Catalogue no. 11-522-XIE

**Statistics Canada International Symposium Series - Proceedings**

# Symposium 2006 : Methodological Issues in Measuring Population Health

2006

Statistics Canada Statistique Canada

Canada

Proceedings of Statistics Canada Symposium 2006
Methodological Issues in Measuring Population Health

# Combining Cycles of the Canadian Community Health Survey

Steven Thomas[1]

## Abstract

The Canadian Community Health Survey consists of two cross-sectional surveys conducted on an alternating annual cycle. Both surveys collect general health information, while the second smaller survey collects additional information on survey specific health issues. Even with the large sample sizes, users are interested in combining the cycles of the CCHS to improve the quality of the estimates, create estimates for small geographical domains, or to estimate for rare characteristics or populations. This paper will focus on some of the issues related to combining cycles of the CCHS including some possible interpretations of the combined result. Possible methods to combine cycles will also be outlined.

KEY WORDS: Cross-sectional Survey; Combining Cycles; Rare Populations.

## 1. Introduction

### 1 Overview

The ***Canadian Community Health Survey*** (CCHS) consists of two cross-sectional sample surveys. These two surveys alternate on an annual basis between a survey with a large sample size and general content and a survey with a smaller sample size and focused content with some general content. The first survey (2001, 2003, 2005…), or the .1 cycle, is designed to collect general health-related information from a sample large enough to provide information for more than 100 health regions in Canada. This requires collecting data from over 130,000 respondents. The second survey (2002, 2004…), or the .2 cycle, focuses on a specific health topic and collects data from a smaller sample of 30,000 respondents to provide information at the provincial level.

Even with such large sample sizes, there are situations where a single occurrence of the survey does not meet the needs of users. The sample size may be sufficient for publishing estimates for the targeted populations of interest but may not be large enough to have the desired statistical power to detect significant differences. In other situations, researchers are interested in pushing the limits of a single cycle of the survey to study subpopulations. This includes populations based on socio-demographic or geographic characteristics, along with studies of rare characteristics, where there simply is not enough sample.

To increase the uses of the data, users are interested in the possibility of combining the different cycles of the CCHS to estimate their parameters of interest. For the most part, data for the same characteristics are collected in all of the .1 cycles and some of the same information is collected by the .2 cycles. Thus, the possibility is there to combine data. However, there are several issues to consider and in some situations, it may not be appropriate to combine. In other situations, where it is deemed appropriate, there will be different methods for combining and the choice of the method will depend on the details of the analysis.

This paper will summarize the different issues that must be considered before attempting to combine surveys with a concentration on the issues related to the CCHS. An overview of the CCHS is given in section 2. In section 3, the possible interpretations of a combined analysis are studied and the methods for doing so are studied in section 4.

[1]Steven Thomas, Statistics Canada, 16th Floor, RH Coats Building, Tunney's Pasture, Ottawa, Ontario, Canada, K1A 0T6

Finally, the issues related to combining cycles of the CCHS are studied through examples in section 5 and a general overview of the major points to consider in combining CCHS cycles is given.

As a final note, this paper should not be considered as covering all possible analyses using a combined dataset. The methods outlined below should be considered only as guidelines for most general types of analyses and there may be exceptions to the ideas outlined below.

# 2 The Canadian Community Health Survey

The Canadian Community Health Survey was created as part of the Health Information Roadmap (CIHI, 1999). The main goal of this survey is to better address the health of the healthcare system and the health of Canadians. To do this, two cross-sectional surveys on an alternating annual cycle are conducted with the size and goal of these surveys varying from year to year.

## 2.1 The .1 Cycles

The first survey, or the .1 cycle, has the main goal of collecting general health information at the Health Region level, a sub-provincial level of geography. The target population for this survey is all persons aged 12 years or older who are living in private dwellings in the ten provinces and three territories during the collection period. Persons living on Indian Reserves or Crown lands, clientele of institutions, full-time members of the Canadian Forces and residents of certain remote regions are excluded from the survey. Over 130,000 respondents are required in each cycle to adequately estimate health characteristics for the age by sex groups of interest. To date, data for three cycles of the .1 survey have been collected and released. Cycle 1.1 collected data from September 2000 to October 2001, 2.1 collected data from January 2003 to December 2003 and cycle 3.1 collected data from Jan 2005 to December 2005.

For all .1 cycles, 3 frames have been used. These frames are an area frame from the Labour Force Survey (LFS), a telephone list frame created by combining telephone directories across the country and a Random Digit Dialling (RDD) frame, which is used in remote areas. This means that data are collected through a mix of telephone and personal interviews. In the 2.1 and 3.1 cycles of the survey, the sampling was controlled so that approximately 50% of the sample came from the area frame and 50% from the list and RDD frames. However, this was not the case in cycle 1.1 where 83% of the sample came from the area frame and 17% from the telephone list and RDD frames.

With every .1 cycle, there have been some slight changes to the sampling plan. First, there has always been the opportunity for interested parties to buy extra sample in order to better target their particular needs. Certain health regions have bought extra sample in order to obtain estimates for sub-health region boundaries. To avoid extraordinary costs, the extra sample was often selected from the telephone frame where it was less expensive to conduct interviews. Second, in every .1 cycle, there has been the opportunity to select optional content to be added to the questionnaire. This allowed the health regions, who chose this option, to obtain more information on issues particular to their area. Finally, there has been sub-sample content where, for a portion of the main sample, additional questions were asked to obtain estimates at a less detailed level of geography. These sub-samples differed from cycle to cycle. For more information on the sub-samples, one can refer to the user guide that accompanies the CCHS data.

For the most part, the questionnaire has remained unchanged between the .1 cycles. This was desired to maintain consistency and comparability between the cycles. However, this has not always been the case. Over time, certain questions have been modified to better collect the information desired. The questionnaire may be modified when it is realized that the question is not being interpreted properly or when concepts change. The data user should note that this does not always mean a change in variable name for the CCHS.

## 2.2 The .2 Cycles

The .2 cycles of the CCHS have varied quite considerably in comparison to the .1 cycles where consistency was important. The .2 cycles have been designed to focus on a particular health topic for a particular target population. There is nothing in these designs to ensure comparability with other cycles. The frames used, the populations targeted, the questionnaires given, and the general survey sampling strategy are all specific to that survey. To date, two cycles have been conducted. For these surveys, estimates were generally required at the provincial level and it required over 30,000 respondents to estimate proportions to the desired level of precision at this level of detail. Estimates were not required for the Territories.

Cycle 1.2 was collected from May 2002 to November 2002 and concentrated on mental health issues. The target population was all persons in the 10 provinces aged 15 years and older with the same exclusions as the .1 cycles. The sample was allocated for provincial estimates with the exception of Nova Scotia and Ontario, where sub-provincial estimates were required. For this cycle only the LFS area frame was used to target this population and information was collected through personal interviews.

Cycle 2.2 was collected from January 2004 to December 2004 and concentrated on nutrition. For this survey, persons of all ages were included in the target population. In each province, two different frames were required to properly target this population, because of the difficulty in targeting children with a single frame. For the most part this was a combination of the LFS area frame with cycle 2.1 respondents that had children in the household. As a special requirement for cycle 2.2, there was also a follow-up questionnaire, where a second dietary recall was collected from approximately 10,000 respondents. Data were collected through personal interviews for the first interview and telephone interviews for the follow-up interviews.

# 3 Combined Analysis

The combination of surveys has been discussed in detail by several authors including Kish (1999). In the general context, the idea is to combine surveys that collect the same information and either together or independently represent the population of interest. The end results are more precise estimates for the common characteristics. This situation arises quite often as different surveys collect the same base information. In the case of repeated surveys, the same questions are often posed to the respondents for each cycle. For these similar surveys, there are different methods to combine the information and the best method for doing so will depend on the details of the analysis. These details are very important and can not be overlooked. The researcher must have a clear picture of the population of interest and the characteristics for which estimates are required.

## 3.1 Defining the Analysis

When a researcher begins his analysis, his first step is to define the analysis in terms of the characteristic of interest and by doing so usually has a population in mind. The two general types of analysis are description and inference. The researcher is either trying to describe the characteristics of a finite population or trying to make inferences about a model or a superpopulation. A finite population is a finite group of individuals at a fixed moment in time. In studying these populations, researchers are mostly interested in descriptive statistics such as means, proportions and totals. For example, the researcher may be interested in the average body mass index for the population. In inference, the researcher is interested in modeling the relationships between characteristics. The parameters of the model are of interest rather than trying to describe some population. An example may be the relationship between smoking and cancer, in which case, the researcher is not interested in confining his research to a particular population at a particular moment in time. For more information on the idea of finite and infinite populations, it is suggested to refer to Binder and Roberts (2006).

In combining information from different surveys, the researcher will have to carefully define the analysis that will be conducted. The methodology for combining generally assumes that the statistics being combined estimate the same thing. For finite populations, there are often slight differences in the populations along with differences in the characteristics of those populations that may make this assumption difficult to assure. In the case of repeated surveys, the populations and their characteristics are different because of changes over time. If the cycles were measuring the exact same thing then repetition of the survey would be redundant. These changes may be minimal for cycles of a survey collected relatively close to one another but may be non-ignorable if the time between surveys is longer. With the intrinsic differences in what is being measured in repeated surveys, any descriptive analysis other than a time-series analysis of moving averages may not make sense, although with a careful description of the analysis, different analyses may be defendable. This is generally not an issue when making inferences, since it is the model that is of interest rather than a particular population. It is assumed that this model will apply to both surveys and any differences can be explained through the model.

# 4 Methods for Combining

There are several methods for combining information collected from different data sources and each method has its advantages and disadvantages. These methods can be broadly divided into two categories as outlined by Binder and Roberts (2006). The separate approach includes composite estimation techniques where estimates are calculated for each survey separately and then combined, while the pooling approach combines the sample data at the micro-data level and the resulting dataset is treated as if it is a sample from one population.

## 4.1 The Separate Approach

With the separate approach, characteristics of interest are estimated with each data source and a weighted average of the estimates is calculated. This composite estimation approach has theoretical attractions but can be time consuming to implement if many estimates are required (Chu, Brick, and Kalton 1999). Under this scenario, assume that two surveys independently collect information on some population parameter $\theta$. Estimates of $\theta$ can be calculated using each data source to have two separate estimates $\hat{\theta}_1$ and $\hat{\theta}_2$. A weighted average of these two estimates can then be calculated as:

$$\hat{\theta}_C = \alpha\hat{\theta}_1 + (1-\alpha)\hat{\theta}_2 .$$

With this model, if $\hat{\theta}_1$ and $\hat{\theta}_2$ are unbiased estimates of $\theta$ then $\hat{\theta}_C$ will also be unbiased for any choice of $\alpha$.

If the estimates do not measure quite the same thing, a combined estimate could be difficult to interpret. This is the inherent problem in combining surveys that cover different moments in time since these surveys are often not measuring the same thing. It is only after careful analysis to ensure that the measurements are the same that the results can be combined and interpreted this way.

This does not mean that the estimates from different surveys or different cycles of the same survey cannot be combined. However, the interpretation described above can not be taken since the assumption that each survey is estimating for the same population parameter $\theta$ can not be made. In this situation, the combined estimate using the separate approach can only be interpreted as and estimate of a weighted average of two different values. The other issue that may arise is when dependent surveys are combined. This becomes an issue since the dependence should be considered in estimating the variance of a combined analysis.

There are many choices for $\alpha$. One choice for $\alpha$ would be .5 which would result in a simple average but this choice would be inefficient. A more attractive option may be to choose $\alpha$ that minimizes the variance of $\hat{\theta}_C$. Here, assuming that the two surveys are independent, $\alpha=V_2/(V_1+V_2)$, where $V_i$ is the sample variance of $\hat{\theta}_i$. Since $V_i$ is unknown, a function of effective sample sizes $n_i*$ may be more applicable, where

$$\alpha=n_1*/(n_1*+n_2*)$$

and

$$n_i*=n_i/D_i$$

where $n_i$ is the sample size and $D_i$ is the design effect of $\hat{\theta}_i$.

With the more efficient methods of calculating $\alpha$, there is the disadvantage that $\alpha$ would have to be calculated separately for each variable included in the analysis. For some analyses, where comparability between the values calculated is important, the same $\alpha$ for all characteristics could be used but would not be as efficient as one specific for the analysis. In this case, a value of $\alpha$ based on a function of original sample sizes may be more appropriate.

## 4.2 The Pooled Approach

The pooled approach for combining consists of combining different surveys at the microdata level. Here, the different datasets are integrated in order to obtain a single dataset which can be then analysed as a single sample. This is an attractive option because of the increased sample size available for analyses. For statistics such as regression parameters, it may make more sense to use a pooled approach to calculate the parameters instead of taking an average of the parameters calculated from the different surveys.

One desirable option with the pooled approach is the rescaling of weights. Assuming that both surveys estimate the same population, an analysis of totals will overestimate the total by a factor of the number of surveys being combined (Korn and Graubard, 1998). Dividing the original sampling weights by this factor to yield new weights $w_i$* would result in an unbiased estimator for the population total but would be inefficient. A more attractive option would be to rescale the weights by a function of variances as described in 4.1. By simply applying the same calculated $\alpha$ values to the original sampling weights for each survey separately, and then combining the original datasets with these weights, the researcher would have a dataset that could be used for analyses.

In the case of linear statistics where $\hat{\theta}_i$ can be expressed as $\hat{\theta}_i = \sum_{i \in S} w_i y_i$ , the estimates would be the same for both the separate and pooled approaches. In the case of a total under the pooled approach, $\hat{Y}_p = \sum_{i \in S} w_i^* y_i$ and under the separate approach:

$$
\begin{aligned}
\hat{\theta}_C &= \alpha \hat{\theta}_1 + (1-\alpha)\hat{\theta}_2 \\
&= \alpha \sum_{i \in S1} w_i y_i + (1-\alpha) \sum_{i \in S2} w_i y_i \\
&= \sum_{i \in S1} \alpha w_i y_i + \sum_{i \in S2} (1\text{-}\alpha) w_i y_i \\
&= \sum_{i \in S1} w_i^* y_i + \sum_{i \in S2} w_i^* y_i \\
&= \sum_{i \in S} w_i^* y_i
\end{aligned}
$$

However, this is not the case for non-linear statistics, in which case the results from the separate approach and the pooled approach will be quite different unless certain criteria are ensured. In the case of a mean, the requirement would be that $\sum_{S1} w_i = \sum_{S2} w_i$ .

The pooled approach has the advantage that once the weights are calculated, the combined dataset can be used for multiple analyses. Unfortunately, the difficulty in using such an approach would be to come up with a set of weights that would be efficient for all variables of interest in all analyses. In fact, the choice of $\alpha$ that is very efficient for one variable may have the opposite effect for other variables of interest. A choice of $\alpha$ based on the average design effects has the advantage that more than one variable would be considered in creating the weights and thus be somewhat optimal for each variable. However, if many variables are of interest in a large dataset, it may be difficult in finding an optimal solution that would be appropriate. In these cases, it may be more appropriate, and just as efficient, to calculate adjusted weights that are a function of sample sizes or to use the simple average of weights.

# 5 Combining Cycles of the CCHS

As outlined above, there are many issues in combining surveys and many options for how to combine them. This section will attempt to clarify these issues in relation to the CCHS through the use of some examples. The first thing to be noted is that it may not be appropriate to combine the general content of cycles of the .2 surveys with cycles of the .1 survey or other .2 cycles. As illustrated in section 2.2, the .2 survey was designed for a specific health topic and the results may not be comparable to other cycles. The possible comparability issues are a result of the mode effects, questionnaire effects, effects from the changes in design and the time effect. The mode effect is an issue since the .2 cycles are usually conducted through personal interviews and the .1 cycles are conducted through a mix of personal and telephone interviews. Questionnaire effects exist since questions are often asked in a different way or a different order. Changes in the design may be an issue since areas that are not covered by the area frame, which are covered by telephone frames will not be included in the sample population. Finally, the time effect may cause temporal differences in the population and the characteristics of the population. For most analyses, it is likely that the characteristics measured are not the same and that the survey populations are not the same. Therefore, many of the assumptions required for combining will not be satisfied.

Another problem with combining the .2 cycles arises when attempting to combine cycle 2.2 data with cycle 2.1 data. These cycles are not independent since the cycle 2.1 sample was used as a frame for cycle 2.2. With the current methodology, the variance would not be able to be estimated because of the assumptions of independence between cycles. Note that this should not be an issue with the .1 cycles. There is the small possibility that the same clusters were visited in multiple cycles but, for the most part, the .1 cycles can be considered independent. With the above issues, it is strongly suggested that the .2 cycles not be used in any combination of CCHS data. As for the other cycles and some of the sub-sample content, there may still be issues in combining but with careful thought and verifications, a combined analysis can be conducted.

## 5.1 Finite Population Example

The first example is to combine the .1 cycles of the CCHS in order to obtain better descriptive statistics for Nunavut. This could include statistics such as rates for smoking or diabetes. With this example, the proportion of respondents deemed as being in poor health according to the health description index will be studied. With each cycle alone, there are sometimes issues where there are not enough observations to obtain estimates of high enough quality. In the case of this index, for cycle 3.1 of the survey, 23 respondents were identified as being in poor health, representing an estimated 3.1% of the population. This estimate was of poor quality with a coefficient of variation of 39.26%. In order to improve the quality of the estimates, a combined estimate may be an option given that the cycles can be considered independent. However, as stated earlier, the first steps in an analysis are to determine the population of interest and decide on the characteristics to be studied. These steps will lead to the best methodology for combining and ignoring this step may lead to misleading interpretations of the results.

The problem in this scenario is determining how to treat the fact that the researcher is combining three different snapshots of an evolving population over a period of 5 years. The fact is that each sample represents a different finite population for a different moment in time. It does not seem appropriate to assume that each sample is selected from the same finite population and therefore, a pooled approach may not be appropriate. There may be some merit in treating each cycle's finite population as a stratum from some larger population but this is not an easy concept. In this scenario, the most straight forward concept is the idea of a moving average of the estimates obtained from the three cycles. 2.2% and 1.6% of the population were estimated as being in poor health from the earlier cycles giving a moving average of 2.3%. Differences in the values are accepted since there is no assumption of equality between cycles. As for the variance, given that the samples are independent, it can be estimated as:

$$\hat{V}\left[(\hat{p}_1 + \hat{p}_2 + \hat{p}_3)/3\right] = \left[\hat{V}(\hat{p}_1) + \hat{V}(\hat{p}_2) + \hat{V}(\hat{p}_3)\right]/9$$

In this case, the combined coefficient of variation is 19.6%, which is a publishable estimate by CCHS standards. A more efficient method may be to take a weighted average. The resulting linear combination of estimates may be difficult to interpret unless it can be assumed that the estimate from each sample is an unbiased estimate of the same population parameter. In which case, the result would be an unbiased estimate of the same parameter. Statistical tests or subject matter knowledge would be required to verify this assumption.

## 5.2 Infinite Population Example

The second example is the combination of health procedure waiting-time information, which was collected in each .1 cycle as sub-sample content from close to 35,000 CCHS respondents. This module of the CCHS collects enough information that generic health procedure analyses can be conducted at the Canada level. However, when the interest is to perform analyses of more detailed medical procedures, there simply are not enough observations. As an example, only 10 individuals were identified in cycle 3.1 as having unacceptable waiting times for cardiac surgery and 8 were identified in cycle 2.1. Each sample alone may not contain enough information for analysis but, given that the samples are independent, the possibility of combining cycles in order to have a larger sample for analysis is there.

The main issue with combining the data is related to the questionnaire design of each cycle of the survey. The biggest changes have been between cycle 1.1 and 2.1 and it is felt here that the questionnaire designs are so different that the cycles may not be measuring the same relationships. Therefore, it is suggested that cycle 1.1 should not be combined with the other cycles. For more information on the changes, the researcher should consult the CCHS user guide that accompanies each cycle of the survey.

Before thinking about how to combine the data, the researcher will have to clarify the expected output. The usual idea in the combination of such datasets is to be able to better estimate the parameters of a regression model. In this case, each sample can be thought of as being independently generated through the same model. The pooled approach of combining may be preferable. By combining at the micro-data level, there will be more sample available to estimate the regression parameters. In this example, the sample size available would be almost doubled by combining the two cycles. The use of the separate approach may be possible. However, the resulting estimates will be different than the pooled approach and the idea of a linear combination of regression coefficients may be difficult to interpret.

In the estimation of a model with the pooled approach, the fact that there are samples coming from different cycles should be considered by including a factor in the model which indicates which cycle the data come from. Interactions between this variable and the key variables in the analysis should also be considered. By including this term in the model along with the interactions, unknown design issues will be considered in the model including changes in collection mode.

In such a situation, questions often arise related to the use of design weights in the estimation of the regression parameters of the model. It is recommended that the design weights be used in the estimation of the model to take into consideration the fact that the data are generated by a survey design that may have an effect on what is being measured. The weighted and unweighted estimates could be compared in this situation to see if the survey design does have an effect. Included in this idea is the use of the supplied bootstrap weights, which is suggested, in the estimation of the variances.

## 5.3 Points to consider when combining cycles of the CCHS

From the above examples, it should be clear that combining cycles of the CCHS is a possibility. The researcher must have it clear in their mind what is desired in combining the data from the different cycles and remember that each of the samples are taken from a different survey population. This can be overcome if the target population can be considered as infinite but in the finite population case, the fact that each population is different must be accounted for in the analysis. A combined result will not clearly represent any of the finite target populations from a particular cycle. If this was the desired result, small area estimation techniques or time-series adjustments may be of interest. The rest of this section will go into detail on some of the points to consider before combining cycles of CCHS data.

The first thing to consider is the questionnaire itself. Often there are changes to the questionnaire in the hope to improve the quality of the responses to the questions. Generally, the concepts are kept the same however, there have been changes and the same variable name does not necessarily mean that the same question was asked. For example, in the tobacco-use module of the CCHS, the question for cycle 1.1 of the survey was “Did you try a nicotine patch to stop smoking?” while in cycle 2.1 of the survey the question was changed to “In the past 12 months, did you try a nicotine patch to quit smoking?” and the variable name remained the same.

It is also important to ensure that the target populations of the survey are the same. The sexual behaviour module of cycle 1.1 of the CCHS was asked to individuals aged 15 to 59 years. In cycle 2.1, this was changed to individuals aged 15-49 to reduce response burden on the upper age category. As well, most of the sexually transmitted disease information, for which this question was asked, was accounted for with the 15-49 year age group.

In some analyses, it is important to ensure that the parameters of interest estimated from the samples being combined are the same. Statistical tests can be used to ensure that there are not statistical differences between the surveys on what is being combined. However, it is unlikely that any statistically significant differences be found given that the sample sizes were small enough that combining was necessary. If this is the situation, it may be preferable to rely on subject matter expertise.

With every .1 cycle of the CCHS there have been changes to the geography used to define the health region boundaries. These changes can be small, where one part of a health region becomes part of another or larger, where the health regions that are defined by the provincial government are completely changed to reflect that government's health mandate. If a combined dataset is required for sub-provincial estimates, it is important that any geographical changes be included. Since it is likely that the most recently defined health regions are of interest, it would be desirable to recode the geography of the previous cycles to this geography. This is not a simple task and Statistics Canada is looking to recode previous data files with new geography in the near future. These data sets would also aid in making any comparisons between cycles where there have been health region boundary changes.

A possible issue for some of the variables is the collection mode effect. The problem is that for some variables, the mode of collection has a serious effect on the responses that are given. For the CCHS, the possible collection modes are by phone and in person and for health related information, people respond differently by these two modes. A study by St-Pierre and Béland in 2004 showed that the most problematic characteristics were self reported height and weight, physical activity index, contact with medical doctors, and self reported unmet health care needs. To remedy this problem, it is ensured, to the highest level possible, that the mix of telephone and personal interviews remains constant. However, because of large buy-in samples, which are usually collected by phone, the mix of telephone and personal interviews can differ from cycle to cycle. It was noted in section 2.1 that the proportion of telephone interviews was quite low for cycle 1.1 making this an issue when combining cycle 1.1 with the other cycles. These differences may make the combined result difficult to interpret.

# 6 Conclusion

The purpose of this document is to show that with careful detailed analysis, it is possible to combine the cycles of the CCHS. The possibilities are there given the rich amounts of similar data that have been collected. However, it is hoped with this document that the researcher realizes that this is not a simple task to be taken blindly. The possibility to combine cycles of the CCHS will come from the fact that the researcher has properly identified the intended result and carefully analysed the data coming from the CCHS to ensure that this result can be obtained. Future work on the CCHS will be required to make sure that the data is appropriate for combining. This will become imperative as continuous collection begins with cycle 4.1 of the survey where the underlying idea is that data collected over any period of time can be combined to create estimates of the population.

# 7 Acknowledgements

The author would like to thank all of the colleagues who helped in the development of this paper. A special thanks to Georgia Roberts, Karla Fox and Michelle Simard for their much appreciated comments.

# 8 References

Binder, D., and Roberts, G. (2006), "Issues Relating to Methods for Analysis of Survey Data", 2006 *Proceedings of the session on Complex Data Structures, Joint Statistical Meetings, to be published.*

Canadian Institute for Health Information (1999), "Health Information Roadmap: Beginning the Journey". (1-895581-32-X).

Chu, A., Brick, J.M., and Kalton, G.(1999). "Weights for Combining Surveys across Time or Space". *Proceedings from the 1999 International Statistical Institute*, 103-104.

Kish, L. (1999). "Cumulating/Combining Population Surveys". *Survey Methodology*, 25, 129-138.

Korn, E. L. and Graubard, B. I. (1998). *Analysis of Health Surveys.* Wiley.

St-Pierre, M. and Béland, Y. (2004). "Mode effects in the Canadian Community Health Survey: a comparison of CAPI and CATI"*, 2004 Proceedings of the American Statistical Association Meeting*, *Survey Research Methods*. Toronto, Canada: American Statistical Association.