Microdata User Guide

National Longitudinal Survey of Children and Youth

Cycle 5

September 2002 to June 2003

Statistics Canada Statistique Canada

Canada

TABLE OF CONTENTS

1.0 Introduction

The National Longitudinal Survey of Children and Youth (NLSCY) Cycle 5 was conducted from September 2002 to June 2003 by Statistics Canada in partnership with Social Development Canada.

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

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

Statistics Canada

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

Special Surveys Division, Statistics Canada
2500 - R, Main Building
Tunney's Pasture
Ottawa, Ontario K1A 0T6

Social Development Canada

Ruth Martin
Assistant Director
Knowledge and Research Directorate
Social Development Canada
3rd Floor 3-042
355 North River Road
Place Vanier Tower B
Ottawa, Ontario K1A 0L1

Telephone: (613) 957-6735 (effective February 25, 2005)
E-mail: ruth.a.martin@sdc-dsc.gc.ca

2.0 Background

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

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

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

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

Survey Population
In Cycle 5, a representative sample of Canadian children aged 0 to 5 years from each of the provinces was surveyed for longitudinal and cross-sectional purposes. The cohort of children aged 8 to 19 years was surveyed for longitudinal purposes.

The Northern cohort of children consists of a census of 5-year-old children who are attending kindergarten in the Yukon. Please refer to Chapter 17.0 of this guide for a detailed explanation of the background of the Survey of Northern Children and a description of the children surveyed in Cycle 5.

Target population
The NLSCY objectives are to produce longitudinal and cross-sectional estimates. Therefore, several populations are targeted in the Cycle 5 sample.

- Cross-sectionally, the Cycle 5 sample represents all children who were 0 to 5 years old on January 1st, 2003.

- Longitudinally, we have three cohorts:

    1) The first cohort represents all children who were 0 to 11 years old in 1994/1995. Those children are now 8 to 19 years old in Cycle 5.

    2) The second cohort represents all children who were 0 to 1 years old in 1998/1999. Those children are now 4 to 5 years old in Cycle 5.

    3) The third cohort represents all children who were 0 to 1 years old in 2000/2001. Those children are now 2 to 3 years old in Cycle 5.

- Cross-sectionally, the census of children living in the territories is restricted in Cycle 5 to 5 year old children who attend kindergarten in the Yukon.

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

Data Release Strategy
The Cycle 5 data are being released in three files: Longitudinal cohort (primary and secondary), Early Child Development cohort (primary and education), and Northern cohort.

- Longitudinal cohort: These are the children from the original longitudinal cohort. These children were aged 0 to 11 years in the first cycle of the NLSCY. For Cycle 5, they are 8 to 19 years old. There are two files for the longitudinal cohort: the primary file contains data collected from the person most knowledgeable about the child or youth and the secondary file contains data collected directly from the youth, aged 10 to 19, on the Self-complete Questionnaires. See Chapter 8.0 for more information about the content of the various questionnaires.

- Early Childhood Development (ECD) cohort includes children who are aged 0 to 5 years in Cycle 5. There are two files for the ECD cohort: the primary file contains data collected from the person most knowledgeable about the child and the education file includes data collected from the child's teacher, for those children attending kindergarten.

- The Northern cohort includes 5-year-olds from the Yukon only. Initially, children from all three territories were to be included in the survey but operational constraints prevented the Northwest Territories and Nunavut from participating in Cycle 5. See Chapter 17.0 for more detail about the Survey of Northern Children.

3.0 Objectives

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

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

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

4.0 Concepts and Definitions

There are many variables and concepts that are critical to the analysis of the National Longitudinal Survey of Children and Youth (NLSCY) data. The following is an explanation of the key concepts in the NLSCY. Derived variables are those that are not asked directly to the respondents but are calculated using information they have provided.

The content for each section of the various questionnaires used in the NLSCY is presented in Chapter 8.0.

The unit of analysis for the NLSCY is the child (0 to 15 years of age) or youth (16 to 19 years of age). See Chapter 5.0 for information on the survey design.

4.1 Definitions

Components
The NLSCY is made up of various components; these are generated by the computer application based on the child's age. The main components are: Adult, Child, Youth, Self-completes, Direct Measures (Peabody Picture Vocabulary Test - Revised (PPVT-R), Who Am I? and Number Knowledge), Math tests and Problem Solving Exercises. These components are described in Chapter 6.0, Data Collection.

An Adult component is generated for all children 0 to 15. For youth 16 and 17 years a modified adult component will be generated (see Survey Instruments for more details). The Adult component is asked of both the person most knowledgeable (PMK) and of the spouse of the PMK if living in the household. No Adult component is generated for Youth 18 or older.

There is a Child component generated for all respondents aged 0 to 15. Children aged 4 and 5 will also receive the Direct measures (PPVT-R, Who Am I? and Number Knowledge). All children 8 to 15, who are in Grades 2 to 10, will also complete a math test.

A Youth component is generated for all youth 16 and older. Youth 16 and 17 years of age will also be required to complete a problem-solving exercise. Parents of youth 18 and older will no longer be the PMK as the Youth will be self-reporting.

Computer-Assisted Interviewing (CAI)
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 complicated flows and edits to be built into the questionnaire which helps with both data quality and ensuring that each respondent answers only the questions appropriate to their situation. The questions are identical whether the interview is conducted using CAPI or CATI. Depending on the composition of the household and the nature of the required components, the interview will be conducted partly or completely by telephone and/or field visit.

Effective Age (**EMMCQ01)**
The age of the child determines, in most cases, the questions that will be asked about him or her. Instead of using the child's actual age, the NLSCY uses a calculated age called effective age. This is done to ensure the child stays in the age group to which he/she is assigned regardless of whether collection takes place before or after the child's birthday. For Cycle 5, the effective age is calculated as 2002 minus year of birth. For example a child born in 1998 would have an

effective age of 4 years old (2002-1998). Note that the actual age of the child at the time of the interview is sometimes different from the effective age.

Imputation
A procedure of completing a response by using values from one or more records on the same file or from external sources (e.g., historical data on respondents, administrative sources, etc.). This procedure is used for the income component of the NLSCY and the Motor and Social Development Score (see Chapter 11.0 for more information).

Actual province
This is the province of residence of the respondent at the time of the interview.

Province of residence at time of selection
This is the province of residence of the respondent at the time the sample was selected.

Early Childhood Development (ECD) Cohort
The ECD cohort is comprised of children aged 0 to 5 years.

Longitudinal Cohort
The longitudinal cohort for Cycle 5 is comprised of children and youths aged 8 to 19 who for the most part were introduced into the survey in 1994/1995 (Cycle 1). It is the intention to follow these children until they are 25 years of age.

4.2 Family Derived Variables

Using the NLSCY data, a child's family may be described in several different ways. Many of the family variables used to describe the NLSCY children were derived from what is known as the relationship grid. As part of the household questionnaire some basic demographic information was collected for all members of the child's household. Using this information it was possible to create an extensive set of variables to describe the child's family situation, including the relationship of each household member to all other members of the household.

The following are some of the family derived variables; the names of the derived variable are given in brackets.

Single-parent Families
Using the relationship grid resulting from the entry-exit component of the questionnaire, a child's single-parent status was derived (EDMCD04).

	Living with two parents	Living with one parent	Not living with a parent	Living independently
ECD Cohort	86.91%	12.85%	0.24%	N/A
Longitudinal Cohort	78.18%	18.4%	0.4%	3.02%

A child's parental status can also be defined in terms of the PMK (EDMPD06A). This variable is not applicable to youth as they are considered to be their own PMK.

	PMK with spouse/partner	PMK without spouse/partner
ECD Cohort	87.19%	12.81%
Longitudinal Cohort	70.40%	16.24%

The two variables describing the child's family are very similar. The small differences are the result of a few cases where the child lived with a parent, but the parent was not selected to be the PMK.

Intact, Step and Blended Families
Children/youth living with two parents are classified as being members of intact, step and/or blended families based on the relationship of these children to the parents.[1]

Intact Family
An intact family consists of a married or common-law couple in which all children are the natural and/or adopted offspring of both members of the couple (EDMCD16).

	Belongs to an intact family
ECD Cohort	78.19%
Longitudinal Cohort	60.50%

Step Family
A step family consists of a married or common-law couple, with at least one child living with them who is the biological or adopted child of one parent but not the other. It should be noted that a child who is the biological child of both parents is said to belong to a step family if at least one of these parents has a step child residing in the household (EDMCD03 and EDMCD15).

	Child is a step child	Child is part of a step family
ECD Cohort	1.97%	8.65%
Longitudinal Cohort	8.45%	9.65%

Blended Families
A blended family consists of a married or common-law couple living with at least two children, one of whom does not share the same natural and/or adoptive parents as the other child(ren). The following are examples of blended families:

- a couple with biological children of the female partner as well as biological children of the male partner (i.e., hers and his)
- a couple with biological children of the female partner as well as children out of the new union (i.e., hers and theirs)

The blended family is a sub-set of the step family (EDMCD14).

	Belongs to a blended family
ECD Cohort	7.29%
Longitudinal Cohort	5.03%

[1] Foster children and children living with only one parent are not included in intact, step, or blended families. In the derivation of intact, step and blended families, if a child was the adopted child of one parent and the biological child of the other parent, then this child was treated like a step child, and thus the family labeled as a step family. In other Statistics Canada publications children of this type are treated as if they were biological children of both parents.

Economic Family (EDMCD01)
For the NLSCY, an economic family is defined as all family members related by blood, marriage, common-law relationship or adoption. Foster children are considered to be part of the economic family. For example, if a woman lives in a household with her spouse and two children as well as her sister and her sister's child, then all of these individuals would be part of one economic family. If a boarder also resided in the household with her child, then this would constitute a second economic family.

Siblings
Siblings include full, half, step, adopted and foster siblings. Only siblings residing in the household have been included in the calculation of the sibling derived variables. In the case of common-law relationships, if both members have brought their own children into the relationship then these children are considered siblings. All siblings living in the household, including adult siblings, are included in the calculation of the sibling derived variables. The sibling derived variables include total siblings, as well as number of older siblings, younger siblings and siblings of exactly the same date of birth (i.e., twins) (EDMCD08, 09, 10 and 11).

Youth
New derived variables for youth 18 and 19 have been created in Cycle 5 to reflect the life changes that are becoming evident in this cohort.

EDMCed21 - Youth living with spouse, common-law partner, boy friend/girl friend (0.21%)
EDMCed22 - Youth have children of their own (0.20%)

The following variables will **not** be available for Youth 18 and 19 years: _DMCD06, _DMCD06A, _DMCD06B, _DMCD06C, _DMCD06D, _DMCD06E, _DMCD06F, _DMCD13, _DMCD18, _DMCD18B, _DMCD19, _DMCD19B, and _DMCD20.

4.3 Person Most Knowledgeable and Spouse

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

In some households, there is no PMK. In cases where the selected child/youth is 16 or 17 and is no longer living with a parent or guardian or for selected youth 18 years or older, there is no PMK in the household.

The following is the breakdown of the relationship of the PMK to the NLSCY children for Cycle 5.

	PMK is biological mother	PMK is step, adoptive or foster mother	PMK is biological father	PMK is step, adoptive or foster father	PMK is not a parent
ECD Cohort	87.87%	0.48%	10.79%	0.13%	0.72%
Longitudinal Cohort *	77.48%	1.4%	6.86%	0.31%	0.6%

*No PMK in 13.36% of cases.

For the majority of cases of the PMK not being a parent, the child had a parent living in the household, but the parent was not selected as the PMK. For the most part this situation occurred when a child had a very young mother living with her own parents, i.e., the child's grandparents, and the grandmother was selected as the PMK.

If the PMK had a partner residing in the household at the time of the interview, then this person was labeled as the spouse. Spouses included both married and common-law partners. Detailed socio-economic information was collected about the spouse/partner in order to describe the family situation of the child

Change in the PMK between cycles
For several reasons, the PMK and his/her spouse could be different people than those designated in the previous cycle. For this reason, analysts should use caution when comparing PMK information from one cycle to the next.

4.4 Respondent

A **cross-sectional respondent** is a child for whom the Adult or Child component was completed. These children represent the population as of January 2003.

A **longitudinal respondent** is a child introduced in a previous cycle for whom the Adult, Child or Youth component was completed. The children introduced in a previous cycle who have died or moved outside of one of the ten Canadian provinces are also longitudinal respondents. They represent children in the reference population who have the same life course (i.e. have died or moved).

A **respondent household** is a household where an Adult component or a Child or Youth component has been completed.

A **respondent child** is a child for whom an Adult component or his/her Child or Youth component has been completed. A respondent household without a complete Adult component can have one respondent child and one non-respondent child.

Please see Chapter 10.0, Survey Methodology – Response Rates, for more information about the definition of a respondent.

4.5 Socio-economic Status

In past cycles of the NLSCY, a measure of socio-economic status (SES) was included. This measure was not available for Cycle 4 and will continue to be unavailable in Cycle 5. The former definition used information about the respondent's employment as classified by the Standard Occupational Classification (SOC 1980). There is now a new coding structure, SOC 1991, and a definition of SES has not been developed using this new classification.

5.0 Survey Methodology - Sample

The National Longitudinal Survey of Children and Youth (NLSCY) is a probability survey designed to provide information about children and youth in Canada. To produce reliable estimates that meet the needs stated by clients, a representative sample of children and youth was selected. This chapter describes the sample selection method and the sample size.

For Cycle 5, the NLSCY sample includes children aged 0 to 5 and 8 to 19 years of age. As in previous cycles, the NLSCY still has the dual objective to produce longitudinal and cross-sectional estimates. However, for Cycle 5, cross-sectional weights are only produced for children aged 0 to 5.

The sample design of the NLSCY is determined to a large degree by the sample design of the Labour Force Survey (LFS). This applies to the stratification, allocation and selection of the sample as well as the estimation methods.[2]

5.1 Labour Force Survey Sample Design

The LFS is a monthly survey that collects labour market data from a national sample of about 60,000 dwellings. The current design was implemented at the end of 1994 following a redesign program that included a reassessment of the survey's principal role as a provider of current labour market information as well as a central vehicle for conducting household surveys within Statistics Canada.

5.1.1 Target Population

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

5.1.2 Stratification

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

Primary Stratification
Provinces are divided into economic regions (ER) and employment insurance economic regions (EIER). ERs are geographic areas of more or less homogeneous economic structure formed on the basis of federal-provincial agreements. They are relatively stable over time. EIERs are also geographic areas, and are roughly the same 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, sub-stratification takes place within these primary strata (see Secondary Stratification). Note that a third set of regions, census metropolitan areas (CMA), is also respected by stratification in the current LFS design, since each CMA is also an EIER.

[2] For a detailed account of the LFS methodology, see Methodology of the Canadian Labour Force Survey, Statistics Canada, Catalogue no. 71-526-XPB.

Types of Areas
The primary strata (ER/EIER intersections) are further disaggregated 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 to allow for the formation of strata in some areas. Urban areas include the largest CMAs down to the smallest villages categorized by the 1991 Census as urban (1,000 people or more), while rural areas are made up of areas not designated as urban or remote.

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

Approximately 1% of the LFS population is found in remote areas of provinces which 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. Some populations, not congregated in places of 25 or more people, are excluded from the sampling frame.

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

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

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

5.1.3 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 1991 Census counts. These clusters are generally a set of one or more city blocks or block-faces. The selection of a sample of clusters (always six or a multiple of six clusters) from each of these secondary strata represents the first stage of sampling in most urban areas. In some other urban areas, census enumeration areas (EA) are used as clusters. In the low density urban strata, a three stage design is followed. Under this design, two towns within a stratum are sampled, and then 6 or 24 clusters within each town are sampled.

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

Within each of the secondary strata in rural areas, where necessary, further stratification is carried out in order to reflect the differences among a number of socio-economic characteristics within each stratum. Within each rural stratum, six EAs or two or three groups of EAs are sampled as clusters.

5.1.4 Dwelling Selection

In all three types of areas (urban, rural and remote areas) selected clusters are first visited by enumerators in the field and a listing of all private dwellings in the cluster is prepared. From the listing, a sample of dwellings is then selected. The sample yield depends on the type of stratum. For example, in the urban area frame, sample yields are either six or eight dwellings, depending on the size of the city. In the urban apartment frame, each cluster yields five dwellings, while in the rural areas and urban EAs, each cluster yields 10 dwellings. In all clusters, dwellings are sampled systematically. This represents the final stage of sampling.

5.1.5 Sample Rotation

The LFS follows a rotating panel sample design, in which households remain in the sample for six consecutive months. The total sample consists of six representative sub-samples or panels, and each month a panel is replaced after completing its six month stay in the survey. Outgoing households are replaced by households in the same or a similar area. This results in a five-sixths month-to-month sample overlap, which makes the design efficient for estimating month-to-month changes. The rotation after six months prevents undue respondent burden for households that are selected for the survey.

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.

5.1.6 Household Members Eligible for the Labour Force Survey

The first month a dwelling is in the LFS, a roster containing information on the household composition is completed. Demographic information including the name, sex, date of birth and education level is obtained for all persons for whom the selected dwelling is the usual place of residence. Labour force information is obtained for all civilian household members 15 years of age or older. Respondent burden is minimized for the elderly (age 70 and over) by carrying forward their responses for the initial interview to the subsequent five months in the survey.

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

5.2 Birth Registry

When a significantly larger number of 1-year-old or 5-year-old children were needed, the Birth Registry was used.

This registry is created by the Health Statistics Division which is maintained by updates provided by the provincial registries to Statistics Canada. Some of the information given on the birth certificate is available on an electronic file. This file contains personal information such as the age of the mother, the birth date, the weight of the child at birth, the postal code and the code for the Census Subdivision (CSD). The remainder of the information is available on microfilm which consists of copies of the birth certificates.

The advantages and the disadvantages of this survey frame are summarized in the following table.

Advantages	Disadvantages
Good coverage. File is already available at Statistics Canada. The reverse record check of the Census of Population uses a similar approach. This allows us to take advantage of experiences gained from that project.	Does not include immigrants. This situation is significant for the 5-year-old children. The address provided is the one at the time of birth. In constructing the frame of 1 year old children, there can be administrative delays as some provinces are late in providing Statistics Canada with their registries. A large amount of time and resources are needed to create a sampling frame (stratification, formation of primary sampling units, capturing the information from the microfilm, merging information …) The name of the child is often incomplete.

In order to reduce collection costs, it was decided to implement a two phase sampling plan. In the first phase, a geographical area was selected, and second, the children born in this area. Many different geographic areas were considered to form the primary sampling units (PSU). None seemed to be appropriate. The CSD had too small a number of births, the Census Division had too many and covered too large an area.

The PSUs were formed by regrouping CSDs. In order to improve the efficiency of the sampling plan, the PSUs were selected with probability proportional to size (number of births). With this scenario, the CMAs would have had a large probability of selection. As the CMAs are relatively dense and two-thirds of children are born in CMAs, it was decided to classify these children into distinct strata.

To summarize, each province was sub-divided into two strata: the stratum of children born in a CMA and the stratum of children born outside of a CMA. For efficiency's sake, simple random sampling (SRS) was used in the CMA stratum. For the stratum of children born outside of a CMA, a PSU was selected using probability proportional to size and then a SRS of children living in the selected PSU's was chosen. Given the small number of births in Prince Edward Island, only an SRS of children was chosen for this province.

5.3 Target Population for Cycle 5

The goal of the NLSCY is to produce both longitudinal and cross-sectional estimates; therefore, it has several different target populations.

5.3.1 Longitudinal Target Populations

Longitudinally, the following populations are represented:

Children aged 0 to 11 in 1994/1995:
In Cycle 1, in 1994/1995, a sample of children aged 0 to 11 was selected. In Cycle 5, those children were between the ages of 8 and 19. Sample reductions were made in the sample in Cycle 2. As a result, only part of the sample is being followed longitudinally. Children dropped between Cycles 1 and 2 can be regarded as Cycle 1 cross-sectional children. It is important to note that, longitudinally, this cohort still represents children aged 0 to 11 in 1994/1995, who were aged 8 to 19 in 2002/2003. This cohort will be followed until the children reach the age of 25.

Children aged 0 and 1 in 1996/1997:
In Cycle 2, in 1996/1997, a longitudinal sample of children aged 0 and 1 was selected. About 2,000 children aged 0 and 2,000 children aged 1 were selected. This cohort was followed for only three cycles (2 through 4), and they are not part of the Cycle 5 sample.

Children aged 0 and 1 in 1998/1999:
In Cycle 3, in 1998/1999, a longitudinal sample of children aged 0 and 1 was selected. About 2,000 children aged 0 and 8,000 children aged 1 were selected. In Cycle 5, those children were 4- and 5-years-old and represent children aged 0 and 1 in 1998/1999. This cohort has been followed for only three cycles (3 through 5).

Children aged 0 and 1 in 2000/2001:
In Cycle 4, in 2000/2001, a longitudinal sample of children aged 0 and 1 was selected. About 2,000 children aged 0 and 2,000 children aged 1 were selected. In Cycle 5, those children were 2- and 3-years-old and represent children aged 0 and 1 in 2000/2001. They will be surveyed for the last time in Cycle 6.

Children aged 0 and 1 in 2002/2003:
In Cycle 5, in 2002/2003, a longitudinal sample of children aged 0 and 1 was also selected. About 2,000 children aged 0 and 2,500 children aged 1 were selected.

5.3.2 Cross-sectional Target Population

The NLSCY cross-sectional estimate, in Cycle 5, covers children aged 0 to 5 on January 1st, 2003. The user should note, however, that this cross-sectional sample is made up of various components:

1) The 4- and 5-year-olds are from the sample of 0- and 1- year-olds selected in Cycle 3.
2) The 2- and 3-year-olds are from the sample of 0- and 1-year-olds selected in Cycle 4.
3) The sample of children aged 0 and 1 were newly selected in Cycle 5.

Cycle 5 is the first cycle in which the children selected in Cycle 1 are not part of the cross-sectional target population. These children are now 8 to 19 years old, so these children and youth, in addition to the children selected in Cycles 3 to 5 as part of the Early

Childhood Development (ECD) program, make up a total sample of children aged 0 to 5 and 8 to 19. Since there are no 6- and 7-year-olds in the Cycle 5 sample, this sample cannot represent one continuous group of children aged 0 to 19 as of January 1st, 2003. Therefore, only the 0- to 5-year-old children are part of the cross-sectional target population.

In addition, the coverage of the cross-sectional sample in previous cycles was deteriorating over time. The children aged 8 to 19 in the sample were selected in 1994, and no updated sample was added to reflect changes in the population during that period. Also, attrition of the sample over five cycles also reduces the representativity of the sample. It was felt that the quality of the coverage of this population was too low to include it in cross-sectional estimation.

5.4 Sample Selection

As described in the section on target population, the Cycle 5 sample consists of several distinct parts. To help the user understand the current composition of the Cycle 5 sample, we should explain briefly how the sample was selected and how it has changed since Cycle 1. Following is a brief history of the sample selection for Cycles 1 to 5.

5.4.1 Cycle 1

Cohort 1
In Cycle 1, in 1994, the initial strategy involved selecting children aged 0 to 11 in each of Canada's 10 provinces. The objective was to be able to produce reliable provincial estimates by age group. The children were then to be followed until they reached the age of 25. Several frames were used to select the initial sample. Households with children in the target population (aged 0 to 11) were selected from the old-design Labour Force Survey (LFS), from the new-design LFS, and from the National Population Health Survey (NPHS). A total of 22,831 responding children made up our longitudinal sample. A breakdown of this total is provided in the Cycle 1 Microdata User Guide.

5.4.2 Cycle 2

Cohort 1
Sample reductions were made between Cycles 1 and 2 on the longitudinal cohort. First, the children from the NPHS were dropped. Then, to reduce the response burden on households with several eligible children, the number of children selected was limited to two per household. Some children were dropped from the sample but 16,903 children remained in our longitudinal sample. These children, all Cycle 1 respondents, were selected for Cycle 2 of the survey.

Cohort 2
A new initiative was introduced to the main survey in Cycle 2: Early Childhood Development. This new initiative focussed on a sample of children aged 0 and 1 to be followed until the age of 5. Two sources were used for the sample: first, we selected children from the LFS; then we added the siblings of Cycle 1 longitudinal children who were already in our sample. A total of 4,153 children were included in our second longitudinal cohort (our first ECD cohort).

Sample Buy-in
Following a request for additional sample, we also added a sample specifically for New Brunswick. The LFS was used for this sample as well. In all, 549 households were selected for a total of 480 responding children. These children were selected for cross-

sectional estimation purposes and were not intended to be followed in subsequent cycles.

5.4.3 Cycle 3

Cohort 1
In Cycle 3, we had the opportunity to convert Cycle 2 non-respondents back into the fold. The same initial sample as in Cycle 2 was used, except for deceased children (12), duplicate cases (3), children who were the wrong age for the survey (3), households that were not traceable in Cycle 2 (2), households that had moved permanently out of the country (52), children on Indian reserves (1) and households that were adamant refusals (112) as recorded in Cycle 2. In all, we excluded 185 Cohort 1 children from Cycle 3, for a longitudinal sample of 16,718.

Cohort 2
Responding children from Cycle 2 as part of the ECD initiative should all have been in the Cycle 3 sample. Unfortunately, an error in assigning cases to interviewers reduced the sample by 164 children. In all, 2,506 of the 2,670 children selected from the LFS were in the Cycle 3 sample. Similarly only responding siblings of Cycle 1 longitudinal children were contacted. A total of 1,483 children made up the second part of the Cohort 2 longitudinal sample.

Cohort 3
The ECD initiative continued in Cycle 3 with a new sample of children aged 0 and 1. The initiative was also interested in "the readiness to learn" aspect of children entering the school system. It was determined that a large sample of 5 year-olds was required to meet these analytical goals. Simultaneously, it was decided that a larger sample of 1 year old children selected in Cycle 3 would meet those objectives once they were 5 years-old in Cycle 5. A sample of about 10,000 would be selected. The LFS is not large enough to provide a sample of 10,000 children as it would require too many rotation groups. Typically, 12 LFS rotation groups, a year's worth of data, generates about 2,000 children aged 0 and 2,000 one-year-olds. Consequently, an additional sampling frame was used to procure a large sample of 1 year-olds. A sample of approximately 2,000 children aged 0 to 11 months was selected using the LFS. A sample of about 8,000 children aged 1 year old was selected using the Birth Registry. Stratification for this second sample is done differently than in the LFS. Because of the overlap between the two frames, certain eligible children were dropped from selection as their household had already been selected for children in the other cohorts. After excluding the children in households already in the survey, we had a sample of 7,944.

The number of 5-year-olds in Cohort 1 was not large enough to meet the analytical goals set out for studying the "readiness to learn" aspect of children entering the school system. To meet these goals, we also used the Birth Registry to select about 7,000 five-year-olds. These children were part of the ECD initiative and were sampled to meet analytical goals for the cross-sectional estimates in Cycle 3 only. After identifying and removing the duplicates from Cohort 1, an extra sample of 7,052 five-year-olds was added.

5.4.4 Cycle 4

Cohort 1
From the initial longitudinal sample of children introduced in Cycle 1, the historical pattern of response to the survey can be varied. Some children and youth have responded to all cycles, some did not respond in the last cycle, but have responded in previous cycles, some did not respond to previous cycles, but did respond to the last cycle. Several combinations of response patterns are possible. The NLSCY strategy was to try to re-

interview as many of the initial cohort as possible. However, it became inefficient to contact households that were unlikely to cooperate or untraceable. It was therefore decided to exclude households after two consecutive cycles of non-response. There were 518 such households excluded in Cycle 4. There were also hard refusals (473), deceased children (7), children who had moved away permanently out of the country (79) and children who had not responded in Cycle 2 and had moved temporarily in Cycle 3 (8). Those children were also considered to be non-respondents for two consecutive cycles. In all, 1,086 children were dropped from the longitudinal cohort sample from Cycle 1, leaving a total of 15,632 selected children.

Cohort 2
The longitudinal children introduced in Cycle 2 are now 4 and 5 years old. Cycle 4 is the last contact cycle for these children. First, to correct the error that occurred in Cycle 3, we added the 164 omitted children to the sample. We also attempted to convert non-respondents from the previous cycle. Only 38 children from the LFS were dropped from the Cycle 4 sample. The reasons for dropping them were age (2), death (2), permanent move (13) and hard refusal (21). A total of 2,632 children from the LFS were included in the Cycle 4 sample.

For the siblings of children introduced in Cycle 1, the situation was more complicated. We had decided in Cycle 2 to limit the number of children surveyed per household to two. However, the addition of siblings contravened that rule for many cases. Consequently, for households in which two children were already being surveyed, we excluded their younger sibling. This reduced the sample by 484 children. In addition we excluded 7 cases that had permanently moved out of the country, 13 hard refusals and 1 death. A total of 978 siblings of Cycle 1 longitudinal children remained in the sample for Cycle 4.

Cohort 3
For children introduced in Cycle 3, only respondents were contacted for this cycle. A total of 1,735 children from the LFS (1 hard refusal was excluded) were contacted again in Cycle 4, along with 6,383 children selected using the Birth Registry.

Cohort 4
The sample of children aged 0 and 1 from the LFS for this cycle was scaled back to the original ECD goals. In all, 5,031 unique households were chosen.

As in the previous cycle, there were not enough five-year-olds (introduced in Cycle 2) to meet the analytical goals of the “readiness to learn” aspect of 5 year-olds entering the school system. The Birth Registry was used once again to select a supplemental sample of 4,399 children.

5.4.5 Cycle 5

Cohort 1
From the initial longitudinal sample of children introduced in Cycle 1, the historical pattern of response to the survey can be varied. Some children and youth have responded to all cycles, some did not respond in the last cycle, but have responded in previous cycles, some did not respond to previous cycles, but did respond to the last cycle. Several combinations of response patterns are possible. The NLSCY strategy was to try to re-interview as many of the initial cohort as possible. However, it became inefficient to contact children and youth that were unlikely to cooperate or untraceable. It was therefore decided to exclude children and youth after two consecutive cycles of non-response where the respondent was under the age of 18. There were 243 such children excluded in Cycle 5. In the cases where the respondent was 18 or 19 years old, households were only excluded after three consecutive cycles of non-response. This

was done to retain more cases in the sample. When the respondent reaches 18 years old, it is their decision to respond to the survey, not the person most knowledgeable (PMK), so the response burden for them is not as great as for the younger children.

Also excluded from the sample in Cycle 5 were hard refusals (192), deceased children and youth (8), children and youth who had moved away permanently out of the country (72) and children and youth on Indian reserves (10). In all, 525 children and youth were dropped from the longitudinal cohort sample from Cycle 1, leaving a total of 15,163 selected children and youth.

Cohort 2
The children selected in Cycle 2 were all part of the ECD initiative; no children from Cohort 2 remain in the Cycle 5 sample.

Cohort 3
For children introduced in Cycle 3, only respondents from Cycle 4 were contacted for this cycle. A total of 6,949 children were contacted again in Cycle 5.

Cohort 4
For children introduced in Cycle 4, only respondents were contacted for this cycle. A total of 4,007 children were contacted again in Cycle 5.

Cohort 5
A total of 4,492 children aged 0 and 1 were selected from the LFS for Cycle 5.

5.4.6 Sample Sizes in Cycle 5

The number of children and youth sampled in Cycle 5 is shown by age and province in the following tables. Note that some children are purely cross-sectional and others are purely longitudinal. Examples of children who were longitudinal only are 5-year-olds introduced in Cycle 4 who had died or moved outside the country. These children were no longer in the target population for the cross-sectional sample, but longitudinally, they still represented the children of the year in which they were selected. For more details on these concepts, please see the chapter on weighting (Chapter 12.0).

There are no 6- and 7-year-old children in the sample for Cycle 5. This is because they are not part of the original survey population from 1994/1995 and they are also not part of the ECD initiative for children aged 0 to 5 years old.

Number of children in the sample by age, Cycle 5

Age	Number of Children	Age	Number of Children
0	2,012	11	1,278
1	2,480	12	1,231
2	1,794	13	1,151
3	2,213	14	1,050
4	2,560	15	1,058
5	4,389	16	1,122
8	1,783	17	1,052
9	1,858	18	1,172
10	1,330	19	1,078
		Total	**30,611**

Number of children in the sample by province, Cycle 5

Province	Number of Children
Newfoundland and Labrador	1,575
Prince Edward Island	900
Nova Scotia	2,053
New Brunswick	1,864
Quebec	5,497
Ontario	8,220
Manitoba	2,333
Saskatchewan	2,413
Alberta	3,077
British Columbia	2,668
Outside the 10 provinces	11
Canada	**30,611**

5.5 Sample Overview

The following graphic gives an overview of the longitudinal sample and the ECD sample.

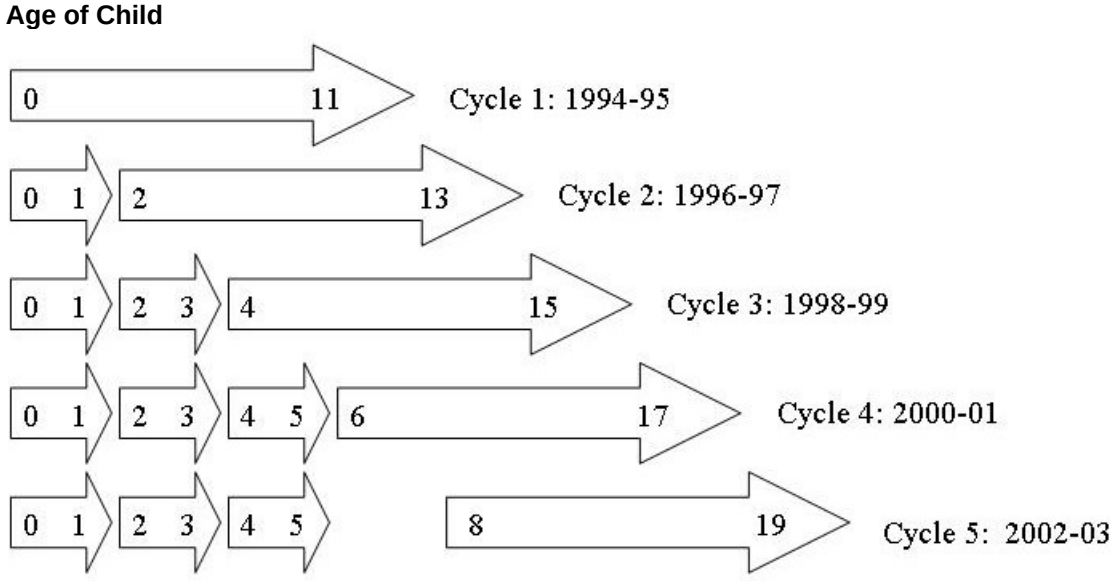

The larger arrows on the right represent the original longitudinal cohort. They were 0 to 11 years old in Cycle 1 and in Cycle 5 they are now 8 to 19 years old.

The smaller arrows represent the four ECD samples. These children were selected when they were 0 or 1 years old and remain in the sample for only three cycles, until they are 4 or 5. The ECD sample introduced in Cycle 2 was done with the survey in Cycle 4, and the ECD sample introduced in Cycle 3 was done with the survey in Cycle 5. The ECD samples introduced in Cycles 4 and 5 will continue in Cycle 6.

In Cycles 3 and 4 an additional cohort of five-year-olds was added to supplement the ECD component. These children are not longitudinal, so they are not represented in the image above.

6.0 Data Collection

Data for Cycle 5 of the National Longitudinal Survey of Children and Youth (NLSCY) were collected between the fall of 2002 and the spring of 2003. They were collected in two main settings: households and schools.

6.1 Household Collection

The survey combines computer-assisted interviewing methods and the use of paper questionnaires. Depending on the composition of the household and the nature of the required components, the interview was conducted partly or completely by telephone and/or field visit. This section provides a brief description of the "Collection tools" or the "Survey instruments", in other words the computer-assisted and paper questionnaire components used in the NLSCY collection.

6.1.1 Entry/Exit Component

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

The person most knowledgeable (PMK) about the child was also identified in this component. This was done once the information about the relationships between household members had been collected.

6.1.2 Child Component

A child component was created for each selected child between 0 and 17 years of age. The person most knowledgeable about the children and youth answered the child component questions. The PMK was usually the child's mother, but it could also be the father, a step-parent or an adoptive parent who lived in the same dwelling. Only the PMK or his/her spouse was permitted to answer the questions in this component.

At the end of this component, the respondent was asked to provide the name, address and telephone number of two people (friends, relatives) who would be able to help us trace the family in two years, when the survey will be repeated.

For the children attending kindergarten, in preparation for data collection in the schools, the PMK was asked to consent to the collection of information from the child's teacher.

List of subjects covered:

- Education
- Health
- Medical and biological information
- Mother's work after the child's birth
- Child's development
- Temperament
- Literacy
- Communication
- Activities
- Behaviour
- Positive behaviour
- Sleep habits

- Motor and social development
- Relationships
- Parenting
- Custody
- Child care
- Expectations (Aspirations)
- Socio-demographic characteristics

Note: For households in which the only child selected was in the 16 to 17 age subroup and was living with his/her parents, only three subjects were covered: Expectations (Aspirations), Custody and Socio-demographic characteristics. If the child was no longer living with his/her parents, the component was not created.

6.1.3 Adult Component

An adult component was created for the PMK and his/her spouse or partner. 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.

List of subjects covered:

- Education
- Labour force
- Income
- Health
- Family functioning
- Neighbourhood safety
- Social support
- Socio-demographic characteristics

Note: For households in which the only child selected was in the 16 to 17 age subgroup and was living with his/her parents, only four subjects were covered: Education, Labour force, Income and Socio-demographic characteristics. If the child was no longer living with his/her parents, the component was not created.

6.1.4 Youth Component

This component is used for 16- to 19-year-olds. 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.

List of subjects covered:

- Moving out of the parental home
- Education
- Labour force
- Career aspirations
- Income
- Health
- Activities
- Social support

6.1.5 Peabody Component (Peabody Picture Vocabulary Test – Revised)

The Peabody Picture Vocabulary Test - Revised (PPVT-R) was used to assess the child's level of receptive vocabulary. It was administered to each selected child aged 4 and 5 years. After obtaining the parent's oral consent, the interviewer asked the child the Peabody component questions directly.

6.1.6 Number Knowledge Component

The Number Knowledge is a direct measure that assesses the child's understanding of the concept of quantities and of the system of whole numbers. This component was administered to children aged 4 and 5 years. After obtaining the parent's oral consent, the interviewer asked the child the Number Knowledge component questions directly.

6.1.7 Control Screen Component

A control screen was created for each questionnaire or paper form required, to remind the interviewer to administer the appropriate child or youth questionnaire. Information (identification code, given name, etc.) was also transcribed from the screen to the paper questionnaire, and the questionnaire's sequence number was entered in the computer.

6.1.8 Who Am I?

The "Who Am I?" questionnaire is a direct measure that assesses child development and learning. This booklet was administered to four- and five-year-olds.

6.1.9 Mathematics Test

The mathematics test (computation exercise) is an objective indicator of the child's academic performance in mathematics. It was administered to children in Grade 2 or above, ranging in age from 8 to 15. It consisted of a set of nine booklets of varying levels of complexity. The level was determined by the child's grade. The test is administered by the interviewer in the child's home.

Level	Usually used for Grade
2	2
3	3
4	4
5	5
6	6
7	7 (Secondary 1 in Quebec)
8	8 (Secondary 2 in Quebec)
9	9 (Secondary 3 in Quebec) (Secondary 1 in Manitoba)
10	10 (Secondary 4 in Quebec) (Secondary 2 in Manitoba) (Level I in Newfoundland and Labrador)

6.1.10 Self-complete Questionnaires – Ages 10 to 19

Respondents between 10 and 19 years of age completed a paper questionnaire on various aspects of their lives. The youth was given the questionnaire during the interview and asked to complete it himself/herself. To ensure confidentiality, the child placed the completed questionnaire in an envelope, sealed the envelope and gave it to the interviewer.

The Self-complete Questionnaires consisted of a set of five booklets, one for each age group. The table below shows the subjects covered by each age-group section in the booklet. The questions for each subject were different for each age group. The booklets are reproduced in Book 2 of the National Longitudinal Survey of Children and Youth, Cycle 5 Survey Instruments 2002-2003.

Topic	Section				
	10 - 11 Booklet #20 E/F	12 - 13 Booklet #21 E/F	14 - 15 Booklet #22 E/F	16 - 17 Booklet #23 E/F	18 - 19 Booklet #24 E/F
Friends and Family	A	A	A	A	-
School	B	B	B	-	-
About me	C	C	C	B	A
Feelings and Behaviours	D	D	D	C	B
My Parent(s)	E	G	G	G	D
Smoking, Drinking and Drugs	G	F	F	D	E
Puberty	F	H	H	-	-
Activities	H	E	E	-	-
Dating / My Relationships	-	H	H	F	C
Health	-	H	H	E	-
Work	-	I	I	-	-

6.1.11 Cognitive Measure for 16- and 17-Year-Olds

The test was administered on a paper questionnaire to be completed by the respondent. It covered reading and mathematics. Two versions of the test were developed based on aptitude as measured by math tests completed by the respondent in previous cycles. Each booklet contained 18 questions designed to measure mathematical aptitude. The mathematics questions dealt with the use of mathematics in everyday activities, such as interpreting graphs and spatial diagrams and solving equations in order to make decisions.

6.1.12 List of Components and Questionnaires for Each Age Group and Interview Type

Ages 0 to 3 years

Interview type	Components	Approximate length of interview
Telephone	Entry/exit	75 minutes
	Adult	
	Child	

Ages 4 and 5 years

Interview type	Components	Approximate length of interview
Telephone and face-to-face interview	Entry/exit	140 minutes
	Adult	
	Child	
	Peabody component	
	"Who Am I?" booklet and its Control Screen component	
	Number Knowledge Component	

Ages 8 and 9 years

Interview type	Components		Approximate length of interview
Telephone and face-to-face interview	Entry/exit		83 minutes
	Adult		
	Child		
	Mathematics test and its Control Screen component	If in Grade 2 or above	

Ages 10 to 15 years

Interview type	Components		Approximate length of interview
Telephone and face-to-face interview	Entry/exit		90 minutes
	Adult		
	Child		
	Mathematics test and its Control Screen component	If in Grade 2 or above	
	Self-complete Questionnaire and its Control Screen component		

Ages 16 and 17 years

Interview type	Components		Approximate length of interview
Telephone and face-to-face interview	Entry/exit		105 minutes
	Adult	Only if youth lives with parents	
	Child		
	Youth		
	Self-complete Questionnaire and its Control Screen component		
	Cognitive Measure for 16-and 17-year-olds and its Control Screen component		

Ages 18 to 19 years

Interview type	Components	Approximate length of interview
Telephone and face-to-face interview	Entry/exit	105 minutes
	Youth	
	Self-complete Questionnaire and its Control Screen component	

6.1.13 Interview Method

The computer-assisted interview was completed by telephone for households in which all the selected children were aged 3 or less and for the Northern collection.

In the remaining households, where the selected children were aged 4 or over, the first few components of the computer-assisted interview were completed by telephone; the rest of the interview, which had both computer-assisted and paper components, had to be completed during a field visit.

6.1.14 Information Kit

Before the collection period, Statistics Canada's regional offices mailed an information kit on the survey to all parents and 16 to 19 year olds.

6.1.15 Collection Period

The collection period was divided into five phases between September 2002 and June 2003.

Wave 1		
Telephone	September 9 to November 15, 2002	Ages 0 to 3 years
Telephone and Personal Interview		
	September 9 to December 23, 2002	Ages 4 and 5 and 8 to 11
	November 1 to December 23, 2002	Ages 18 and 19
Wave 2		
Telephone and Personal Interview		
	January 3 to March 31, 2003	Ages 12 to 17
Wave 3		
Telephone	April 7 to June 13, 2003	Ages 0 to 3
Northern Collection		
Telephone	January to March, 2003	Age 5 years

6.1.16 Collection Personnel (Training, Supervision and Control)

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

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

CATI Training:
Classroom training for CATI interviewers was a one day course covering the following topics.

- Overview of the survey and its components
- Entry/Exit component
- Child Component
- Adult Component
- Tracing
- Overview of survey methodology

CAPI Training:
Classroom training for CAPI interviewers was a three day course covering the following topics.

DAY 1	Overview of the survey and its components "Contact/Demo" component Child Component (aged 1 year) Self-complete Questionnaires and Cognitive Measure (aged 10 to 17)
DAY 2	Child Component (aged 12) Math tests Tracing Youth Component (aged 18 and 19)
DAY 3	Non-response Child Component (aged 5) Direct Measures Interviewing young children Non-biased interviewing

6.2 School Collection

This phase of data collection took place in the schools between April and June 2003. For children in the sample who were attending kindergarten, the PMK was asked to consent to the collection of information from the child's teacher. Teachers were asked to complete the questionnaires and mail them back to Statistics Canada in the envelopes provided. The kindergarten teacher questionnaire was not administered in private schools.

The kindergarten teacher's questionnaire dealt with the child's academic performance and behaviour at school, the teacher's methods of instruction and the atmosphere in the classroom.

7.0 Data Processing

The main outputs of the National Longitudinal Survey of Children (NLSCY), Cycle 5 are "clean" data files. This chapter presents a brief summary of some of the processing steps involved in producing these files.

The processing of the NLSCY Cycle 5 data was done using the Generalized Processing Environment. This is a generic system that follows a series of steps to “clean” a file from beginning to end. The main steps were:

- Clean up
- Age and gender edits
- Relationship edits
- Pre-edit
- Longitudinal edits
- Flow edits
- Coding
- Consistency edits
- Derived variables

7.1 Computer Generated Edits

As discussed earlier, all of the information for the household (except for the 10 and 11, 12 and 13, 14 and 15, 16 and 17 and 18 and 19 year old Self-complete Questionnaires and the Kindergarten Teacher’s Questionnaire) was collected in a face-to-face or telephone interview using a computer-assisted interviewing (CAI) application. As such, it was possible to build various edits and checks into the questionnaire for the various household CAI components in order to ensure high quality of the information collected. Below are specific examples of the types of edits used in the NLSCY computer-assisted interviewing application:

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

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

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

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

7.2 Data Capture

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

- the Self-complete Questionnaires for 10- to 19-year-olds
- the Who Am I? for 4- to 5-year-olds
- the Math tests for 8- to 15-year-olds who are in grades 2 to 10
- the Cognitive Measure for 16- and 17-year-olds
- the Kindergarten Teacher's Questionnaire

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

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

7.3 Clean Up

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

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

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

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

The longitudinal file also contains 377 records that were created for some longitudinal children for whom no data was collected in this cycle. These are children who are now deceased or who have moved out of the country, but who will be kept on the longitudinal file for weighting

purposes. For these records, all variables except for the longitudinal weight (EWTCW01L) have been set to ”Not stated”.

7.4 Age and Gender Edits

In this step, verification of all age variables is conducted. A comparison to the previous cycle is done for the date of birth and the reported age. The age of all reporting children and youths is compared with the previous cycle. 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 establishes the relationship between the members of the household and creates the family derived variables. This step performs a standard set of edits against the relationship information entered for all members of a given household; some inconsistencies are corrected automatically by an application using a set of rules, while others are flagged for manual review and recoding. A related set of derived variables is produced through the relationship edits.

7.6 Pre-edits

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

7.7 Longitudinal Edits

Editing was also performed to ensure consistency between cycles. The longitudinal edits consist of carrying forward data from the previous cycle for questions that remain consistent over time and that are only asked during the first interview. For example, if the child’s birth weight was reported in Cycle 4, it was not asked again to the respondent in Cycle 5. However, the value does appear on the Cycle 5 file.

Flags were set for inconsistencies between cycles in the derived-variable step.

7.8 Flow Edits

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

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

7.9 Coding of Open-ended Questions

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

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

Career aspirations questions were asked in the Youth Questionnaire for 18- and 19-year-olds:

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

How they are recorded
The interviewer recorded, in words, the answer provided by the respondent. At Head Office, these written descriptions were converted into industry and occupation codes which describe the nature of the respondent's work. Similar information was collected for the spouse/partner and codes assigned to describe the nature of the work.

How they are coded
The open-ended questions were coded using various standard classifications. Occupation questions were coded using the 1991 Standard Occupational Classification codes (SOC) and the industry questions were coded using the 1997 North American Industry Classification System (NAICS). Grouped versions of these codes are available on the data file (ELFPcD7A and ELFPcD8A for the PMK, and ELFScD7A and ELFScD8A for the spouse/partner).

7.10 Consistency Editing

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

For example, in the Socio-demographic Section, for children who were not born in Canada, question ESDCQ2B asks what year they first immigrated to Canada. There was a consistency edit which compared the year of immigration to the child's year of birth. If the year of immigration was before the year of birth then the year of immigration was set to "Not stated" in the edit.

A data file was produced for the 10- to 19-year-old respondents. For questions that did not apply to an age group, the variables were set to "Not applicable" codes (6, 96, 996…).

Another data file was produced for the Kindergarten Teacher's Questionnaire respondents. Records on the Kindergarten Teacher's Questionnaire data file were removed if the parent or school board had not consented to participate in the survey. Also, the parent-reported grade and the teacher-reported grade of the child were compared. In cases where there was a discrepancy, the teacher-reported grade was accepted as correct.

7.11 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 has been carried out. Imputation is the process whereby missing or inconsistent items are "filled in" with plausible values. For the NLSCY, imputation was carried out for household income, PMK income, youth income, and motor and social development. See Chapter 11.0 for more details on imputation.

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

All imputation flag variables on the NLSCY data file have an "I" as the sixth character of the variable name. For example, the imputation flag variable for the income of the PMK would be named EINPcI1A.

7.12 Derived Variables

Combining Items
A number of variables have been derived by combining questions on the questionnaire in order to facilitate data analysis. For example, in the Labour Force Section, one of the questions is on the number of weeks worked while in the Adult Education Section, there is a question on whether they are currently going to school. The combination of these two questions forms a derived variable based on the actual situation of work and study.

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

Derived Variable Name
All derived variables on the NLSCY data file have a "D" as the fifth character of the variable name. For example, the name of the derived variable for the “Primary care arrangements” is ELFPD51.

7.13 Standard Coding Structures

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

Refusals
During a CAI interview, the respondent may choose to refuse to provide an answer for a particular item. The CAI system has a specific function key that the interviewer presses to indicate a refusal. This information is recorded for the specific item refused and transmitted back to Head Office.

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

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

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

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

In some cases a single question or series of questions was not applicable. For example, the question on number of hours per week the child is cared for in a daycare centre (ECRCQ1G1) is only applicable for children for whom this type of care is used (ECRCQ1G = 1). Otherwise there will be a code 996 for this question.

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

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

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

- As part of the CAI interview, the interviewer was permitted to enter a “Refusal” or “Don't know” code, as described above. When this happened, the CAI system was often programmed to skip out of this particular section of the questionnaire. In the case of a “Refusal”, it was assumed that the line of questioning was sensitive and it was likely that the respondent would not answer any more questions on this particular topic area. In the case of a “Don't know“, it was assumed that the respondent was not well enough informed to answer further questions. As part of the NLSCY processing system, it was decided that all of these subsequent questions should be assigned a “Not stated” code. A not stated code means that the question was not asked of the respondent. In some cases it is not even known if the question was applicable to the respondent.

- In some cases a questionnaire was not started or it was started but ended prematurely. For example, there may have been some kind of an interruption, or the respondent decided that she/he wished to terminate the interview. If there was enough information collected to establish the household as a responding household, then all remaining unanswered questions on the questionnaire (and on questionnaires that had not yet been started) were set to “Not stated”. The one exception was that if it was known that a certain section or a certain questionnaire was not applicable, then these questions were set to “Not applicable”.

- Some paper questionnaires were mailed back partially complete. If there was enough information collected about the respondent, then all remaining incomplete items on the questionnaire were set to “Not stated”. The one exception was that if it was known that a certain section or a certain questionnaire was not applicable, then these questions were set to “Not applicable”.

- Another situation in which not stated codes were used was as a result of consistency edits. When the relationship between groups of variables was checked for consistency, if there was an error, often one or more of the variables was set to “Not stated”.

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

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

7.14 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 straight-forward for the user.

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

Format for Variable Names
The **first** character of the variable name refers to the NLSCY cycle:

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

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

The **fourth** character refers to the collection unit or the unit to which the variable refers. There are nine possibilities[3].

“**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 youths
“**C**” if the variable refers to the Self-complete Questionnaire
“**T**” if the variable refers to the Kindergarten Teacher’s Questionnaire
“**W**” if the variable refers to a weight
“**M**” if the variable refers to the mother

The **fifth, sixth, seventh and eighth** characters of the variable name (for example ESDCQ2B and ELFPcD78) could refer to the following:

e the lower case letter refers to the NLSCY Cycle in which the variable first appeared on the file or the cycle in which changes to a previously asked question were made. For example, in the variable name ELFPcD78, "c" indicates the variable was new in Cycle 3. In subsequent cycles, new variables will also be

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

identified using the lower case letter to represent the cycle. New variables in Cycle 4 will contain a "d", in Cycle 5 an "e", etc. 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 5, the variable was treated as new and contains an "e".

Q refers to the variable for a question that was asked directly on one of the NLSCY questionnaires.

S refers to a score calculated for one of the scales used on the questionnaire.

D means the variable was derived from two or more questions that were asked on the questionnaire or coded variables.

I means the variable is a flag created to indicate that an item has been imputed.

Z means the variable is a flag created to indicate an inconsistency in reported data between the current and previous cycles.

nnx refers to the question or variable identification. Generally “nn” is a sequential number assigned to the variable; and “x” is a sequential alphabetic indicator for a series of variables of a similar type.

7.15 Acronym Names for the Questionnaire Sections

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

Acronym	Variable	Collected or Derived From:
GE	Geographic	Sample information.
HH	Household	Dwelling characteristics.
MM	Variables collected as part of the household roster.	Basic demographic variables for each household member. These variables are included on the NLSCY data file for the child, youth, the PMK and the spouse/partner.
DM	Demographic - derived to explain the living arrangements of the child.	Information from the household roster and relationship grid.
SD	Socio-demographic	Child on the Child Questionnaire and for the PMK and spouse/partner on the Adult Questionnaire.
HL	Health	PMK and spouse/partner on the Adult Questionnaire, for the child on the Child Questionnaire and for the youth on the Youth Questionnaire.
CH	Adult Chronic Conditions	PMK and spouse/partner in the Health Section of the Adult Questionnaire.
RS	Restriction of Activities	PMK and spouse/partner in the Health Section of the Adult Questionnaire.
DP	Depression Scale	Parent Questionnaire (this scale was administered to the PMK).
ED	Education	Children aged 4 to 15 years on the Child Questionnaire, for the PMK and spouse/partner on the Adult Questionnaire and for the youth on the Youth Questionnaire.
LF	Labour Force	PMK and spouse/partner on the Adult Questionnaire and for the youth on the Youth Questionnaire.
IN	Income	Household income and personal income of the PMK collected on the Adult Questionnaire and for the youth on the Youth Questionnaire.
FN	Family Functioning	Adult Questionnaire (section asked to the PMK or spouse/partner).
MD	Medical/Biological	Child Questionnaire (0 to 3 years)
TM	Temperament	Child Questionnaire (3 to 35 months).
LT	Literacy	Child Questionnaire (0 to 9 years).
AC	Activities	Child Questionnaire (0 to 15 years) and for the youth on the Youth Questionnaire.
BE	Behaviour	Child Questionnaire (0 to 11 years).
MS	Motor and Social Development	Child Questionnaire (0 to 47 months).
RL	Social Relationship	Child Questionnaire (4 to 9 years).
PR	Parenting Style	Child Questionnaire (0 to 15 years).
CR	Child Care	Child Questionnaire (0 to 13 years).

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

7.16 Examples of Variables Names

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

Variable Name Refers to:

Variable Name	Refers to:
ELFSQ2	Q2 in the Labour Force Section for the spouse/partner
E	Cycle 5 variable
LF	Labour Force Section
S	Spouse/partner
Q	An item asked directly on the questionnaire
2	The second question from the Labour Force Section of the spouse/partner questionnaire

Variable Name	Refers to:
EPRCS03	a positive interaction score on the parenting scale for a 2 to 15 year old child
E	Cycle 5 variable
PR	Parenting Section
C	Child
S	A score
03	The identification number of the item

7.17 Final Processing Files

The following files were created for the NLSCY, Cycle 5:

- Early Child Development Cohort 0 to 5 year olds (Adult, Child and Household)
- Longitudinal Cohort 8 to 19 year olds (Adult, Child, Youth and Household)
- Self-complete Questionnaires
- Education (Kindergarten Teachers)
- North (5-year-olds in the Yukon)

8.0 Content of the Survey

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

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

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

8.1 Survey Components

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

Household This is the first part of the interview. The household roster asks for basic demographic information for each household member and their relationship to everyone else in the household.

Adult Questions asked about the person most knowledgeable (PMK) and spouse. For children aged 16 and 17, not all the sections in the adult component are asked. The adult component is completed once even if there are two children in the household. No adult component is generated for youths aged 18 years and older.

Child Questions about the selected child asked to the PMK. A child component is completed for each selected child aged 0 to 5 and 8 to 17. The only sections of the Child Questionnaire asked about youth aged 16 and 17 are the Aspirations and Expectations section, Custody and the Socio-Demographics section.

Youth Questions asked about the selected youth, if he/she is aged 16 to 19 years old. In this section respondents answer questions about themselves in a computer-assisted interview (CAI)

Self-completes Respondents aged 10 to 19 answer questions about themselves in a paper questionnaire.

Direct assessments Several direct assessments are done with the children and youth; these are described in Chapter 16.0.

Kindergarten Teacher Questionnaires

For children who are in kindergarten, a questionnaire is sent to their teacher.

8.2 Demographic Variables

The demographic variables are collected in the household questionnaire. As part of the household questionnaire some basic demographic information (e.g., age, gender, and marital status) is collected for all members of the child's household. The relationship grid is also completed as part of this questionnaire i.e., the relationships of everyone in the household to all the other members of the household. Using this information it is possible to create an extensive set of variables to describe the child's family situation. Most of these derived variables are critical to the analyses of NLSCY data and are described in Chapter 4.0.

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

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

8.3 Adult Questionnaire

Education (Parent)

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

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

Labour Force

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

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

A complete description is recorded for the current or most recent job. Industry and occupation coding was carried out using the North American Industry Classification System (NAICS) 1997 and the 1991 Standard Occupational Classification codes.

Labour Force Derived Variables

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

ELFPcD5A / ELFScD5A:	NAICS code for PMK's / spouse's current job
ELFPcD6A / ELFScD6A:	SOC91 code for PMK's / spouse's main job
ELFPcD7A / ELFScD7A:	Standard industry code for current job - grouped
ELFPcD8A / ELFScD8A:	Standard occupation code for current job - grouped

Income
In the Income section of the survey, the sources of income and the income are collected for each household. There are also a few questions which ask for the perceptions of the PMK or the PMK's spouse regarding how well they think 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 11.0 for a detailed explanation of how income is imputed.

Two derived variables (EINHD04A and EINHD05A) have been created to compare the household income to the low income cut offs (LICO)[4]. LICOs are used to distinguish "low income" family units from "other" family units. A family unit is considered "low income" when its income is below the cut off for its family size and its community. The variable EINHD03A gives the value of the LICO by geographic area.

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

Chronic Conditions
PMKs and their spouses are asked whether or not they have any long-term conditions (e.g., allergies, asthma, and high blood pressure). A derived variable (ECHPdD01 or ECHSdD01) indicates that the respondent answered "Yes", they have at least one of the long-term conditions.

Restriction of Activities
The PMKs and their spouses are asked a series of questions about whether or not their activities are restricted at home, work, school, etc. A derived variable (ERSPdD01 or ERSSdD01) is also created stating whether or not the PMK or spouse reported an activity restriction.

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

Depression Scale
A Depression Scale (EDPPS01) was administered to the PMK as part of the Adult Questionnaire.

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

[4] For more information about Statistics Canada's low-income measures, please see Low income cut offs from 1994 - 2003 and low income measures 1992 - 2001, Catalogue no. 75F0002MIE2004002.

Neighbourhood Safety
This section gathers information about the respondent's satisfaction with his/her neighbourhood as a place to raise children, including perception of the extent of danger and problems, and of social cohesion or "neighbourliness". Two scales are created in this section: Neighbourhood Safety Score (ESFHdS5), indicating the degree of perceived neighbourhood safety and the Neighbours Score (ESFHS6), indicating the degree of neighbour cohesiveness.

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

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

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

8.4 Child Questionnaire

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 varied depending upon the age of the child, with more information being collected for the older children who have had greater school experience.

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

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

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

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

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

For a child four or five 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 (EHLCcD2A). The HUI3 is a generic health status index that is able to synthesize both quantitative and qualitative aspects of health. The index, developed at McMaster University's Centre for Health Economics and Policy Analysis, is based on the Comprehensive Health Status Measurement System (CHSMS). It provides a description of an individual's overall functional health, based on eight attributes: vision, hearing, speech, mobility (ability to get

around), dexterity (use of hands and fingers), cognition (memory and thinking), emotion (feelings), and pain and discomfort.

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

Medical/Biological

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

For each child under two, the nature of the delivery, general health of the child at birth and the use of specialized services following the birth were 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 were derived variables created for this section that should be noted. Two variables were derived to indicate the gestational age of the child. EMDCD06 gives the gestational age in days and EMDCD07 indicates if the child was born prematurely (gestational age 258 days or less), in the normal range (gestational age 259 to 293 days) or late (gestational age 294 days or later).

A variable was derived (EMDCD08) to indicate if the child was of normal birth weight (2,500 grams), moderately low birth weight (1,500 to 2,499 grams) or very low birth weight (< 1,500 grams).

Work After Birth

These questions are asked to determine the time interval after which mothers returned to work following the birth of a child and the extent to which these mothers participated 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 about 30 items covering five domains of development.

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

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

The NLSCY is using the ASQ's for children aged 3 to 47 months, inclusive. In consultation with the publisher, Statistics Canada has converted the questionnaires so they could be asked as part of the CAI application. The gross motor portion of the ASQs was dropped in Cycle 5 as this concept is covered in other portions of the survey.

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

Temperament
This section measures the temperament of young children by asking the parent about the degree of difficulty their child presents them. 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 children do certain activities with their parents, such as tell stories, sing songs and teach new words. These questions were adapted from the Early Childhood Longitudinal Study in the United States.

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

Communications
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 three year olds. Four and five year olds are asked a similar question as part of the Health Status Index in the Health Section.

Activities
This section measures the child's participation in various non-school activities and the amount of household responsibility taken on by 10 and 11 year olds at home. The latter questions are used to create the Home Responsibilities Score (EACCS6), indicating the degree of home responsibilities. The section will give some sense of how the child spends his/her time, of personal interests, as well as the degree of interaction with peers.

Several questions are included for children 4 and 5 years old and 8 and 9 years old to determine how often parents get to do certain activities with their children, such as eating a meal, playing a game and doing chores together. When there is a spouse/partner in the household, these questions are asked about both the PMK and his/her spouse/partner.

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

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

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

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

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

The questions in this section asked 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 three years; the questions vary by the age of the child. Three scores (EMSCS01, EMSCS02 and EMSCdS03) are derived from these questions.

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

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

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

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

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

Child Care
This section provides basic information about the methods of care currently provided for the child while the parents are working or studying, plus some information on previous care. Concepts measured include both the amount of time spent by the child in child care and the methods of care used for each child. In addition, information is obtained on the number of changes in child care arrangements that the child has experienced and the reason(s) for changes in the past 12

months. The section also identifies whether or not a child care centre is profit or non-profit, whether home care is licensed or unlicensed, and the ratio of caregivers to children.

Socio-demographics
These questions gather socio-demographic information on the selected child. Such information as ethnicity and country of origin is collected.

8.5 Youth Questionnaire (Ages 16 to 19)

Parent Report

Aspirations and Expectations (Ages 16 and 17)
These questions are included to assess parental aspirations and expectations for their youth, and parental views on their youth's school experiences. Providing help with school work, discussing school experiences and future educational plans has been linked to school success.

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

Youth Report (Computer-assisted Interviewing, Ages 16 to 19)

Moving Out of the Parental Home (Ages 18 and 19)
There are numerous transitions that a youth goes through from adolescence to adulthood. Undoubtedly, one of these major transitions is when the youth leaves the parental home for the first time to live independently in their own residence.

This section is designed to gather information on how many times the 18 and 19 year old respondents have left home. Some of these youth will be living away from home for school or work either permanently or temporarily. It was felt that information should be collected on this transition because of its importance in the movement from childhood to adulthood.

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

Youth Education (Ages 16 to 19)
This section looks at the youth's education experience. The section is divided into four parts:

1) school leavers (those who are not in school and have not graduated from high school),
2) school finishers (those who are not in school and have graduated),
3) currently in school (for youth still in high school), and
4) post-secondary (for youth who are attending a post-secondary education institution).

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

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

In Cycle 5, an integrated education section was developed. Because there would be many respondents of different ages and similar education statuses, the youth were streamed to the correct questions based not on age but education status.

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 and 19)
This section is new for the youth aged 18 and 19 in Cycle 5. The questions are similar to those asked of the youth aged 16 and 17. However, there are more questions taken from the adult labour force section due to the increased age of the respondents. Some of the youth may be working as their main activity and the questions need to reflect this possibility.

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

Youth Career Aspirations (Ages 18 and 19)
These questions are new in Cycle 5 for youth aged 18 years and older. 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.

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

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

For youth reporting a desired future career, occupation coding was carried out using the 1991 Standard Occupational Classification codes. From this information, the variable EASYED03 was created.

Youth Income (Ages 16 to 19*)*
The youth Income section asks the youth about their income from various sources in the last 12 months. These questions are similar to those asked of the parents. The degree to which youth make autonomous decisions may be measured through their spending behaviors.

In Cycle 5, a series of questions concerning payment of housing/shelter expenses were included for youth aged 18 and 19. At this age, many youth may be moving out of the parental home for the first time. Determining whether they are paying for shelter is important data to collect in order to assess how youth adjust to financial responsibilities.

Youth Health (Ages 16 to 19)
This section asks about the youth’s general health, injuries, chronic conditions and restriction of activities. These questions are similar to the child and adult health questions. There are

also some questions relating to the sleep patterns of the youth. Sleep is an important indicator of the youth's attitude toward their bodies and how they take care of themselves. The amount of sleep reported can be used to help understand if youth are successfully balancing the demands of work, school, volunteering, sports, etc.

Youth Health (Ages 18 and 19)
This section is new in Cycle 5 for youth aged 18 and older. It includes questions that ask about height, weight, questions relating to weight loss and weight gain. These questions are concerned with body image and the degree to which youth take steps to enhance the way they look.

Youth Activities (Ages 16 to 19)
Adolescence can be a time of high involvement in a variety of activities that are not school related. It is important to measure these activities to understand how this involvement can contribute to good outcomes. This section includes questions about physical activities, literacy activities, television watching, computer use and community involvement. These questions have been adapted from the questions asked of younger adolescents. The 16 and 17 year olds and the 18 and 19 year olds may receive different questions in this section depending on their age.

Questions are included for the youth aged 16 and 17 about the youth's access to a vehicle and whether or not 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.

Neighbourhood (Ages 16 and 17)
Neighbourhood factors have been shown to influence child and youth outcomes in a variety of domains (e.g., school achievement, behaviour, emotional and social functioning, motor and social development). These effects increase as children move through the life course, increasing their interactions and exposure to extra-familial environments. This has been evidenced in the academic literature, as well as by research conducted using NLSCY data (Boyle and Lipman, Kohen et. al, Offord and Lipman).

Youth Social Support (Ages 18 and 19)
This section is asked of the 18 and 19 year olds in Cycle 5 as part of the CAI application. These questions gather information on the youth's social support network by asking about how many close friends they have and the degree to which they are influenced by them.

Some of the questions in this section are similar to those asked of the younger youth in the Self-completes. In addition, the Social Support Scale from the adult CAI component was added to the 18 and 19 year olds questionnaire as a more age-appropriate measure.

8.6 Self-completes (Ages 10 to 19)

The objective of these questionnaires is to collect information directly from the youth on a variety of aspects of his/her life to supplement information obtained from the parent. The questionnaire also collects information from the youth on subjects about which only the youth could reliably report. For 16 to 19 year olds, some information is still collected on the Self-completes even though these youth are reporting their own information in the CAI portion of the interview. We felt that youth may be more comfortable answering sensitive questions on a paper questionnaire rather than face-to-face with an interviewer.

Friends and Family (Ages 10 to 17)
The objective of this section is to determine how well the youth feels he/she gets along with others.

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

The Friends Scale (EFFCS01) 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 Furnman and Buhmester's Network of Relationships Inventory.

School (Ages 10 to 15)
This section asks about the youth's attitude towards school, how well he/she is doing at school, the importance of good grades, feelings of safety and acceptance at school, perception of the teacher with respect to fairness and providing extra help. For 14 and 15 year olds, there is a series of questions about school based extra-curricular activities, such as sports or drama. These questions have been modified by the Project Team from the Western Australia Child Health Survey, Northwest Territories Health Attitudes, Knowledge and Behaviours Study, Marsh Self-Description Questionnaire, and the World Health Organization (WHO) Survey on Health Behaviours in School Children.

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

About Me (Ages 10 to 19)
These questions are used to determine the youth's overall self-esteem. A score is calculated (EAMcS02) based on the answers to these questions.

For youth aged 12 to 19, additional questions are asked about youths' 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". This measures the degree to which the youth relates to other people at home, school and 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 scale was developed by Dr. Reuven BarOn and Dr. James D.A. Parker. This measure was selected because it assesses the respondent's social, personal, and emotional abilities, as opposed to their behaviours.

Youth aged 14 to 19 are also asked about painful events, such as a break-up with a boyfriend/girlfriend or death of someone close to them.

Feelings and Behaviours (Ages 10 to 19)

Behaviour Checklist (Ages 10 to 15)
This section replicates the behaviour checklist used in the parent-report CAI Child Questionnaire. It provides indicators of the following behaviours: conduct disorder, hyperactivity, inattention, physical aggression, indirect aggression, emotional disorder, anxiety, and prosocial behaviours. Scores for these behaviours are also created.

Risky Behaviours (Ages 10 to 19)
These questions about risky behaviours, such as staying out all night without permission, are also replicated from the Child Questionnaire. The questions are expanded for the older age groups to capture behaviours that may become more common as the youth get older. These questions were adapted by the Project Team from the National Longitudinal Survey of Youth

at Ohio State University, Western Australia Child Health Survey and from questions provided by Dr. Richard Tremblay from the University of Montreal.

Suicide (Ages 12 to 19*)*
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. A score (EHTCbS1b) is calculated based on these questions.

My Parent(s) (Ages 10 to 19)
This section aims to capture the youths' relationship with their parents/guardians from several different angles. Questions are geared to uncover the amounts of understanding, fairness and affection received from each parent/guardian as well as conflict resolution practices and parental supervision. The youth's impression of their parents'/guardians' relationships and conflict resolution skills is also addressed.

My Parents and Me (Ages 10 to 15)
Three scales are created using these questions:

1) Parental Nurturance (EPMCcS1),
2) Parental Rejection (EPMCbS2b) and
3) Parental Monitoring (EPMCcS3).

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 (EPMCdS4), and
2) Conflict Resolution Scale – Father (EPMCdS5).

Puberty (Ages 10 to 17)
Puberty is an important marker of physical development. This section asks the youth about key physiological indicators and their perceptions of their own puberty. These questions were provided by Dr. Richard Tremblay from the University of Montreal.

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

Smoking, Drinking and Drugs (Ages 10 to 19)
This section asks questions to determine if the youth has used cigarettes, alcohol or drugs and the extent of usage. The behaviours have been correlated with negative behaviours and outcomes, such as delinquent behaviours and poor school performance. The questions vary by age.

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

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

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

Questions on driving under the influence of drugs and alcohol and being a passenger in a car with a drunk driver are included for 16 to 19 year olds. These are important risk-taking activities in this age group. The questions have been adapted from the North Carolina Evaluation of School-Based Health Centers.

Activities (Ages 10 to 15)
The objective is to determine the youth's extent of participation in activities outside of school hours and use of free time. Activities include sports, arts, dance or music, Guides or Scouts, jobs and volunteering. Reading for pleasure, using a computer and watching television are also covered. Generally, the activities are also covered on the CAI parent-report Child Questionnaire for children under 10.

Literacy Activities (Ages 14 and 15)
These questions ask about how often youth engage in literacy activities outside of school, such as using a library or reading. These questions are similar to those asked of the PMK for younger children and of the 16 to 19 year olds in the CAI questionnaire.

Health (Ages 12 to 17)
Youth are asked to report on their height and weight, symptoms of stress, use of seatbelts and helmets, healthy eating and dating. The questions vary with age.

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

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

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

Work and Money (Ages 12 to 15)
Youth are asked about their work during the school year and those aged 14 and 15 are asked about work last summer. The 14 and 15 year olds are asked more detailed questions about their job(s), such as hours worked and pay. They are also asked about whether work reduces the amount of time they spend studying.

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

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

Dating (Ages 12 to 19)
This section asks youth about their experiences with a boyfriend/girlfriend and their sexual activity. The question about sexual behaviour on the 12 and 13 year old questionnaire was modified from the Youth and Aids Survey. Questions are also asked about contraceptive use and, for the 16 to 19 year olds, reasons for abstaining from sex or reasons for not using birth control. These questions were designed by the Project Team in consultation with experts from youth surveys such as the 1992 British Columbia Adolescent Health Survey and the Minnesota Adolescent Health Survey.

9.0 Validation of the Survey Scales

9.1 Validation of Scale Data

9.1.1 Scale Definition

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

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

9.1.2 Scales and Calculations

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

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

9.1.3 Evaluation of Scale Data

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

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

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

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

9.2 Factor Analysis

9.2.1 Factor Analysis for Scales

The factor structure of each scale was determined based on data from the first cycle. The factor structure imposed on the scales already used in the first cycle and repeatedly utilized in the second, third, fourth and fifth cycles of the survey was the result of analysis of data from the first cycle. This analysis was redone for the majority of the scales for Cycle 4 and it is these results that are summarized below. For the results from the Cycle 1 factor analysis, please refer to the Cycle 1 Microdata User Guide.

1. The sample of respondents for each scale (and age group, if the scale used different questions for different age groups), was randomly divided into two half-samples. This was done to find out whether different samples would yield the same results.

2. Principal component analysis was carried out separately on each half-sample to find out how many factors should be extracted in the subsequent factor analysis. In principle, the same number of factors as was found in the literature was expected. In practice, however, some scales showed a different number of factors because in some cases factors combined while in others new factors emerged.

3. Factor analysis was done on each half-sample and the factor structure and loading of each factor were compared across the half-samples.

4. In the factor analysis, the items for each child in the appropriate age group were used, multiplied by the child's normalized weight. An individual's statistical weight is normalized by dividing his/her weight (_WTCW01C)[5] by the average weight for all individuals. Thus, the sum of the normalized weights is equal to the sample size.

5. Once the factor structures were analysed and the items included in each factor were determined, scores were calculated. To produce the scores, one was subtracted from each item so that the lowest possible score would be zero (0). A score of zero indicates that the child has no problems for all factors in the Behaviour Scale except for the prosocial factor, where a score of zero indicates the absence of prosocial behaviour. Some items were imputed. The imputed values were computed by a procedure (the SAS PRINQUAL procedure) that determines which of the possible values for an item is the most plausible for an individual in view of his/her response profile, the response profiles of others in the sample, and the number of factors included in the analysis.

6. The score for each factor on the scale was derived by totaling the values of the items that made up that factor (including imputed values). The score was set to "missing" if too many of the values of any items included in the factor were unreported. A value may be missing if the parent refused to answer or did not know the answer to the item.

[5] 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 EPRCS01 in Cycle 5 begins with the letter "E" on the microdata file and here is referred to as _PRCS01.

9.2.2 Data Transformation Using Optimal Scaling

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

9.2.3 Factor Analysis Using Weighted Data

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

9.3 Calculation of Scores and Item Imputation

9.3.1 Calculation of Scores for Each Factor

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

9.3.2 Example of Factor Score Computation

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

_PRCQ04	How often do you get annoyed with your child for saying or doing something he/she is not supposed to?
_PRCQ08	Of all the times you talk to your child about his/her behaviour, what proportion is praise?
_PRCQ09	Of all the times you talk to your child about his/her behaviour, what proportion is disapproval?
_PRCQ13	How often do you get angry when you punish your child?
_PRCQ14	How often do you think the kind of punishment you give your child depends on your mood?
_PRCQ15	How often do you feel you have problems managing your child in general?
_PRCQ18	How often do you have to discipline your child repeatedly for the same thing?

The answer categories for these items were of two types:

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

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

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

9.3.3 Negative Loading

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

9.3.4 Non-response Code

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

When the score was being calculated for each factor there was a possibility that one or more of the items making up the score had a non-response code (“Don't know”, “Refusal” or “Not stated”). If the number of items with a non-response code was above a certain threshold, the factor score was set to “Not stated”. Generally this threshold value was set at 10% of the items.

9.3.5 Raw Items

It should be noted that in addition to the scores, the raw items for each scale are included on the microdata file. This will allow researchers to consider alternate factor structures if desired. For the raw items the original values (in the 1 to 5 range for the Parenting Scale) have been retained before any rescaling or reversal of values took place.

9.4 Reliability Measures for Scales

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

9.4.1 Cronbach's Alpha

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

9.4.2 Interpretations of Cronbach's Alpha

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

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

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

9.4.3 What is a Satisfactory Level of Reliability?

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

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

9.5 Parent-reported Scales

The remainder of this chapter provides and 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. For convenience, the scales are listed in alphabetical order. The table below provides a brief summary of the NLSCY scales followed by individual descriptions of each scale.

9.5.1 Behaviour Scale

The objective of the Behaviour Scale is to assess aspects of the behaviour of children two years of age and older.

Initially, an attempt was made to measure the following behaviours for children aged 2 and 3:

- hyperactivity,
- emotional disorder,
- anxiety,
- physical aggression,
- inattention,
- prosocial behaviour,
- separation anxiety and
- opposition

For children between 4 and 11 years of age, an attempt was made to measure similar behaviours to the 2 and 3 year olds; separation anxiety and opposition behaviours were omitted, while indirect aggression, conduct disorder and property offences were added.

Theoretical Constructs
Below are the theoretical constructs that were used for the factor analysis. The actual scales that emerged from the analysis vary from these constructs.

Separation Anxiety (Ages 2 and 3 years)
Includes items _BECQ6CC, _BEC6QDD1, _BECQ8LL1, _BECQ8PP1 and _BEC8TT1 from Achenbach's Child Behavior Checklist (CBCL).

Opposition (Ages 2 and 3 years)
Includes items _BECQ6G, _BECQ6R1, _BECQ8E1, _BECQ8T1, _BECQ8Z1 and _BECQ8J1 drawn from Achenbach's CBCL.

Conduct Disorder (Ages 2 to 11 years)
Includes items _BECQ6G, _BECQ6X, _BECQ6AA, _BECQ6FF, _BECQ6JJ and _BECQ6NN from the Ontario Child Health Study (OCHS).

Hyperactivity (Ages 2 to 11 years)
Includes items _BECQ6B, _BECQ6I, _BECQ6P and _BECQ6W from the OCHS and _BECQ6QQ and _BECQ8HH from the Montreal Longitudinal Survey. In previous cycles, item _BECQ6N was included in this construct. A decision was made to drop this item from Cycle 4 and all future cycles as respondents found it to be too repetitive.

Emotional Disorder (Ages 2 to 11 years)
Includes items _BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM and _BECQ6RR from the OCHS. Anxiety includes NLSCY items taken from OCHS emotional disorder items (_BECdQ6F, _BECQ6Q, _BECQ6V and _BECQ6CC). In previous cycles the items _BECQ6Y and _BECQ6II were included. A decision was made to remove both items from Cycle 4 and all future cycles.

Indirect Aggression (Ages 2 to 11 years)
Includes items _BECQ6J, _BECQ6R, _BECQ6Z, _BECQ6LL and _BECQ6TT from Lagerspetz, Bjorngvist and Peltonen of Finland.

Physical Aggression (Ages 2 and 3 years and ages 8 to 11 years)
Includes items _BECQ6X from the Montreal Longitudinal Survey and _BECQ6G, _BECQ6AA and _BECQ6NN from the OCHS.

Inattention (Ages 2 to 11 years)
Includes items _BECQ6P from the OCHS and _BECQ6QQ from the Montreal Longitudinal Survey.

Prosocial behaviour (Ages 8 to 11 years)
Includes items _BECQ6A, _BECQ6H, _BECQ6M, _BECQ6GG and _BECQ6OO from the OCHS and _BECQ6D, _BECQ6U, _BECQ6BB, _BECQ6SS and _BECc6UU from the Montreal Longitudinal Survey; the last four items are from a scale devised by K. Weir and G. Duveen. In Cycles 1 through 3, these items were asked of all children aged 4 to 11. In Cycle 4, all 4 and 5 year olds were excluded from this scale and were asked the questions in the positive behaviour section instead.

Factor Analysis for the Behaviour Scale
The following indicates the items that were included on the questionnaire to measure these various constructs of behaviour. A complete factor analysis was carried out for the Behaviour Scale to assess the psychometric properties of this scale for the NLSCY population. As part of this analysis, the items that loaded into each construct or factor were compared to the expected result described below. The results of this analysis are presented later on in this section.

Analysis of Children Aged 2 and 3 years, based on Cycle 4 data
In the sample there were 7,122 children aged 2 and 3 years. The group was split into two sub-samples of 3,477 and 3,645 individuals, and the analysis for this age group was performed separately for each sub-sample. The non-response rate for most items was about 2.0%. Some individuals were excluded from the analysis that produced the factors. The exclusion criteria were as follows: individuals with 10% or more items coded "missing" or refused were not included in the analysis. After the criteria were applied, there were 3,413 and 3,565 individuals left in the sub-samples to be analysed.

FACTOR	SCORE	ITEMS
Hyperactivity - inattention	_BECdS01	_BECQ6B, _BECQ6I, _BECQ6P, _BECQ6W, _BECQ6QQ, _BECQ8HH
Emotional disorder - anxiety	_BECdS03	_BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM, _BECQ6RR
Physical aggression - opposition	_BECS04	_BECQ6G, _BECQ6X, _BECQ6NN, _BECQ6R1, _BECQ8E1, _BECQ8T1, _BECQ8Z1, _BECQ8J1
Separation anxiety	_BECS05	_BECQ6CC, _BECQ6DD1, _BECQ8LL1, _BECQ8PP1, _BEC8TT1

Cronbach's Alpha for Children Aged 2 and 3 years, based on Cycle 4 data
Cronbach's alpha (raw value) was computed with SAS using normalized weighted data.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Hyperactivity - inattention (_BECdS01)	0.739	_BECQ6P	0.684
Emotional disorder - anxiety (_BECdS03)	0.600	_BECQ6MM	0.518
Physical aggression - opposition (_BECS04)	0.716	_BECQ8Z1	0.677
Separation anxiety (_BECS05)	0.584	_BECQ6DD1	0.460

Analysis of Children Aged 4 to 11 years, based on Cycle 4 data
There were 13,765 children in the 4 to 11 age group. Two sub-samples of 6,830 and 6,935 were created for analysis. The item non-response rate was approximately 3.5% for most of the 47 items involved in the analysis. Individuals were excluded from the analysis if there were 10% or more items coded "missing," or refused. After this criteria was applied 6,574 and 6,681 individuals remained in the sub-samples to be analysed.

FACTOR	SCORE	ITEMS
Hyperactivity - inattention	_BECdS06	_BECQ6B, _BECQ6I, _BECQ6P, _BECQ6S, _BECQ6W, _BECQ6QQ, _BECQ8HH
Emotional disorder - anxiety	_BECdS08	_BECdQ6F, _BECQ6K, _BECQ6Q, _BECQ6V, _BECQ6MM, _BECQ6RR, _BECQ6CC
Physical aggression - conduct disorder	_BECdS09	_BECQ6G, _BECQ6X, _BECQ6AA, _BECQ6FF, _BECQ6JJ, _BECQ6NN
Indirect aggression	_BECS10	_BECQ6J, _BECQ6R, _BECQ6Z, _BECQ6LL, _BECQ6TT

Cronbach's Alpha for Children Aged 4 to 11 years, based on Cycle 4 data
Cronbach's alpha (raw value) was computed with SAS using normalized weighted data.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Hyperactivity - inattention (_BECdS06)	0.815	_BECQ6P	0.774
Emotional disorder - anxiety (_BECdS08)	0.736	_BECQ6V	0.684
Physical aggression - conduct disorder (_BECdS09)	0.772	_BECQ6AA	0.714
Indirect aggression (_BECS10)	0.766	_BECQ6LL	0.711

Note: The scores for these factors could not be computed in, 510, 501, 501, and 1,031 cases respectively because of unreported values.

Analysis of Children Aged 6 to 11 years, based on Cycle 4 data
There were 7,687 children in the 6 to 11 age group. Two sub-samples of 3,751 and 3,936 were created for analysis. The item non-response rate was approximately 3.8% for the items involved in the analysis. Individuals were excluded from the analysis if there were 10% or more items coded "missing," or refused. After this criteria was applied 3,407 and 3,578 individuals remained in the sub-samples to be analysed.

FACTOR	SCORE	ITEMS
Prosocial behaviour	_BECdS07	_BECQ6A, _BECQ6D, _BECQ6H, _BECQ6M, _BECQ6U, _BECQ6BB, _BECQ6GG, _BECQ6OO, _BECQ6SS, _BECc6UU

Cronbach's Alpha for Children Aged 6 to 11 years, based on Cycle 4 data
Cronbach's alpha (raw value) was computed with SAS using normalized weighted data.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Prosocial behaviour (_BECdS07)	0.831	_BECQ6SS	0.804

Analysis of Children Aged 8 to 11 years, based on Cycle 4 data
There were 4,399 children in the 8 to 11 age group. Two sub-samples of 2,186 and 2,213 were created for analysis. The item non-response rate was approximately 2.8% for the six items involved in the analysis. Individuals were excluded from the analysis if there were 10% or more items coded "missing," or refused. After this criteria was applied 2,081 and 2,093 individuals remained in the sub-samples to be analysed.

FACTOR	SCORE	ITEMS
Property offences	_BECdS11	_BECQ6C, _BECQ6E, _BECQ6L, _BECQ6T, _BECQ6DD, _BECQ6PP

Cronbach's Alpha for Children Aged 8 to 11 years, based on Cycle 4 data
Cronbach's alpha (raw value) was computed with SAS using normalized weighted data.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Property offences (_BECdS11)	0.623	_BECQ6C	0.544

NOTE: The scores for these factors could not be computed in 225 cases because of unreported values.

9.5.2 Depression Rating Scale

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

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

Items Included in the Depression Rating Scale
The Depression Rating Scale includes 12 questions, each of which contains four response categories. In order for the lowest score value to be zero, the value for each question was reduced by one in calculating the score. As well, the answer categories were reversed for questions having a negative loading (_DPPQ12F, _DPPQ12H, and _DPPQ12J). The total score (_DPPS01) may therefore vary between 0 and 36, a high score indicating the presence of depression symptoms.

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 done based on data from the first cycle.

Analytical Results, based on Cycle 1 data
In analysing this scale, unweighted data were used. The sample size was 13,439 PMK's. However, once the observations containing mostly missing values were eliminated, the analysis dealt with only 13,140 PMK's. The non-response rate for the various questions in the rating scale was roughly 2.0%, whereas for the total score, a non-response rate of 2.2% was obtained. There was no imputation for the variables in this rating scale.

Cronbach's Alpha for Depression Rating Scale, based on Cycle 1 data
In spite of the possibility of extracting more than one factor from the Depression Rating Scale, a single-factor analysis was used since the interest was in developing a global depression index. Following the analysis, the 12 variables of the scale were all kept as components of this factor since all 12 loading values met the established threshold.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Depression (_DPPS01)	0.820	_DPPQ12D	0.790

9.5.3 Family Functioning Scale

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

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

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

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

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

Analytical Results, based on Cycle 1 data
In analysing this scale, unweighted data were used. The non-response rate for the different variables was between 1.3% and 1.4%, whereas for the total score, a non-response rate of 1.9% was obtained. There was no imputation for the variables in this scale.

Cronbach's Alpha for Family Functioning Scale, based on Cycle 1 data

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Family functioning (_FNHS01)	0.880	_FNHQ01L	0.870

Distribution of Values for the Family Functioning Scale
When the values for the factor score for the Family Functioning Scale are examined for the NLSCY children, the distribution that is observed is not a continuous one. In fact, the most common score is 12. This is a result of the fact that there are 12 items in the scale and four possible rescaled values (0 to 3). Many respondents had a rescaled score of one for every item in the scale and thus an overall score of 12. This means that the respondent answered "Agree" to all of the items in the scale which were positive and "Disagree" to all of the negative items, as opposed to the more extreme answers of "Strongly agree" or "Strongly disagree."

9.5.4 Home Responsibilities Scale

Objectives and Overview
The object of the activities scale is to measure the child's participation in home responsibilities.

This set of questions about responsibilities is from the Home Observation for Measurement of the Environment-Short Form questionnaire in the National Longitudinal Survey of Youth, Ohio State University.

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

Analytical Results, based on Cycle 1 data
In the Cycle 1 sample there were 3,434 children aged 10 and 11 years. They were divided into two sub-samples of size 1,705 and 1,729 and an analysis was done on each sample. The non-response rate for the five items was 1.3%. Individuals with missing values were excluded from the analysis that was conducted for the purpose of constructing the factor. After these exclusions, the sub-samples contained 1,680 and 1,709 individuals respectively. No imputation took place. As a result of factor analysis, one factor was identified: the activities factor (_ACCS6). Items _ACCQ6A to _ACCQ6E loaded into the factor.

Scale Score
To produce the score, one was subtracted from each item so that the lowest score would be zero (0). The values for each item were reversed so that a high score would indicate a high degree of home responsibilities. The final score was derived by totalling the values of all items with non-missing values. The score ranges from 0 to 15. A score of zero indicates the respondent does not participate in home responsibilities.

Once the factor structures were analysed and the items included in the factor was determined, the score was calculated. No imputation was done on the values. If any values were missing the final score was set to missing. A value may be missing if the child refused to answer or did not know the answer to the question.

Cronbach's Alpha for Home Responsibilities Scale, based on Cycle 1 data

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Home responsibilities (_ACCS6)	0.778	_ACCQ6B	0.705

Note: The final activities score could not be calculated for 45 (1.3%) individuals, due to missing values for the items comprising this factor.

9.5.5 Motor and Social Development Scale

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

The Motor and Social Development Scale
The Motor and Social Development (MSD) Scale was developed by Dr. Gail Poe of the United States National Center for Health Statistics. The MSD Scale consists of a set of 15 questions that vary by the age of the child, which measure dimensions of the motor, social and cognitive development of young children from birth through three years. 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 recent versions of the National Child Development Survey in England.

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

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

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

Although there were different sets of questions depending on the age in months of the child, differences were observed when comparing score within these age bands. For example, there was a specific set of questions for children aged 4 to 6 months. It was found that children who were 6 months old had scores that were on average higher than those 4 months old. Therefore a decision was made to produce standardized scores. These scores, calculated for each age in months, would make it possible to compare scores across ages. All children, aged three years or less, received a standardized score based on Cycle 1 data and a standardized score based on the Cycle 5 data.

Standardized Scores based on Cycle 5 norms
Each child aged 3 to 47 months was assigned a standard score. This standardization was done by one month age groups. For each month age group the mean and standard deviation of the raw score was found and were used to produce a normalized score with a mean of zero and a standard deviation of one. This score was adjusted such that the mean MSD score was 100 and the standard deviation was 15. Therefore children who are 3 months old have an average MSD score of 100, children who are aged 4 months have an average MSD score of 100, ... and children aged 47 months have an average MSD score of 100.

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

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

Standardized Scores based on Cycle 1 norms
A second standardized score (_MSCS02) was calculated for all children 0 to 47 months. This score was calculated in the same way as mentioned above, except that the norms were derived using the data from Cycle 1, and then applied to the Cycle 5 data.

Overall there are no major differences found when comparing the scores found using Cycle 5 norms and the scores found using Cycle 1 norms. The score calculated using Cycle 1 norms should be used in order to compare scores over cycles. This score is available for all cycles of data.

The Motor and Social Development Scale questions have remained unchanged throughout the five 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.

9.5.6 Neighbourhood Safety Scale

Objectives and Overview
To gather information on the respondent's satisfaction with his/her neighbourhood as a place to raise children, including perception of the extent of danger and problems, and of social cohesion or "neighbourliness". Recent 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. The information on the parent's perception of the neighbourhood is supplemented by the interviewer's observation of several aspects of the block where the respondent lives.

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

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

The Cycle 5 scale questions are the same as the Cycle 1 questions with the exception of _SFHQ05C where there has been a small wording change. Also, the questions that made up the neighbourhood problems scale in Cycle 1 (ASFHQ07A to ASFHQ07F) have not been included in Cycle 5.

Analytical Results, based on Cycle 4 data
In the sample there were 31,744 children. They were divided into two sub-samples and analysis was done on each sub-sample. Individuals with missing values were excluded from the analysis conducted for the purpose of constructing the factor. After these exclusions the sub-samples contained 15,720 and 16,024 individuals respectively. As a result of the factor analysis, two factors were identified: the neighbourhood safety factor and the neighbours factor. The items that comprised each factor are described in the following table:

FACTOR	SCORE	ITEMS
Neighbourhood safety	_SFHS5	_SFHQ05A, _SFHQ05B, _SFHQ05C
Neighbours	_SFHS6	_SFHQ06A, _SFHQ06B, _SFHQ06C, _SFHQ06D, _SFHQ06E

Scale Score
If too many values were missing the final score was set to missing. To produce the final scores, one was subtracted from each item so that the lowest score would be zero (0). All the score values were reversed. The final score was derived by totalling the values of all items with non-missing values. A score of zero indicates the following for the two neighbourhood scales:

- a low degree of neighbourhood safety
- a low degree of neighbourhood cohesiveness

Cronbach's Alpha for Neighbourhood Safety Scale, based on Cycle 4 data
Cronbach's alpha for these factors are given in the table below (calculated using SAS).

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Neighbourhood safety (_SFHS5)	0.701	_SFHQ05B	0.544
Neighbours (_SFHS6)	0.883	_SFHQ06C	0.848

Note: Scores could not be calculated for 3,571 (11.2%) and 6,534 (20.6%) individuals due to missing values.

9.5.7 Parenting Scales

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

The questions from the Child's Questionnaire used to measure these aspects of parenting are identified in the following paragraphs. A complete factor analysis was done on the parenting scales to evaluate the psychometric properties of these scales for the NLSCY population. The make-up of each factor obtained during the analysis was compared to that which had been indicated in the literature. The analytical results are presented later in this section.

Questions _PRCQ01 to _PRCQ18 and _PRCQ21 to _PRCQ24 on positive interaction, ineffectiveness and on coherence were provided by Dr. M. Boyle of the Chedoke-McMaster Hospital, McMaster University, based on the work of Dr. Ken Dodge (Vanderbilt University) and an adaptation of the Parent Practices Scale of Strayhorn and Weidman. (For children ages 0 to 23 months, only questions _PRCQ01 to _PRCQ07 were asked.)

Calculation of Parenting Scores
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 value would be zero (0). For each of the four factors, a score of zero indicates:

- the absence of positive interaction for the positive interaction factor;
- the absence of ineffective interaction for the ineffective factor;
- the absence of consistent parenting for the consistency factor;
- the existence of punitive interaction or aversion producing practices for the ineffective parenting factor.
- a low degree of parent-child conflict (12 to15 years only)

Analytical Results for Children Aged 0 to 23 months, based on Cycle 4 data
There were 4,008 children in the sample for the age group 0 to 23 months. The group was split into two sub-samples of 1,987 and 2,021 individuals, and the analysis for this age group was performed separately for each sub-sample. The non-response rate for the seven items ranged from 1.82% to 2.07%. Individuals were excluded from the analysis that produced the factors when 10% of the data for that factor was missing. After the criterion was applied, there were 1,922 and 1,943 individuals left in the sub-samples to be analysed. No imputation was done. The factor analysis derived two factors for this age group: positive interaction (_PRCS01), and ineffective (_PRCS02). The items making up each factor are listed below.

FACTOR	SCORE	ITEMS
Positive interaction	_PRCS01	_PRCQ01, _PRCQ02, _PRCQ03, _PRCQ06, _PRCQ07
Ineffective	_PRCS02	_PRCQ04, _PRCQ05

Cronbach's Alpha for Children Aged 0 to 23 months, based on Cycle 4 data
Cronbach's alpha (raw value) was computed with SAS using normalized weighted data.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Positive interaction (_PRCS01)	0.661	_PRCQ07	0.587
Ineffective (_PRCS02)	0.339	N/A - only two items included	N/A - only two items included

Note: The scores for these factors could not be computed in 163 and 145 cases respectively because of unreported values.

Analytical Results for Children Aged 2 to 11 years, based on Cycle 4 data
There were 21,777 children in the sample for the age group 2 to 11 years. The group was split into two sub-samples of 10,784 and 10,993 individuals, and the analysis for this age group was performed separately for each sub-sample. The non-response rate for each of the 21 items ranged from 3.07% to 3.69%. Individuals missing 10% or more of the items were excluded from the analysis. After the criteria were applied, there were 10,321 and 10,554 individuals left in the sub-samples to be analysed. The factor analysis derived four factors for this age group: positive interaction (_PRCS03), ineffective (_PRCS04), consistency (_PRCS05) and punitive/aversive (rational) parenting (_PRCS06). The items making up each factor are listed below.

FACTOR	SCORE	ITEMS
Positive interaction	_PRCS03	_PRCQ01, _PRCQ02, _PRCQ03, _PRCQ06, _PRCQ07
Ineffective	_PRCS04	_PRCQ04, _PRCQ08*, _PRCQ09, _PRCQ13, _PRCQ14, _PRCQ15, _PRCQ18
Consistent	_PRCS05	_PRCQ10, _PRCQ11, _PRCQ12*, _PRCQ16*, _PRCQ17*
Rational	_PRCS06	_PRCQ21, _PRCQ22*, _PRCQ23, _PRCQ24*

* Indicates that the item value was reversed when computing the score.

Cronbach's Alpha for Children Aged 2 to 11 years, based on Cycle 4 data
Cronbach's alpha (raw value) was computed with SAS using normalized weighted data. In general, Cronbach's alphas computed by SAS are lower than those produced by SPSS.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Positive interaction (_PRCS03)	0.778	_PRCQ02	0.715
Ineffective (_PRCS04)	0.701	_PRCQ13	0.654
Consistent (_PRCS05)	0.664	_PRCQ12	0.576
Rational (_PRCS06)	0.551	_PRCQ22	0.359

Note: The scores for these factors could not be computed in 902, 1,074, 1,488 and 1,016 cases respectively because of unreported values.

Conflict Resolution Scale for Children Aged 12 to 15 years
The conflict resolution score (_PRCbS09) was created for children aged 12 to 15. The following items were used in the factor analysis: _PRCb30A, _PRCb30B, _PRCb30C, _PRCb30D, _PRCb30E, _PRCb30F, _PRCb30G, and _PRCb30H.

The factor structure of this scale was determined based on data from Cycle 3. To produce the scores, one was subtracted from each item so that the lowest possible score value would be zero (0). For this factor, a high score indicates a high level of conflict between parent and child.

Analytical Results for Children Aged 12 to 15 years, based on Cycle 4 data
There were 4,155 children in the sample for the age group 12 to 15 years. The group was split into two sub-samples of 2,090 and 2,065 individuals, and the analysis for this scale was performed separately for each sub-sample. The non-response rate for the eight items ranged from 4.52% to 5.12%. In total 262 cases had one or more missing values and were excluded from the analysis. These cases were given a missing value for the overall score since no imputation was completed. Items _PRCb30A and _PRCb30H were reversed in the calculation of the score. All values were recoded from 1 to 5 to 0 to 4. The final score ranges from 0 to 30 with a high score indicating a higher degree of parent-child disagreements. The Cronbach's alpha value for the score is 0.745.

9.5.8 Social Support Scale

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

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

In Cycle 5, this section is asked of all PMK's with children/youth less than 16 years of age and includes the following items: _SPHQ01A, _SPHQ01B, _SPHQ01C, _SPHQ01D, _SPHQ01E, _SPHQ01F, _SPHQ01H, and _SPHQ01I.

Changes to Social Support Section
In Cycle 4 the following changes were made to the Social Support section and these changes were kept for Cycle 5:

- The original six items used in Cycle 1 and Cycle 3 were kept; however, items _SPHQ02A to _SPHQ02D used in Cycle 1 were replaced by the F-COPES items.

- Two additional questions from the above mentioned social integration sub-scale (items _SPHQ01H and _SPHQ01I) were also added. The questions on social integration are significant because they assess one's feeling of belonging to a group that shares similar interests, concerns, and activities thus measuring another factor of social support.

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

Analytical Results, based on Cycle 4 data
There were 30,325 children in the sample for the age group 0 to 16 years. The group was split into two sub-samples of 15,001 and 15,324 individuals, and the analysis for this scale was performed separately for each sub-sample. The non-response rate for the eight items averaged about 2.6%. In total 2,033 cases had more than one missing value and were excluded from the analysis. These cases were given a missing value for the overall score since no imputation was done. Items _SPHQ01A, _SPHQ01E, _SPHQ01F and

_SPHQ01I were reversed in the calculation of the score. All values were recoded from 1 to 4 to 0 to 3. The final score ranges from 0 to 24 with a high score indicating a higher degree of social support. The factor structure imposed was the same as that used in Cycle 1.

This analysis was done using normalized weighted data.

Cronbach's Alpha for Social Support Scale, based on Cycle 4 data

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Social support (_SPHS01)	0.877	_SPHQ01F	0.853

9.5.9 Temperament Scale

This scale was used in the NLSCY to measure the various aspects of the temperament of young children from ages 0 to 3 (difficult, unpredictable, dull, not adaptable, persistent/unstoppable, irregular and affectively negative). This section was completed by the PMK. The scale was developed by Dr. John Bates of the University of Indiana and was originally known as the Infant Characteristics Questionnaire (ICQ). Dr. Jo-Anne Finegan of Toronto's Hospital for Sick Children created a revised version of the ICQ to be used for three year olds. The Temperament Scale has been used in large-scale studies and is considered by specialists to be the best available measure for use in household surveys.

This particular scale went through many alterations from Cycles 1 through 3. After the validation of Cycle 1 results, only the difficult construct was used in Cycle 2. In Cycle 3, the not adaptable construct was again included in the survey. The following section will show specifically what variables were kept, discarded and reintroduced and in what cycles these changes were made. This following section will outline which questions have been removed and added in each cycle, according to the age group.

In Cycle 4, it was decided to ask the temperament questions of children aged 0 to 2 years rather than 0 to 3 years as had been done in all previous cycles. In addition three new items were added to the temperament scale (see description by age breakdown listed below). These changes were made based on factor breakdown and reliability analysis undertaken by J. Douglas Willms[6]. No additional changes were made for Cycle 5.

Changes to Questions for Children Aged 3 to 11 months

Action Taken	Variable Name
Questions removed from Cycle 2 and 3	_TMCQ2, _TMCQ3, _TMCQ4, _TMCQ10, _TMCQ15, _TMCQ16, _TMCQ18, _TMCQ23
Questions added in Cycle 4	_TMCQ15, _TMCQ23a

[6] J. Douglas Willms is a Professor at the University of New Brunswick and the Director of the Canadian Research Institute for Social Policy.

Changes to Questions for Children Aged 6 to 11 months

Action Taken	Variable Name
Questions removed from Cycle 2 and 3	_TMCQ13, _TMCQ24
Questions removed from Cycle 3 only	_TMCQ11, _TMCQ12, _TMCQ14
Question removed from Cycle 2 but reintroduced in Cycle 3	_TMCQ27
Questions added in Cycle 3 for the first time	_TMCQ25, _TMCQ26
Questions reintroduced in Cycle 4	_TMCQ15, _TMCQ23

Changes to Questions for Children Aged 1 to 3 years

Action Taken	Variable Name
Questions removed from Cycle 2 and 3	_TMCQ2, _TMCQ3, _TMCQ4, _TMCQ10, _TMCQ13, _TMCQ15, _TMCQ18, _TMCQ21, _TMCQ22, _TMCQ23, _TMCQ24, _TMCQ28, _TMCQ32
Questions removed from Cycle 2 but reintroduced in Cycle 3	_TMCQ25, _TMCQ26, _TMCQ27, _TMCQ29, _TMCQ30, _TMCQ31
Questions reintroduced in Cycle 4	_TMCQ15, _TMCQ23a

As previously mentioned, the Temperament Scale has been used in large-scale studies and is considered by specialists to be the best available measure for use in household surveys. Unfortunately, when used in the context of the NLSCY data, the scale proved to be problematic. Subsequently, there is no Cronbach's alpha value available.

There are three explanations as to why the Temperament Scale did not fare so well. Firstly, data analysis shows that the distributions of these items are highly skewed and multi-modal. Secondly, the focus group study of the temperament questions found that some people found it difficult to understand the concept of the "average" child, which is included in the anchor and in the introduction to the Temperament Scale in the survey. Finally, exploratory factor analysis undertaken by Statistics Canada found that the items in the scale load strongly onto difficult, but that the remaining theoretical factors are unstable with regard to their loading (remaining items do not load on expected theoretical concepts). For these reasons this scale has not been created since Cycle 2. However, the questions making up the scale are still included on the release file for researchers to conduct their own analysis.

Children Aged 3 to 5 months
For children aged 3 to 5 months, the scale is made up of questions _TMCQ01, _TMCQ05 to _TMCQ08, _TMCQ14, _TMCQ15, _TMCQ17, _TMCQ19, _TMCQ20, _TMCQ23 and _TMCQ33 and is intended to measure the extent to which the child is difficult, not adaptable, unpredictable and dull.

Children Aged 6 to 11 months
For children aged 6 to 11 months, the scale is made up of questions _TMCQ01, _TMCQ05 to _TMCQ09, _TMCQ15, _TMCQ17, _TMCQ19, _TMCQ20, _TMCQ23, _TMCQ25, _TMCQ26, _TMCQ27 and _TMCQ33. This expanded list of questions measures the same four aspects of temperament as for children 3 to 5 months old.

Children Aged 1 and 2 years
For children aged 1 and 2 years, questions _TMCQ01 to _TMCQ08, _TMCQ11, _TMCQ12, _TMCQ15, _TMCQ17, _TMCQ19, _TMCQ23a, _TMCQ25, _TMCQ26, _TMCQ27, _TMCQ29, _TMCQ30, _TMCQ31 and _TMCQ33 should theoretically measure the degree to which the child is difficult, irregular, not adaptable, affectively negative and persistent/unstoppable.

Meanings of Ratings for Specified Behaviours
The respondent, in most cases a parent, is required to answer each question in the scale by assigning a rating between 1 and 7. For all questions except _TMCQ14, a one means that the child has a favourable response or usually exhibits the specified behaviour, while a seven indicates that the child reacts negatively or seldom displays the behaviour in question. If the child is in the middle, a four is assigned. In question _TMCQ14, the meanings of the ratings are reversed.

9.6 Child-reported Scales

9.6.1 General-self Scale

Objectives and Overview
The objective of the General-self Scale is to measure the child's overall self-esteem. The self-esteem scale was expanded each year to include the oldest cohort. This means that by Cycle 5 the items making up this scale are asked of all youths aged 10 to 19 years.

In Cycle 2 and subsequent cycles, the factor scores were derived based on the factorial structure identified in Cycle 1. Below is a description of the items that were included on the questionnaire to measure these scales. The analysis used to construct the scale and the results of the analysis where based on Cycle 1 data.

Questions _AMCQ01A to _AMCQ01D on overall self-esteem were taken from the General-self Scale of the Marsh Self Description Questionnaire developed by H.W Marsh.

Once the factor structures were analysed and the items included in each factor were determined, the scores were calculated. No imputation was done for missing values. If any values were missing, the final score was set to missing. To produce the final scores, one was subtracted from each item so that the lowest score would be zero (0). The final score was derived by totalling the values of all items with non-missing values. A score of zero indicates a lack of general self-esteem for the General-self Scale.

Analytical Results, based on Cycle 1 data
In the sample there were 3,434 children aged 10 and 11 years. They were divided into two sub-samples of size 1,705 and 1,729 and analysis was done on each sample. The non-response rates for the eight items ranged from 14.0% to 15.8%. Individuals with missing values were excluded from the analysis conducted for the purpose of constructing the factor. After these exclusions, the sub-samples contained 1,371 and 1,413 individuals respectively, for analysis purposes. As a result of factor analysis, the general self factor was identified. The items that make up this factor are described below.

FACTOR	SCORE	ITEMS
General-self	_AMCS02	_AMCQ01A, _AMCQ01B, _AMCQ01C, _AMCQ01D

Cronbach's Alpha for General-self Scale, based on Cycle 1 data
Cronbach's alpha coefficients (raw values) were calculated with SAS, using the normalized weighted data.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
General-self (_AMCS02)	0.728	_AMCQ01C	0.629

Note: Scores could not be calculated for 555 individuals (16.2%), due to missing values.

9.6.2 Behaviour Scale

Objectives and Overview
This section replicates the behaviour checklist included in the Child's Questionnaire completed by the PMK. All children aged 10 to 15 years answer these questions in the self-complete portion of the survey. It is intended to provide indicators of the following behaviours: conduct disorder, hyperactivity, inattention, physical aggression, indirect aggression, emotional disorder, anxiety and prosocial behaviours. For a more detailed explanation refer to Section 9.5.1 of this chapter. All analysis presented below was done in Cycle 1.

Analytical Results, based on Cycle 1 data
In the sample there were 3,434 children aged 10 and 11 years. They were divided into two sub-samples of size 1,705 and 1,729 and analysis was done on each sub-sample. The non-response rates for the eight items ranged from 13.6% to 16.7%. Individuals with missing values were excluded from the analysis conducted for the purpose of constructing the factor. After these exclusions, the sub-samples contained 1,352 and 1,398 individuals respectively, for analysis purposes. As a result of imposed factor analysis, six factors were identified: hyperactivity - inattention, prosocial behaviour, emotional disorder - anxiety, physical aggression - conduct disorder, indirect aggression and property offences. The items that comprised each factor are described in the following table.

FACTOR	SCORE	ITEMS
Indirect aggression	_FBCS01	_FBCQ01J, _FBCQ01R, _FBCQ01Z, _FBCQ1LL, _FBCQ1TT
Emotional disorder - anxiety	_FBCdS02	_FBCQ01F, _FBCQ01K, _FBCQ01Q, _FBCc01V, _FBCQ1CC, _FBCQ1MM _FBCQ1RR
Physical aggression - conduct disorder	_FBCS03	_FBCQ01G, _FBCQ01X, _FBCQ1AA, _FBCQ1FF, _FBCQ1JJ _FBCQ1NN
Hyperactivity - inattention	_FBCdS04	_FBCc01B, _FBCQ01I, _FBCQ01P, _FBCQ01S, _FBCQ01W, _FBCQ1HH _FBCQ1QQ
Prosocial behaviour	_FBCS05	_FBCQ01A, _FBCQ01D, _FBCQ01H, _FBCQ01M, _FBCQ01U, _FBCQ1BB, _FBCQ1GG, _FBCQ1OO, _FBCQ1SS, _FBCc1UU
Property offences	_FBCS07	_FBCQ01C, _FBCQ01E, _FBCQ01L, _FBCQ01T, _FBCQ1DD _FBCQ1PP

Cronbach's Alpha for Behaviour Scale, based on Cycle 1 data
Cronbach's alpha for these factors are given below (calculated using SAS):

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Indirect aggression (_FBCS01)	0.728	_FBCQ1LL	0.657
Emotional disorder - anxiety (_FBCdS02)	0.760	_FBCQ1MM	0.717
Physical aggression - conduct disorder (_FBCS03)	0.738	_FBCQ1AA	0.678
Hyperactivity - inattention (_FBCdS04)	0.751	_FBCQ1QQ	0.717
Prosocial behaviour (_FBCS05)	0.766	_FBCQ1SS	0.741
Property offences * (_FBCS07)	0.623	_FBCQ01C	0.544

* These results are based on parent-reported scale analysis not done on child-reported property offences.

Note: The scores for these factors could not be computed in 566 (16.5%), 597 (17.4%), 585 (17.0%), 621 (18.1%), 587 (17.1%) and 225 (17.2%) cases respectively because of unreported values.

9.6.3 Depression Scale

General Information about the Depression Scale
The depression scale used to measure PMK depression was also used for children aged 12 and 13 years in Cycle 2 and again for those aged 16 and 17 years in Cycles 4 and 5. The factor structure that was used for the PMK scale was also imposed on the youth scale.

In order to produce the score, one was subtracted from each item so that the lowest score would be zero (0). The final score was derived by totalling the values of all items with non-missing values. As well, the answer categories were reversed for questions having a negative loading (_FBCd10F, _FBCd10H, and _FBCd10J). The total score (_HTCbS1B) may therefore vary between 0 and 36, a high score indicating the presence of depression symptoms.

FACTOR	SCORE	ITEMS
Depression	_HTCbS1B	_FBCd10A, _FBCd10B, _FBCd10C, _FBCd10D, _FBCd10E, _FBCd10F, _FBCd10G, _FBCd10H, _FBCd10I, _FBCd10J, _FBCd10K, _FBCd10L

For more information on the creation of this scale and related analysis, refer to the parent-reported Depression Scale discussed earlier in this chapter (Section 9.5.2).

9.6.4 Friends Scale

Objectives and Overview
The object of the Friends Scale is to measure how well the child feels he/she gets along with his/her peers. In order to understand how the factorial structure was determined in Cycle 1, below is a description of the items that were included on the questionnaire in Cycle 1 to measure peer relations, the analysis used to construct the scale and the results of the analysis.

Score Calculation
Once the factor structures were analysed and the items included in the factor was determined, the score was calculated. No imputation was done on the values. If any values were missing the final score was set to missing. A value may be missing if the child refused to answer or did not know the answer to the question.

To produce the score, one was subtracted from each item so that the lowest score would be zero (0). The final score was derived by totalling the values of all items with non-missing values. The score ranges from 0 to 16. A score of zero indicates the respondent does not have a lot of friends and does not have positive relations with other children.

Analytical Results for Friends Scale, based on Cycle 1 data
In the Cycle 1 sample there were 3,434 children aged 10 and 11 years. They were divided into two sub-samples of size 1,705 and 1,729 and analysis was done on each sub-sample. The non-response rates for the four items ranged from 10.9% to 11.5%. Individuals with missing values were excluded from the analysis conducted for the purpose of constructing the factor. After these exclusions, the sub-samples contained 1,508 and 1,529 individuals respectively, for analysis purposes. No imputation took place. As a result of factor analysis, one factor was identified: the friends factor (_FFCS01). All items, _FFCQ01 to _FFCQ04, loaded into the factor.

Cronbach's Alpha for Friends Scale, based on Cycle 1 data
Cronbach's alpha coefficients (raw values) were calculated with SAS.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Friends (_FFCS01)	0.779	_FFCQ04	0.779

Note: Scores could not be calculated for 397 (11.6%) individuals due to missing values. The analysis presented here was based on Cycle 1 data.

9.6.5 My Parents and Me Scale

Objectives and Overview
This section was part of the Self-complete Questionnaire given to children in the 10 to 15 age group. The objective was to complement the Parenting Section on the Child's Questionnaire completed by the PMK by gathering information directly from the child regarding his/her perception of his/her relationship 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.

The scale that was used was also 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. This information will complement the constructs measured in the parent-completed Child's Questionnaire (positive child-parent interaction, ineffective child-parent interaction, and consistent child-parent interaction, aversive and non-aversive parent management techniques.)

The objective of the My Parents and Me Scale is to measure the child's perception of his/her relationship with his/her parents and parental supervision. Below is a description of the items that were included on the questionnaires to measure family relations for children aged 10 to 15 years, the analysis used to construct the scale and the results of the analysis.

Questions _PMCcQ1A to _PMCcQ1Q were taken from the Western Australia Child Health Survey. In addition to these questions, questions _PMCcQ1R and _PMCdQ1T were also used. The scale 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.

To construct the My Parents and Me Scale for the NLSCY, a factor analysis was conducted to test the theoretical construct. In the factor analysis the items were multiplied by the child's normalized weight. An individual's statistical weight is normalized by dividing his/her weight (_WTCW01C) by the average weight of all individuals. Consequently, the sum of the normalized weights is equal to the sample size.

Analytical Results, based on Cycle 3 data
In the sample of children aged 10 to 15 years there were 5,539 children. The sample was divided into two sub-samples and an analysis was done on each sub-sample. Individuals with missing values were excluded from the analysis conducted for the purpose of

constructing the factor. After these exclusions, the two sub-samples contained 2,509 and 2,584 individuals respectively.

As a result of the factor analysis, three factors were identified for ages 10 to 15 years: the parental nurturance factor, the parental rejection factor and the parental monitoring factor. The items that comprised each factor are described in the following table.

FACTOR	SCORE	ITEMS*
Parental nurturance	_PMCcS1	_PMCcQ1A, _PMCcQ1D, _PMCcQ1K, _PMCcQ1H, _PMCcQ1I, _PMCcQ1M, _PMCcQ1Q
Parental rejection	_PMCbS2b	_PMCcQ1C, _PMCcQ1G, _PMCcQ1J, _PMCcQ1L, _PMCcQ1O, _PMCcQ1P, _PMCcQ1R
Parental monitoring	_PMCcS3	_PMCcQ1B, _PMCcQ1F, _PMCcQ1N, _PMCcQ1E, _PMCdQ1T

Cronbach's Alpha for My Parents and Me Scale, based on Cycle 3 data
Cronbach's alpha coefficients (raw values) were calculated with SAS, using the normalized weighted data.

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Parental nurturance (_PMCcS1)	0.88	_PMCcQ1M	0.855
Parental rejection (_PMCbS2b)	0.73	_PMCcQ1R	0.680
Parental monitoring (_PMCcS3)	0.57	_PMCdQ1T	0.459

9.6.6 Neighbourhood Structure Scale

Objectives and Overview
To gather information on the respondent's satisfaction with his/her neighbourhood, including perception of the extent of danger and problems, and of social cohesion or "neighbourliness".

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.

Analytical Results, based on Cycle 4 data
In the sample there were 2,057 youths aged 16 and 17 years. They were divided into two sub-samples and analysis was done on each sub-sample. Individuals with missing values were excluded from the analysis conducted for the purpose of constructing the factor. After these exclusions the sub-samples contained 1,008 and 1,049 individuals respectively. As a result of the factor analysis, two factors were identified: the neighbourhood safety factor and the neighbours factor. The items that comprised each factor are described in the following table.

FACTOR	SCORE	ITEMS
Neighbourhood structure	_ACYdS01	_ACYd13A to _ACYd13D, _ACYd13F, _ACYd13G

Scale Score
If too many values were missing the final score was set to missing. To produce the final scores, one was subtracted from each item so that the lowest score would be zero (0). All the score values were reversed. The final score was derived by totaling the values of all items with non-missing values. A score of zero indicates a low perception of neighbourhood structure.

Cronbach's Alpha for Neighbourhood Structure Scale, based on Cycle 4 data
Cronbach's alpha for these factors are given in the table below (calculated using SAS).

FACTOR	CRONBACH'S ALPHA (RAW)	ITEM THAT LOWERS CRONBACH'S ALPHA THE MOST IF IT IS EXCLUDED	CRONBACH'S ALPHA IF THE ITEM IS EXCLUDED
Neighbourhood structure (_ACYdS01)	0.729	_ACYd13G	0.665

Note: Scores could not be calculated for 495 (24.1%) individuals due to missing values.

9.6.7 Conflict Resolution Scales for Youth

Analytical Results, based on Cycle 4 data
Two conflict resolution scores were created for youth aged 16 and 17 years. One score relates to the relationship between the youth and their mother (_PMCdS4) and the other score refers to the relationship between the youth and their father (_PMCdS5). A high score indicates an elevated number of disagreements between youth and parent. The following items were used in the factor analysis.

FACTOR	SCORE	ITEMS
Conflict resolution - Mother	_PMCdS4	_PMCdQ6C, _PMCdQ6D, _PMCdQ6E, _PMCdQ6F, _PMCdQ6G, _PMCdQ6H, _PMCdQ6I, _PMCdQ6J, _PMCdQ6K, _PMCdQ6L
Conflict resolution - Father	_PMCdS5	_PMCdQ9C, _PMCdQ9D, _PMCdQ9E, _PMCdQ9F, _PMCdQ9G, _PMCdQ9H, _PMCdQ9I, _PMCdQ9J, _PMCdQ9K, _PMCdQ9L

Analytical Results, based on Cycle 4 data - Mother
There were 1,856 youth in the sample aged 16 and 17 years. The group was split into two sub-samples of 907 and 949 individuals, and the analysis for this scale was performed separately for each sub-sample. The non-response rate for the included items was approximately 26%. In total 530 cases had one or more missing values and were excluded from the analysis. These cases were given a missing value for the overall score since no imputation was completed. All values were recoded from 1 to 5 to 0 to 4. The final score ranges from 0 to 40 with a high score indicating a higher degree of parent-child disagreements. The Cronbach's alpha value for the score is 0.840.

Analytical Results, based on Cycle 4 data - Father
There were 1,856 youth in the sample aged 16 and 17 years. The group was split into two sub-samples of 907 and 949 individuals, and the analysis for this scale was performed separately for each sub-sample. The non-response rate for the included items was approximately 32%. In total 593 cases had one or more missing values and were excluded from the analysis. These cases were given a missing value for the overall score since no imputation was completed. All values were recoded from 1 to 5 to 0 to 4. The final score ranges from 0 to 40 with a high score indicating a higher degree of parent-child disagreements. The Cronbach's alpha value for the score is 0.885.

9.6.8 Emotional Quotient

Developed by Dr. Reuven BarOn and Dr. James D.A. Parker, the Emotional Quotient Inventory Youth Version (EQ-i:YV) is a measure of emotional intelligence. This measure comprises five major dimensions: intrapersonal, interpersonal, adaptability, stress management, and general mood. When compared to other possible measures, there are several reasons why this instrument was preferred and, eventually chosen. First, the questions are generally very positive and are short and simple. They also address the respondent's social, personal, and emotional "abilities", as opposed to their behaviours. This measure was used for youth aged 10 years and older in the Self-complete Questionnaires.

Prior to calculating the Emotional Quotient (EQ) score, the response category values were reversed for three items: _EQYEQ03, _EQYEQ08 and _EQYEQ13. Then one was subtracted from each of the items to permit a score of zero. Once these two steps had been completed the values were summed for each of the dimensions and five scores were created. The lowest scores for a particular scale represents the negative end of the Emotional Intelligence continuum, while the higher scores represent the positive end. For example, 33 on total EQ would mean that the individual is much more emotionally intelligent 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 table below presents interpretive guidelines for the standard scores. Standard scores for the EQ-i:YV set the mean values at 100 and each standard deviation at 15; however, you will notice there are 10 points around the mean values for differentiating between the descriptors in the table below.

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

FACTOR	SCORE	ITEMS
Intrapersonal (RA)	_EQYES06	_EQYEQ01, _EQYEQ06, _EQYEQ11
Interpersonal (ER)	_EQYES07	_EQYEQ02, _EQYEQ07, _EQYEQ12
Stress management (SM)	_EQYES08	_EQYEQ03*, _EQYEQ08*, _EQYEQ13*
Adaptability (AD)	_EQYES09	_EQYEQ04, _EQYEQ09, _EQYEQ14
General Mood (GM)	_EQYES10	_EQYEQ05, _EQYEQ10, _EQYEQ15

* Indicates that the values were reversed.

FACTOR	SCORE VARIABLES
Emotional quotient 4 (EQYES04)	_EQYES06 to _EQYES09
Emotional quotient 5 (EQYES05)	_EQYES06 to _EQYES10

Below are brief definitions of what is measured by the 5 composite scales and the 15 subscales (only the 5 composite scales appear in the EQ-i:YV while these and the 15 subscales appear in the EQ-i which are bulleted below under each of the composite scales):

1) **Intrapersonal competencies - self-awareness and self-expression:**
   These competencies include the following sub-components that govern our ability to be aware of ourselves, to understand our strengths and weaknesses, and to express our thoughts and feelings nondestructively.
   - Self-regard: The ability to be aware of, understand and accept ourselves.
   - Emotional self-awareness: The ability to be aware of and understand our emotions.
   - Assertiveness: The ability to express our feelings and ourselves nondestructively.
   - Independence: The ability to be self-reliant and free of emotional dependency on others.
   - Self-actualization: The ability to set goals and the drive to achieve them.

2) **Interpersonal competencies - social awareness and interpersonal relationship:**
   These competencies include the following sub-components that govern our ability to be aware of others' emotions, feelings and needs, and to be able to establish and maintain cooperative, constructive and mutually satisfying relationships.
   - Empathy: The ability to be aware of and understand how others feel.
   - Social responsibility: The ability to identify with and feel part of our social group.
   - Interpersonal relationship: The ability to establish and maintain mutually satisfying relationships with others.

3) **Stress management competencies - emotional management and regulation:**
   These competencies include the following sub-components that govern our ability to manage emotions so that they work for us and not against us.
   - Stress tolerance: The ability to effectively and constructively manage our emotions.
   - Impulse control: The ability to effectively and constructively control our

emotions.

4) **Adaptability competencies - change management:**
These competencies include the following sub-components that govern our ability to manage change, by realistically and flexibly coping with the immediate situation and effectively solving problems as they arise.
   - Reality-testing: The ability to validate our feelings and thinking with external reality.
   - Flexibility: The ability to cope with and adapt to change in our daily life.
   - Problem-solving: The ability to generate effective solutions to problems of a personal and social nature.

5) **General mood - self-motivation:**
General mood is a facilitator of emotionally and socially intelligent behavior and includes the following sub-components that govern our ability to be optimistic, positive and sufficiently self-motivated to set and pursue our goals.
   - Optimism: The ability to have a positive outlook and look at the brighter side of life.
   - Happiness: The ability to feel content with ourselves, others and life in general.

For further information, see the following:

1) BarOn, R. (2004). The Bar-On Emotional Quotient Inventory (EQ-i): Rationale, description, and summary of psychometric properties. In Glenn Geher (Ed.), *Measurement of emotional intelligence*: *Common ground and controversy*. Hauppauge, NY: Nova Science Publishers, pp. 111-142.

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

9.6.9 Ages and Stages

The Ages and Stages component was generated for all children 3 to 47 months. In Cycle 4 this component had been completed using paper questionnaires but a decision was made to incorporate it into the Cycle 5 application to increase response rates and facilitate data capture.

The questions included in the Ages and Stages questionnaires 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. In Cycle 4, a fifth category, Gross Motor Skills was included, but a decision was made not to incorporate this in Cycle 5 as many of these skills are covered by the Motor and Social Development questions. For more information on this measure please refer to Chapter 8.0.

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

9.7 Cycle 5 Scales

9.7.1 Parent-reported Scales

Score Variable	Scale Name	Universe
_DPPS01	Depression Rating Scale	PMK of children 0 to 15 years
_FNHS01	Family Functioning Scale	PMK or spouse of children 0 to 15 years
_SFHS5	Neighbourhood Safety Scale	PMK or spouse of children 0 to 15 years
_SFHS6	Neighbours Scale	PMK or spouse of children 0 to 15 years
_SPHS01	Social Support Scale	PMK or spouse of children 0 to 15 years
_ACCS6	Home Responsibilities Scale	PMK of children 10 to 13 years

Score Variable	Behaviour Scale	Universe
_BECdS01	Hyperactivity - inattention	PMK of children 2 to 3 years
_BECdS03	Emotional disorder - anxiety	PMK of children 2 to 3 years
_BECS04	Physical aggression - opposition	PMK of children 2 to 3 years
_BECS05	Separation anxiety	PMK of children 2 to 3 years
_BECdS06	Hyperactivity - inattention	PMK of children 4 to 11 years
_BECdS07	Prosocial behaviour	PMK of children 8 to 11 years
_BECdS08	Emotional disorder - anxiety	PMK of children 4 to 11 years
_BECdS09	Physical aggression - conduct disorder	PMK of children 4 to 11 years
_BECS10	Indirect aggression	PMK of children 4 to 11 years
_BECdS11	Property offences	PMK of children 8 to 11 years

Score Variable	Motor and Social Development Scale	Universe
_MSCS01	Raw Score	PMK of children 0 to 47 months
_MSCS02	Standardized Score - based on Cycle 1 norms	PMK of children 0 to 47 months
_MSCdS03	Standardized Score - based on Cycle 5 norms	PMK of children 3 to 47 months

Score Variable	Parenting Scale	Universe
_PRCS01	Positive interaction	PMK of children 0 to 23 months
_PRCS02	Ineffective	PMK of children 0 to 23 months
_PRCS03	Positive interaction	PMK of children 2 to 11 years
_PRCS04	Ineffective	PMK of children 2 to 11 years
_PRCS05	Consistent	PMK of children 2 to 11 years
_PRCS06	Rational	PMK of children 2 to 11 years
_PRCbS09	Conflict Resolution Scale	PMK of children 12 to 15 years

9.7.2 Child-Reported Scales (Self-complete)

Score Variable	Scale Name	Universe
_FFCS01	Friends Scale	Children 10 to 17 years
_AMCS02	General-self Scale	Children 10 to 19 years

Score Variable	Behaviour Scale	Universe
_FBCS01	Indirect aggression	Children 10 to 15 years
_FBCdS02	Emotional disorder - anxiety	Children 10 to 15 years
_FBCS03	Physical aggression - conduct disorder	Children 10 to 15 years
_FBCdS04	Hyperactivity - inattention	Children 10 to 15 years
_FBCS05	Prosocial behaviour	Children 10 to 15 years
_FBCS07	Property offences	Children 10 to 15 years

Score Variable	My Parents and Me Scale	Universe
_PMCcS1	Parental nurturance	Children 10 to 15 years
_PMCbS2b	Parental rejection	Children 10 to 15 years
_PMCcS3	Parental monitoring	Children 10 to 15 years

9.7.3 Youth-reported Scales (Self-complete)

Score Variable	Scale Name	Universe
_ACYdS01	Neighbourhood Structure Scale	Youth 16 to 17 years
_HTCbS1B	Depression Scale	Youth 16 to 17 years

Score Variable	Conflict Resolution Scale	Universe
_PMCdS4	Conflict resolution – Mother	Youth 16 to 17 years
_PMCdS5	Conflict resolution – Father	Youth 16 to 17 years

Score Variable	Emotional Quotient	Universe
_EQYES04	Emotional quotient 4	Youth 10 to 19 years
_EQYES05	Emotional quotient 5	Youth 10 to 19 years

10.0 Survey Methodology - Response Rates

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 labeled 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.

The first section of this chapter deals with total non-response. For any survey, response rates are used to measure the effectiveness of the data collection process and are also a good indicator of the quality of the estimates produced. The National Longitudinal Survey of Children and Youth (NLSCY) can produce both cross-sectional and longitudinal estimates. As a result, total non-response can be expressed both cross-sectionally and longitudinally, which is done in the first sections of this chapter. In Cycle 5, for the first time since the beginning of the survey, cross-sectional weights are only calculated for children aged 0 to 5 and not for all ages covered by the survey. For more details about this, see Chapter 12.0 on weighting. As a result, the data collected for the 8- to 19-year-olds are used only for longitudinal purposes. For this reason, in this chapter, cross-sectional response rates will be given only for children aged 0 to 5. Longitudinally, response rates are presented for the cohorts introduced in Cycles 1, 3 and 4. These response rates are useful to help assess the quality of the longitudinal estimates which can be produced.

The second part of the chapter describes partial non-response. Response rates for different portions of the questionnaire and for the direct measure components will be presented.

10.1 Definitions

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

A **respondent household** is a household for which an adult component or a child/youth component of the survey is completed. Therefore, a respondent household that does not have a completed adult component may still have one respondent child and one non-respondent child.

A **respondent child** is a child whose child component or adult component of the survey is completed.

A **non-respondent child** is an in-scope child for which neither the child component nor the adult component is complete. Non-response can occur either because the child and the guardian(s) refused to do the survey, the child could not be traced, or because the interviewer was unable to complete the interview.

A **respondent youth 16 or 17 years old** is a youth whose child component or youth component or adult component of the survey is completed.

A **non-respondent youth 16 or 17 years old** is an in-scope youth for which neither the child component, the youth component nor the adult component is complete. Non-response can occur either because the youth and the guardian(s) refused to do the survey, the youth could not be traced, or because the interviewer was unable to complete the interview.

A **youth, 18 years old or above**, is a **respondent** if his/her youth component is completed. Note that there is no adult component for youth 18 years old or above. On the other hand, a **youth 18 years old and over** is a **non-respondent** if his/her youth component is not completed.

An **out-of-scope child** is a child that is selected into the sample, but is not part of the target population. A child may be out-of-scope either because he/she is deceased, residing outside of Canada, an inmate of an institution, or living on an Indian reserve. In contrast, an **in-scope child**

is a child that 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 target population when first selected in the survey. On the other hand, it is also possible to have children who are cross-sectionally in-scope but longitudinally out-of-scope. This situation occurs when children were introduced to a longitudinal cohort but, for some reasons, had the wrong age. If the age of these children is covered cross-sectionally (0 to 5 years old), these children get a cross-sectional weight.

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

The **refusal rate** is defined as the number of children or youth who refused over the number of in-scope children or youth.

A **to-be-traced child/youth** is a child/youth that needs to be searched for, usually because the child/youth has moved from the listed address.

The **to-be-traced rate** is defined as the number of cases to be traced over the total sample size (in scope and out of scope).

A non-respondent child/youth is considered **not traced** if attempts at tracing a 'to-be-traced' child/youth are unsuccessful.

The **successful trace rate** is defined as the number of successfully traced children/youth over the number of children/youth that require tracing.

10.2 Cross-sectional Response Rates

The Cycle 5 sample is the aggregation of children introduced in four different cohorts, namely Cycles 1, 3, 4 and 5. The first cohort, selected in Cycle 1, consists of children whose age at the time of selection ranged from 0 to 11 years. In Cycle 5, these children are now 8 to 19 years old. Children newly introduced in Cycle 3 were 0 and 1 year old at the time. In Cycle 5, they are now 4 and 5 years old. In Cycle 4, children aged 0 and 1 were also introduced to NLSCY. In Cycle 5, these children are 2 and 3 years old. Finally, in Cycle 5, we have added new 0 and 1 year old children to the survey. Consequently, the Cycle 5 sample has children aged 0 to 5 and 8 to 19 years of age. As previously mentioned in Cycle 5, only the children 0 to 5 years old get a cross-sectional weight. For this reason, the cross-sectional response rates are only given for the children aged 0 to 5 years. Longitudinally, response rates are shown for children introduced in Cycles 1, 3 and 4 and are given in Section 10.3 of this chapter.

Table 1: Unweighted Cross-sectional Response Counts for Children, NLSCY, Cycle 5, Canada

Effective Age	Sample Total	Out-of-scope	In-scope	Respondents	Non-respondents	To-be-traced	Traced
0 and 1	4,492	98	4,394	3,252	1,142	474	421
2 and 3	4,007	12	3,995	3,476	519	719	549
4 and 5	6,949	42	6,907	6,188	719	1,144	995
Total	**15,448**	**152**	**15,296**	**12,916**	**2,380**	**2,337**	**1,965**

Table 2: Unweighted Cross-sectional Response Rates for Children, NLSCY, Cycle 5, Canada

Effective Age	Response rate (%)	Non-response rate (%)	To-be-traced rate (%)	Successful trace rate (%)
0 and 1	74.01	25.99	10.55	88.82
2 and 3	87.01	12.99	17.94	76.36
4 and 5	89.60	10.40	16.46	86.98
Total	**84.44**	**15.56**	**15.13**	**84.08**

Tables 3 and 4 present the provincial response counts and rates, respectively. The results show that the Atlantic provinces have high response rates and high successful trace rates, whereas Ontario has the lowest response rate, and the lowest successful trace rate.

Table 3: Unweighted Cross-sectional Response Counts for Children by Province, NLSCY, Cycle 5

Province	Sample Total	Out-of-scope	In-scope	Respondents	Non-respondents	To-be-traced	Traced
Newfoundland and Labrador	771	5	766	688	78	110	100
Prince Edward Island	477	3	474	415	59	63	57
Nova Scotia	1,012	3	1,009	884	125	188	166
New Brunswick	923	3	920	792	128	135	110
Quebec	2,634	24	2,610	2,199	411	362	324
Ontario	4,327	60	4,267	3,489	778	621	480
Manitoba	1,215	13	1,202	1,028	174	196	163
Saskatchewan	1,203	14	1,189	1,006	183	200	178
Alberta	1,473	14	1,459	1,254	205	252	214
British Columbia	1,409	10	1,399	1,160	239	207	171
Other	4	3	1	1	0	3	2
Canada	**15,448**	**152**	**15,296**	**12,916**	**2,380**	**2,337**	**1,965**

Table 4: Unweighted Cross-sectional Response Rates for Children by Province, NLSCY, Cycle 5

Province	Response rate (%)	Non-response rate (%)	To-be-traced rate (%)	Successful trace rate (%)
Newfoundland and Labrador	89.82	10.18	14.27	90.91
Prince Edward Island	87.55	12.45	13.21	90.48
Nova Scotia	87.61	12.39	18.58	88.30
New Brunswick	86.09	13.91	14.63	81.48
Quebec	84.25	15.75	13.74	89.50
Ontario	81.77	18.23	14.35	77.29
Manitoba	85.52	14.48	16.13	83.16
Saskatchewan	84.61	15.39	16.63	89.00
Alberta	85.95	14.05	17.11	84.92
British Columbia	82.92	17.08	14.69	82.61
Other	100.00	0.00	75.00	66.67
Canada	**84.44**	**15.56**	**15.13**	**84.08**

Tables 5 and 6 provide the response counts and rates respectively, for children that required tracing. These tables show the impact that tracing had on response rates by age group. Note that interviewers were instructed not to trace children aged 0 and 1. Although, this instruction was not followed in all cases, and some households responded after being traced, there is a low response rate for this age group.

Table 5: Unweighted Cross-sectional Response Counts for To-be-traced Children, NLSCY, Cycle 5, Canada

Effective Age	To-be-traced Total	Out-of-scope	In-scope	Respondents	Non-respondents
0 and 1	474	22	452	117	335
2 and 3	719	11	708	451	257
4 and 5	1,144	25	1,119	836	283
Total	**2,337**	**58**	**2,279**	**1,404**	**875**

Table 6: Unweighted Cross-sectional Response Rates for To-be-traced Children, NLSCY, Cycle 5, Canada

Effective Age	Response rate (%)	Non-response rate (%)
0 and 1	25.88	74.12
2 and 3	63.70	36.30
4 and 5	74.71	25.29
Total	**61.61**	**38.39**

The following table provides counts for in-scope children by province and age group, while Table 8 provides counts for respondents by province and age group. Using these counts, Table 9 displays the cross-sectional response rates by province and age group.

Table 7: Unweighted Cross-sectional Counts of In-scope Children by Province and Age, NLSCY, Cycle 5

Province	Effective Age			
	0 and 1	2 and 3	4 and 5	Total
Newfoundland and Labrador	130	124	512	766
Prince Edward Island	130	106	238	474
Nova Scotia	239	236	534	1,009
New Brunswick	200	205	515	920
Quebec	736	708	1,166	2,610
Ontario	1,346	1,264	1,657	4,267
Manitoba	364	301	537	1,202
Saskatchewan	341	329	519	1,189
Alberta	456	347	656	1,459
British Columbia	452	375	572	1,399
Other	0	0	1	1
Canada	**4,394**	**3,995**	**6,907**	**15,296**

Table 8: Unweighted Cross-sectional Counts of Respondent Children by Province and Age, NLSCY, Cycle 5

Province	Effective Age			
	0 and 1	2 and 3	4 and 5	Total
Newfoundland and Labrador	96	104	488	688
Prince Edward Island	95	97	223	415
Nova Scotia	181	212	491	884
New Brunswick	149	170	473	792
Quebec	554	613	1,032	2,199
Ontario	986	1,089	1,414	3,489
Manitoba	274	267	487	1,028
Saskatchewan	242	298	466	1,006
Alberta	342	309	603	1,254
British Columbia	333	317	510	1,160
Other	0	0	1	1
Canada	**3,252**	**3,476**	**6,188**	**12,916**

Table 9: Unweighted Cross-sectional Response Rates by Province and Age, NLSCY, Cycle 5

Province	Effective Age – Response Rates (%)			
	0 and 1	2 and 3	4 and 5	Total
Newfoundland and Labrador	73.85	83.87	95.31	89.82
Prince Edward Island	73.08	91.51	93.70	87.55
Nova Scotia	75.73	89.83	91.95	87.61
New Brunswick	74.50	82.93	91.84	86.09
Quebec	75.27	86.58	88.51	84.25
Ontario	73.25	86.16	85.33	81.77
Manitoba	75.27	88.70	90.69	85.52
Saskatchewan	70.97	90.58	89.79	84.61
Alberta	75.00	89.05	91.92	85.95
British Columbia	73.67	84.53	89.16	82.92
Other	N/A	N/A	100.00	100.00
Canada	**74.01**	**87.01**	**89.59**	**84.44**

Tables 7 and 8 shown above always used the province at the time of selection to give the response counts and response rates (Table 9). However, between the sample selection and the data collection, some children have moved. Table 10 gives the number of responding children by age and the province at the time the data were collected. Note that the province used in this table is the variable EMMCQ01 in the master file.

Table 10: Unweighted Cross-sectional Counts of Respondent Children at the time of collection by Province and Age, NLSCY, Cycle 5

Province	Effective Age			
	0 and 1	2 and 3	4 and 5	Total
Newfoundland and Labrador	96	105	481	682
Prince Edward Island	95	97	224	416
Nova Scotia	181	209	494	884
New Brunswick	149	172	465	786
Quebec	554	615	1,036	2,205
Ontario	987	1,081	1,430	3,498
Manitoba	274	263	478	1,015
Saskatchewan	241	294	461	996
Alberta	342	321	615	1,278
British Columbia	333	319	504	1,156
Other	0	0	0	0
Canada	**3,252**	**3,476**	**6,188**	**12,916**

10.3 Longitudinal Response Rates

In a longitudinal survey, the longitudinal response rate shows how much attrition has taken place since the children were first introduced to 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 takes all of the components of attrition into account. Specifically, neither the actual number of children in non-responding households, nor the number of children from non-responding households in the Labour Force Survey, is known. Therefore, the denominator needed to determine the ratio (and hence the rate) is also unknown. The solution used to overcome this problem is described in the next sections.

10.3.1 Children Selected in Cycle 1

Table 11 below shows the longitudinal response rates by province and Table 12 the longitudinal response rates by age for the children introduced in the first cycle. Readers must be aware of the following elements when using these tables:

- The denominator for all the percentages shown in these tables is the number of responding children in Cycle 1 who were followed in Cycle 2. Note that since the sample size has been reduced from Cycle 1 to Cycle 2, not all Cycle 1 responding children are considered longitudinal.

- A unit is considered a longitudinal respondent in a given cycle if it has a longitudinal weight in that cycle. This definition is slightly different from what was used in the Microdata User Guides for the previous cycles. This definition implies that some units might not have reported data but are still considered as respondents. This situation occurs for the children who have died or moved out of the country. These children do not have reported data but still have a longitudinal weight since they represent other children in the population in the same situation.

- Since the longitudinal weights for the children introduced in Cycle 1 always represent the population of children at the time of selection, the province used in Table 11 is the province in Cycle 1. Of course, this might differ from their current province of residence. In theory, the year of birth and the gender of a child should not change from cycle to cycle. However, for different reasons, a small number of children have some discrepancies in these variables. When this is the case, the latest known information was used.

- In order to retain the highest possible number of children from the original cohort, attempts are made to convert children who did not respond in a previous cycle. It is then possible to have children who have not responded in a particular cycle but have responded in the subsequent cycle(s).

- The “Respondents in all cycles” columns show the number of children who have responded in all cycles. A weight for these children who have responded in all cycles is available in the master file. See Chapter 12.0 on weighting for more details about this.

Table 11: Unweighted Longitudinal Response Rate for Children Selected in NLSCY, Cycle 1, by Province

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

Since the introduction of the Cycle 1 cohort, Ontario has shown the highest attrition rate. On the other hand, the Atlantic provinces and Manitoba have the highest retention rates. Overall, about 2/3 of the number of Cycle 1 longitudinal respondents have responded in all cycles, with provincial variations from 62.5 % in Ontario to 72.5% in Newfoundland and Labrador.

Table 12: Unweighted Longitudinal Response Rate for Children Selected in NLSCY, Cycle 1, by Age

Age in Cycle 1	Number of Respondents in Cycle 1	Respondents in Cycle 2		Respondents in Cycle 3		Respondents in Cycle 4		Respondents in Cycle 5		Respondents in all cycles	
		Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1	Number	% of Cycle 1
0 and 1	4,052	3,740	92.3	3,638	89.8	3,229	79.7	3,157	77.9	2,793	68.9
2 and 3	2,916	2,662	91.3	2,585	88.6	2,337	80.1	2,229	76.4	2,008	68.9
4 and 5	2,666	2,433	91.3	2,372	89.0	2,104	78.9	1,982	74.3	1,769	66.4
6 and 7	2,393	2,170	90.7	2,103	87.9	1,872	78.2	1,719	71.8	1,557	65.1
8 and 9	2,451	2,238	91.3	2,158	88.0	1,890	77.1	1,786	72.9	1,592	65.0
10 and 11	2,425	2,225	91.8	2,141	88.3	1,878	77.4	1,650	68.0	1,417	58.4
Total	**16,903**	**15,468**	**91.5**	**14,997**	**88.7**	**13,310**	**78.7**	**12,523**	**74.1**	**11,136**	**65.9**

The general trend seen from this table is that the response rate decreases as the age increases. This situation stands out in particular for children who were 10 and 11 in Cycle 1. These youths, aged 18 and 19 in Cycle 5 have the lowest response rate.

Table 13 below provides, by province and age in Cycle 1, the number of longitudinal children who responded in Cycle 5.

Table 13: Number of Longitudinal Children from Cycle 1 who responded in Cycle 5, by Province and Age

Province	Effective Age						
	0 and 1	2 and 3	4 and 5	6 and 7	8 and 9	10 and 11	Total
Newfoundland and Labrador	161	116	108	114	129	127	755
Prince Edward Island	84	56	60	57	50	57	364
Nova Scotia	229	146	159	116	127	126	903
New Brunswick	183	148	142	114	107	98	792
Quebec	642	434	364	299	309	313	2,361
Ontario	813	556	479	424	452	380	3,104
Manitoba	266	194	154	127	146	117	1,004
Saskatchewan	242	184	174	126	153	123	1,002
Alberta	283	199	170	173	162	175	1,162
British Columbia	254	196	172	169	151	134	1,076
Canada	**3,157**	**2,229**	**1,982**	**1,719**	**1,786**	**1,650**	**12,523**

10.3.2 Children Selected in Cycle 3

The longitudinal response rates for children introduced in Cycle 3 are presented by province in Table 14. In Cycle 3, all the newly introduced longitudinal children were 0 or 1 years old, so the response is not shown by age. Again, readers should be aware of the following elements when using these tables:

- The denominator for all the percentages shown in these tables is the number of responding children in Cycle 3.

- A unit is considered a longitudinal respondent in any given cycle if it has a longitudinal weight in that cycle. This definition is slightly different from what was used in the Microdata Users Guides for the previous cycles. This definition implies that some units might not have reported data but are still considered as respondents. This situation occurs for the children who have died or moved out of the country. These children do not have reported data but still have a longitudinal weight since they represent other children in the population in the same situation.

- Since the longitudinal weights for the children introduced in Cycle 3 always represent the population of children at the time of selection, the province used in Table 14 is the province in Cycle 3. This province might be different from their current province of residence in Cycle 5. In theory, the year of birth and the gender of a child should not change from cycle to cycle. However, for different reasons, a small number of children have some discrepancies in these variables. When this is the case, the latest known information was used.

- For the children newly introduced in Cycle 3, only the respondents to a given cycle are followed in the next cycle. There is no attempt to convert non-respondents from a previous cycle. Consequently, there is no distinction made between respondents to a given cycle (Cycle 5) and respondents to all cycles.

Table 14: Unweighted Longitudinal Response Rate of Children Selected in NLSCY, Cycle 3, by Province

Province	Number of Respondents in Cycle 3	Respondents in Cycle 4		Respondents in Cycle 5	
		Count	% of Cycle 3	Count	% of Cycle 3
Newfoundland and Labrador	566	526	92.9	500	88.3
Prince Edward Island	270	241	89.3	227	84.1
Nova Scotia	599	533	89.0	490	81.8
New Brunswick	595	523	87.9	480	80.7
Quebec	1,332	1,186	89.0	1,049	78.8
Ontario	1,943	1,657	85.3	1,421	73.1
Manitoba	643	554	86.2	501	77.9
Saskatchewan	611	532	87.1	477	78.1
Alberta	748	646	86.4	592	79.1
British Columbia	667	593	88.9	536	80.4
Canada	**7,974**	**6,991**	**87.7**	**6,273**	**78.7**

The Atlantic Provinces, led by Newfoundland and Labrador, have a higher response rate than the rest of the population. Ontario has the lowest response rate and the lowest retention rate.

10.3.3 Children Selected in Cycle 4

The longitudinal response rates for children introduced in Cycle 4 are presented by province in Table 15. In Cycle 4, as in Cycle 3, all the newly introduced longitudinal children were 0 or 1 year old, so the response is not shown by age. Again, readers should be aware of the following elements when using these tables:

- The denominator for all the percentages shown in these tables is the number of responding children in Cycle 4.
- The definition of a longitudinal respondent is the same as the one used for units introduced in Cycles 1 and 3.
- Since the longitudinal weights for the children introduced in Cycle 4 always represent the population of children at the time of selection, the province used in Table 15 is the province in Cycle 4. Again, this province might be different from their current province. In theory, the year of birth and the gender of a child should not change from cycle to cycle. However, for different reasons, a small number of children have some discrepancies in these variables. When this is the case, the latest known information was used.

Table 15: Unweighted Longitudinal Response Rate of Children Selected in NLSCY, Cycle 4, by Province

Province in Cycle 4	Number of respondents in Cycle 4	Respondents in Cycle 5	
		Count	% of Cycle 4
Newfoundland and Labrador	122	102	83.6
Prince Edward Island	105	96	91.4
Nova Scotia	233	209	89.7
New Brunswick	198	163	82.3
Quebec	672	580	86.3
Ontario	1,221	1,048	85.8
Manitoba	289	257	88.9
Saskatchewan	316	287	90.8
Alberta	322	287	89.1
British Columbia	364	306	84.1
Canada	**3,842**	**3,335**	**86.8**

10.4 Partial Non-response

The previous sections on non-response dealt with the issues of representation of the sample for cross-sectional or longitudinal estimation of individuals based on their household characteristics. These types of non-sampling errors can usually be dealt with effectively by adjusting the sample weight to reflect the more accurate picture of the population measured. Other types of non-response are also measured in this survey and are usually not corrected through an adjustment of survey weights.

Although a person may provide enough information to qualify as a respondent, some of the questions (variables) may still be not answered, resulting in partial non-response. Some reasons for this are (in no particular order) unwillingness to answer sensitive questions, respondent fatigue, accidental skipping of parts of the questionnaire, or operational difficulties. Usually, the nature of partial non-response depends on the subject matter. For instance, the Motor-Social Development component, 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, while the “Not stated” category usually indicates that the respondent was not asked for the information. This is true for computer-assisted response capture but not for self-completed paper questionnaires. For the latter, “blank responses” are categorized as “Not stated” even though the respondent may have seen the question.

(Note: The "Not applicable" category is not a non-response but a valid skip of a particular piece of information for a particular respondent.)

The "Not stated" response category is often the result of proxy responses where the person providing the information does not know the true response (e.g., a person reporting for his/her spouse). In other cases, the respondent may not understand the question. However, because of the type of application used to collect the NLSCY data, two types of situations emerge that contribute to the number of "Not stated" cases: either the interview was never completed in its entirety – partial response, because the survey could not be completed before the end of the collection period, or because the respondents avoided the interviewers (tacit refusal) – or the respondent refused to complete a specific portion of the questionnaire or one of its components.

The NLSCY questionnaire is divided into cohesive blocks of information, such as heath, labour force or education. Other sections are collected in a different manner, which requires specific conditions or administration procedures. Cognitive assessments are collected in person, directly from the child, sometimes on the computer, sometimes on a paper test. Self-complete questionnaires are paper questionnaires given to young people to complete on their own in private. When this information is missing, it often results in a block of missing information categorized as "Not stated".

The following sections will explore the issue of component non-response for a number of these situations. Typically, non-response analysis is performed on the sample population and represents characteristics of the sample units, so unweighted response rates are presented. This analysis is intended to inform researchers who use these variables in their analysis of sources of error not remedied by the survey weights. With the number of variables in the NLSCY files, an exhaustive analysis would not be possible, so researchers may find in their own analysis that component non-response may affect the results in other ways. Researchers who suspect that component non-response may cause differential results in their estimates should take corrective measure to compensate for the non-response or detail the impact of component non-response in their findings.

10.4.1 Measurements in Early Childhood Development

Non-response to the Cycle 5 Peabody Picture Vocabulary Test - Revised

Children born in 1997 and 1998 were eligible for the Peabody Picture Vocabulary Test - Revised (PPVT-R) in Cycle 5 of the NLSCY. In responding households, there were 6,192 such children, of whom 5,538 completed the PPVT-R. Also, the PPVT-R scores were estimated for 37 children who partially completed the PPVT-R, so these children will also be considered respondents. Thus, the overall response rate to the PPVT-R was 5,574 of 6,190 children, or 90.1%.

The response rate was 89.4% for males and 90.8% for females. It was 90.5% for children born in 1997 and 89.2% for children born in 1998. The differences between the sexes and between the years of birth are both statistically significant.

The variables that we found were most highly associated with non-response are described in the sections below. They are ranked in descending order of the strength of their correlation with response. Children with the following characteristics were more likely to be PPVT-R respondents:

- Only children who understood English or French well enough to follow instructions were given the test. By design those who did not (130 in all) were excluded from testing. Similarly, children who were out of the jurisdiction during the collection period (2) were not given the PPVT-R. This does not imply that

these children's skills were deficient in any way; simply that it was not feasible or appropriate to apply the PPVT-R assessment to them.

- Children from the five easternmost provinces, particularly Prince Edward Island (224 children, 96% response rate) had higher response rates.

- The 613 area code (Eastern Ontario) also had a higher response (95%).

Children with the following characteristics were less likely to be PPVT-R respondents:

- Children with PMKs who said in Cycle 4 that their race/colour was best described as Native/Aboriginal were less likely to take the PPVT-R.
    - Twenty-six children whose PMK spoke Cree had a particularly low response rate (46%).

Aboriginal PMK (unweighted)	Number eligible	Response rate
Yes	170	72%
No	5,907	91%
Not stated, Don't know or Refusal	113	79%

- The response rates were generally higher to the east of the Ontario-Québec border. Among the telephone area codes, the 416 area (Toronto) had a lower response rate (82.7% for 168 children) and the 613 area (Eastern Ontario) had a higher response rate (95.3% for 212 children).

Province (unweighted)	Number eligible	Response rate
Newfoundland and Labrador	481	93.6%
Prince Edward Island	224	96.0%
Nova Scotia	494	92.7%
New Brunswick	465	92.0%
Quebec	1,036	92.7%
Ontario	1,432	86.6%
Manitoba	478	87.0%
Saskatchewan	461	89.4%
Alberta	614	87.8%
British Columbia	505	90.5%
Other	2	0.0% *

* Not tested because they were out of the jurisdiction during the collection period.

- Children whose PMK has a low level of education, particularly less than elementary (60 children, 68% response rate).
    - We found that the response rate increased steadily as the PMK's level of education increased.

PMK education (unweighted)	Number eligible	Response rate
No schooling or some elementary	60	68%
Completed elementary or some secondary	646	86%
Completed secondary, some post-secondary or completed community college, technical college, CÉGEP or nursing training	3,904	91%
Completed university or teacher's college, Master's, doctorate, medical degree, or other education/training not listed above.	1,486	94%
Not asked, Not stated, Refusal	94	28%

- In Cycle 5, the responding adult was asked, "Does a physical condition or mental condition or health problem reduce the amount or kind of activity [the child] can do at home?"
    - Children with a physical, mental or health problem that often reduces activities at home had lower response rates (73 kids, 64% response rate).

Activity restriction (unweighted)	Number eligible	Response rate
Yes, sometimes	100	88%
Yes, often	73	64%
No	5,980	94%
Not asked	37	3%

- Children whose household was in the lowest income adequacy category at the time they were sampled in Cycle 3 (210 children, 72% response rate).
    - Response rates increase as income category increases and peak just before the highest level of income adequacy (note that income adequacy for Cycle 5 was not available at the time of our assessment and Cycle 3's measurement had a stronger correlation than Cycle 4's measurement).

Income adequacy in 1998-1999 (unweighted)	Number eligible	Response rate
Lowest	210	72%
Lower middle	814	86%
Middle	1,891	90%
Upper middle	2,314	93%
Highest	961	91%

- Responses to the question in Cycle 5 "If you tell [your child] will get punished if [he/she] doesn't stop doing something, and [he/she] keeps doing it, how often will you punish [him/her]?" were shown to be related to non-response.

Frequency of punishment (unweighted)	Number eligible	Response rate
Never, or less than half the time	884	84%
About half the time or more (includes N/A)	5,199	92%
Not asked, Not stated, or Refusal	107	37%

Non-response to the Cycle 5 "Who Am I?" Assessment

Children born in 1997 and 1998 were eligible for the "Who Am I?" assessment in Cycle 5 of the NLSCY. In responding households, there were 6,192 such children, of whom 5,162 completed the Who Am I?, for a response rate of 83.4%. A small difference was noted between the response rates of girls (83.8%) and boys (82.9%).

We noted a large difference when we looked at the effective age of the children. Younger children were less likely to complete the "Who Am I?" assessment. The response rate was 77.2% for children born in 1998, compared with 86.9% for children born in 1997.

While two children could not be assessed because they did not live in the 10 provinces at the time of the interview, the remaining children are shown in the following table by province. The rates vary greatly, from a high of 89% in Newfoundland and Labrador to a low of 73.4% in Manitoba.

Province (unweighted)	Number eligible	Response rate
Newfoundland and Labrador	481	89.0%
Prince Edward Island	224	83.9%
Nova Scotia	494	84.8%
New Brunswick	465	81.9%
Quebec	1,036	83.0%
Ontario	1,432	84.8%
Manitoba	478	73.4%
Saskatchewan	461	82.6%
Alberta	615	80.8%
British Columbia	504	87.7%
Other	2	0.0% *

* Not tested because they were out of the jurisdiction during the collection period.

Non-response to the Cycle 5 "Number Knowledge" Assessment

Children born in 1997 and 1998 were eligible for the "Number Knowledge" (NK) assessment in Cycle 5 of the NLSCY. In responding households, there were 6,192 such children, of whom 5,580 completed the NK assessment, for a response rate of 90.1%. A difference was noted between the response rates of girls (90.8%) and boys (89.4%).

A difference was also noted when we looked at the effective age of the children. Younger children were less likely to complete the NK assessment. The response rate was 88.9% for children born in 1998, compared with 90.8% for children born in 1997.

While two children could not be assessed because they did not live in the 10 provinces at the time of the interview, the remaining children are shown in the following table by province. The rates vary moderately, with the best rates in the East. Prince Edward Island had the highest rate, at 96%, and Manitoba the lowest, at 86.4%, marginally below Ontario's 86.8%.

Province (unweighted)	Number eligible	Response rate
Newfoundland and Labrador	481	94.0%
Prince Edward Island	224	96.0%
Nova Scotia	494	92.9%
New Brunswick	465	92.5%
Quebec	1,036	92.3%
Ontario	1,432	86.8%
Manitoba	478	86.4%
Saskatchewan	461	90.0%
Alberta	615	87.8%
British Columbia	504	90.7%
Other	2	0.0% *

* Not tested because they were out of the jurisdiction during the collection period.

10.4.2 Cognitive Assessments for Older Cohorts

Non-response to the Cycle 5 mathematics tests

The NLSCY mathematics tests are made of 20 computational questions answered in the home by respondents aged 8 to 15. Test respondents were non-institutionalized students in the ten provinces and were either ungraded (grades based on effective age) or in Grades 2 to 10. The component response rate was 81% in Cycle 5.

Note that for this non-response analysis, a small number of individuals had values imputed (mostly with previous-cycle information) for the following Cycle 5 characteristics used in the analysis: grade, province of educational institution, mathematics progress, diagnosis of learning disability, diagnosis of mental handicap, hoped-for future education, presence of condition/health problem, and school type. Such variables are indicated by an asterisk (*****). Also, the results are based on raw data and may differ slightly from those found in the release data.

Our analysis showed that lower response had the stronger correlation with the following variables.

- **For Grade level*:** while ungraded children probably have other reasons for their low response rate (reasons that may also explain why they were ungraded), these ungraded children tend to have lower mathematical ability.

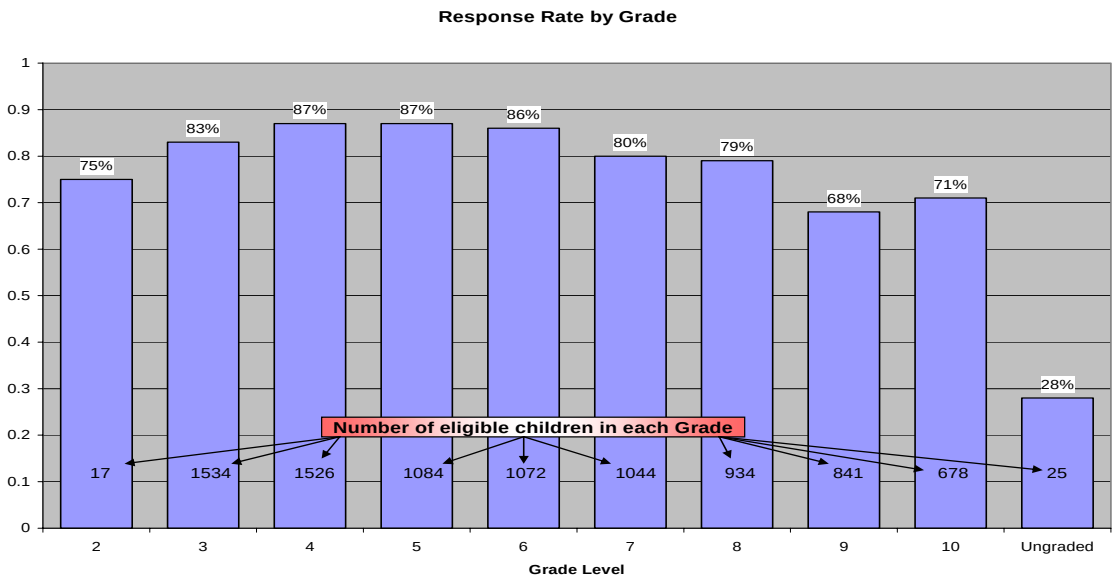

- **For province of educational institution*:** Quebec, which generally has the highest overall mathematics test scores, also has the highest response rate. Alberta continues to have a lower rate of response (as in Cycle 4 for both the Cognitive Measure and the math test).

Province	Number eligible	Response rate
Newfoundland and Labrador	445	82%
Prince Edward Island	243	84%
Nova Scotia	608	83%
New Brunswick	561	83%
Quebec	1,704	90%
Ontario	2,265	79%
Manitoba	709	77%
Saskatchewan	688	76%
Alberta	907	73%
British Columbia	778	81%

- **Past good performance on the Math Assessment:** Those who did well in the Cycle 4 mathematics test were more likely to respond to the Cycle 5 mathematics assessment. (Note that anyone who attempted less than five items in the mathematics test is considered to be a non-respondent to the test). In fact, the Cycle 4 math raw score is a better predictor of Cycle 5 math test non-response than the Cycle 5 question "How is [the child] doing in [mathematics] at school this year?". This reinforces our previous findings that non-response to the test is often related to the child's ability to perform. When doing domain comparison, one should investigate the level of non-response to ascertain its impact on the results.

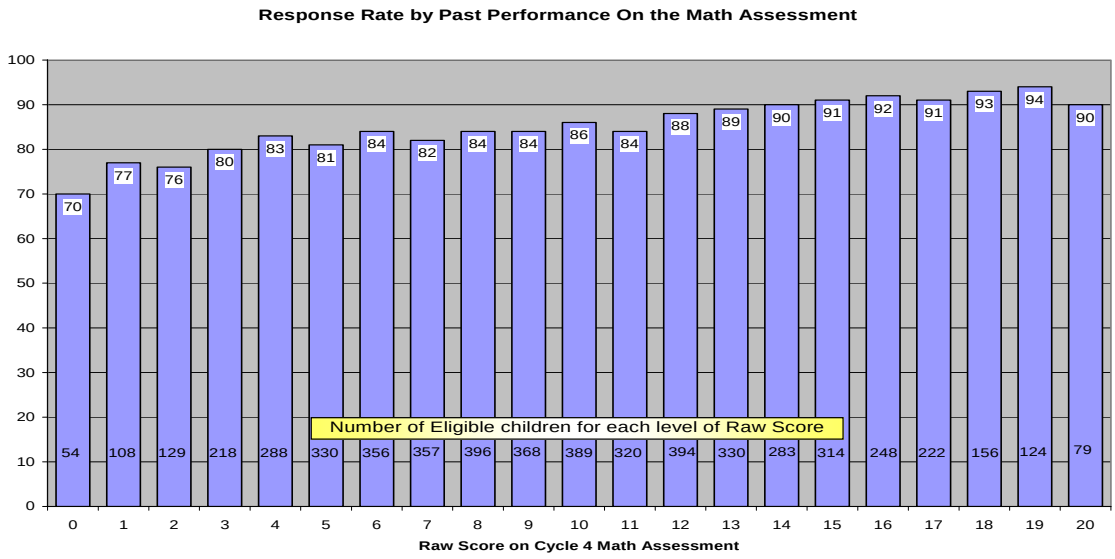

  - o There were 1,866 children not eligible for test in Cycle 4, of whom 80% responded in Cycle 5. These would mostly be children who were not yet in Grade 2 or who were home-schooled in Cycle 4.

  - o There were 1,095 test non-respondents in Cycle 4, of whom only 68% respondent in Cycle 5, and there were 484 cases of total non-response from Cycle 4 re-contacted for Cycle 5, only 63 of whom completed their assessment. Clearly, a predisposition not to respond may be inherent for this sub-group.

- **Non-response by Cycle 4 interviewer observation of task persistence.** The response rate is quite low for the 53 children whom the interviewers classified as "Refusals" for this question. Again we see at the bottom of the table that the 5% of Cycle 5 survey respondents who were not in Cycle 4 had a low response rate on the mathematics test.

Child's response during activities: Task persistence (unweighted)	Number eligible	Response rate
Persists with task	6,062	86%
Attempts tasks … briefly	450	80%
Attempts tasks … after much encouragement	252	77%
Refuses	53	47%
Don't know	70	73%
Not stated	1,537	70%
Not in Cycle 4	484	63%

- **Non-response by how the child is doing in mathematics at school*.** The response rate is higher for children who are doing well in mathematics according to their parents. Interestingly, 73% of children were considered to be above average, but only 5% were rated below average.

Mathematics progress (unweighted)	Number eligible	Response rate
Very well	3,825	85%
Well	2,481	82%
Average	2,107	74%
Poorly	442	74%
Very poorly	53	60%

- **Non-response by diagnosis of learning disability*.** The non-response rate of eligible learning disabled children is double that of other eligible children. The statistical significance of the relationship between having a learning disability and response rate may be slightly greater than the significance of the relationship between mental handicap and response. Of the 49 children with a mental handicap, 41 were reported to have a learning disability as well.

Learning disability status (unweighted)	Number eligible	Response rate
Learning disabled	373	61%
Not learning disabled	8,535	82%

- **Non-response by diagnosis of mental handicap*.** Only 9 of 51 eligible mentally handicapped children responded to the mathematics test.

Mental handicap status (unweighted)	Number eligible	Response rate
Mentally handicapped	51	18%
Not mentally handicapped	8,857	81%

- **Non-response by how far the PMK hopes child will go in school*.** The response rate to the Cycle 5 cognitive measure was highest for children whose parent hoped the child would go to university. For children whom parents hoped would attend unspecified post-secondary education, the response rate was average. The response rates of the remaining categories were not significantly different from one another (Primary or High School, or Community College, CEGEP or Nursing School, or Trade, Technical or Vocational School, or Business College, or other).

Hoped-for education (unweighted)	Number eligible	Response rate
University	5,428	85%
Unspecified post-secondary	1,101	81%
Other	2,258	74%

- **Non-response by presence of physical or mental condition or health problem*.** Children with a physical condition or mental condition or health problem that reduces the amount of activity they can do at home have a lower rate of response to the Cycle 5 mathematics test.

Reduces amount of activity (unweighted)	Number eligible	Response rate
Yes, sometimes	255	73%
Yes, often	124	48%
No	8,529	82%

- **Non-response by school type*.** Eight children who were otherwise eligible for the test did not attend school and so were not given the test. One child was in an institution and was not tested. Of 83 eligible home-schooled children, only 4 responded. In previous cycles, home-schooled children were not eligible for the test because it was to be administered at school. Although they ultimately could be tested in the home, the application still did not generate a test component.

School type (unweighted)	Number eligible	Response rate
Public school	7,325	82%
Catholic school, publicly funded	1,128	80%
Private school	347	86%
Taught at home	83	4%
Other	25	48%

Non-response to the Cycle 5 Cognitive Measure for 16- and 17-year-olds

In Cycle 5, the Cognitive Measure for 16- and 17-year-olds in the NLSCY consisted of 18 mathematical questions with an emphasis on spatial relationships and everyday problem-solving. This test was administered in the home.

Of the approximately 1,772 youths who had a chance to take the Cognitive Measure, 1,404 attempted at least 4 items and so were considered respondents to the test. The response rate of 79.23% compares favourably to the 68.52% response rate for this component in Cycle 4.

Of the main demographic variables, sex had the greatest effect on non-response. Females had a response rate of 83%, whereas males had a response rate of only 75%. There was no significant difference in response rates between different ages or school grades. Significantly different response rates were measured for different geographical areas, however. For example, while the response rate was 91% for the 45 youths from Montréal, it was only 65% for the 40 youths from Calgary.

The variables that we found were most highly associated with non-response are described in each section below. They are ranked in descending order of the significance of their correlation with response. Before this ranking was performed, however, the categories marked with an asterisk (*) were redistributed among the other (response) categories.

Non-response by Cycle 5 Self-complete Questionnaire completion
For the purposes of this table, a fully completed questionnaire is one with item responses to all the questions that are applicable to all respondents.

Self-complete Questionnaire (unweighted)	Number eligible	Response rate
Fully completed	796	95%
Partially completed	682	92%
No questions answered	294	8%

Non-response by Cycle 4 math test Item Response Theory (IRT) score (DMACS03)

Cycle 4 non-respondents were less likely to respond to the Cycle 5 Cognitive Measure. Those who took the math test in Cycle 4 and scored in the top 81% had a particularly high response rate to the Cognitive Measure.

Math test Item Response Theory score (unweighted)	Number eligible	Response rate
541 or above	1,050	88%
540 or below	222	79%
Test not taken	375	61%
Was not in Cycle 4	125	56%

Non-response by Cycle 4 interviewer observation of task persistence (DOBCDQ1A)

The response rate is quite low for the 32 youths whom the interviewers classified under "Refuses". The 117 youths among Cycle 5 survey respondents who were not in Cycle 4 also had a low response rate to the Cognitive Measure .

Child's response during activities: Task persistence (unweighted)	Number eligible	Response rate
Persists with task	1,179	86%
Attempts tasks … briefly	68	75%
Attempts tasks … after much encouragement	27	67%
Refuses*	32	34%
Don't know or Refusal*	17	59%
Not stated*	324	69%
Not in Cycle 4*	125	56%

Non-response for those given special help

For the purposes of this table, a child is considered to have received special help when “Yes” is the most recent response to the following question (which was asked in the first three cycles): "Does [your child] receive special help because a physical, emotional, behavioural, or some other problem limits the kind or amount of school work [he/she] can do?".

Receives special help? (unweighted)	Number eligible	Response rate
Yes	110	58%
No	1,662	81%

Non-response by frequency of book-reading
The response rate is lowest for youths who do not read books at all. However, the highest response rates do not seem to come from those who read the most books.

Excluding for school or work, how often do you read books? (unweighted)	Number eligible	Response rate
Daily	351	90%
Weekly	351	87%
Monthly	275	88%
Several times a year	125	95%
Once or twice a year	173	94%
Never	322	78%
No response*	175	6%

Non-response by Cycle 4 parental hopes for child's academic attainment (DEDCQ18B)
Youths whose parents hoped they would go to university had the highest response rate on the Cycle 5 Cognitive Measure. (Surprisingly, this question from Cycle 4, had a stronger association with Cognitive Measure non-response than a similar question from Cycle 5.)

How far do you hope [your child] will go in school? (unweighted)	Number eligible	Response rate
Primary or secondary	134	67%
Community college, CEGEP or nursing school	193	74%
Trade, technical, or vocational school, or business college	191	77%
University	1,029	86%
Other	28	79%
No response*	197	59%

Non-response by attendance at day camp in summer prior to Cycle 1 (AACCQ8A)
In Cycle 1, the question "Last summer, did [your child] attend a day camp or recreational or skill-building activity that ran for half days or full days (e.g., music program, reading program, athletic program)?" was asked. The response to this question was strongly associated with the Cycle 5 Cognitive Measure response.

Attendance (unweighted)	Number eligible	Response rate
Yes	522	85%
No	1,220	77%
No response*	30	67%

Non-response by propensity to steal outside the home (DFBCQ1PP)
Respondents to the Cycle 4 Self-complete Questionnaire were more likely to respond to the Cycle 5 Cognitive Measure if they indicated that they did not steal outside the home.

I steal outside my home (unweighted)	Number eligible	Response rate
Never	1,210	87%
Sometimes or often	108	69%
In Cycle but no response*	329	61%
Not in Cycle 4	125	56%

Non-response by frequency of PMK alcohol consumption
In the table below, the frequency of PMK alcohol consumption is that for the most recently available cycle.

Alcohol consumption (unweighted)	Number eligible	Response rate
Every day or 4 to 6 times a week	74	77%
2 or 3 times a week	190	89%
Once a week	266	85%
2 or 3 times a month	223	78%
Once a month	213	80%
Less than once a month	481	74%
Never	325	77%

10.4.3 Self-complete Components

Non-response to the Self-complete Component for 10- and 11-year-olds

The Self-complete Component for 10- and 11-year-olds in the NLSCY is a short paper booklet comprising questions on topics such as education, parents and emotions. 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. We consider responding questionnaires to be those which have at least one response to any of the questions. Of those who answered at least one question in the booklet (the smallest number of items answered was in fact 11), the mean item response rate to items on the questionnaire is about 95%.

Approximately 2,189 of the 10- and 11-year-olds were eligible to take the Self-complete Component. Of those, 1,838 answered at least the minimum number of items (actually 11) of the 152 questions that are applicable to all youths. The remaining youths are counted as component non-respondents for the purposes of this analysis. Thus, the component response rate is 84.0%.

The response rate was 84.5% for males and 85.5% for females, which is not a significant difference. For children born in 1991, the response rate was 85.4%, but for those born in 1992 it was 84.6%.

The variables that we found were most highly associated with non-response are described in the sections below. They are ranked in descending order of the strength of their correlation with response.

Non-response by number of items attempted in the Cycle 5 math tests
Children who attempted at least two items in the Cycle 5 math test were most likely to respond to the Cycle 5 Self-complete Questionnaire.

Number of attempted items (unweighted)	Number eligible	Response rate
0 or 1	253	40%
2 to 19 (English booklet)	830	94%
2 to 19 (French booklet, levels 2 to 5)	229	71%
2 to 19 (French booklet, level 6 to 10)	102	81%
20 (English booklet)	650	94%
20 (French booklet)	125	77%

Non-response by child component status
About half of the child component non-respondents (who were also Cycle 5 respondents) were respondents to the Self-complete Questionnaire.

Status (unweighted)	Number eligible	Response rate
Answered all key questions	2,121	85%
Did not answer all key questions	68	40%

Non-response by Regional Office
The response rate was high in the Atlantic Region, particularly Nova Scotia (98%). However, Québec, particularly Sherbrooke (54%) had a low response rate.

Regional Office (unweighted)	Number eligible	Response rate
Atlantic (Newfoundland and Labrador, Prince Edward Island, Nova Scotia and New Brunswick)	436	94%
Québec	435	70%
Ontario	543	79%
Prairies (Manitoba, Saskatchewan and Alberta)	586	91%
Pacific (British Columbia)	189	85%

Non-response by frequency of an adult reading with the child in Cycle 2 (BLTCQ07)
Those children with whom an adult read daily in Cycle 2 were more likely to respond to the Cycle 5 Self-complete Questionnaire.

Reading frequency (unweighted)	Number eligible	Response rate
Less than daily	661	78%
Daily or more	1,398	87%
No response	76	76%
Not in Cycle 2	54	86%

Non-response for the variable BBECQ6M (related to peacemaker characteristic)
The variable BBECQ6M gives the responding adult's response to the question "How often would you say that [the child], if there is a quarrel or dispute, will try to stop it?". Children who never tried to stop a quarrel were less likely to respond to the Cycle 5 Self-complete Questionnaire.

Will try to stop a quarrel (unweighted)	Number eligible	Response rate
Never	656	78%
Sometimes or often	1,372	87%
Not asked or no response	161	85%

Non-response by marital status of the spouse/partner of the PMK in Cycle 1
The response rate to the self-complete form is lower for children whose PMK's spouse was unmarried in Cycle 1.

Spouse,s marital status (unweighted)	Number eligible	Response rate
Married	1,661	85%
Common-law	259	74%
No spouse	269	86%

Non-response to the Self-complete Component for 12- and 13-year-olds

The Self-complete Component for 12- and 13-year-olds in the NLSCY is a short booklet comprising questions mostly of a private nature on topics such as misbehaviour, suicide and dating. 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. Of those who answered at least one item in the booklet, the mean item response rate is about 97%.

Approximately 1,956 of the 12- and 13-year-olds had a chance to take the Self-complete Component. Of those, 1,543 answered at least one (actually, at least 62) of the 189 items that are applicable to all youths. The remaining youths are counted as component non-respondents for the purposes of this analysis. Thus, the component response rate is 78.9%.

The response rate was 78.3% for males and 79.5% for females, which is not a significant difference. For children born in 1989, the response rate was 78.5%, but for those born in 1990 it was 79.2%.

The variables that we found were most highly associated with non-response are described in the sections below. They are ranked in descending order of the strength of their correlation with response.

Non-response by number of items attempted in Cycle 5 math tests
Children who did not attempt any items in the Cycle 5 math test were least likely to respond to the Cycle 5 Self-complete Questionnaire. 97% of children who answered at least 6 items correctly in the French math booklet responded to the Self-complete Questionnaire.

Number of items attempted (unweighted)	Number eligible	Response rate
0	391	49%
English	1,186	84%
French (5 correct items or less)	74	85%
French (6 correct items or more)	305	97%

Non-response by Regional Office
Response to the self-complete form varied greatly by regional office. The response rate was low across the Prairies, with a particularly low rate in Regina (48%). The province of Québec had a high response rate, particularly in the Montréal metropolitan area (99%).

Regional office (unweighted)	Number eligible	Response rate
Atlantic (Newfoundland and Labrador, Prince Edward Island, Nova Scotia and New Brunswick)	437	73%
Québec	360	93%
Ontario	480	86%
Prairies (Manitoba, Saskatchewan and Alberta)	506	63%
Pacific (British Columbia)	173	88%

Non-response by ADMCD10 (number of younger siblings in household in Cycle 1)
The number of younger siblings is negatively correlated with probability of response to the Cycle 5 Self-complete Questionnaire. The number reported in Cycle 1 has a higher correlation than the number reported in subsequent cycles, which implies that it is the number of siblings who are no more than about 4 or 5 years younger that is important. In a previous analysis, we had found that children with siblings of a similar age (less than about 4 years' difference) tend to do worse on the math tests, which may be a related effect.

Number of younger siblings (unweighted)	Number eligible	Response rate
0	1,048	85%
1	827	72%
2	78	72%
3	3	0%

Non-response by Cycle 4 math test Item Response Theory Score
The IRT score from the Cycle 4 math test is positively correlated with likelihood of response to the Cycle 5 Self-complete Component. Also, children who responded to the Cycle 4 math test were more likely to respond to the Cycle 5 Self-complete Component.

Number correct in Cycle 4 math test (unweighted)	Number eligible	Response rate
217 to 382	181	74%
383 to 677	1,292	84%
Cycle 4 math test component non-respondent	370	68%
Not in Cycle 4	113	64%

Non-response by activity restriction
The variable DHLCdQ5B in Cycle 4 is the answer to the question "Does a physical condition or mental condition or health problem reduce the amount or the kind of activity [the child] can do at home?". Children without such a condition or problem had the highest rate of response to the Cycle 5 Self-complete Questionnaire.

Activity restriction (unweighted)	Number eligible	Response rate
Yes, sometimes	52	54%
Yes, often	37	73%
No	1,692	81%
Item non-response	62	72%
Not in Cycle 4	113	64%

Non-response by DBECQ6GG
The variable DBECQ6GG from Cycle 4 is the response to the question "How often would you say that [the child] spontaneously helps to pick up objects which somebody has dropped?".

Helps to pick up objects (unweighted)	Number eligible	Response rate
Never	285	71%
Sometimes	1,025	81%
Often	452	85%
Item non-response	81	72%
Cycle 4 non-respondent	113	64%

Non-response by Cycle 4 self-complete status
For the purposes of this table, a child was considered to be a respondent to the Cycle 4 Self-complete Questionnaire if the variable DFFCQ01 has a valid answer. Response to the Cycle 4 Self-complete Component is associated with response to the Cycle 5 Self-complete Component.

Cycle 4 self-complete status (unweighted)	Number eligible	Response rate
Cycle 4 self-complete respondent	1,502	83%
Cycle 4 self-complete non-respondent	341	66%
Cycle 4 non-respondent	113	64%

Non-response to the Self-complete Component for 14- and 15-year-olds

The Self-complete Component for 14- and 15-year-olds in the NLSCY is a short booklet comprising questions mostly of a private nature on topics such as feelings, parents and puberty. 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. Of those who answered at least one item in the booklet, the mean item response rate is about 97%.

Approximately 1,709 of the 14- and 15-year-olds had a chance to take the Self-complete Component. Of those, 1,427 answered at least 50 of the 219 items that are applicable to all youths. The three youths who answered between 2 and 19 items and the 279 who answered no items are counted as component non-respondents for the purposes of this analysis. Thus, the component response rate is 83.5%.

It may be of interest that the response rate was 81% for males and 86% for females. For children born in 1987, the response rate was 82%, but for those born in 1988 it was 85%.

The variables that we found were most highly associated with non-response are described in the sections below. They are ranked in descending order of the strength of their correlation with response.

Non-response by Cycle 5 math test Classical score (EMACS02)
Response to the Self-complete Component is highly correlated with response to the math test. There is also a significant difference between the Self-complete Component response rates of the high scorers and the low scorers on the math test.

Classical scale score (unweighted)	Number eligible	Response rate
366 to 617	821	94.0%
618 to 854	363	97.5%
Attempted 4 items or less (usually none)	525	57.3%

Non-response by Cycle 4 math test Classical score (DMACS02)
The difference in self-complete response rates between children with high Cycle 4 math test scores and those with low scores is moderately significant. The difference in self-complete response rates between the math test respondents and the math test non-respondents is substantial, however.

In addition to the Cycle 4 and 5 math tests, other components with response rates that are highly correlated with Cycle 5 self-complete response rates are (with the most significant first) the Cycle 4 self-complete form, the Cycle 3 self-complete form, the Cycle 5 child component (asked of the parent), the Cycle 3 reading test, and the Cycle 3 math test. For brevity, tables for these components have been omitted.

Classical scale score (unweighted)	Number eligible	Response rate
287 to 550	1,016	87.4%
557 to 729	269	91.8%
Cycle 4 component non-respondent	335	69.9%
Cycle 4 non-respondent	89	65.2%

Non-response by province of school/institution
The response rates of 14 and 15 year-olds in Alberta and Saskatchewan (particularly in Regina) were quite low. Prince Edward Island had the highest rate (96%).

Province (unweighted)	Number eligible	Response rate
British Columbia	175	83%
Alberta	189	72%
Saskatchewan	119	70%
Manitoba	122	77%
Ontario	429	88%
Québec	300	90%
New Brunswick	108	91%
Nova Scotia	108	81%
Prince Edward Island	55	96%
Newfoundland	104	82%

Non-response by Cycle 1 income adequacy (AINHD07)
The probability of response to the Cycle 5 self-complete form increases with Cycle 1 household income adequacy. Income adequacy in other cycles and socio-economic status (SES) are also associated with response but the correlation is weaker. In fact, the more recent the cycle, the weaker the correlation for both SES and income adequacy.

Income adequacy (unweighted)	Number eligible	Response rate
Lowest to middle range	900	80.0%
Upper middle	591	86.1%
Highest	218	90.8%

Non-response by variable AHLCQ48D from Cycle 1
The number of times that the PMK saw or talked on the phone about the child's physical and mental health with a public health nurse or a nurse practitioner is inversely associated with Cycle 5 self-complete response.

Number of times (unweighted)	Number eligible	Response rate
0	1,447	85%
1	185	82%
2 to 10	48	67%
No answer but in Cycle 1	29	59%

Non-response by number of older siblings in Cycle 1 (ADMCD09)
Whether or not the child has any older siblings in the household seems to have an effect on the probability of response to the Cycle 5 self-complete form. Using data from Cycle 1 gives the strongest correlation. Whether the sibling is in the household does not seem to be as important as having older siblings. Older siblings may have moved out by the time we collected information in Cycle 5, thus devaluing this information by Cycle 5.

Number of older siblings (unweighted)	Number eligible	Response rate
0	775	87%
1	636	81%
2 or more	298	79%

Non-response by variable DSCCQ13 from Cycle 4 Self-complete Instrument
The item DSCCQ13 in Cycle 4 stated "If I have problems at school, my parents are ready to help" and respondents were given six response categories from which to choose.

My parents are ready to help (unweighted)	Number eligible	Response rate
All the time	896	86%
Most of the time	212	91%
Some of the time, rarely or never	87	92%
No problems at school	111	93%
No answer but in Cycle 4	314	72%
Cycle 4 non-respondent	89	65%

Non-response by variable DPMCCQ1M from Cycle 4 Self-complete Instrument
Children who answered "Never", "Rarely" or "Sometimes" to the Cycle 4 question DPMCCQ1M had the highest response rates to the Cycle 5 Self-complete Instrument.

My parents speak of the good things I do (unweighted)	Number eligible	Response rate
Never or Rarely or Sometimes	284	93%
Often or Always	997	86%
No answer but in Cycle 4	339	72%
Cycle 4 non-respondent	89	65%

Non-response by variable CAMCQ01C from Cycle 3 Self-complete Instrument
Children who least agreed with the statement "A lot of things about me are good" in Cycle 3 were most likely to respond to the Cycle 5 self-complete form.

A lot of things about me are good (unweighted)	Number eligible	Response rate
False or mostly false	33	100%
Sometimes false / sometimes true	157	90%
Mostly true	414	83%
True	863	86%
No answer but in Cycle 3	172	60%
Cycle 3 non-respondent	70	84%

Non-response to the Self-complete Component for 16- and 17-year-olds

The Self-complete Component for 16- and 17-year-olds in the NLSCY is a short booklet comprising questions mostly of a private nature on topics such as truancy, drug use and relationships. These are self-administered questionnaires that the youth completes in private, away from both parents and interviewer. Questionnaires are returned in a sealed envelope to the interviewer during the visit. Of those who attempted the booklet, the mean item response rate is about 98%.

Approximately 1,772 of the 16- and 17-year-olds were eligible for the Self-complete Component. Of those, 1,476 answered at least 33 of the 106 items that are applicable to all youths. (Two youths answered 8 items and the remainder answered none.) Thus, the component response rate is 83.3%.

Of the main demographic variables, sex had the greatest effect on non-response. Females had a response rate of 86.2%, compared with only 80.4% for males. There was no significant difference in response rates between different ages or school grades. Significantly different response rates by geographical area are the 67.5% response rate of the 40 youths from Calgary and the 78.9% response rate of the 144 youths from Manitoba.

The variables that we found were most highly associated with non-response are described in the sections below. They are ranked in descending order of the strength of their correlation with response.

Non-response by Cycle 5 Cognitive Measure score
Since the Self-complete Questionnaire was distributed at about the same time as the Cognitive Measure, it is not surprising that the response rates for the two instruments are highly correlated. However, there was no significant difference between the Self-complete Component response rates of the high scorers and the low scorers on the Cognitive Measure. Of the youth who took the Cognitive Measure, almost all responded to the Self-complete Questionnaire.

Cognitive Measure score (unweighted)	Number eligible	Response rate
Test score above mean	698	98.6%
Test score less than or equal to mean	706	97.9%
Cognitive Measure non-respondent	368	26.4%

Non-response by youth component status
The response status to the youth component is highly correlated with self-complete response.

Youth component status (unweighted)	Number eligible	Response rate
Answered at least 50% of the key questions	1,599	91%
Otherwise	173	9%

Non-response by Cycle 4 math score
Although youths with higher Cycle 4 math scores were more likely to respond to the Cycle 5 Self-complete Component, the difference is not significant. The difference in self-complete response rates between the math test respondents and the math test non-respondents is substantial, however.

Cycle 4 math score (unweighted)	Number eligible	Response rate
Cycle 4 math score equal to or above mean	577	90.5%
Cycle 4 math score below mean	695	88.8%
Cycle 4 math test component non-respondent	375	68.5%
Cycle 4 non-respondent	125	64.0%

Non-response by Cycle 4 self-complete status
For the purposes of this table, a child was considered to be a respondent to the Cycle 4 Self-complete Questionnaire if the variable DFBCQ1AA has a valid answer.

Cycle 4 self-complete status (unweighted)	Number eligible	Response rate
Cycle 4 self-complete respondent	1,319	89%
Cycle 4 self-complete non-respondent	328	67%
Cycle 4 non-respondent	125	64%

Non-response by task persistence during Cycle 4 activities
Youth who displayed persistence during the Cycle 4 activities were more likely to respond to the Cycle 5 Self-complete Questionnaire.

Task persistence (unweighted)	Number eligible	Response rate
Persists with task	1,179	89%
Attempts task briefly	68	82%
Attempts task after much encouragement	27	85%
Refuses	32	41%
In Cycle 4 but not reported	341	75%
Cycle 4 non-respondent	125	64%

Non-response by Cycle 3 self-complete status
For the purposes of this table, a child was considered to be a respondent to the Cycle 3 Self-complete Questionnaire if the variable CFFCQ01 has a valid answer.

Youths' non-response tendency is quite persistent, although it does weaken over time. Compare the "Non-response by Cycle 4 self-complete status" above.

Cycle 3 self-complete status (unweighted)	Number eligible	Response rate
Cycle 3 self-complete respondent	1,533	85%
Cycle 3 self-complete non-respondent	196	71%
Cycle 3 non-respondent	43	72%

Non-response by use of day care centre in summer before Cycle 1
Children who were in a child care arrangement in the summer before Cycle 1 were more likely to respond to the Cycle 5 self-complete form. This relationship was particularly strong for children who were in daycare centres in Cycle 1 (when they were aged about 8 or 9).

Used daycare centre (unweighted)	Number eligible	Response rate
Yes	49	92%
No, a different arrangement was used	1,061	86%
No child care arrangement used	632	79%
Not stated	30	80%

Non-response by how far parent expects child will go in school
Youths whose parents expect them to become highly educated according to the Cycle 5 child component have the highest response rate. (This component comprises questions asked of the PMK about the child.).

Expected extent of education (unweighted)	Number eligible	Response rate
Not graduate high school	33	67%
Secondary graduation	221	79%
Technical, trade or vocational school	133	83%
Community college, CEGEP or apprenticeship	388	85%
Bachelor's degree	695	88%
Master's, PhD, or more than one bachelor's degree	89	89%
Post-secondary, unspecified studies	97	94%
Refusal, Not stated or Missing	116	82%

Non-response by frequency of assigned homework in Cycle 4
Children who are used to doing homework seem to be more likely to fill out the Self-complete Questionnaire.

How often child is assigned homework (unweighted)	Number eligible	Response rate
Daily	822	88%
Less than daily or never	751	82%
Not applicable, Don't know or Not stated	74	72%
Cycle 4 non-respondent	125	64%

Non-response by Cycle 1 socio-economic status (AINHD08)
Socio-economic status (SES) was not calculated past Cycle 3 because the obsolete 1980 Standard Occupational Classification was no longer derived for the parents. However, we can still see that SES is strongly associated with response. The measure of SES used in the NLSCY is a combination of the following five variables: years of education for the PMK and spouse, occupational status of the PMK and spouse, and household income. Specifically, each of these five (or three, for single-parent families) input variables is standardized and summed.

Cycle 1 socio-economic status (unweighted)	Number eligible	Response rate
Low /medium (SES < 0.2)	1,230	81%
High (0.2 <= SES)	528	88%
Not stated	14	86%

Non-response to the Self-Complete Component for 18- and 19-year-olds

The Self-complete Component for 18- and 19-year-olds in the NLSCY is a short booklet comprising personal questions about subjects like feelings, criminal behaviour and drug use. Of those who attempted the booklet, the mean item response rate is about 98.6%.

Approximately 1,711 of the 18- and 19-year-olds had a chance to take the Self-complete Component. Of those, 1,466 answered at least one of the 83 items that are applicable to all youths. Thus, the component response rate is 85.7%.

Of the main demographic variables, sex had the greatest effect on non-response. Females had a response rate of 89.3%, compared with only 81.9% males. There was no significant difference in response rates between different years of birth. The most significantly different response rate by geographical area was the 94.5% response rate of the 109 youths from Newfoundland and Labrador.

The variables that we found were most highly associated with non-response are described in the sections below. They are ranked in descending order of the strength of their correlation with response.

Non-response by youth component status
The youth component response status is highly correlated with self-complete response. Other components with response rates that are strongly associated with Cycle 5 self-complete response are, in order of significance, the Cycle 4 Cognitive Measure, the Cycle 4 self-complete form, the Cycle 5 adult component, the Cycle 3 math test, the Cycle 3 self-complete form, the Cycle 2 self-complete form, and the Cycle 3 reading test.

Youth component status (unweighted)	Number eligible	Response rate
All key questions answered	1,576	91%
Answered at least 50% of key questions	39	79%
Otherwise	96	7%

Non-response by diagnosis of mental handicap
Ten of the twelve 18- and 19-year-olds who had been reported in at least one of the first five cycles to have a mental handicap were non-respondents to the Cycle 5 self-complete form.

Mental handicap (unweighted)	Number eligible	Response rate
Yes	11	9%
No	1,700	86%

Non-response by Cycle 1 impulsiveness (ABECQ6S)
In Cycle 1, the question “How often would you say that [the child] is impulsive, acts without thinking?” was asked of the responding adult. Children who were said to be often impulsive were less likely to respond to the Cycle 5 self-complete form.

Impulsive (unweighted)	Number eligible	Response rate
Never	784	85%
Sometimes	765	88%
Often	143	75%
No response	19	84%

Non-response by number of younger siblings of child in household in Cycle 1 (ADMCD10)

Youths with no younger siblings in their household have the highest response rates on the Cycle 5 self-complete form. This correlation is slightly stronger when the number of younger siblings in Cycle 1 is used (rather than the number in later cycles).

Number of younger siblings (unweighted)	Number eligible	Response rate
None	850	90%
One or more	861	81%

Non-response by ARLCQ08

Variable ARLCQ08 is the response of the PMK to the question “During the past 6 months, how well has [the child] gotten along with his/her parent(s)?”. Those children who had gotten along with their parent(s) least well were also least likely to respond to the Cycle 5 Self-complete Component.

How well? (unweighted)	Number eligible	Response rate
Very well / quite well	1,486	87%
Pretty well / not too well / not well at all	208	78%
No response	17	82%

11.0 Imputation

The definition of a respondent, as given in Chapter 10.0, is a child or youth, who has at least one of the child/youth component or the adult component completed. For the respondents, there exist many cases of partial non-response. In some cases this may be for an entire component or only for certain questions. Imputation is the process whereby missing or inconsistent items are "filled in" with acceptable values. In the National Longitudinal Survey of Children and Youth (NLSCY), imputation is carried out for certain variables in the adult and youth Income Section and for the Motor and Social Development Scale.

Imputation flags have been included on the NLSCY file so that users will have information on the extent of imputation and what specific items have been imputed on what records. All imputation flags on the NLSCY data file have an "I" as the sixth character of the variable name. For example, the name of the imputation flag for household income (EINHEQ03) is EINHEI02.

11.1 Adult Income Imputation

Several income questions were asked during the NLSCY household interview. Information on income, broken down into three sources, was asked for the person most knowledgeable (PMK), his/her spouse, as well as a total amount for other household members who were 15 years old or older. Those three income sources are: wages and salary, self-employment net income, and Employment Insurance benefits. Information on income broken down into four sources was also asked at the household level. Those four income sources were: Child Tax Benefit/National Child Benefit, social assistance, child and spousal support and other sources.

Income is a sensitive topic. As a result, some respondents refused to provide answers to the detailed income questions. Among those, some provided an estimate of their total household income or an estimate of their income using ranges. Finally, for those who responded, amounts declared in the Income Section were sometimes incoherent with answers provided in the Labour Force Section (for example, an individual might have reported working in the past 12 months according to answers provided in the Labour Force Section but no wages or self-employment income were reported in the Income Section). Income imputation was carried out to fill in the holes resulting from partial 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 ages 18 or 19 years old and those who have youth ages 16 or 17 years old living independently. Of the 20,239 eligible households at least one income variable was imputed for a total of 5,130 households, which corresponds to 5,822 children or youth.

Imputation was carried out for each of the income sources. Imputation was done using a nearest neighbour approach. This method first identifies a respondent to the Income Section (a donor) who has similar characteristics as the individual or household with incomplete income data (the recipient). Once the nearest neighbour has been identified, the income amounts reported by the donor are used to impute the missing income amounts for the recipient. Two types of imputation were done. First the three sources of personal income for the PMK and the spouse were imputed. The remaining variables were imputed at the household level.

Household level imputation was done in one of three ways. For households that provided an estimate of household income, this estimate was used to help determine the donor. For households that provided an estimate of household income in ranges, the value of the range was used to help determine the donor. If there was no additional income information, then only other household variables, for example province, were used. The imputation flags provide information on how the imputation was done.

Imputation flag = 1 Estimated total was used to determine the donor
Imputation flag = 2 Income range was used to determine the donor
Imputation flag = 3 No additional income information was used to determine the donor

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

Variable	Overall Imputation Rate	Imputation Rate for Longitudinal File (Ages 8 to 19)	Imputation Rate for Early Childhood Files (Ages 0 to 5)
EINPC1AA PMK income from wages and salaries	10.2%	7.5%	12.5%
EINPC1AB PMK income from self-employment	10.0%	7.8%	11.9%
EINPC1AC PMK income from Employment Insurance benefits	10.5%	6.9%	13.6%
EINPED04 Total personal income for PMK	14.3%	10.2%	17.9%
EINSC1AA Spouse income from wages and salaries	15.4% *	11.2% *	18.8% *
EINSC1AB Spouse income from self-employment	13.6% *	10.5% *	16.1% *
EINSC1AC Spouse income from Employment Insurance benefits	12.1% *	7.9% *	15.5% *
EINSED04 Total personal income for spouse	18.7% *	13.8% *	22.6% *
EINHED3P Total personal income from other household members 15 years old and older	19.4% *	13.6% *	39.8% *
EINHE1AD Household income from child tax benefits	13.5%	9.8%	16.7%
EINHE1AE Household income from social assistance	9.9%	7.1%	12.4%
EINHE1AF Household income from child and spousal support	9.9%	7.0%	12.5%
EINHE1AG Household income from other sources	11.2%	8.0%	14.0%
EINHEQ03 Total household income	24.1%	18.4%	29.1%

* Households where there was no spouse and/or no other members 15 year old and older were not included in the calculation of the imputation rate.

Note: Due to a technical problem with the questionnaire, most households were not asked about income for other household members.

11.2 Youth Income Imputation

Information on income, broken down in five sources, was asked as part of the youth component. The youth were asked their income from odd jobs, employers, scholarships or bursaries, parents and other sources.

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

Income was imputed for all respondents who were 18 or 19 years old, but only for some of the respondents who were 16 or 17 years old. Of the 1,761 youth who were 16 or 17 years old and who responded to the survey, only 1,599 had information from the youth component. Only these youth were considered for imputation. This is possible as youth who are 16 or 17 years old who only completed the child and/or adult component are considered a respondent. There was not enough information about the youth who did not complete the youth component to be able to impute income for them. From those who completed the youth component a total of 742 youth ages 16 to 19 years old had at least one source of income imputed.

Imputation was carried out for each of the five sources of income. Imputation for most cases was done using a nearest neighbour approach. This method first identifies a respondent to the Income Section who has the same characteristics as the individual with incomplete income data. Once the nearest neighbour has been identified, the missing sources of income are copied to the recipient record. The imputation assigned incomes to cases where sources of income were not reported by the respondent, or where the amounts reported in the Income Section did not agree with what was reported in the Labour Force Section. In the remaining cases, where only one of the four income sources was missing, and there was a total income provided in ranges, a plausible value was imputed.

If the youth provided an estimate of their total income in ranges, this value was used to help determine the donor. The imputation flags provide information on how the imputation was done.

Imputation flag = 1 Income range was used to determine the donor
Imputation flag = 2 Plausible value imputation
Imputation flag = 3 No additional income information was used to determine the donor

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

Variable	Imputation Rate (%)
EINYEQ1A Income from odd jobs	10.1
EINYDQ1B Income from employers	11.9
EINYEQ1E Income from scholarships or bursaries	1.5
EINYEQ1C Income from parents	7.1
EINYDQ1D Income from other sources	3.9
EINYED01 Total youth income	23.1

11.3 Motor and Social Development Imputation

Originally, to obtain the raw Motor and Social Development (MSD) score for a child (variable **EMSCS01**), all 15 applicable questions had to be answered either "Yes" or "No". However, it was noted that a large proportion of the incalculable records had only one or two missing responses among the 15 questions. Missing is defined as “Don’t know”, “Refusal” or “Not stated” and was primarily “Don’t know”. We decided that we could obtain a reasonably accurate score making use of the 13 or 14 valid responses and sensibly imputing for the missing items.

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

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

Naturally, to have 13 or 14 questions in common, all potential donors had to be in the same age in months range as the child to be imputed. For example, an eight month old child missing EMSCQ21 had potential donors age 7 to 9 months who were asked the same 15 questions (EMSCQ12 to EMSCQ26) and had the same pattern of “Yes” and “No” responses for EMSCQ12 to EMSCQ20 and EMSCQ22 to EMSCQ26.

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, 520 additional MSD scores were obtained by having at least one response imputed; 453 had exactly one response imputed and 67 had exactly two responses imputed. This represents 7.7% of all eligible children.

The imputation flag variables EMSCIS1A and EMSCIS1B identify which MSD questions were imputed. A value of zero means that no imputation was done for the MSD questions.

12.0 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 for a given date. 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 5 is a combination of samples selected in Cycles 1, 3 and 4 (1994, 1998 and 2000) and a new sample selected in Cycle 5 (2002). The sample selected in Cycle 2 has already been interviewed three times in Cycles 2, 3, and 4 and is no longer in the survey.

Based on its design, the NLSCY aims to satisfy two separate objectives. It is possible to produce longitudinal estimates, cross-sectional estimates or both for children aged 0 to 5. For children and youth aged 8 to 19 only longitudinal estimates can be produced. As a result, each unit in the survey has one, two or three weights associated with it. This is described below.

The **longitudinal weighting of Cycle 5 respondents** represents the original population at the time of selection. These weights are computed using all Cycle 5 respondents to represent the population at the time of their original selection. These weights were also calculated in the first four cycles. Therefore, there is a longitudinal weight for children who were selected in Cycle 1, which represents children ages 0 to 11 in January 1995, a longitudinal weight for children who were selected in Cycle 3, which represents children ages 0 and 1 in January 1999 and a longitudinal weight for children who were selected in Cycle 4, which represents children ages 0 and 1 in January 2001.

In Cycle 5, a second longitudinal weight, known as the “funnel” weight, was calculated for the Cycle 5 longitudinal respondents who were respondents in all cycles. The **longitudinal weighting of all-cycle respondents** represents the same population as the weighting of Cycle 5 respondents for the sample selected in Cycle 1, i.e. children ages 0 to 11 in January 1995.

In the previous cycles, a cross-sectional weight, which represented the population at the time of collection, was calculated for all the responding units. In the same way in Cycle 5, the new units that were selected would represent children ages 0 and 1 in January 2003, the children ages 0 and 1 who were selected in Cycle 4 would represent children ages 2 and 3 in January 2003, and so on. If we wanted to calculate a cross-sectional weight for all the children and youth in Cycle 5, the children who were 0 to 11 years old in Cycle 1 would represent children and youth ages 8 to 19 years old in January 2003. Unfortunately, because of the time delay between selection and Cycle 5 collection, this sample no longer properly represents children and youth ages 8 to 19 years old in January 2003. Since children and youth who have immigrated to Canada after 1994 did not have the chance to be represented in the sample chosen in 1994, the actual sample of 8 to 19 year olds in the NLSCY does not cover this part of the Canadian population of children and youth 8 to 19 years old. As well, at each cycle, units are lost because of non-response. Due to the combined effect of non-response at every cycle this sample is no longer representative. Therefore, for Cycle 5, **cross-sectional weighting** was only done for children ages 0 to 5 years old, that is to say, those introduced in Cycles 3, 4 and 5.

Three weight variables are available for the NLSCY. It should be noted, as discussed above, that all units do not have a value for all three weights.

- EWTCW01C is the cross-sectional weight
- EWTCW01L is the standard longitudinal weight
- EWTCdW1L is the funnel weight

12.1 Longitudinal Weighting of Cycle 5 Respondents

The NLSCY weighting strategy is based on a series of adjustments applied to an initial weight. In theory, each child's initial weight is approximately equal to the inverse of his/her probability of selection. For example, for households selected from the Labour Force Survey (LFS), the initial weight is the LFS subweight. For longitudinal children, the initial weight is determined using the weight computed for the cycle in which they were selected. The final weight, either cross-sectional or longitudinal, is obtained by applying various adjustments to the initial weight.

This section explains the adjustments made to the initial weights and the procedures for weighting the longitudinal samples.

12.1.1 Definition of a Longitudinal Respondent

As defined in Chapter 4.0, a longitudinal respondent is a child who was introduced in a previous cycle and whose adult component or child or youth component is complete. For youth 18 years old and above, the youth component must be completed in order to consider the youth a respondent. Children who were introduced in a previous cycle and died or moved outside Canada's ten provinces are also longitudinal respondents. They represent similar children in the reference population.

12.1.2 Weighting Method for the Longitudinal Samples

The first step in computing the longitudinal weight of Cycle 5 respondents is to determine the initial weight. An important adjustment in the weighting process is the non-response adjustment. The weight of respondents is adjusted for non-respondents using the characteristics of all responding children and youth.

To represent the characteristics of non-respondents more accurately, the most recent information, which is from Cycle 4, is given preference. However, some Cycle 5 respondents did not respond in Cycle 4. Consequently, they have no initial weight for that cycle. The first step, then, is to determine that weight. It will have to be based on the child’s weight in an earlier cycle.

Using the initial weight, two steps are required to obtain the longitudinal weight for the children selected in Cycles 1, 3 and 4. These adjustment factors are applied to the initial weight to produce the final longitudinal weight.

12.1.3 Determining the Initial Weight

Cycle 5 respondents may or may not have responded in Cycle 4. There were 826 Cycle 5 respondents who did not respond in Cycle 4. These respondents will be referred to as **converted respondents**. To use the Cycle 4 information to model non-response, each Cycle 5 respondent must have a Cycle 4 weight. Converted respondents do not have such a weight.

These 826 Cycle 5 respondents have a Cycle 4 weight of 0, while other respondents have a non-zero weight. The sum of the weights represents and must always represent the Canadian population. If a weight is assigned to each of these respondents, the weight of the units who responded in both Cycles 4 and 5 will have to be reduced. Some of the weight of the Cycle 4 and 5 respondents is transferred to the units converted in Cycle 5. The sum of the weights of all units in Cycle 4 does not change. The adjustment is computed for each age-province combination.

The initial weight of the **children introduced in Cycle 1** is defined as:

For the Cycle 4 and 5 respondents:

$$\text{initial weight}_{\text{Cycle 5}} = \frac{\text{Cycle 4 weight before post-stratification} \times \sum_{\text{respondents Cycles 4 and 5}} \text{Cycle 1 initial weight}}{\sum_{\text{respondents Cycles 4 and 5}} \text{Cycle 1 initial weight} + \sum_{\text{converted respondents}} \text{Cycle 1 initial weight}}$$

For the converted respondents in Cycle 5:

$$\text{initial weight}_{\text{Cycle 5}} = \frac{\text{Cycle 1 weight} \times \sum_{\text{respondents Cycles 4 and 5}} \text{Cycle 4 weight before post-stratification}}{\sum_{\text{respondents Cycles 4 and 5}} \text{Cycle 1 initial weight} + \sum_{\text{converted respondents}} \text{Cycle 1 initial weight}}$$

For **children introduced in Cycle 3 and Cycle 4**, there are no converted respondents. All children sampled in Cycle 5 were respondents in Cycle 4. Therefore, the initial weight is the Cycle 4 weight before post-stratification.

12.1.4 First Adjustment: Non-response Adjustment

The initial weights reflect the attrition (non-response) observed in Cycles 1 to 4. The Cycle 5 attrition must also be taken into account. Therefore, an adjustment factor is calculated, which reflects the Cycle 4 characteristics of the respondents and non-respondents.

Using Cycle 4 variables, homogeneous response groups (HRG) are created. The HRG method involves grouping individuals with the same likelihood of response. Then an adjustment factor is computed for each HRG. That factor is defined as follows:

$$\text{Non-response adjustment} = \frac{\sum_{\text{Respondents + non-respondents}} \text{adjusted weights in the HRG}}{\sum_{\text{Respondents}} \text{adjusted weights in the HRG}}$$

Certain constraints (size of the adjustment factor and minimum size of each HRG) are imposed when the HRGs are formed so that reasonable, reliable adjustment factors can be obtained.

There are separate sets of HRGs for the children introduced in Cycles 1, 3 and 4. The three sets are required because the samples do not necessarily have the same non-response behaviour. There is every reason to believe that this behaviour varies with the number of times the individual has been interviewed. The non-response adjustment model needs to take that into account.

12.1.5 Second Adjustment: Post-stratification Adjustment

The second adjustment factor ensures consistency between the estimates produced by the survey and Statistics Canada's population estimates. This method is called post-stratification. For the sample of children selected in Cycle 1, the target population is the population of all children between the ages of 0 and 11 at the beginning of 1995. The post-stratification adjustment of that sample ensures consistency between the sum of the weights and the January 1995 population estimate for each province-age-sex combination. For the sample of children selected in Cycle 3, the population estimates for January 1999 are used, and for children introduced in Cycle 4, the estimates for January 2001 are used.

12.2 Longitudinal Weighting of All-cycle Respondents – Children Introduced In Cycle 1

The longitudinal weighting of all-cycle respondents or funnel weighting is very similar to the longitudinal weighting of Cycle 5 respondents. In fact, the all-cycle respondents are a subset of the Cycle 5 respondents. The general approach is the same: an initial weight is multiplied by a non-response adjustment factor and a post-stratification adjustment factor. However, the initial weight and the creation of the homogeneous response groups are different.

12.2.1 Determining the Initial Weight

When computing the longitudinal weighting of respondents to all cycles, there are no converted respondents. The initial weight is therefore the Cycle 4 longitudinal weight of respondents to all cycles before post-stratification. This weight takes into account attrition from Cycles 1 to 4.

12.2.2 First Adjustment: Non-response Adjustment

Once again, non-response adjustment is based on the creation of HRGs. The HRGs are created using Cycle 4 variables. The HRGs are different from the HRGs created while weighting for all longitudinal respondents introduced in Cycle 1 since Cycle 5 respondents who didn’t answer in one of the previous cycles have a different non-response mechanism than respondents to all cycles. The adjustment factor is computed for each HRG.

12.2.3 Second Adjustment: Post-stratification Adjustment

As in the case of the weighting of Cycle 5 responding children, an adjustment factor is computed to ensure consistency between Statistics Canada’s population estimates and the NLSCY's estimates. The target population is again children aged 0 to 11 in January 1995. The adjustment is computed for each age-sex-province combination.

12.2.4 Comparison of the Longitudinal Weights of Cycle 5 Respondents and All-cycle Respondents

Some 20 variables were used to compare the estimates made using each set of longitudinal weights. The proportions for each combination were compared. No significant difference was observed for the variables considered. No matter which set of weights was

used, the conclusions were the same. For more details on choosing the right set of weights for the type of analysis being performed, see Chapter 18.0, Analytic Issues.

12.3 Cross-sectional Weighting for 0 to 5 Year Old Children

Cross-sectional weighting for 0 to 5 year old children involves representing the population at the time of collection, i.e., in January 2003. The cross-sectional sample consists of children introduced in Cycles 3 and 4 as well as children aged 0 and 1 who were first surveyed in Cycle 5.

In the following paragraphs, the adjustment factors that, when multiplied by the basic weights, produce the weights for the cross-sectional sample are described. Those adjustment factors vary depending on the cycle in which the child was first introduced.

First, cross-sectional weights were calculated independently for the children selected in 1998, 2000 and 2002. After that, each of those components represented its own target population. Those target populations are entirely separate, so they can be combined without further adjustments. The last step, post-stratification, ensures consistency between the survey estimates and Statistics Canada's population estimates.

12.3.1 Definition of a Cross-sectional Respondent

As defined in Chapter 4.0, 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 ten provinces are out-of-scope. They are not in the January 2003 target population.

12.3.2 Cross-Sectional Weight of Children First Surveyed in Cycle 5

Children selected from the Labour Force Survey sample
For children selected from the LFS sample, the weighting strategy is similar to the approach taken in previous cycles.

Adjustment 1: Adjustment for the number of rotation groups
The LFS sample is composed of six rotation groups. Each group is a representative sub-sample of the LFS target population. In the NLSCY, we used 15 rotation groups, but only 12 rotation groups to represent children aged 0 and 12 rotation groups to represent 1-year-old children. Hence the first adjustment is 6/12 for all 0- and 1-year-old children. The adjusted weight is obtained by multiplying the LFS weight by 6/12.

Adjustment 2: Non-response adjustment
In surveys such as the NLSCY, some households fail to provide responses for a variety of reasons (refusal, special circumstances, language problems, temporary absence). To compensate for the reduction in sample size due to non-response, the weights of respondents are increased. This adjustment is made by multiplying the subweights of respondent households by the following factor:

$$\text{Adjustment factor} = \frac{\sum_{\text{Sample}} \text{adjusted weights of households in the NLSCY stratum}}{\sum_{\text{Respondents}} \text{adjusted weights of households in the NLSCY stratum}}$$

In this equation, the adjusted weight is the weight obtained after Adjustment 1. A different factor is computed for each of the strata defined by the LFS specifically for non-response. Those strata are defined using the following information: province, economic region, census metropolitan area, urban or rural area, apartment frame, special region or not. Each stratum has a response rate of at least 60% and a minimum of 50 children. Strata that are too small or have response rate under 60% were combined until the above requirements were met.

Adjustment 3: Adjustment for households with more than one economic family
Some households contain more than one economic family. In such cases, one economic family must be chosen at random before a child can be selected. This adjustment is the inverse of the family's selection probability.

Adjustment 4: Adjustment for households with more than one in-scope child
In Cycle 5, only one child was surveyed in new households. Consequently, when there was more than one in-scope child in a family, one of them was selected at random. This adjustment compensated for the selection process.

This was the last adjustment made for these children before the weights were integrated.

12.3.3 Weighting of Children Selected in Cycles 3 and 4

It was not necessary to apply all the adjustments described in the preceding section, since that had already been done in Cycles 3 and 4. The initial weight was the same as the one used for longitudinal weighting of Cycle 5 respondents introduced in Cycle 4. Some children introduced in Cycle 3 were cross-sectional respondents but longitudinally out-of-scope. Since none of these children have longitudinal weights in Cycle 4, an initial weight must be computed for them. Two adjustments, as described below, were required.

Initial weight of children introduced in Cycle 3
Some children introduced in Cycle 3 have the particularity of being longitudinally out-of-scope due to their age. Nonetheless, they are still part of the Cycle 5 cross-sectional target population. Since there is Cycle 4 information available for those children, it is not necessary to process them as converted respondents. A new Cycle 4 non-response adjustment is computed to take into account these children and then apply to the Cycle 3 initial weight. The resulting weight is the Cycle 4 initial weight. Where w3 represents the Cycle 3 initial weight, we have:

$$\text{Adjustment factor} = \frac{\sum_{\text{Respondent s}} \text{w3} * \text{non-response adjustment} + \sum_{\text{Out-of-scope}} w3}{\sum_{\text{Respondent s + Out-of-scope}} \text{w3}}$$

as the adjustment factor. The Cycle 4 initial weight is therefore the Cycle 3 initial weight times this new adjustment factor.

Adjustment 1: Non-response adjustment
This factor augmented the basic weight to compensate for non-response. The adjustment used at this stage was computed with cross-sectional respondents. Out-of-scope children were not included in the numerator or the denominator.

Adjustment 2: Adjustment for interprovincial migration
The purpose of the second adjustment was to minimize the impact of rare types of interprovincial migration. Some children selected in 1998 or 2000 moved to another

province after the first interview. In some cases, that created excessive weights for the new province of residence. For example, the weight of a child selected in Ontario was much larger than the weight of a child selected in Prince Edward Island. If a child selected in Ontario moved to Prince Edward Island and kept his/her initial weight, it would have a huge impact on the estimates for Prince Edward Island. This type of migration is very rare in the target population. Accordingly, it is unreasonable to assume that the child who moved from Ontario to Prince Edward Island represented a large number of children who did likewise in the target population. This would be a very unusual event. As a result, the weight of such children was adjusted downward.

12.3.4 Integration of Weights

The two weight computation methods described above were used to produce estimates for their respective target populations. These target populations are completely separate, contrary to the previous cycles. Consequently, an adjustment factor did not need to be computed to compensate for the overlap of multiple populations. A final adjustment was also needed to ensure that the weights produced estimates consistent with population estimates derived from other data.

Post-stratification adjustment
The weights computed to that point were post-stratified to ensure that the national and provincial estimates were consistent with the population estimates for children aged 0 to 5 in January 2003. For Cycle 5, the post-strata were defined by province, age and sex. This adjustment factor was computed for each post-stratum. It was defined as the ratio of the population estimates to the sum of the post-stratum weights.

This adjustment was the final step in weighting the cross-sectional sample for Cycle 5 of the NLSCY.

13.0 Data Quality and Coverage

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

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

13.1 Sampling Error

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

Sampling error can be estimated using the sampling variance. For more details on calculating the estimated sampling error, see Chapters 14.0 and 15.0.

13.2 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 aren't known, it is very difficult to measure bias.

13.3 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. 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 Ph. D. and 5 with a Ph. D. There are many more chances for a non-Ph. D. to falsely report having a Ph. D. 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 Ph. D.'s 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 NSLCY 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 13.4. 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.

13.4 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.

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

13.5 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, don’t necessarily retain these properties on the survey file. For example, we see many incomes reported as exact multiples of $10,000, and many heights reported in exact feet. In the population, the number of households with income $19,501 to $20,500 is probably comparable in size to the number of households with income $20,501 to $21,500. The survey results would show a very different picture with the first group many times larger than the second due to respondent approximation of income.

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

13.6 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.

13.7 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".

13.8 Language of Interview

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

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

13.9 Conflicting Information

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

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

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

The data from the current cycle can also conflict with what has been collected in past cycles. For example, for some children, a parental separation was reported at Cycle 1, but at Cycle 5 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.

13.10 Data Quality for Body Mass Index

13.10.1 Body Mass Index

Body Mass Index (BMI) is a standardized scale to measure body mass. The release variable name for the respondent's BMI score is DBMIES01. A BMI score is calculated by dividing weight by height squared:

$$BMI = \frac{\text{Weight in Kilograms}}{(\text{Height in Meters}) \times (\text{Height in Meters})}, \text{ or}$$

$$BMI = \frac{(\text{Weight in Pounds}}{(\text{Height in inches}) \times (\text{Height in inches})) \times 703}$$

$$BMI = \frac{(\text{Weight in Kilograms}}{(\text{Height in centimeters}) \times (\text{Height in centimeters})) \times 10{,}000}$$

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

PMK reported for 2- to 11-year-olds - DHLCQ03B (height) and DHLCQ04A (weight)
Self-complete for 12- to 17-year-olds - DHTCbQ01 (height) and DHTCbQ02 (weight)

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

13.10.2 Body Mass Index – Centers for Disease Control

The CDC have proposed a set of percentile ranges to classify BMI scores as either: underweight, normal, at risk of overweight or overweight. These percentile ranges are age-specific by sex, and are based on American height and weight data. The CDC cutoffs are based on the person's age broken down into one month intervals. Consequently, in processing the NLSCY data, the age in months variable (DMMCdQ1B) was used to derive the cutoffs. The percentile ranges proposed by the CDC can potentially be used for 0- to 20-year-olds. The release name for this variable is DBMIES03.

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

13.10.3 Body Mass Index – Tim Cole, et al.

A set of international BMI cutoffs for 2- to 18-year-olds were proposed by Tim Cole, Mary Bellizzi, Katherine Flegal, and William Dietz in the British Medical Journal (Establishing a standard definition for child overweight and obesity worldwide: international survey. *BMJ* 2000;320:1–6). These cutoffs classify BMI scores into three categories: normal, overweight or obese. Similar to the CDC cutoffs, these categories are age and sex specific. In contrast to the CDC cutoffs, the international cutoffs are in half-year intervals for age and were based on studies from six nationally representative datasets of body mass indices in childhood (United States, Brazil, Great Britain, Hong Kong, Netherlands, Singapore). The release name for this variable is DBMIES02.

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

13.10.4 Body Mass Index – Data Quality

There are several issues that affect the quality of the BMI scores. First, there is a higher rate of non-response for the BMI variables as it is necessary that both the height and weight variables contain valid responses in order to calculate a score. Second, the data collected for height and weight are based solely on estimates provided by the parent or the youth rather than on accurate clinical measurements. The result of this method of collection is less accurate values for height and weight and correspondingly less accurate values for the BMI. Typically, a respondent will round the values of height or weight that they report, which leads to higher BMI values than would otherwise have been calculated based on clinical measurements. For example, a PMK will report the child as being 5' (feet) tall instead of 5'2"(inches) or 4'10", or maybe they will report that the child's weight is 110 pounds instead of 113 or 108. These small errors in estimated height and weight can translate into a much larger degree of error in the BMI resulting in a change in classification from 'overweight' to 'normal' or 'obese' depending on how height and/or weight was rounded.

13.11 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.

13.12 Person Most Knowledgeable

At each cycle, one adult in the household is identified as the person most knowledgeable about the child. The PMK answers the Child component, giving information about the child's health, education, behaviour, etc. The child's characteristics are measured indirectly as reported by the PMK. From cycle to cycle, the PMK can change. For a given child, perhaps the mother was the PMK at Cycle 1, then the father at Cycle 2 and Cycle 3, and then the mother again at Cycle 4 and Cycle 5. Observed changes in the child's characteristics may be due, in part, to having different people answer these questions. The child's environment is not truly different, but the person answering the questions has changed and, naturally, has a different perspective. The variable EDMCD06 shows the relationship of the PMK to the child.

13.13 Total 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. In Chapter 10.0, cross-sectional and longitudinal response rates were computed for various domains:

provinces, age groups, and others. More details on the weighting procedure and how it attempts to adjust for total non-response are given in Chapter 12.0.

In this section, this potential bias due to non-response is discussed in more detail.

Non-response models using variables that are most closely related to the likelihood of responding were used during weighting. This technique adjusts the sampling weights to correct for the potential bias due to non-response. However, since not every survey variable can be considered when doing this modelling, it does not guarantee that there is no bias due to non-response.

The homogeneous response groups (HRG) were created such that the weights of respondents will compensate for the non-respondents having certain similar characteristics (e.g., number of parents in the household, income range, geography). Still, within any given HRG, the non-respondents may differ from the respondents in other important unobserved or unknown ways.

Non-response cumulates over time. As we have fewer and fewer participants, the estimated sampling error increases, but the potential for bias also increases. After many cycles, it would be highly improbable that the participants who continue to co-operate are a random sub-sample of the Cycle 1 respondents. As I have said, with the survey weights, we attempt to control this as much as possible, but the potential for non-response bias is still present.

13.14 Partial Non-response

Just because a respondent agrees to participate in the survey, it doesn't necessarily mean we have their answers to every single NLSCY question that applies to them (see Section 10.3). In some cases, a respondent may not answer a particular question or even an entire module or component.

Item non-response can occur when the respondent feels the question matter is too sensitive or the respondent may legitimately not recall or know the answer. On the Self-complete Questionnaire, items may be simply overlooked or skipped because instructions were not followed.

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 cooperate and answer the Child and Adult component, but her spouse may refuse to do his Adult 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

Whenever there is partial non-response, we are left with a “hole” in the dataset.

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, re-weighting 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 some measures to account for it.

Statistics Canada does impute missing values for a few of the questions in the NLSCY: adult income, youth income, household income; adult labour force and Motor and Social Development items (see Chapter 11.0 for more details). Missing values for all other questions will have a

"Don't know", "Refusal" or "Not stated" code on the final data file. It is up to each data user to deal with partial non-response in a manner that is appropriate to the research being undertaken.

A short workshop on techniques for dealing with non-response for the NLSCY is being developed. It is planned that this workshop will be presented at interested Research Data Centres in 2005.

13.15 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, 3, 4, and 5, but not Cycle 2.

If data from every cycle are crucial, the analyst can limit himself to children without cycle non-response and use the longitudinal weights for this group, variable EWTCWd1L. This weight is available for children introduced in Cycle 1.

13.16 Coverage

A sample is selected from a frame. That frame is intended to represent the target population. For the NLSCY, the sample is intended to represent the population for both longitudinal and for cross-sectional purposes, at the time of collection. However, when the frame does not represent the population accurately, there may be coverage errors. If the frame does not contain all the units in the target population, there is under-coverage. If the frame contains units that are not in the target population, there is over-coverage. More detail on the sampling frames can be found in Chapter 5.0.

Coverage problems may arise at the time of sample selection or during subsequent contacts.

13.16.1 Coverage Issues

1) **Respondents from the Labour Force Survey**
   Most of the children surveyed in the NLSCY were selected from households that had taken part in the Labour Force Survey (LFS). The LFS does not cover every child living in the ten provinces. For example, children living on Indian Reserves are not part of the LFS target population. These exclusions are listed in Chapter 5.0. In addition, this selection method leads to some problems related to coverage.

   The first problem stems from the fact that only LFS respondents were considered in forming the NLSCY sample. Some households that were LFS non-respondents had eligible children. These households were excluded from the NLSCY sample, which could lead to coverage errors.

   The second problem relates to the fact that only households that had children at the time of the LFS were included in the NLSCY sample. If the household membership did not include eligible children at the time of the LFS, they were out-of-scope for the NLSCY. These households may change between the time of the LFS interview and when the NLSCY started interviewing.

   Thirdly, the LFS underwent its last redesign in 1994. Over time, the number of occupied dwellings in certain areas changes and the efficiency of the survey design gradually deteriorates. This has a small impact on all surveys using the LFS as a frame, including the NLSCY.

Lastly, the LFS, though generally a good vehicle for social surveys, was designed to measure labour force characteristics, not necessarily measure characteristics of Canadian children. There isn't any real problem per se. We are only noting that the sample design, though near-optimal and very cost efficient, is not optimal for a survey targeting children.

2) Respondents from the Birth Register

In Cycle 3, the Birth Register was used to select a large sample of one-year-olds. There are also some limitations of the Birth Register frame. One problem stems from the fact that the Birth Register does not have full coverage of every child born in 1997 who resides in Canada at sample selection time. Some births may not have been reported or were reported after the sample was selected and one-year-olds from immigrant families would not be on the Birth Register and therefore had no chance of being selected.

Another problem is due to the length of time between birth and sample selection. Some children may no longer be in the target population because they have died or have moved out of the country. Nevertheless, these potential errors are generally minor, and their impact is negligible.

3) Non-uniform coverage by month of birth

For the children aged 8 to 19 at Cycle 5, we have a fairly uniform distribution of sample by month of birth. Likewise, for children born in 1999 or later, all birth months are nearly equally represented – with the minor exception that the very youngest children born in the last few months of 2002 are slightly underrepresented.

This distribution of sample by month of birth is not at all uniform for children born in 1997 or 1998. There is a coverage anomaly with respect to month of birth in that we have no children in the sample born in January, February, March or April of 1997. The children born in the latter two-thirds of 1997 are weighted to represent all 1997 born children. This is a consequence of a change in how children are classified by age. In Cycle 3, the age at the time of the interview was the key to determining content. It made sense, at the time of Cycle 3 collection in 1998/1999, to avoid introducing two-year-olds – the target population was 0- and 1-year-old children. At Cycle 4 and Cycle 5, the child's year of birth was the key to determining contact. The child's actual age at the time of the interview, and hence specific month of birth was no longer crucial.

For children born in 1998 (considered 4 years old at Cycle 5), the sample contains nearly ten times as many children born in January as born in December. Children born in January, February and March are overrepresented relative to the younger children born that same year.

4) Non-uniform coverage by age in months

There are no children aged 34, 35, or 36 months at the time of the Cycle 5 interview. This is a consequence of how the sample was collected in Cycle 5. At the first wave of collection, starting in September 2002, only those children born January to August in 1999 or 2000 were included in this wave. In the second wave in April 2003, children born September to December in 1999 or 2000 were interviewed. As a consequence, most children born in 2000 were 25 to 32 months old at the time of interview and most children born in 1999 were 37 to 44 months old at the time of interview. There is a gap for the 33 to 36 months old range.

5) Coverage and sample attrition

Even if we could exclude any potential error at the time of the first collection, another type of error can surface when households are contacted for a second, third or more

times. This error is due to non-response. As noted in the previous section, some households with specific characteristics are less likely to respond in each succeeding cycle. A good example would be households with low-income. Although non-response adjustment can compensate for the loss of these respondents in the short term, the accrued loss of low-income households may lead to biased estimates when the non-respondents can no longer be compensated by similarly profiled respondents.

13.16.2 Coverage and Immigration

For the age 0 to 5 cohort at Cycle 5, children in previous cycles were selected to represent their respective 1999 and 2001 target populations. By assigning cross-sectional weights to these same children to represent the 2003 population, we introduce two potential coverage problems: international immigration and inter-provincial migration.

1) International immigration

Between the time when the sample was selected and when the Cycle 5 collection started, a period of four or two years had elapsed, depending on in which cycle the sample was selected. Immigrant children who arrived in Canada during that period were not eligible for selection. The following table attempts to show the differences between the cross-sectional and longitudinal target populations. The *Total Number of Children* column is the total number of children aged 2 to 5 in January 2003. In other words, this is the size of the cross-sectional target population for this age group. The *Longitudinal Target Size* column is the sum of the number of 0- and 1-year-old children present in January 1999 and the number of 0- and 1-year-old children in January 2001. Basically, this is the size of our longitudinal target population.

Table 13.15: Difference Between Estimates of Total Population for the 2 to 5 Year Age Group by Province

Province	Total Number of Children	Longitudinal Target Size	Difference	Relative Difference (%)
Ontario	575,535	534,634	40,901	7.11
Alberta	151,065	145,507	5,558	3.68
Quebec	301,997	293,858	8,139	2.70
Prince Edward Island	6,059	5,929	130	2.15
British Columbia	168,632	165,071	3,561	2.11
Manitoba	52,911	51,808	1,103	2.08
New Brunswick	30,365	29,981	384	1.26
Nova Scotia	37,811	37,379	432	1.14
Newfoundland and Labrador	20,087	20,050	37	0.18
Saskatchewan	44,628	44,789	-161	-0.36
Canada	**1,389,090**	**1,329,006**	**60,084**	**4.33**

The differences are fairly large for Ontario (over 7%) and overall (4.3%). The negative difference for Saskatchewan means that there are fewer children aged 2 to 5 in January 2003 than there were aged 0 to 1 in January 1999 or age 0 to 1 in January 2001.

Since recent immigrants and children born in Canada may have different characteristics, we run the risk of bias.

Also, this table does not tell the whole story. A certain number of children aged 2 to 5 die or leave the ten provinces between when they are selected and the Cycle 5 collection. These children cannot contribute to the cross-sectional estimates at Cycle 5 and such children are not counted in the first column. Our cross-sectional under-coverage is at least 4.3% and in reality, is slightly larger because 60,084 is a count of net immigration, and ideally, we would use a count of gross immigration. Gross immigration is calculated as the cross-sectional target population size minus the longitudinal target population size plus the death and emigration count.

2) Inter-provincial migration

The second source of error is inter-provincial migration. The cross-sectional sample is intended to represent the population of children for each province at the time of collection. The province of residence may be different from the province at the time of selection if the child has moved. The weight represents the province at time of selection and can potentially be much larger than the weights of the other children in the province of residence at the time of collection. When this occurs the initial sampling weight is modified. This deviation from the original sample design impacts the known probability of selection associated with the individual that has jumped stratum.

The problem can be described using this scenario. Children, who were selected in provinces with low probability of selection, such as Ontario and Quebec, moved to a small province with a much higher probability of selection. Such children might dominate the estimation because of their excessively large sampling weight. To address this situation, outlier detection techniques have been used to determine which cases required a modification to the initial weight. Fewer than ten records were affected.

Alternately, children who migrated from small provinces to large provinces had much smaller sampling weights than children originally selected in the large provinces. Their impact on the provincial estimates is therefore reduced (perhaps becoming insignificant) compared to what it would have been if they had stayed in their original province of selection. No action is taken when the weights are much smaller that the mean weight of the current province of residence.

13.17 Indicators of Sampling Error

This section aims to provide users some idea of the level of sampling error associated with estimated proportions for some key domains. The indicators are based on the average standard error of several estimates. Users should note that sample size plays a big part in determining the sampling error. To achieve a 50% reduction in the standard error associated with an estimate requires four times the sample size.

The other factors that influence sampling error are the sample design and the variability in the population of the characteristic being measured.

Table 13.16.1: Estimated Standard Error Associated with a Proportion of 50%, by Age Group and Province of Selection

Province	Age					
	8 to 9	10 to 11	12 to 13	14 to 15	16 to 17	18 to 19
Newfoundland and Labrador	4.6	5.5	5.5	5.3	5.0	5.1
Prince Edward Island	6.1	7.5	6.9	7.4	8.2	7.2
Nova Scotia	4.0	5.2	4.9	5.6	5.6	5.4
New Brunswick	4.2	4.7	4.6	5.2	5.4	6.0
Quebec	2.7	3.4	3.7	4.1	3.8	4.0
Ontario	2.4	3.1	3.3	3.4	3.3	3.6
Manitoba	3.9	4.8	5.4	6.1	5.5	6.3
Saskatchewan	3.7	4.4	4.4	5.2	4.7	5.2
Alberta	3.5	4.4	4.6	4.5	4.6	4.6
British Columbia	4.0	4.9	5.2	5.3	5.3	5.6
Canada	**1.3**	**1.7**	**1.8**	**1.8**	**1.7**	**1.9**

Table 13.16.2: Estimated Standard Error Associated with a Proportion of 50%, by Age Group and Province of Residence, Cross-sectional Weights

Province	Age		
	0 to 1	2 to 3	4 to 5
Newfoundland and Labrador	5.5	5.3	4.5
Prince Edward Island	5.4	5.3	4.0
Nova Scotia	4.2	3.9	2.8
New Brunswick	4.3	4.1	2.8
Quebec	2.8	2.6	1.9
Ontario	2.0	1.9	1.6
Manitoba	3.4	3.5	2.6
Saskatchewan	3.5	3.2	2.6
Alberta	3.1	3.1	2.3
British Columbia	3.2	3.2	2.5
Canada	**1.2**	**1.1**	**0.9**

Ninety-five percent confidence intervals can be constructed from the standard errors in these tables as 50% ± 1.96 * standard error. For example, for Canada for children aged 4 to 5, the 95% confidence interval is 48.2% to 51.8% and in Quebec for the same age range, the confidence interval becomes 46.3% to 53.7%.

For estimated measures of sampling error associated with a wider range of domains and proportions, users are encouraged to consult the Excel sampling variance spreadsheets. For measures of sampling error for other types of estimates, users are encouraged to use the

Bootstrap weights. These are described in Chapter 15.0.

13.18 Conclusion

Data quality is affected by various sources of error. Efforts are made at all steps (interviewer training, collection monitoring, processing, weighting, etc.) to reduce the potential for errors.

Data users are encouraged to consider how sampling and non-sampling errors may affect the variables they are attempting to analyse.

14.0 Guidelines for Tabulation, Analysis and Release

This chapter of the documentation outlines the guidelines to be adhered to by users tabulating, analysing, 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.

14.1 Rounding Guidelines

In order that estimates for publication or other release derived from the National Longitudinal survey of Children and Youth (NLSCY) microdata files correspond to those produced by Statistics Canada, users are urged to adhere to the following guidelines regarding the rounding of such estimates:

a) Estimates in the main body of a statistical table are to be rounded to the nearest hundred units using the normal rounding technique. In normal rounding, if the first or only digit to be dropped is 0 to 4, the last digit to be retained is not changed. If the first or only digit to be dropped is 5 to 9, the last digit to be retained is raised by one. For example, in normal rounding to the nearest 100, if the last two digits are between 00 and 49, they are changed to 00 and the preceding digit (the hundreds digit) is left unchanged. If the last digits are between 50 and 99 they are changed to 00 and the preceding digit is incremented by 1. For example, an estimated total of 21,352 would be rounded to 21,400.

b) Marginal sub-totals and totals in statistical tables are to be derived from corresponding unrounded components and then are to be rounded themselves to the nearest 100 units using normal rounding.

c) Averages, proportions, rates and percentages are to be computed from unrounded components (i.e., numerators and/or denominators) and then are to be rounded themselves to one decimal using normal rounding.

d) Sums and differences of aggregates (or ratios) are to be derived from their corresponding unrounded components and then are to be rounded themselves to the nearest 100 units (or the nearest one decimal) using normal rounding.

e) In instances where, due to technical or other limitations, a rounding technique other than normal rounding is used resulting in estimates to be published or otherwise released which differ from corresponding estimates published by Statistics Canada, users are urged to note the reason for such differences in the publication or release document(s).

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

14.2 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.

When producing simple estimates, including the production of ordinary statistical tables, users must 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.

14.3 Calculating Survey Estimates

Unit of Analysis
The NLSCY files have been set up so that the child is the unit of analysis. The weight that can be found on each record (EWTCW01C for the cross-sectional sample, EWTCW01L for the longitudinal sample of children introduced in Cycle 1, Cycle 3 and Cycle 4, and EWTCWd1L for the longitudinal sample of children introduced in Cycle 1 who were respondents to all five cycles) is a "child" weight. Estimates of parents or families cannot be made from the NLSCY microdata file.

14.3.1 Categorical Estimates

Categorical estimates are estimates of the number, or percentage of the surveyed population possessing certain characteristics or falling into some defined category. An estimate of the number of persons possessing a certain characteristic may also be referred to as an estimate of an aggregate.

Examples of Categorical Questions:

Q: Was (the child) born before, after or on the due date?
R: Before / After / On due date

Q: Compared to other babies in general, would you say the (the child's) health at birth was:
R: Excellent / Very good / Good / Fair / Poor

Estimates of the number of children with a certain characteristic can be obtained from the microdata file by summing the final weights of all records possessing the characteristic(s) of interest. These estimates may be cross-sectional or longitudinal. Proportions and ratios of the form $\hat{X} / \hat{Y}$ are obtained by:

a) summing the final weights of records having the characteristic of interest for the numerator ($\hat{X}$),
b) summing the final weights of records having the characteristic of interest for the denominator ($\hat{Y}$), then
c) dividing the numerator estimate by the denominator estimate ($\hat{X} / \hat{Y}$).

14.3.2 Quantitative Estimates

Quantitative estimates are estimates of totals or of means, medians and other measures of central tendency of quantities based upon some or all of the members of the surveyed population. They also specifically involve estimates of the form $\hat{X} / \hat{Y}$ where $\hat{X}$ and $\hat{Y}$ are each survey estimates.

An example of a quantitative estimate is the average number of days of care received by babies who required special medical care following birth. The numerator is an estimate of the total number of days for which babies required special care. The denominator is the number of babies who required special care at birth.

Examples of Quantitative Questions:

Q: For how many days, in total, was this care received?
R: |_|_|_| days

Q: What was the child's weight at birth in pounds and ounces?
R: |_|_| pounds and |_|_| ounces

Estimates of quantities can be obtained from the microdata file by multiplying the value of the variable of interest by the final weight for each record, then summing this quantity over all records of interest. For example, to obtain an estimate of the total number of days of special care received by infants who were born prematurely multiply the number of days for which special care was received by the final weight for the record, then sum this value over all records for which the child was born prematurely.

To obtain a weighted average of the form $\hat{X} / \hat{Y}$, the numerator ($\hat{X}$) is calculated as for a quantitative estimate and the denominator ($\hat{Y}$) is calculated as for a categorical estimate. For example, to estimate the average number of days spent in special care by premature babies,

a) estimate the total number of days ($\hat{X}$) as described above;
b) estimate the number of children ($\hat{Y}$) in this category by summing the final weights of all records for babies which were premature; then
c) divide estimate a) by estimate b) ($\hat{X} / \hat{Y}$).

14.4 Guidelines for Statistical Modeling

Sample Design
As mentioned earlier, the NLSCY is based upon a complex sample design, with stratification, multiple stages of selection, and unequal probabilities of selection of respondents. Using data from such complex surveys presents challenges to analysts because the survey design and the selection probabilities affect the estimation and variance calculation procedures that should be used. In order for survey estimates and analyses to be free from bias, the appropriate survey weights should be used whenever possible. For more details about the survey weights, refer to Chapter 12.0.

Variance Estimates
While many analysis procedures found in statistical packages allow weights to be used, the meaning or definition of the weight in these procedures may differ from that which is appropriate in a sample survey framework, with the result that while in many cases the estimates produced by the packages are correct, the variance estimates that are calculated are poor. Users should estimate design-consistent variances using the Bootstrap weights and tools described in Chapter 15.0.

14.5 Release Guidelines

Data users must not release or publish any estimate that would allow the identification of a specific respondent or reveal any individual's responses. For this reason, estimates (for example, the cells in a cross-tabulation) should have at least five contributing respondents.

Apart from the above requirement, all estimates can be considered releasable, but before releasing and/or publishing any estimate from the NLSCY, users should first determine the quality

of the estimate. This means that the standard error associated with the estimate must be calculated (see Chapter 15.0). Users should also consider how non-sampling errors discussed in Chapter 13.0 may affect the estimate.

Also, the number of children who contribute to the calculation of the estimate should be determined. If this number is small, the standard error associated with the weighted estimate is expected to be large, and the estimate is likely unreliable for most purposes. When considering proportions, one can certainly infer from the survey that a certain characteristic is rare, but the true rate cannot usually be determined from the survey data. For example, one can use the NLSCY to estimate that 1 out of 1,500 children have some specific health problem, but the true rate may be twice this estimate or half this estimate. In such instances, releasing a less exact estimate (i.e., the rate is estimated to be less than 0.5%) would be preferable as this is the level of precision that the survey can legitimately claim.

15.0 Variance Estimation

The National Longitudinal survey of Children and Youth (NLSCY) is a probabilistic survey for which a sample has been selected to represent the target population. A given variability is inherent in the sampling process. As described in Chapter 13.0, survey estimates based on a sample are not exact; they have a certain sampling error associated with them. The amount of sampling error depends on the sample size, the variability in the population of the characteristic being measured, the sample design, and the response rate achieved. Two statistical terms exist for the sampling error associated with a given estimate: the sampling variance, which can be obtained by the methods described in this chapter; and the standard error, which is simply the square root of the sampling variance. These measures of sampling error are themselves estimates, but for convenience, we will use the word sampling variance to mean the estimated sampling variance and standard error to mean the estimated standard error.

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.

15.1 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. This section aims to explain these terms with the aid of some examples.

Unfortunately, the term variance is used for two very different things. In the first instance, the variability of a characteristic in the population is called the **population variance**. For example, consider the population of all 10 year-old boys in Canada, there is some variability in their heights in centimetres. This variability is called the population variance.

The mathematical definition of the population variance is:

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

where y_i is the value for person i ;

$\overline{Y}$ is the average of y_i ; and

N is the size of the population.

The population variance is simply a characteristic of the population. Its value does not depend at all on the survey sample. For example, some populations have little variability (e.g. the height of 10 year old boys in the population); others have a large variability (e.g. the distance employees live from their place of employment).

The other variance term is the **sampling variance.** The sampling variance of an estimate is the variability in the estimate due to the estimate coming from a sample and not a census. Returning to the case of 10 year-old boys in Canada, let us draw a sample from this population. From this sample, we calculate a survey estimate, the mean height. Associated with this estimate is the sampling variance of the estimate, which unlike the population variance does depend on the sample. If we take a different sample, we will get a different value for the sampling variance. In general, larger samples yield smaller sampling variances. So, by taking a larger sample of 10 year-old boys in Canada, we can reduce the sampling variance associated with the estimate of their mean height. For some simple designs, like simple random sampling, a straightforward

formula exists to calculate the sampling variance associated with an estimate. For more complex designs, other techniques, such as the use of Bootstrap weights described in this chapter, are needed to obtain a value for the sampling variance. It is the sampling variance that we discuss in this chapter.

The square root of the population variance is called the **standard deviation**.

The square root of the sampling variance associated with an estimate is called the **standard error** of the estimate.

As mentioned in Chapter 15.0, the population variance is one thing that affects the sampling variance of the estimate. If the population is very homogenous in the characteristic we are interested in, then we only need a small sample to get a reliable estimate. If, on the other hand, the population is very heterogeneous, we need a larger sample to capture the population variability. In other words, there is more sampling error when there is more potential for different samples to yield very different results.

Note that the population variance (or standard deviation) is not usually known. One can obtain an estimate of the standard deviation from the survey using the following formula.

$$\text{Estimated Standard Deviation} = \sqrt{\frac{\sum_i w_i (y_i - \bar{y})^2}{(\sum_i w_i) - 1}}$$

where w_i is the weight associated with observed value y_i.

This estimate of the standard deviation can be obtained from statistical software. For example, the SAS option VARDEF = WFD with PROC MEANS (or PROC UNIVARIATE or other procedures) and the appropriate WEIGHT statement will give the estimate of the standard deviation.

However, it should be noted that the estimated standard deviation alone is not a measure of the quality of an estimate. One must calculate the standard error associated with the estimate.

15.2 Coefficient of Variation

Another measure of sampling error is the coefficient of variation (CV). The coefficient of variation is defined as the standard error of the estimate divided by the estimate itself, that is:

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

As opposed 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 example, CVs of provincial totals may be more useful indicators of sampling error than the corresponding sampling variances.

Although CVs are useful for indicating the quality of estimates such as totals, there are some pitfalls that users should be aware of when using CVs to examine the quality of proportions. A few potential problems are outlined below.

Issue 1: Coefficients of variation for very small (or very large) proportions
Since the standard error of a proportion p is the same as the standard error of the proportion $(1-p)$, the CVs of p and $(1-p)$ may differ substantially because the denominators are p and $(1-p)$ respectively. One can imagine a scenario when p is very small giving a very large CV for p, but the CV of $(1-p)$ being excellent. Let's use the following example to illustrate. Suppose we have calculated the standard error of the estimates p and $(1-p)$ as 0.0475, for a value of p of 0.95. The CV for the value of p, 0.95 would be:

0.0475 / 0.95 = 0.05 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%. This is a very poor CV.

Issue 2: Applying the Quality Guidelines in the case of proportions
The quality guidelines can be a very useful tool when determining the quality of an estimate. However, in the case of proportions, the guidelines should be intelligently applied, as the following examples show.

Example 1: An estimated proportion of 0.5 with a 99% confidence interval of 0.1 to 0.9 falls into the **marginal** category, using the previously published quality guidelines (the CV is 31%). But the confidence interval is so large, that the estimate is not really giving us much information.

Example 2: Suppose that we have 27,000 sampled individuals of which 44 have a characteristic we are interested in studying. Using the survey weights, we calculate that 0.16% of the population has this characteristic, with a corresponding CV of 34%. A CV of 34% is classified as **unacceptable** by applying the previously published quality guidelines. Although 0.16% is a very small proportion, if we construct its 99% confidence interval we get (0.02%, 0.30%). This means we have considerable confidence that the true rate is less than, say 0.5%. Of course, data users should be cautious as the true rate could be 0.05% or 0.25% or even 0.30%. Depending on the goal of the research, maybe a statement that the estimate is smaller than 0.5% is meaningful. Therefore, blindly applying the quality guidelines and throwing away an estimate because of its high CV may not be appropriate.

Consequently, if users wish to use CVs as a measure of sampling error when dealing with proportions, they are strongly encouraged to calculate the CVs for both the proportions p and $(1-p)$. CVs associated with proportions, particularly proportions greater than 0.5, can be misleading as the above examples illustrate. As well, users are encouraged to construct confidence intervals as a measure of the sampling error when dealing with estimated proportions.

Research is currently taking place to find better alternatives than the CVs for the extreme proportions (close to 0 and close to 1). However, for now, a solution which will meet all needs has yet to be found. Consequently, users must use caution if they wish to use CVs for proportions.

15.3 Importance of the Sampling Variance

The sampling variance of an estimate is one way to characterize the quality of the estimate. An estimate with a high sampling variance is considered unreliable (i.e., not precise). An estimate with low sampling variance is usually more reliable, although still subject to non-sampling errors

such as response errors, which are discussed in Chapter 13.0. Stating survey estimates without corresponding estimates of the sampling error can be very misleading. Also, the sampling variance is required for statistical tests such as hypothesis tests, which determine if two estimates are statistically significantly different.

15.4 Sampling Variance Calculation

It would be difficult to derive an exact formula to calculate the sampling variance for the NLSCY due to the complex sample design, non-response adjustments, treatment of out-of-scope units, and the post-stratification. A very good way to approximate the sampling variance is to use the Bootstrap method. The idea behind the Bootstrap 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 12.0). 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. For the NLSCY, a set of 1,000 Bootstrap weights is available. The sampling variance calculation using these 1,000 Bootstrap weights involves calculating the estimates with each of these 1,000 weights and then calculating the variance of these 1,000 estimates.

Two tools, both making use of the Bootstrap weights, have been developed to help users calculate the sampling variance and the CV for their estimates. These tools are:

- An Excel spreadsheet, with a Visual Basic interface, that enables users to retrieve approximate sampling variances for proportions across a large number of domains (e.g., by age and by province).
- Macros to calculate the sampling variance, using the Bootstrap weights.

The choice of tool to use depends on the type of analysis and the level of precision required.

In previous cycles, a third variance approximation tool was available: CV look-up tables. Using a representative design effect (the design effect compares the variance of estimators from the NLSCY sample design to those from a simple random sample) users were able to obtain CVs for some domains, by age cohort or by province. These tables are not available for Cycle 5 since the Visual Basic interface and Bootstrap macros are more flexible and more accurate.

15.4.1 Spreadsheet with Approximate Sampling Variances for Proportions

A set of spreadsheets is available to users to (approximately) calculate the sampling variance 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.

15.4.2 SAS and SPSS Macros to Calculate the Sampling Variance Using the Bootstrap Weights: Bootvar 3.0

SAS and SPSS macros have been developed to calculate the sampling variance using the Bootstrap weights. From the methods available, this replication method is the best approach to obtain a very good estimate of the sampling variance. For any domain, the sampling variance of estimates such as: totals, ratios, difference of ratios, linear and logistic regression coefficients, can be calculated using this method. The sampling variance calculated using this method takes into account the sample design and also, the specificities of the variable of interest. Finally, as opposed to the spreadsheet, the user is not restricted to pre-defined domains.

This method has many advantages but requires more work from the researcher. The sampling variance calculation using these macros is more time consuming than using the spreadsheet. The user must first become familiar with the macros before using them. However, these macros have been developed in such a way that they are easy to use. The researcher must have access to the macros, to the data files and to the Bootstrap weight files. Access to these tools is possible in the Statistics Canada Research Data Centres (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, this method is strongly recommended to calculate the sampling variance of any estimate which must be published, as this method provides a more exact measure of the sampling variance.

Again, details on how to use these programs can be found in separate documentation that accompanies the programs and Bootstrap weights.

Each set of weights available for Cycle 5 (cross-sectional, longitudinal weights for those introduced at Cycle 1, funnel weights for those introduced at Cycle 1, longitudinal weights for those introduced at Cycle 3, and longitudinal weights for those introduced at Cycle 4) has a corresponding file with Bootstrap weights.

15.4.3 Other Methods

Other commercially available software can properly calculate the sampling variance from the Bootstrap weights provided. SUDAAN (setting design = BRR) and WesVar are two such software available in the RDCs. The results are equivalent to those obtained from Bootvar.

To calculate the sampling variance for estimates not included in Bootvar, analysts may wish to write their own programs implementing the Bootstrap method. However, this is not a trivial matter.

15.4.4 Taylor Linearization and Other Techniques

The Bootstrap weight files contain variables indicating the primary sampling unit (PSU) and stratum from which the individual was selected. Some existing software packages (such as Stata, SUDAAN or SAS) have procedures that calculate sampling variance estimates using design information (stratum and PSU) and the survey weights. The technique is known by several names: Taylor Linearization or Binder or robust variance estimation. The problem with using these procedures with the NLSCY data is that they require at least two PSUs per stratum, and the NLSCY very often does not satisfy this requirement. Although collapsing of strata is possible, its effectiveness at this point is

unclear, as a thorough comparison of sampling variances obtained this way to Bootstrap sampling variances has not yet been done. 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 can calculate the sampling variance of estimates using built-in procedures (e.g., PROC UNIVARIATE in SAS). However, many of these routines do not take into account the sample design (e.g., stratification), which means the variance is usually underestimated. Therefore, these procedures are not recommended since they can lead to erroneous conclusions.

16.0 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 for Cycle 5 (collection in 2002-2003) 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 had a history of providing advice to the NLSCY Team.

Two early childhood assessments were selected to be administered to children four and five years of age in addition to the Peabody Picture Vocabulary Test - Revised (PPVT-R); the Who Am I? and the Number Knowledge Test. The Ages and Stages Questionnaires were also added in Cycle 4 to help assess the development level of children from 4 to 47 months old. These questionnaires are described in Chapter 8.0, Content of the Survey.

The NLSCY conducts direct assessments of older children these assessments are also described in this chapter.

16.1 The Peabody Picture Vocabulary Test - Revised

The PPVT-R was designed to measure receptive or hearing vocabulary and can be used for any age group, up to adult. The test was developed by Lloyd and Leota Dunn, at the University of Hawaii, and has been widely used in large-scale data collections as well as assessments. A French adaptation of the PPVT-R was developed by the test's authors and Claudia M. Thériault at St. Thomas University in Fredericton, New Brunswick. The French test is called the Échelle de vocabulaire en images de Peabody (EVIP).

For the NLSCY, the PPVT-R was used to measure school readiness for children in the four to five year age group. Verbal parental consent was required before the test was administered. If permission was granted, the interviewer then administered the test to the child in the home. The child looked at pictures on an easel and identified the picture that matched the word the interviewer read out.

A total raw score was calculated for each child who completed the PPVT-R by computing correct responses. A standardized score was also assigned to each child. Standard scores allow for comparisons of scores across age groups. Obviously a five year-old would be expected to perform better on the PPVT-R than a four year-old and thus have a higher score. The standard score takes into account the child's age.

Standard scores for a test are usually based on the distribution of scores obtained from the entire population. In the absence of scores for the entire population a representative sample distribution, called the norm sample, is more often used. Each cycle the NLSCY yields a representative sample of children. Any of those samples would be a viable option to estimate the distribution of scores measured in the population. While each sample is selected probabilistically, albeit for different reference period, experts in the field of cognition might disagree as to whether

differences between the estimated distributions from one sample to the other reflects a true population difference over time or simply resulting from sampling error.

For robustness and to facilitate longitudinal analysis, we have elected to use not one but all normative samples for each of the cycles of collection in the NLSCY. With 5 cycles of data, the samples each represent an expected outcome for children aged 4 and 5 over the years 1996 to 2002. Each year a new sample of cases will be added to the file, making the norm sample larger. Because the bulk of cases are now static, it is not expected to change the overall distribution significantly.

For Cycle 5[7] the norms are based on 28,214 records from Cycles 1 to 5 of the NLSCY main survey with PPVT-R raw scores. Some records with zero PPVT-R raw scores were excluded from the data used to create the norms: 1 record from Cycle 2, 49 records from Cycle 3, and 6 records from Cycle 4. These zero scores were probably incomplete tests so they are not reliable and would underestimate the true measure of ability (particularly in Cycle 3). To obtain the norms, each record was weighted by its cross-sectional weight divided by the average cross-sectional weight of records from the same cycle. The PPVT-R individuals in the norm sample were assigned standard scores so the mean of the standard scores was 100 and the standard deviation was 15 for all age groupings. This standardization was done by 2 month age groups. 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).

16.1.1 Psychometric Properties of Scores

This section addresses the quality of the test itself as it applies to the survey population, as compared to the original population for which the test was developed. We find that the test still provides a reasonable assessment of the child's ability, and we outline the reasons below.

The Raw Score
One of the main advantages of the test in a survey context is that it is tailored to the child's age and performance so that not all of the questions need to be asked to determine the ability level. Based on the age of the respondent a starting question is selected, and the test proceeds with increasingly difficult questions. When the respondent appears to be answering at random – at least six out of the last eight questions are missed - the test stops and a score is derived based on the rank of the last question and the number of incorrect answers.

Questions are ranked by increasing order of difficulty and are designed to be equally spaced on the "difficulty scale". Originally the PPVT-R test was calibrated by using a representative sample of about 5,000 English speaking children. Similar efforts were undertaken to calibrate the French version. The difficulty items were calculated using the Rasch model. In the language of Item Response Theory (IRT), this is known as the one-parameter logistic model.

Since the calibration test was done some years ago, it is natural to expect some drift of the difficulty items, as the language itself evolves and some words become more or less common. To verify whether this is the case an IRT analysis of the items was done in Cycle 4, and derived scores based on the new difficulties of the items were created. For

[7] Note that for cycles 1, 2 and 3 a different norm sample was used. While some slight variation will exist, these are well within sampling error that results from using different samples as the norm sample. Alternate norm PPVT-R scores have been calculated for these cycles using the cumulative norm sample and will be made available later in 2006.

some of the items we did find some deviation from what was to be expected in the original test.

However, no systematic deviations were found in the differences for the measured outcome. Consistently no overestimating or underestimating of the child's ability was measured in any portion of the test, and the scores derived by using the IRT were consistent with the raw scores. The correlation coefficient between the two scores from the tests was 94% for the English version, and 96% for the French version, which is high by any standard. Therefore we are confident that the raw scores can be used as they are.

For a number of children (127) the test was not completed in the field as per the directives of the instrument and no score could be assigned by the application. These cases resulted in a score of zero despite a significant number of test questions being answered for some cases. By and large most cases were dropped before achieving any information about the child's ability level. In fact, 94 cases stopped after only one plate (i.e., set of images) and a further 6 only went up to 5 plates before abandoning the test. We were able to derive approximate raw scores for the remaining 27 children. For the rest of the respondents we used the score that was produced by the application using the original rules of the PPVT-R test.

The Standardized Score

As described earlier, the standardized score is determined by using population distributions for each age. Strictly speaking we can never know the population distribution, since applying the test to the whole population is not feasible. One way to deal with this is to use the sample that we have within an age group as representative of the population in that age group, and derive the necessary percentiles.

However, certain sample limitations exist that need to be addressed before the score can be standardized. By inspecting the percentiles for different ages, we would expect an increasing trend in the ability measurement with age. While there is an overall increasing trend, for a lot of ages the trend is reversed. This is due to the fact that within each age group the sample is not large enough, and a lot of noise is introduced due to poor representation.

The better approach is to use the percentiles from the sample as a starting point, and smooth the progression with age until we are satisfied that we have a "natural" progression. We used the progression of the original percentiles from the PPVT-R handbook as an example of what degree of smoothing should be expected. Then we used the resulting points as the percentiles for standardization. We should note here that even though features of the norms were similar, the percentiles drifted upwards over the years, which, according to the experts, can be expected.

The test is usually applied to children whose effective age is 4 or 5 (note that the children's real age may include 3 and 6 year-olds if they are assessed early before their fourth birthday or after their sixth in the next calendar year).

Final Note

The PPVT-R scores used in the NLSCY are a valid measurement of ability. There is however non-response, which should be handled on a case by case basis, when doing analysis to minimize the potential for biased estimates. For more information about non-response, please see Chapters 10.0 and 13.0.

16.2 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 four and five year old children.

The assessment was developed by Dr. Robbie Case from the University of Toronto, with colleagues, including Yukari Okamoto at the University of California at Santa Barbara. The assessment is constructed based on Dr. Case's theory of central conceptual structures for explaining the development of children's thought. Before his death in May 2002, Dr. Case was adapting the test for the NLSCY. Following Dr. Case's death, Yukari Okamoto assisted the NLSCY team in completing the adaptations of the assessment for the survey.

Theoretical Background
According to Dr. Case's theory, four developmental levels can be distinguished in children's understanding of numbers: pre-dimensional, uni-dimensional, bi-dimensional, and integrated bi-dimensional. Some degree of mastery of each level is required prior to continuing on to the next. Typically the four levels are attained at the ages of 4, 6, 8, and 10. The pre-dimensional level assesses the ability to count by rote and to quantify small sets, using concrete objects. This knowledge is important for the uni-dimensional level where children deal with changes in quantity without objects than can be touched or seen. The uni-dimensional level assesses children's knowledge of the number sequence and ability to handle simple arithmetic problems. To solve the items, children must rely on a "mental counting line" in their heads. This "line" integrates their understanding of numbers and quantities. This assessment measures the essential prerequisites for successful school learning.

For the purpose of the NLSCY, we are aiming to assess children's understanding at the first three levels – also referred to as levels 0, 1, and 2. Dr. Case felt that a child between the ages of four and five was unlikely to complete questions higher then level 2 therefore the top level of difficulty was omitted from the NLSCY assessment.

Assessment Description
In consultation with Dr. Case and Dr. Okamoto the test was revised for the NLSCY. The assessment has been made continuous with three levels; some items were revised or dropped to make the test somewhat shorter. The original version of the test was discontinuous (i.e. the child had to pass sufficient items at any one level to go to the next level). Since it is accepted that we cannot expect a child to do well at a level without also doing well at the preceding level, it is sensible to stop administering the assessment after a certain number of missed items. The test was also programmed into the computer application, so that the stopping rule was automatically applied. The interviewer asked the child the question and then entered the answer. The application determines whether or not the child answered correctly.

The test is composed of 22 questions. 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[8]. Requiring children to pass two such items before they get a point increases confidence that children have the knowledge required by the item.

Children are not permitted to use a pencil and paper to answer the questions, which are given orally. Instead, the children must rely on a "mental counting line", which integrates the child's understanding of numbers and quantities. Children do have access to the various manipulative aids such as chips and a number card to help solve the problems.

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

The administration of the Number Knowledge assessment should take approximately 10 to 15 minutes.

Scoring
An "Age Equivalent Score" (DKNCdS01) is derived based on the child's responses. The "Age Equivalent Score" assigns a point for each of the three levels passed and then the points are totalled (a maximum of 1 point for each level completed can be assigned). Passing a level means passing a certain number of items from that level – for instance, for the pre -dimensional level, 3 out of 5 items must be correct. A child failing to answer any questions at the first level will get the minimum (zero), while a child who answers all the questions of all three levels correctly receives the maximum (three).

Level 1 represents the proportion of correct responses for the pre-dimensional level. There are 5 items in this level. To reach the age equivalent of this level the child must achieve a proportion of at least 0.6 (i.e., get 3 out of 5 correct responses). Level 2 represents the proportion of correct responses for the uni-dimensional level. There are 7 items in this level. To reach the age equivalent of this level the child must achieve a proportion of at least 0.6 (i.e., get 4 out of 7 correct responses). Level 3 represents the proportion of correct responses for the bi-dimensional level. There are 8 items in this level. To reach the age equivalent of this level the child must achieve a proportion of at least 0.6 (i.e., get 5 out of 8 correct responses).

Evaluation of the Assessment
Analysis was conducted on the Number Knowledge data to validate this assessment. The analyses included: comparing age equivalent scores to the child's age, comparison with the Who Am I? and an analysis of the items and of 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 Chapters 10.0 and 13.0, for more information on non-response.

16.3 Who Am I?[9]

The purpose of the Who Am I? 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 four and five year old children.

The assessment was developed by Dr. Molly de Lemos and her colleagues at the Australian Council for Educational Research (ACER). The NLSCY team worked closely with Dr. de Lemos to make some modifications to the assessment for the NLSCY (mainly dropping the drawing task) and to enhance the administration and scoring procedures for the NLSCY context.

Theoretical Background
The Who Am I? instrument assesses the developmental level of young children from the ages of 3 years to 7 years. 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 maturation, culture, experiences, and language skills.

[9] For more information about the Who Am I?, please see "Patterns of Young Children's Development: An International Comparison of Development as Assessed by *Who Am I?*" by Molly de Lemos (R-02-5E). This research paper was published by the Applied Research Branch, Human Resources Development Canada.

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 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 (EWIcdS02) is composed of a circle, cross, square, triangle and diamond which the child attempts to copy. The symbols scale (EWIcdS03) is composed of printing their name, printing some letters, numbers, words and a sentence. 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 themselves. The drawing scale is not used in the NLSCY due to time constraints. Dr. Molly de Lemos was consulted before the drawing scale was dropped for the NLSCY.

The assessment consists of an appealing booklet in which the child completes the tasks as the assessor turns the pages and gives instructions. The booklet takes about 10 minutes to complete and is scored in head office. The child completes as much as he/she is able but is encouraged to produce at least a scribble for each task.

Scoring

In addition to the two scales retained in the NLSCY, there is a combined total score (EWIcdS01). As mentioned, in the NLSCY, the drawing scale is not included and will not be discussed here. Each sub-scale (copying and symbols) is composed of four levels. The scorer uses detailed scoring instructions to determine the child's level for each task. Finally, the total of the copying and symbols score gives a general overview of the child's developmental level.

For the NLSCY, the Who Am I? assessment is hand-scored by trained individuals at Statistics Canada. These individuals have been trained to recognize signs of each level in a child's responses. Scorers who cannot make a decision on a child's level because the work does not fit clearly into one level are asked to make a judgment about the child's level based on the score on other items. Scoring was done by a small number of people and was subject to quality control procedures. The head of the scoring team also met with a staff member from ACER to consult on scoring procedures.

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 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 recommended by the author was used for dealing with cases in which the child makes no attempt. This involved assigning a score based on the score to other items. For example, if a child had a score of 4 on the square and did not attempt the diamond then a score of 3 would be applied to the diamond.

Dr. de Lemos felt that imputation was necessary to make the NLSCY data more consistent with data collected with the Who Am I? in other studies. In most cases, the Who Am I? is administered by the child's teacher or an ACER researcher trained in child development. The NLSCY uses lay interviewers who only have a short time, in the interview setting, to develop rapport with the child. This made it harder for the interviewers to convince the children to attempt the more difficult items. The imputation rules attempt to adjust the scores to better reflect the child's developmental level.

Evaluation of the assessment
Analysis was conducted on the Who Am I? data to determine whether this assessment was valid. The analyses included: comparing age equivalent scores to the child's age, comparison with the Number Knowledge, comparison with Who Am I? in other studies, an analysis of the items and of non-response. As the scoring procedures were being developed the NLSCY team consulted with Dr. de Lemos.

All 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 Chapters 10.0 and 13.0, for more information on non-response.

16.4 Mathematics Computation Exercise

The Mathematics Computation Exercise administered to the child is a shortened version of the Mathematics Computation Test of the standardized Canadian Achievement Tests, Second Edition (CAT/2). The CAT/2 is a series of tests designed to measure achievement in basic academic skills.

The CAT/2 mathematical operations test measures the student's ability to do addition, subtraction, multiplication and division operations on whole numbers, decimals, fractions, negatives and exponents. Problem solving involving percentages and the order of operations are also measured. The short version of the test developed for the purposes of the NLSCY consists of 20 questions at each level, except levels 9 and 10 which have 15 questions.

Two substantial changes occurred in Cycle 5. While no changes were made to the existing items of the tests, a desire to link the cognitive development of 16- and 17-year-old youths to the previously administered math tests for lower aged children meant a slight modification to our overall strategy. In order to bridge the two instruments without changing the tests for grades 2 to 8, we added 5 items from grades 7 and 8 to a new version of the grade 9 test (formerly the

grades 9/10 test), and we added a number of "mathematical" items from the cognitive assessment to the last math tests for youths in grade 10. (Note that the pre-existing items for the math tests for grades 9/10 are still all included). The next change in Cycle 5 has to do with the collection period. The math assessment for children from grades 2 to 10 (under the age of 16) were administered much earlier than in previous cycles. An evaluation of the impact of this change has been done on the classical scores (Canadian Test Centre scores). Our initial results show that there is a substantial effect.

This is an unfair assessment of individuals, however, as their ability is directly related to the learned curriculum. Our analysis, aimed at correcting this imbalanced measurement, was used to model the progressive cognitive growth on the classical score for estimating the mis-measurement of abilities over time. This is an estimated learning growth model based on all the children tested over time. A similar approach was used when creating the IRT scores to adjust for this effect. The final scores have been corrected for this mis-measurement to allow more reasonable comparisons of results over the different cycles. Please note that another correction was required to adjust for the overall early completion of the assessments as compared to previous cycles. So that learning gains and trajectories will be more consistent in the longitudinal context as well.

To account for the tests being performed over several months instead of a specific period, we adjusted the children's scores using a model based on multiple regressions to predict growth in abilities. Learning gains are not expected be a linear function over this period because of the Christmas break, affecting the rate at which kids learn information. However, this period is also one where very few interviews were conducted, so its effect could easily be ignored for simplification.

We could have complicated the adjustment greatly at this point by looking at the multivariate dynamics of growth in children's math abilities, but found that a simple corrective factor had a reasonable effect on the scores consistent with past test performances. While some error in prediction is expected, it is certainly well within the other errors associated with sampling and measuring (sampling error, simple response variance, etc.).

Because of how the sample is distributed in the field for collection, children in levels 2 to 6 were assessed mostly before Christmas, while older kids were spread out more evenly over the collection period. We calibrated all the scores to a standard date (May 1, 2003). This corrective measure was mostly an increase over the derived Canadian Testing Centre (CTC) score based on the raw score but it could also be a reduction for kids tested after the reference date. For the date-of-test adjustment in the CTC scores, we grouped the kids into two groups (grades 2 to 6 and grades 7 to 10). The growth rate was 0.12127 per day for the older kids and 0.21299 per day for the younger kids. For the IRT score adjustment, we used one adjustment factor for all kids (0.1214 points per day). The formula used to adjust the ability scores is given by:

Revised score = (Original Score) + (Number of days to 01/05/03) * Daily Growth Rate

You will notice that this corrected score is no longer discrete because of this adjustment. Analysts can refer to Cycle 4's table showing the **relation between the raw scores and scaled scores** if they wish to collapse the revised score into the original CTC approximate scale values.

Scoring

Each child who took the mathematics test was given a raw (gross) score, a scaled score referred to as the classical scaled score and an IRT scaled score. The raw (gross) score is obtained simply by adding the number of correct answers. The classically derived scale score and the IRT scaled score are described as follows.

The approach of the Item Response Theory was used successfully in Cycle 2 to derive scores for the reading comprehension tests. The IRT is a measurement system commonly used in

psychometric and educational testing. The IRT test scoring involves calculating either the most likely or the expected value of the ability of the examinee. The probability of a correct response to a question is assumed to be a certain logistic function of the examinee's ability. This probability is an S-shaped curve over the range of abilities. Its shape depends on the difficulty of the question, and sometimes also the discriminating power of the question (in the two-parameter IRT model) and the chance of a hypothetical no-ability examinee guessing correctly (in the three-parameter model, for multiple-choice questions).

Unlike the approach of the classical theory, the IRT makes it possible to scale the scores without preset population standards. Using common test items linking grades, standards are estimated from the entire population of children taking the test for this cycle. Scores are derived ranking each child within a level, then the scores are vertically scaled to reflect the progression of scores throughout all the levels. In order to ensure comparability from year to year, each sample from each cycle must represent equivalent populations.

The three-parameter logistical model was chosen for the math tests. The three-parameter model takes into consideration both the difficulty and the discrimination of the item and also considers the pseudo-guessing component. In this way, the IRT takes into consideration the pattern of responses. Two children with the same raw (gross) score will not have the same scaled score unless they answered exactly the same way. For example, a child who only answered the 5 easiest questions correctly would have a lower scaled score than the one who only answered the 5 hardest questions correctly.

The IRT scores were also corrected for the different collection period. A growth projection model was used to adjust the score to reflect a more standard reference period for all children and youth tested. Only one final IRT score was released; it represents the best comparative measurement within Cycle 5 and comparisons between cycles.

This score differs from the other scaled score reported for the math test as it provides a greater precision in the estimates of test performance. Unlike the other reported scale score, this score is not referenced to an external population of expected performance but is instead measured against the expected performance of the current population as estimated by all test subjects. Past rank test analysis performed using both methods of scoring showed no significant difference between the two measurements.

The classical scaled score is derived from standards (norms) established by the Canadian Test Centre in 1992. The CTC developed these standards from a sample of Canadian children from all 10 provinces (however, the test has been developed in English only and so the sample represents only the English schools), which is referred to as the normative sample. The children from the normative sample received the complete test. The scaled scores are units of a single scale with equidistant intervals that cover all of the grade levels. The scale was developed using a Thurstone procedure derived from the classical testing theory.

The fact that a short test was used for children in the NLSCY sample meant that it was not possible to directly associate the CTC scaled scores with the raw (gross) scores obtained in the survey. For this reason, the CTC normative sample was used to calculate the percentile rank for each raw (gross) score on our shortened version of the test. For example, using level 6, we find in the short test a percentile rank of 0.94% for a raw (gross) score of 1. On the complete test, the percentile ranks of 0.55% and 0.99% correspond to raw (gross) scores of 3 and 4 and to scaled scores of 315 and 319 respectively. After linear interpolation, we obtain a scaled score of 318 for the gross score of 1 on the short version of the test.

The 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 (EMACeS02) has been produced. The final scores have been corrected for this mismeasurement for cross-sectional comparisons. However, another correction was required to adjust for the overall early completion of the assessments as compared to previous cycles, so that learning gains and trajectories will be more consistent in the longitudinal context.

16.5 Cognitive Measure (16- and 17-Year-Olds)

The Cognitive Measures for the 16- and-17-year-olds were introduced in Cycle 4 when the oldest cohort reached this age milestone. 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 youths at this juncture.

Strategy and Revision
A pre-existing instrument had already been developed and tested for the Youth in Transition Survey (YITS), whose main role was to assess these abilities among other factors. 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.

Time and cost considerations are always at the forefront when developing any new components for this survey and the cognitive instrument was no exception. While we had a large enough bank of questions, seemingly covering a large breadth of abilities, it was felt that the length of the assessment would be too long. To minimize this and to reduce the response burden on the breadth of abilities being measured, a two-level assessment was introduced. Using the YITS pilot study results on the psychometric properties of the items, we constructed two ability-level tests having a substantial overlap of content that would allow us to calibrate and score both tests on the same scale. Historical information about the youths being administered the assessment were used to pre-identify the level to be administered during the actual collection.

A number of risks were minimized by using this strategy; we tried to position this option with enough leeway that an assessment was always possible even if the level had been misdiagnosed with a resulting loss of precision for the individual measurement. However, the items' psychometric properties had never been re-calibrated on this age group, and this would have an impact on the efficacy of these assessments if those properties had changed significantly. Our true evaluation of these assessments was always intended to be done on the whole population during the full application of the instrument.

The two tests, one for higher-ability youths and a slightly easier one for lower-ability youths, were assigned based on data from previous cycles. Past responses were used to pre-select the respondents into the high-ability group and the low-ability group.

Methodology for Scoring
To obtain the Cognitive Measure score, the three-parameter model from the Item Response Theory was used. For free-response items, the pseudo-chance parameter (to model guessing in multiple-choice items) was fixed at a value of zero. The Maximum Likelihood Estimation (MLE) of the three item parameters (discrimination, difficulty and pseudo-chance) followed by the *Expected A Priori* (EAP) estimation of the Cognitive Measure score was performed in an iterative process until the Cognitive Measure score converged sufficiently. To make the ability score consistent with

other NLSCY measures, each ability estimate used in the parameter estimation was weighted by the survey weight adjusted for component non-response. The statistical software SAS was used to perform these computations.

As a final step, a lower bound was placed on the Cognitive Measure score, which raised the scores of the bottom 76 respondents. Removing or lowering this bound decreased the correlation of the Cognitive Measure score with all of the math test scores from the first three cycles of the NLSCY. These 76 respondents, overly low scores appear to be an artifact of guessing or the low-stakes nature of the test.

Treatment of attempted items
Before the iterative estimation process began, raw (gross) scores for each item were calculated. Incorrect answers were scored as zero and correct answers were scored as 1, which is standard for the IRT. Partially correct answers were scored as marks received divided by the maximum possible mark; e.g., 1/2 for half-marks.

Treatment of items with no response
Items without a response ("unanswered items") can be either "omitted" items or "not-reached" items. Omitted items are those that the examinee probably saw but did not answer. In the scoring of the Cognitive Measure, an unanswered item earlier in the test than the last item attempted was considered to be an omitted item. The first item after the last item attempted was also considered to be an omitted item. In this case, the respondent probably saw the question, decided that it was too difficult, and stopped taking the test.

Usually, examinees omit items because they do not know the correct answer. Therefore, an omitted free-response item was given a zero mark. An omitted multiple-choice item was given a mark of one divided by the number of choices.

Not-reached items are those that the examinee probably did not see. These items do not provide any information about the ability of the respondent. In the Cognitive Measure scoring, all items up to and including the item after the last attempted item were considered to be reached items. Items not reached by a respondent were ignored in the estimation of that respondent's ability.

Note: The Cognitive Measure is not timed. The respondent can take as much time as needed to complete the test.

Scores (EMAydS01)
The scores presented on the data file have a mean of 0 and a standard deviation of 1. This scale is standard in IRT.

Linking the Cognitive Assessment to the Math results
We are currently developing a strategy to scale the cognitive score on the same continuum as the math scores. This would allow analysts to use the cognitive measurement (which is largely based on a mathematical construct) as an extension of the mathematical knowledge learned in previous cycles. The new value for this variable (EMAydS02) will be released as soon as the development and our evaluation have been completed.

17.0 Survey of Northern Children, 2003

17.1 Introduction

The Survey of Northern Children (SNC) was conducted by Statistics Canada in conjunction with the National Longitudinal Survey on Children and Youth (NLSCY), Cycle 5 between January and March 2003, with the cooperation and support of Social Development Canada.

The SNC is a census survey of children who were born in 1997 and who were enrolled in kindergarten in the Yukon in September 2002. The survey was conducted at the same time as Cycle 5 of the NLSCY, using the same interviewers as well as a similar survey instrument. It should be noted that the survey was renamed the Survey of Northern Children when the longitudinal panel, which had been established in 1994 at the outset of the NLSCY, was dropped. Now as a cross-sectional survey, the target children in this survey will be replaced by a new sample of 5 year old children in the next cycle.

The development plan for the SNC was based on the assumption that the ministries of education for the three territories would provide Statistics Canada with enrolment lists, including the phone number of the parent or guardian. In 2002/2003, a sample frame was made available from school boards within the Yukon only. Nunavut and the Northwest Territories were unable to provide sample frame information before the start of the collection period and, as a result were not covered in Cycle 5.

17.2 Background

National Longitudinal Survey on Children and Youth, Cycles 1, 2 and 3
The collection of information about Canadian children living in the territories began with Cycle 1 of the NLSCY in 1994 and 1995. Cycle 2 data were collected between 1996 and 1997, and Cycle 3 collection took place between the fall of 1998 and the spring of 1999.

Both the NLSCY and the National Population Health Survey (NPHS) were interested in gathering data for residents of the Yukon and Northwest Territories, and so the two surveys were combined and shortened in Cycles 1, 2 and 3 in order to keep the interview to a reasonable length. Interviews were conducted by interviewers who had been hired and trained by the Yukon and Northwest Territories Bureaus of Statistics. Households were contacted by telephone or visited in person, depending on their location, and interviewers used paper and pencil questionnaires.

Parents provided the information in all cycles of these northern surveys. A number of sections which are on the main NLSCY questionnaires, such as temperament, family functioning, and child care, were excluded from the northern survey while other, more relevant questions were added.

National Longitudinal Survey on Children and Youth, Cycle 4
In 2000/2001, data for the territories were collected by Statistics Canada interviewers for the first time, simultaneously with the collection of the NLSCY in the provinces. The sample for the SNC was based on the responding children from the previous cycle.

Data collection was done from the regional offices, and interviewers were able to use the computer-assisted telephone interviewing (CATI) application. This meant that data could be processed jointly with the NLSCY and thus, more quickly.

Summary of Cycles 1 to 4
Tracking of respondents in the North can be difficult due to the relatively high mobility of the population. In longitudinal terms, this results in serious coverage issues since children who no longer reside in the North are no longer considered part of the sample. Many respondents were

lost over the four cycles due to this type of attrition and it became evident following the completion of Cycle 4 that data quality, particularly the longitudinal aspect, was a serious issue. Data for Cycle 4 could only be made available for the Yukon and Northwest Territories. The Government of Nunavut recommended that the dataset not be released, given the very low response rates.

17.3 Establishing the Survey of Northern Children

Difficulties with collection in the North, compounded by poor response rates, lack of knowledge about the sample design and problems encountered with data processing, have all served to jeopardize the accuracy of both longitudinal and cross-sectional estimates for the first four cycles of collection.

In an effort to maintain the momentum provided by the progressive cycles of the NLSCY, while at the same time creating a database of quality information about children, data collection in the North underwent a thorough review in which a number of alternate collection options were considered. What was required was a new sampling frame that would provide current, reliable respondent information.

In 2002, discussions began with each of the territories to obtain contact information about 5 year old children from school board registration records. Our experience with the SNC in the Yukon has shown that files from the Ministry of Education work extremely well as a frame. Unfortunately, as previously mentioned, Nunavut and the Northwest Territories were unable to provide sample frame information before the start of the collection period, and so data are not available for those territories.

Objectives
The difficulties that became apparent while conducting Cycles 1 to 4 of the NLSCY in the North led to the decision to focus our efforts on one specific age group for which an up-to-date frame could be constructed, i.e., children aged 5 years. The first five years of life are critical to a child's development. The early years shape long-term outcomes related to scholastic achievement, to employment success, to health, to quality of life and to the ability to adapt.

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

Content
In general, the SNC application was similar to that of the NLSCY, Cycle 5 for the provinces with a few exceptions.

1) Only questions that applied to 5 year old children were asked of the person most knowledgeable (PMK) about the child or the PMK's spouse;

2) Since data collection was done by telephone using a CATI application, direct measures used in the NLSCY were not part of the interview for the SNC collection. Therefore, the following NLSCY tests were not included:
    - Peabody Picture Vocabulary Test - Revised
    - Number Knowledge Test
    - Who Am I?

3) Finally, in order to reduce respondent burden as well as remove questions that were not relevant, the following modules were removed from the application:
    - Neighbourhood Safety

- Mother's Work after Birth
- Medical and Biological Information
- Sleep
- Child Care
- Custody

For further details on concepts and definitions, data collection methods and data processing methods for the SNC, which were identical to those in the NLSCY, please refer to the appropriate chapters in the present Microdata User Guide. It is important to note that the income variables for the North were not imputed.

Survey Methodology
The target population of the SNC consisted of children who were born in 1997 and who were enrolled in kindergarten in the Yukon. Since this was a census survey, there was no need for a sample design, sample selection or sample allocation.

The Ministry of Education in the Yukon provided a list of 349 children who came from 343 households. Information was not available for two of these children, and another two were considered to be out-of-scope. These four children were not included in the weighting process. Of the 345 children remaining on the sample file, 301 responded to the SNC.

Weighting
The SNC is a census survey, not a probability sample, which means that the weighting is quite simple. In addition, we were required to produce only cross-sectional weights since the SNC is not longitudinal.

Cross-sectional weights were calculated for all children who responded to the SNC. Again, since this was a census survey, the initial cross-sectional weight assigned to all children was one.

Non-Response Adjustment
In order to calculate the cross-sectional weight, a simple non-response adjustment was made to adjust the weights of the respondents to account for the selected children who did not respond.

The only available information for non-respondent children is that which is provided on the enrolment lists from the school boards. Thus, it is hard to make inferences about these children. The only derived variable that proved significant is whether the child lived in Whitehorse or elsewhere. This creates two different groups. The non-response adjustment (NRA) factor is calculated as follows:

$$\text{Non - response adjustment} = \frac{\sum_{\text{Respondents + non-respondents}} 1 \text{ in the group}}{\sum_{\text{Respondents}} 1 \text{ in the group}}$$

Post-stratification
The target population for this survey is 5 year old children who were enrolled in kindergarten in the Yukon in September 2002. This is different from targeting all 5 year old children (including those not enrolled in schools) in the Yukon. For this reason, the cross-sectional weights will not be post-stratified to conform to known population totals since those totals would include all children. There are 345 children on the file and so the weights should add up to 345. This survey was a census, which means that the record for each child is counted once and has equal weight in the survey results.

For more details about the weighting process, please refer to the appropriate chapters in the present Microdata User Guide.

Data Quality
The estimates derived from this survey are based on a census of children, so therefore there is no sampling error. Errors that are not related to sampling may occur at almost every phase of a survey operation. These are non-sampling errors.

In this section some of the non-sampling errors that occurred in the survey are discussed.

Overall Response Rates
In the SNC, a respondent is defined as a child who has at least one complete adult component or child component.

Of the children who were part of the SNC Yukon sample, 87.25% responded.

Table 1: Cross-sectional Response

Territory	Number of Children in Sample	Number of Respondents	Cross-sectional Response Rate
Yukon	345	301	87.25 %

There were many reasons why some households did not respond to the survey. In some cases, the interviewers were unable to trace the household, or to make contact with a selected household during the collection period. In other cases, the household refused to participate in the survey.

Partial Response Rates
According to the definition given above, only one component needed to be completed in order for a household to be considered as a responding household. Almost all respondents completed both components.

Table 2: Component Response

Territory	Number of Respondents	Both Components Completed	Only Child Component Completed	Only Adult Component Completed
Yukon	301	299	2	0

Guidelines for Tabulation, Analysis and Release
The sample design used for the North is a census. When producing simple estimates, including the production of ordinary statistical tables, users must apply the proper population weight. Even for this survey where a census was performed, a weight exists to correct for under-coverage error caused by families who did not respond to the survey. If proper weights are not used, the estimates derived from the data file cannot be considered to be representative of the targeted population, and will not correspond to those produced by Statistics Canada. In effect, the weight assigned to each child reflects the number of children represented by a particular respondent.

Special attention should be given to estimates from this survey where a census of all children was done. While sampling error is not relevant to the estimates produced, non-sampling error can take on a more significant importance. The weights provided on the file are actually for coverage correction due to the number of non-responding cases. Non-sampling errors reflect other reasons for having an imperfect estimator. While both sample and census surveys can exhibit these types of error, a census estimator is often more vulnerable as it is expected to reflect accurate information.

18.0 Analytic Issues

This chapter provides users with an overview of the various analytic issues that should be considered when analysing National Longitudinal Survey of Children and Youth (NLSCY) data. Some of the points mentioned in this chapter have already been explored in greater detail in previous chapters. The purpose of this chapter is to highlight some of the key data analysis issues.

18.1 How a Complex Sample Design Affects Analysis

As described in Chapter 5.0 on Survey Methodology - Sample and Chapter 12.0 on Weighting and Treatment of Non-response, the children in the NLSCY sample were selected using a complex sample design in order to meet various client needs and address certain operational constraints. To make effective use of the data and produce reliable estimates, it is important to use sample weights, either longitudinal or cross-sectional, depending on the type of analysis.

18.2 Unit of Analysis

In the NLSCY, the unit of analysis is always the child. While some household data were collected, no estimates can be produced at the household level; all estimates must be at the child level. For example, the number of children living in single-parent households can be estimated but **not** the number of single-parent households.

18.3 Type of Analysis

With the NLSCY, users have their choice of different types of analysis: longitudinal analysis, cross-sectional analysis or even treating the NLSCY data as if they were from a repeated survey. Each type of analysis involves different issues.

Longitudinal Analysis
In any longitudinal analysis, the population represented when using the longitudinal weights is always the population at the time of the child's initial selection. With the survey now in its fifth cycle, the analyst must deal with the problem of sample erosion, in other words, whether or not the sample can still be considered to be representative of the longitudinal population given the cumulative effect of non-response. In addition, since the children selected in Cycle 1 have now responded five times, conditioning bias may be present. All of these problems are explained later in this chapter (Section 18.7).

In addition, two sets of longitudinal weights are available, variables EWTCW01L and EWTCWd1L. The first set comprises longitudinal children who responded in Cycle 5, whether or not they responded in all previous cycles. The second set of longitudinal weights, called "funnel" weights, computed for children introduced in Cycle 1, applies to children who have responded at every cycle.

For example, if an analyst were interested in the characteristics at Cycle 5 of children introduced in Cycle 1, but not interested in their data from Cycles 2, 3 or 4, then the first set of longitudinal weights would be most 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. 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
As mentioned before in the Microdata User Guide, cross-sectional weights for Cycle 5 are only calculated for children aged 0 to 5, namely those newly introduced in Cycles 3 to 5. For those children now aged 2 or 3 who were selected in Cycle 4, and those children

now 4 or 5 who were selected in Cycle 3, the sample at Cycle 5 was not updated to take into account immigration. Consequently, children who came to Canada after the sample was selected for their respective age are excluded from the sample. The impact of this situation is discussed in the Chapter 13.0 on Data Quality and Coverage.

Repeated Analysis
A repeated analysis is an analysis in which a single age or age group is studied at different points in time. For example, with the NLSCY, one year-olds can be studied in 1994-1995, 1996-1997, 1998-1999, 2000-2001 and 2002-2003. Depending on the domain being analysed, some of the longitudinal or cross-sectional issues already mentioned may arise. Other issues affecting analysis should also be considered. First, for a particular age group, the number of children in the sample may vary substantially from one cycle to another. This is particularly true for one year-olds and five year-olds. Also, when choosing a domain, the analyst must determine whether the samples are independent or not.

For example, consider one year-olds selected in Cycles 1 through 4. By the very nature of the sample selection process, all samples of one year-olds are independent. By contrast, a study of 11 year-olds would not yield independent sampled units: the 11 year-olds in Cycle 1 come from the original sample of children aged 0 to 11; in Cycle 2, the 11 year-olds are actually the nine year-olds from Cycle 1; the 11 year-olds in Cycle 3 are the seven year-olds selected in Cycle 1; and in Cycle 4 they are the five year-olds selected in Cycle 1. Thus, all these “samples” of 11 year-olds come from the same initial sample. In fact, many of these children may be siblings measured at different points in time, sharing similar if not identical household characteristics. Consequently, they are statistically dependent. In computing the variance, the analyst must be sure to measure the covariance between the four samples of 11 year-olds.

Note that for Cycle 5, since cross-sectional weights were calculated only for children aged 0 to 5, this type of analysis can only be done with children of these ages.

18.4 Non-response

Like any other survey, the NLSCY is subject to non-response. There are two main types of non-response: total and partial. For the detailed definitions of these concepts, see Chapter 10.0.

Total Non-response
Total non-response is the complete lack of data for a selected and eligible child due to factors such as refusal to take part in the survey or inability to trace the child. Total non-response is taken into account and corrected within the sampling weights. This process is described in detail in Chapter 12.0. The greater the amount of total non-response, the more difficult it is to adjust the weight of responding units to compensate for the non-response. Moreover, from a longitudinal standpoint, with the first set of longitudinal weights described above, a child who responded in Cycle 5 may not have responded in a previous cycle. The analyst must account for this in the data analysis. See Chapter 10.0 of this Microdata User Guide for the cross-sectional and longitudinal response rates.

Partial Non-response
All questions need not be answered at Cycle 5 for the child or youth to be considered a respondent at Cycle 5; a child can be deemed a respondent even if some questions or even entire components are not answered. For a given child, there are many sources of information. Depending on the child’s age, the information may come from one or more of the following sources: the child, a parent or a teacher. Thus, even if one of these sources failed to provide the desired information, the child may still be considered a respondent. In particular, the income variables and some Motor and Social Development questions

may be imputed. Most other missing values are set to "Not stated". The analyst must therefore determine the extent of non-response in the variables of interest before proceeding with the analysis. Refer to Chapter 10.0 for the response rates for some components of the survey, or consult the codebook for the detailed answers to all the questions.

For more detailed information on how to deal with partial non-response, a workshop has been developed. The slides of this workshop are available in the Research Data Centres. As well, the workshop can be presented by a methodologist from Statistics Canada upon request.

18.5 Data Processing

The NLSCY contains a huge amount of information. While every effort is made to ensure that the data are of high quality, every reported value cannot be checked since this would prevent the timely release of survey data. Consequently, editing focuses on key variables: some online editing is performed during collection; additional editing is performed after collection during data processing. While some errors may remain after processing, efforts are made to keep these to a minimum.

18.6 Coverage of the Cross-sectional and Longitudinal Samples

The goal of any survey is to be able to produce reliable estimates that are representative of the target population identified in the survey's planning phase. After each cycle, a quality evaluation is conducted to determine whether the sample is still representative of the population it is supposed to represent. The findings of that study are presented in Chapter 13.0.

Briefly, two factors in the NLSCY reduce the coverage of the sample over time. First, from a longitudinal standpoint, cycle non-response decreases the size of the originally selected sample. Since it has been shown that the characteristics of non-respondent units are different from those of respondent units, a significant amount of non-response by units with specific characteristics will tend to reduce the sample’s overall coverage.

For the cross-sectional sample, coverage is reduced by an additional factor: changes in the Canadian population due to immigration. The combined effect of these two factors makes it impossible to produce cross-sectional weights for children 8 to 19 years old. Even though the cross-sectional weights are available for children aged 0 to 5, top-ups were not conducted for 2 to 5 year old children to factor in immigration.

Sample sizes vary greatly from cycle to cycle for certain ages
The survey goals and therefore the sample design have changed from one cycle to another. When analysing domains with specific age groups, the analyst must keep in mind that the sample size may vary substantially from cycle to cycle. As a result, the precision of the estimates will also be different from one cycle to another. This is particularly true for one year-olds and five year-olds. For more details, see Chapter 5.0.

Sample reduction between Cycles 1 and 2
Sample reductions were made between Cycles 1 and 2. First, those children who were also in the National Population Health Survey (NPHS) were dropped. Secondly, to reduce the response burden to large households, the number of children selected within a household was limited to two. This meant that some children were removed from the longitudinal sample. Specifically, only 16,903 of the 22,831 children who responded in Cycle 1 were selected for the Cycle 2 sample. For further details, see Chapter 5.0.

18.7 Conditioning Bias

In a longitudinal survey, the same respondents are re-contacted at predetermined intervals. This can result in the behaviour and responses of longitudinal respondents being affected over time; changes that are a consequence of being revisited by the same survey.

18.8 Combining Data Over Time

When analysing small domains or subgroups of the target population, the sample size may be very small. To decrease the number of small sample units, the analyst may decide to combine data from several cycles. For example, in a study of 10 year-olds with a specific medical condition in Prince Edward Island, there may not be a large enough sample to yield precise estimates. Consequently, the analyst may decide to combine data from Cycles 1 to 4. When combining data over time, the analyst must be aware that this new domain of estimation is not a known or intuitive population but rather a set of children with specific characteristics at different points in time. The time of collection should be used to determine whether or not the cycle that the child belongs to has an effect on analysis. Other issues, such as which survey weight to use and whether or not the units in the domain are independent must also be considered.

18.9 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 and WesVar take into account the sample design in calculating the variance, using the Balanced Repeated Replication (BRR) method. The creation of BRR weights differ in theory from the creation of Bootstrap weights, but the variance estimator is the same. As a result, the NLSCY Bootstrap weights can be used with these applications.

There is a stand alone version of SUDAAN, as well as an integrated version with SAS. The latter gives the flexibility to use the SUDAAN procedures within SAS. With its nine procedures, SUDAAN can produce estimates of means, ratios and totals, independence tests, linear, log-linear and logistic regressions and survival tests. SUDAAN can read SAS and SPSS files, as well as other common types of files.

WesVar uses a "point and click" approach, which makes it easy to learn. The types of analysis are more limited than SUDAAN, but WesVar still allows the variance estimation of means, ratios and totals, independence tests, as well as linear and logistic regressions. WesVar can read SAS, SPSS and other common type of files, but the application requires an additional step to create a WesVar file before proceeding with the analysis. Bootstrap weights can be used with other applications which offer the required programming environment and the desired analytical tools. SAS and SPSS macros have been developed by the NLSCY to use the Bootstrap weights to produce variance estimates based on the sample design. For more details on variance calculation, refer to Chapter 15.0.

18.10 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 15.0 on Variance Estimation and Chapter 14.0 on Guidelines for Tabulations, Analysis and Release.

18.11 Understanding the Difference Between "Not stated", "Don’t know", "Refusal" and "Valid skip"

All questions in the NLSCY do not apply to all children. When working with NLSCY data, a question that was not intended for a particular child will have the response "Valid skip". For a question that was intended for an individual and no answer was provided, "Not stated", "Don’t know", or "Refusal" appears in the data file. When analysing particular populations, the analyst should take care to ensure that the questions of interest are applicable. In the case of examining 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.

18.12 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 15.0 on Variance Estimation.

Appendix I
Guidelines for Researchers and Analysts Using the National Longitudinal Survey of Children and Youth

Some analysts and researchers using the National Longitudinal Survey of Children and Youth (NLSCY) database have expressed a need for guidelines to help them plan their analyses and report their findings. The purpose of this document is to respond to those requests.

This document is made up of two main sections. The first section concerns the research proposal. It gives the reader recommendations on different methodological aspects to consider when submitting a research proposal using the NLSCY as a source of data. The second section concerns research papers and reports. It deals with recommendations on what to consider when writing a paper using the NLSCY data. Many elements included in the section on preparing a research proposal are also found in the section on writing a paper. These two components can be used together, or as separate documents.

Before you submit a research proposal for review:

Methodological considerations

Before undertaking any analysis using the NLSCY data, researchers and analysts should first familiarize themselves with the complexity of the NLSCY and the resulting implications for analysis. The purpose of this document is to facilitate their work by clearly identifying the key methodological issues to be considered when using NLSCY data.

This document identifies several important methodological elements to be considered when submitting a research proposal. A companion document specifies the elements to consider when submitting a paper for review. Authors are encouraged to use these documents to ensure that they have addressed the relevant elements before submitting their research proposal or their paper.

The NLSCY data can be used in many ways. The main objective of the NLSCY is to allow inferences to be made about a population, using a probabilistic sample. This document has been written with this objective in mind. When NLSCY data are used with objectives other than making inferences about the population (for example a case study), some of the elements described in this document might not apply. However, for such cases, caveats provided by the author will help to put the analytical framework into perspective for the reader.

Elements of the Analytical Framework

There are six main elements to be considered in preparing a research proposal or paper using the NLSCY. These include:

1) Data sources,
2) Factors affecting the analysis,
3) Variables,
4) Type of analysis,
5) Variance estimation, and
6) Methods of analysis.

1) Data sources

All sources of data to be used in the analysis can be specified as follows:

a Specify the main source of data to be used in the analysis.

- NLSCY
- Other (specify)

b Indicate what other sources of data, if any, will be used in the analysis and whether these data will be included as raw data or in tabular form.

c If the analysis is to be limited to a subgroup or domain, provide a description of the subgroup or domain; e.g., age groups, provinces, variables with certain characteristics, etc.

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
  - Re-weighting
  - Reported as a value
  - Ignored, analysis to be done with the respondents only.
- Analysis of characteristics of non-respondents versus respondents has to be done to identify possible biases.

c Data limitations
- Provide the sample sizes, overall and for all sub-domains, where this information is known. Sample sizes will be needed that are sufficiently large both to respect confidentiality and to give reliable estimates.
- Indicate if any other limitations are foreseen with the use of the NLSCY in your project.

3) Variables
- Provide a preliminary list of the variables in the NLSCY file to be used in the analysis.
- Indicate both predictor and outcome variables to be considered, to the extent that this is known.

Note that extensive information about variables can be learned before accessing the master files by studying questions in the questionnaires (on the Statistics Canada web site) or examining variable lists in the data files (via the Data Liberation Initiative at university libraries).

4) Type of analysis

a Indicate the kind of study planned, whether longitudinal, cross-sectional, or both. Note that if both kinds are included in the analysis, the target population may differ from one type to the other.

b Specify the kind of survey weights to be used, whether longitudinal, cross-sectional or both. Note that if estimates of both cross-sectional and longitudinal populations are to be analysed, make sure to use the appropriate weights for each analysis.

c If survey weights were not to be used, include an explanation of why not. Note that it is unlikely that the use of survey weights is irrelevant to the analysis.

5) Variance estimation

Various methods are available for estimating precision when making inferences, including the measurement of the variances and/or coefficients of variation (CV). The research proposal should include some indication of the approach to be used, if possible. Options include:

- Approximations using the CV look-up tables (available for the first 4 cycles)
- Use of the NLSCY Excel Interface with CVs for many domains of interest
- Use of the bootstrap weights with the Bootvar program, SUDAAN, or some other program that incorporates the bootstrap weights
- No estimation of variance or coefficient of variation. Note that this would imply that no statistical inferences are being made.
- Use of other software (specify software: _____). Note that very few software programs are capable of handling the complex survey design of the NLSCY when estimating the variance.

6) Methods of analysis

a Present a description of planned analytical methods.

b Describe the statistical techniques to be used to determine whether the estimates are statistically significant.

c Plan to include confidence intervals based on appropriate variance calculation in the analysis.

Before you submit a paper for review:

Methodological considerations

Before undertaking any analysis using the NLSCY data, researchers and analysts should first familiarize themselves with the complexity of the NLSCY and the resulting implications for analysis. The purpose of this document is to facilitate their work by clearly identifying the key methodological issues to be considered when using NLSCY data.

This document identifies several important methodological elements to be considered when submitting a paper for review. A companion document specifies the elements to consider when submitting a research proposal. Authors are encouraged to use these documents to ensure that they have addressed the relevant elements before submitting their research proposal or their paper.

The NLSCY data can be used in many ways. The main objective of the NLSCY is to allow inferences to be made about a population, using a probabilistic sample. This document has been written with this objective in mind. When NLSCY data are used with objectives other than making inferences about the population (for example a case study), some of the elements described in this document might not apply. However, for such cases, caveats provided by the author will help to put the analytical framework into perspective for the reader.

Elements of the Analytical Framework

There are six main elements to be considered in preparing a research proposal or paper using the NLSCY. These include:

1) Data sources,
2) Factors affecting the analysis,
3) Variables,
4) Type of analysis,
5) Variance estimation, and
6) Methods of analysis.

1) Data sources

All sources of data to be used in the analysis can be specified as follows:

a Specify the main source of data to be used in the analysis.
- NLSCY
- Other (specify)

b Indicate what other sources of data, if any, were used in the analysis and whether these data were included as raw data or in tabular form.

c If the analysis was limited to a subgroup or domain, provide a description of the subgroup or domain; e.g., age groups, provinces, variables with certain characteristics, etc.

d Specify the cycle or cycles of the NLSCY were used.

2) Factors affecting the analysis

The paper should include a description of factors that restricted or affected the analysis:

a Description of the target population
- Provide a clear definition of the target population of the NLSCY.
- If the target population differed from the NLSCY definition, include a statement about the potential impact on the analysis.
- If comparative sources were used, include a statement about how their target populations differed from the NLSCY population.

b Treatment of 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
  - Re-weighting
  - 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 sub-domains.
- Ensure that the sample sizes used in the report are sufficient both to respect confidentiality and to give reliable estimates.
- Indicate if there are any other limitations with the use of the NLSCY in your project (eg., with the variables used).

3) Variables

- Describe the variables in the NLSCY file that were used in the analysis.

4) Type of analysis

a. Indicate the kind of study, whether longitudinal, cross-sectional, or both. Note that if both kinds were included in the analysis, the target population may differ from one type to the other.

b Specify the kind of survey weights used, whether longitudinal, cross-sectional or both. If estimates for both cross-sectional and longitudinal populations were reported, ensure that the appropriate weights were used for each analysis.

c If survey weights were not used, include an explanation of why not. It is unlikely that the use of survey weights is irrelevant to the analysis.

5) Variance estimation

Describe the method of estimating precision when making inferences, including the measurement of the variances and/or coefficients of variation (CV) used:

- Approximations using the CV look-up tables (available for the Cycles 1 to 4)
- Use of the NLSCY Excel Interface with CVs for many domains of interest
- Use of the bootstrap weights with the Bootvar program, SUDAAN, or some other program that incorporates the bootstrap weights
- No estimation of variance or coefficient of variation was done. Note that this would imply that no statistical inferences can be made in the paper.
- Use of other software (specify software: _____). Note that very few software programs are capable of handling the complex survey design of the NLSCY when estimating the variance.

6) Methods of analysis

a. Present a description of all analytical methods used.

b Describe the statistical techniques used to determine whether the estimates were statistically significant.

c Include confidence intervals based on appropriate variance calculation.

Summary

A reviewer of your paper, who has access to the same data as you do, should be able to reproduce perfectly your results and reach the same conclusions, given the methodology you have used is sound and well explained in your paper.

Appendix II
Partial Non-response Analysis

Family income, environment and behaviour of young children eight years later: Findings from Cycle 5 of the National Longitudinal Survey of Children and Youth

The Daily, 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 1994/95 and 10 to 13 in Cycle 5 in 2002/03 for whom data are available for each of the predictor and outcome variables under study in the report.

In 1994/95, 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 2002/03, 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 1994/95 child anxiety, 1994/95 punitive parenting, and 2002/03 maternal depression. Respondents in the Prairie Region had a significantly lower response rate for 1994/95 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 1994/95 punitive parenting practices than those in the Atlantic Region and those in Quebec (94% compared with 99% and 98% respectively). No other differences were statistically significant for this variable. Finally, respondents in Ontario had a significantly lower response rate for

2002/03 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 1994/95 (aged 2 to 5) and 2002/03 (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 1994/95	92%	93%	92%	91%	93%	95%	93%	92%	92%	92%
Aggressive behaviour score 2002/03	77%	78%	76%	74%	78%	78%	78%	77%	75%	77%
Anxiety score 1994/95[1]	94%	94%	94%	93%	94%	**97%**	95%	93%	**92%**	94%
Anxiety score 2002/03	76%	77%	76%	75%	77%	77%	78%	77%	74%	76%
Prosocial behaviour score 1994/95	87%	86%	88%	86%	88%	91%	88%	87%	88%	86%
Prosocial behaviour score 2002/03	76%	76%	76%	75%	76%	76%	78%	77%	73%	75%
Self esteem score 2002/03	81%	81%	82%	79%	82%	83%	82%	82%	78%	84%
Family functioning score 1994/95	98%	98%	98%	97%	98%	99%	99%	98%	97%	98%
Family functioning score 2002/03	96%	96%	96%	94%	96%	98%	96%	95%	97%	96%
Maternal depression score 1994/95	98%	98%	98%	99%	98%	98%	97%	99%	97%	98%
Maternal depression score 2002/03[2]	95%	95%	95%	92%	96%	**98%**	96%	**93%**	97%	97%
Punitive parenting score 1994/95[3]	96%	96%	96%	95%	97%	**99%**	**98%**	96%	**94%**	95%
Punitive parenting score 2002/03	78%	79%	78%	74%	79%	81%	77%	80%	76%	80%
Nurturing parenting score 2002/03[4]	74%	74%	73%	**68%**	**75%**	74%	76%	73%	71%	76%
Parental monitoring score 2002/03	76%	77%	75%	70%	77%	77%	77%	76%	72%	77%

Notes

Bold print indicates statisitcally 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