI
HBMIeS02	2	N	SELFBMI2	SELFBMI2
HBMIeS03	1	N	SELFBMI3	SELFBMI3
HEQYeS06	2	N	RA	RA
HEQYeS07	2	N	ER	ER
HEQYeS08	2	N	SM	SM
HEQYeS09	2	N	AD	AD
HEQYeS10	2	N	GM	GM
HEQYeS04	2	N	EQ4	EQ4
HEQYeS05	2	N	EQ5	EQ5
HFFcS01	2	N	FRFS01	FRFS01
HAMcS02	2	N	ABMS01	ABMS01
HFBcS01	2	N	BEHS01	
HFBCdS02	2	N	BEHS02	
HFBcS03	2	N	BEHS03	
HFBCdS04	2	N	BEHS04	
HFBcS05	2	N	BEHS05	
HFBcS07	2	N	BEHS07	
HPMCcS1	2	N	PARS01	
HPMCbS2B	2	N	PARS02	
HPMCcS3	2	N	PARS03	
HPMCdS4	2	N		PARS04
HPMCdS5	2	N		PARS05
HHTCbS1b	2	N		DEPS01
HWTCw01L	12.4	N		
HWTCwD1L	12.4	N		

Source: Statistics Canada, National Longitudinal Survey of Children and Youth.