2006 Census Internet Mode Effect Study

Chantal Grondin[1], Limei Sun[1]
[1]Statistics Canada, 100 Tunney's Pasture Driveway, RHC-15, Ottawa, ON, Canada, K1A 0T6

Abstract
For the first time in 2006, Canadian households had the option of responding to the Census via the Internet. Almost one in five households chose to report their data with this new collection mode. A study was undertaken to check for the presence of an Internet mode effect in the data. Different means were used to assess the presence of a mode effect. However, since the decision to use the Internet is not random—in fact, characteristics of Internet reporters and Paper reporters are different—comparing both groups could only be done after adjusting as much as possible for the differences in characteristics. This was done using the Propensity Score method. The findings of the 2006 Census Internet Mode Effect Study are given in the paper.

Key Words: Mode effect, Propensity score, Internet collection, Census

1. 2006 Canadian Census of Population

1.1 Overview of the Census
Approximately 13 million dwellings and a population of more than 31 million were enumerated in the 2006 Census. The population was counted at their usual place of residence in Canada regardless of their location on Census Day. Self-enumeration was used to count most of the population, while the remainder was counted via interviews or the use of collective dwelling administrative records. Two main types of questionnaires in both official languages were used to collect the majority of the Census data: Form 2A, the short questionnaire with eight person questions, was distributed to four out of every five households (80%); Form 2B, the long questionnaire, with 53 person questions (including the eight questions from the short form) and 8 dwelling questions, was distributed to the remaining 20% of households.

Offering the Internet option in the 2006 Census is consistent with Canadian Government Policy, as more and more government services are now accessible online. The Internet option also turned out to be a solution to a number of important issues with the Census: 1) There have been privacy concerns on the part of respondents regarding local enumerators, 2) recruitment of a large temporary workforce for data collection was always a challenge for the Census, and 3) resources needed for data entry were very important before 2006 but could slightly be reduced because of the Internet option in 2006. Overall, the quality of the 2006 Census data has increased in certain regards, especially due to lower item non response and invalid rates, and to lower rates of manual coding. Moreover, we have seen an important difference in the follow-up rates for Internet reporters compared to paper reporters, the former being much lower than the latter.

1.2 Features of the Internet questionnaire
In order to minimize the mode effect and to facilitate the integration of data received from different response modes, the electronic versions of the questionnaires corresponded as closely as possible to the paper versions in terms of question wording, instructions and presentation of response choices. Determined efforts were made to adhere to the paper form while incorporating many Internet questionnaire standards and conforming, as much as possible, to the guidelines for presentation of federal government Internet sites. The **short and long online questionnaires** used an interactive multi-page design. With this design, the questionnaire was presented screen-by-screen, each of them displaying a question or group of questions related to a common theme (questions on labour for example).

On both paper and Internet questionnaires, respondents were asked to report their date of birth. An additional screen was added in the Internet version to **confirm the age** of respondents. Confirming the age was particularly important in the long questionnaire as persons younger than 15 were neither presented with the same questions nor subjected to the

same online edits as those 15 or over. In fact, no edits were performed for questions related to marital/common-law status, and no questions were asked related to mobility, education, labour market and income for people younger than 15. The Internet questionnaire also had an **additional screen to confirm income** for people 15 years old or over. This additional screen was important because the income part of the questionnaire was very complex, with thirteen sub-questions. This confirmation screen provided respondents the opportunity to verify their responses and make corrections if needed.

Internet standards were followed, including **check boxes** when multiple responses were possible and **radio buttons** when only one response was allowed. Other electronic features were used, such as the **drop-down menus** for selecting the day and month of birth or for selecting provinces/territories, as well as **automated skip patterns** to reduce response burden by ensuring respondents were not presented with irrelevant questions.

1.3 Online edits in the Internet application

When implementing online edits in an application, one has to keep in mind that complex edits can increase the burden on respondents to provide precise responses and therefore increase the time required to complete the questionnaire. The potential benefits in data quality obtained from complex edits need to be weighed against the increase in respondent burden to achieve an appropriate balance. For the 2006 Census, it was decided to keep the online edits relatively simple; they were performed on one question at a time and no consistency edits were performed between questions. The decision to go this route was based in part on the results of usability tests and qualitative studies done with different versions of the Internet application. These tests were helpful in identifying which edits seemed more appropriate and in line with our goal to stay as close as possible to the paper questionnaire.

In the end, the Internet application included four types of **online edits** or **validation messages**. **Non-response** messages appeared when respondents had not answered a question. **Partial response** messages appeared when respondents provided only a partial response to a question, for example, if they omitted the city name from their address. **Invalid response** messages appeared for numerical responses when respondents entered numbers outside the range established for a question. Finally, **amount verification** messages appeared for questions related to money amounts, when the response appeared unusual. This type of message asked respondents to verify that they have entered the correct amount, for example, "*Please verify the amount you entered for part (f), if correct leave as is*". All these messages followed the same approach. When respondents clicked the *Next* button, the information on the current page was validated and, if necessary, the application displayed the same screen again, noting any problems at the top of the page in red text, for example, "*Please answer Question 5 for John Doe*." The question and field requiring attention appeared in red, and a red arrow highlighted the missing response to assist the respondent, who could then either fill in the missing information or continue to the next screen. If the respondent chose to move on without making any changes, the next screen was presented. If the respondent added or changed any information, the responses were validated again. This approach was consistent with the Common Look and Feel guidelines prescribed for Canadian government web sites in that pop-up windows should not be used within pages to convey information to respondents.

2. Internet Mode Effect Study: Methods and Results

A study was done to compare data reported on Internet to data reported on paper to examine if there seemed to be a mode effect. The definition of mode effect used in this study is any sign that the data reported on Internet would have been different had it been reported on paper, whether it be due to the respondent or to the Internet application as such. It is important to study mode effects because when changes are observed between Censuses, one needs to be able to distinguish between real changes and changes due to the introduction of a new collection mode.

The Internet Mode Effect Study was done using data from private dwellings who received the long form (although some results are given for short forms). The main interest was to compare Internet versus paper questionnaires. The variables used correspond as much as possible to what was received from collection, although editing was sometimes done to some variables. We used variables before any imputation was done to them, as we wanted to study the data as it was reported. The results are always unweighted since the goal was to compare two groups of reporters (Internet versus paper), not to do inference on the entire population.

The results of this study are presented in four parts. First, we look at questionnaire **follow-up and non response rates** by question. Second, we examine some unexpected effects of validation messages. Third, we look at some other

situations that resulted from the use of the Internet application. In the most important and last part of this study, we compare distribution of answers for Internet reporters and paper reporters, and identify the largest differences. Because part of these differences in answers are due to the difference in characteristics of people using each mode, we then use the **Propensity Score method** to account for that, and redo the comparisons of answers after adjusting the data. We also discuss the limitations of this method and give some results.

2.1 Follow-up and item non response rates

2.1.1 Follow-up rates for edit failure

The first notable impact when comparing the questionnaires received from Internet versus paper questionnaires was the difference in follow-up rates. After data capture, questionnaires were transmitted to data processing where completion edits were performed. These edits were based on a score strategy. The questionnaires for households exceeding a predefined threshold value were forwarded to follow-up. As can be seen in the following table, Internet questionnaires had much lower follow-up rates than paper questionnaires, especially for the long form.

Table 1: 2006 Census Follow-up Rates, By Response Mode and Type of Questionnaire

Response mode	Follow-up rate (%) by type of questionnaire	
	Short Forms	Long Forms
Internet	2.5	5.7
Mail	5.6	39.1

Source: Statistics Canada, 2006 Census

The lower Internet follow-up rates are a direct result of the overall lower number of partial or invalid answers, as well as lower non-response to each question, mainly due to the presence of validation messages. Another factor that could have had an impact on follow-up rates may be linked to characteristics of people who chose to report their data on the Internet. Indeed, Internet reporters are more likely to have a higher education level, and this might be associated with a better capacity to respond to the questionnaire.

2.1.2 Non response rates

The following graph shows average non response rates for groups of questions related to a same theme, by response mode.

Figure 2: Average Non Response Rates for Groups of Questions, by Response Mode

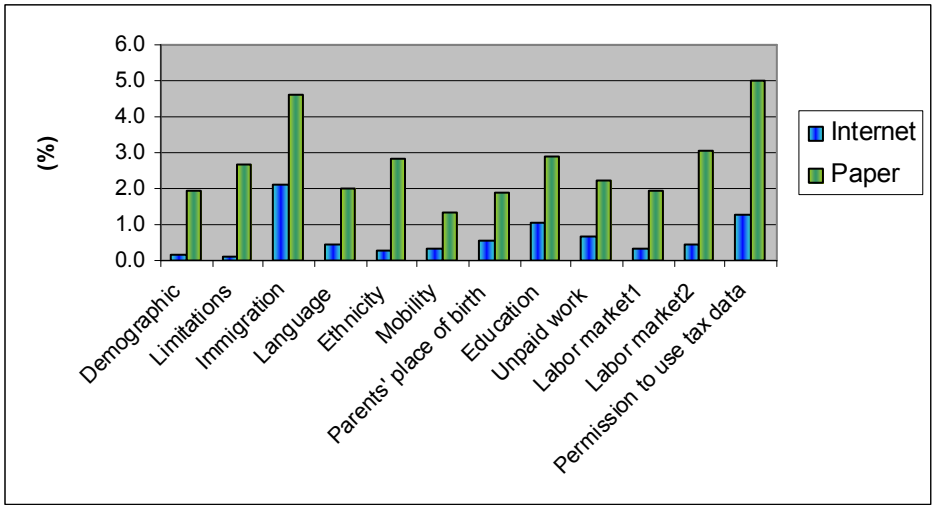

Source: Statistics Canada, 2006 Census

We can see that non response rates for Internet questionnaires are systematically lower than those for paper questionnaires and that is true for every single question without exception. This difference is primarily due to the online

edits or validation messages, and also to a lesser extent to the online help and explanations of why each question was asked.

It is clear that validation messages have had a positive effect in reducing non response (or invalid responses). However we also know from qualitative studies that some Internet reporters, when faced with validation messages, feel that they absolutely must provide an answer in order to be able to continue. This can in fact become problematic in certain cases, especially when respondents don't know the answer to a question. We will examine some examples in the next sub section.

2.2 Unexpected effects of validation messages

The Census questionnaire has two questions on mobility, one asking where the respondent lived 1 year ago, the other asking the same question about 5 years ago. The difficulty with these questions often comes from recalling the postal code (or zip code) of a previous address or from having to respond on behalf of someone else in the household. We looked at the answers of every respondent who said they lived in a different city 1 year or 5 years ago (migrants), and compared the reported postal code to the one of the current address. In theory, they should be different (with a few exceptions in some rural areas).

We see from Table 2 that the proportion of migrants who reported their current postal code as their previous one is higher among Internet reporters than among paper reporters. This is especially true for the 5-year mobility question. Although this phenomenon is not only related to the use of Internet, it is possible that validation messages which are prompting for answers when none are provided had the effect of increasing this behaviour for Internet reporters.

Table 2: Proportion of Migrants Who Reported the Same Postal Code Before and After Migration

Response mode	**1-year mobility**	**5-year mobility**
Internet	4.3%	13.2%
Paper	3.6%	9.5%

Source: Statistics Canada, 2006 Census

To further study the effect of validation messages, we also looked at the reported day of birth of Internet reporters alone. We were able to split Internet reporters into two groups: those whose reported day of birth was valid from the start, and those whose response was first missing (and hence received a validation message) before entering a valid response. For these two groups, we looked at the distribution of the day of birth as is shown in Table 3.

Table 3: Distribution of Day of Birth for Internet Reporters Who Gave a Valid Initial Answer Versus Those Who Received a Validation Message

Day of Birth	**Initial answer was valid**	**Received a validation message before entering a valid answer**
1st of the month	3.6%	7.3%
2nd to 28th of the month	From 3.2% to 3.6%	From 2.5% to 3.8%

Source: Statistics Canada, 2006 Census

As we can see, for the group of Internet reporters who gave an valid initial answer, the frequency of day of birth "1" is very similar to any other day (from the 2nd to the 28th) in the month. However, for Internet reporters who received a validation message, implying that they first had left the answer blank, the frequency of day of birth "1" is much larger compared to any other day in the month. This seems to support the hypothesis that when respondents do not know the answer to a question and get a validation message, they feel compelled to report any valid answer, in this case, to select day of birth "1" which was the first choice in the drop-down menu.

2.3 Some Technology-Related Effects

As part of this study, we were able to link the 2006 Census database to the 2001 one. This means that for a given respondent, we had access to the data reported to both Censuses. We used this information to study concordance rates between answers reported in 2001 and in 2006 for questions where no change was expected in time. This allowed us to identify a problem with the use of drop-down menus for some questions. More particularly, the place of birth question made use of a drop-down menu from which Canadian-born could select their province/territory of birth. Response error

can sometimes occur with these types of menus. Indeed, once a selection is made in a drop-down menu, one must click outside of the menu (or use the tab key) to move the browser's focus away from the drop down list. If the scroll wheel on the mouse is used while the drop-down menu is still "selected"—for example in an attempt to scroll down to the next question after responding—the response will be changed, sometimes without the respondent noticing. The result in such situations is that for this small fraction of respondents, their answer tended to be moved to responses lower (often the last) on the list. By using the linkage file, we could see that the concordance rates between the 2001 and 2006 answers for the place of birth were lower for Internet reporters compared to paper reporters. Understanding that drop-down menus could have generated response errors, we also looked at the responses for other questions using such menus. It turned out that the province/territory of highest degree or diploma was also affected. Edit and imputation strategies were put in place to detect and remedy the problem. In other situations where a drop-down menu was used but the next question was immediately visible without the need for scrolling, we found no evidence of this kind of problem.

Another technology-related effect was observed with the question on number of hours worked. Although the number of responses affected by this was quite small, we noticed that Internet reporters were much more likely to have reported 1 hour worked compared to paper reporters. We know from our experience with the Census questionnaire that paper reporters occasionally enter a "0" in the text box for number of hours worked instead of checking the "None" box. We believe that the same has happened with the Internet reporters, but for them, a validation message for the text box said that valid values had to be between 1 and 168. It seems likely that some of the "1 hour" answers from Internet reporters were reported for people who did not work, but were not clear on how to indicate this.

2.4 Comparing Internet and Paper answers distribution

2.4.1 Unadjusted comparisons

We produced frequency distributions for every 2006 Census question by response mode (Internet versus paper). Statistically speaking, the distributions of answers of Internet reporters and paper reporters were always different. With such a large sample of observations, any small difference is likely to be considered statistically different. Also, because the non response rates were on average 2 percentage points lower for Internet compared to paper, it affected the answer distributions. But more importantly, the characteristics of people reporting data using each mode are different.

We identified the largest differences between the Internet and paper answers distribution (after excluding non respondents). We could see that a lot of the largest observed differences were related to immigration or ethnicity, or to language. It is not surprising because we know that immigrants are more likely to be using Internet to respond to the Census compared to Canadian born, and that people who speak French are less likely to use Internet than those who speak English. So to be able to take into account the difference in characteristics of Internet and paper reporters when comparing answers from these two modes, we used a method called the Propensity Score Method. This method allows removing a large part of the bias due to the difference in characteristics of the two groups from these comparisons. This method and its limitations are examined in the next sub section.

2.4.2 Propensity Score Method and adjustment

In randomized experiments, two treatment groups can often be compared to one another because their units are likely to be similar. However in our case, trying to compare Internet reporters' answers to paper reporters' answers is potentially misleading because the units who chose to use the Internet mode differ from the ones who chose to respond on a paper questionnaire. We had to find a way to group units from both response modes in such a way as to make the comparison more meaningful. To this effect, balancing scores are found to be helpful because they are a function of observed covariates such that their conditional distribution is the same for units who received both treatments (or who used each response mode in our case). One such balancing function is called the propensity score, that is, the propensity towards exposure to a treatment (Internet) given a set of observed covariates or characteristics.

With the propensity score method, each respondent is attributed a propensity score, which represents the conditional probability of assignment to a particular treatment (in our case the probability of using Internet) given a vector of covariates. An estimate of the propensity score was obtained using a logistic regression model. The variables that are related to the chances of using Internet were entered in the model as explanatory variables. In our model, the variables used to explain the probability of using Internet were: geography, urban/rural indicator, census family type, household size, number of children in certain age groups, language spoken at home, knowledge of non official languages, age, age

of Person 1 in the household[1], visible minority status, sex, highest level of education, citizenship and family income categories.

Once we had a predicted probability of using Internet for all respondents, we divided them into 10 sub groups[2] of equal size based on the similarities of these probabilities. People within each sub group having a similar probability, they are hence more comparable. Comparisons of answers could be done within each sub group to see if there seemed to be a mode effect. However we chose to standardize the distribution of Internet respondents within each sub group to make it match the distribution of paper reporters (considered here to be the reference group) within these same sub groups. This way, we could compare the standardized answers for the entire Internet group to the ones of the paper group. This standardization was done simply by computing the proportion of paper reporters within the 10 sub groups as well as the proportion of Internet reporters. Dividing the first by the second, we get the adjustment factor which must be used to standardize the Internet distribution to make it comparable to the paper distribution.

2.4.2.a Limitations

As in any study, there are limitations that we have to be aware of. In the case of the Census, there are at least three. The first limitation is a **proxy effect**. It has to do with the fact that the Census data are collected in clusters (households), and often, there is only one person responding for the entire cluster. The decision to use Internet to report the data may have a lot to do with the characteristics of the person actually reporting the data, but not necessarily with other people in the household. Note that in the Census, we don't know who reported the data for whom. So although we say we are modeling the probability of using Internet, we are really modeling the probability of having data reported on Internet. Other studies have shown that Internet users are usually younger and more educated than non Internet users. But what does this mean in the Census if a young educated person responds for a household in which there are other older and less educated people? The data associated to these other people will also appear as being reported by Internet but their characteristics are not in line with what is expected from Internet users. Consequently, this situation brought a lot of variability into the modeling of characteristics of Internet reporters and resulted in a poor fit of the model. To overcome this limitation, we tried as much as possible to use household level (or higher level) variables in the model, although we also had to use some person level variables.

The second and non negligible limitation with this study could be called the **chicken and egg dilemma**. It comes from the fact that the only data available to model the probability of using Internet is data that was actually collected using each mode, therefore data which could potentially be tinted by a mode effect. For example, we know that there is a mode effect with the number of hours worked, Internet reporters being more likely to have entered "1 hour" compared to paper reporters. The number of hours could then show up as a characteristic that has an impact on the probability of using Internet. But this characteristic is not "causing" people to use the Internet, it is actually the other way around. The higher frequency of "1 hour" answers is a mode effect. If this characteristic was to be used to model the probability of using Internet, then most likely we would lose track of this mode effect when comparing the adjusted distribution of answers, that is to the distribution after the propensity score adjustment is done. The consequence of this second limitation is that we must be very careful in the variables we keep in our model. A variable that shows a significant impact on the probability of using Internet can either be a true characteristic of Internet users, or a mode effect. The choice of variables is then influenced by what we know (from other studies) are related to the chances of using Internet, plus what common sense tells us. Once a variable is entered, it is under the assumption that there is no mode effect related to it. When comes the time to compare the distribution of answers for Internet reporters and paper reporters after the Propensity Score adjustment, we will only be able to compare variables that were not already used in the model.

The third limitation has to do with the very large sample size of the Census database. When working with very large samples and trying to compare estimates, every small difference will show up as being statistically significant. So instead of trying to identify all differences between the Internet and paper groups, we put our efforts into identifying the largest ones. We hence started by assuming that there was no mode effect. We identified the largest unadjusted differences between the two modes (before using the Propensity Score method) and then, compared them with the adjusted differences using the Propensity Score method. The results are presented in the next sub section.

1. The age of person 1 was taken as a proxy for the age of the person actually responding to the questionnaire. We don't know that it is necessarily person 1, but we believe it is in a lot of cases.
2. Research done by Cochran (1968) has shown that 5 sub groups are often sufficient to remove 90% of the bias

2.4.2.b Results

Table 4 presents variables and their specific response categories that showed the largest (unadjusted) differences between Internet and paper proportions. The way this table should be interpreted is the following. Let us look at the third line of the table. We see that the proportion of Internet reporters who fall in the age category "60 years or over" is 11.1%, while the corresponding proportion for paper reporters is 28.7%, hence an unadjusted difference of 17.7 percentage points. After the Propensity Score adjustment, where characteristics of each group of reporters are taken into account, this difference (which we call the adjusted difference) is reduced to 0.3 percentage points. This indicates that a large part of the unadjusted difference was due to a difference in characteristics between Internet and paper reporters. Once these characteristics were taken into account with the Propensity Score method, the new adjusted difference became very small.

The results from Table 4 are interesting in that they generally show an important reduction of the differences between Internet and paper proportions after the Propensity Score adjustment. However one exception stands out for those who reported $0 as their income tax, as can be seen on the first line of the table. The adjusted difference is almost unchanged. This is clearly a mode effect. This question is the last one among the income questions, and is the only one which doesn't have "yes" and "no" checkboxes. It is easy to think that Internet reporters who did not know the amount they paid in income tax simply entered a zero instead of leaving the field blank to avoid getting a validation message. On paper, it was easy just to leave the field blank.

Table 4: Unadjusted and Adjusted Largest Differences in Proportion Between Internet and Paper Reporters

THEME	VARIABLE	Response	INTERNET proportion	PAPER proportion	INTERNET - PAPER	
					Unadjusted	Adjusted
Income	Income tax amount	0	54.6	34.7	19.9	19.8
	Income from pension plan	No	95.1	83.9	11.2	2.6
Demographics	Age group	60+	11.1	28.7	-17.7	-0.3
Immigration and ethnicity	Year of immigration	1950-1969	12.1	29.6	-17.5	-2.8
		2000-2006	25.4	13.1	12.4	1.7
	Visible minority	White	76.3	87.6	-11.3	-1.5
Mobility	Residence 5 years ago	Same Address	51.3	67.7	-16.4	-5.0
		Same City	26.1	16.9	9.2	3.6
Labour market	When last worked	<2005	41.2	55.3	-14.2	2.1
	Hours worked	None	30.9	43.6	-12.7	-2.3
		36-45	35.1	27.2	7.9	0.8
Unpaid work	Number of hours for unpaid child care	None	54.7	65.6	-10.9	-3.4
Limitations	Activity limitation	No	89.9	83.2	6.7	0.2
Education	University diploma	No	69.4	78.5	-9.1	-1.4
	College diploma	No	67.5	73.9	-6.4	-2.8
	High school diploma	Yes	74.1	65.7	8.3	0.5
Language	Mother tongue	NOL[1]	24.5	16.7	7.8	1.3
		French	18.1	25.3	-7.3	0.8
	Home language	French	17.1	24.3	-7.2	0.8
	Work language	French	16.3	23.0	-6.7	-1.2
	Speaks a non official language	No NOL	70.3	79.2	-8.9	-1.1

1. Non official language.

Source: Statistics Canada, 2006 Census

We also see from Table 4 that the adjusted differences for mobility questions are still relatively high. It is possible that what we see here is related to the postal code problem that we have seen earlier (Table 2). The question on unpaid work for child care also shows a relatively high adjusted difference. There is no explanation for now as to why this is like that. More research might help shed light on this situation.

3. CONCLUSION

Allowing Canadians to use Internet to respond to the Census is the way of the future. This first experience has proven that there are multiple benefits to having this new response mode. A lot of savings were done in terms of questionnaire follow-up. With regards to validation messages, the results indicate that in general, they are effective in obtaining answers to questions that respondents might otherwise have overlooked and in having respondents correct errors they inadvertently committed. Along with the automated skips of non-applicable questions, these messages result in a general perception among respondents that the electronic questionnaire is "intelligent". This responds in part to the high expectations the general public has with regard to electronic questionnaires.

We have also seen that prompting for responses for questions that are more difficult to answer could have amplified certain unwanted behaviours, like reporting a valid but incorrect answer when the true response is not known. As well, some technology-related problems, like the use of drop-down menus and the choice of valid range of answers for certain questions have been identified in this study. These technical problems have already been taken care of in the 2011 Census Internet application. Nonetheless, more research should be conducted to better understand how respondents react to such validation messages. If Internet is more and more used in the future, it is essential that we understand how people react to the application and its different features.

As for the Propensity Score method, although there were important limitations to its use, it gave interesting results. A lot of the initially large observed differences were considerably reduced by adjusting the Internet distribution to the paper distribution by sub groups based on their propensity score. Also, some differences which were almost unchanged by the adjustment were clear indication of mode effects.

We have identified very few mode effects in the 2006 Census. For those few that were found, some required extra processing steps in order to reduce their impact, others affected very small number of observations, and some had no impact at all. Overall, we can say that using Internet in the 2006 Canadian Census was a major enhancement.

As for keeping the Internet questionnaire as close as possible to the paper version, it is not clear whether this was the best option. There is no consensus among researchers on the best methods to design questionnaires in a mixed mode survey. One school of thought is to use the same questionnaire format in each mode, while the second school of thought is to take advantage of the full potential of each mode (Dillman; De Leeuw, 2005; Yiptong, 2007). More empirical studies should be undertaken to better understand the impact of each option.

References

De Leeuw Edith D. "To mix or not to mix data collection modes in surveys", Journal of Official Statistics, Volume 21, number 2, pp 233-255, 2005.

Dillman Don A. and Bowker Dennis K. The Web questionnaire challenge to methodologists. Dimensions of Internet Science, edited by Ulf-Dietrich Reips and Michael Bosnjak. Pabst Science Publishers.

Laroche D. and Grondin C. "Impact of Online Edits and Internet Features in the 2006 Canadian Census", United Nations Statistical Commission and Economic Commission for Europe, Work Session on Statistical Data Editing, April 2008.

Laroche D. "2006 Census Internet Data Collection", Technical Report presented at Statistics Canada's Advisory Committee on Statistical Methods, April 30th-May 1, 2007.

Rosenbaum P.R. and Rubin D. B. "Reducing Bias in Observational Studies Using Subclassification on the Propensity Score", Journal of the American Statistical Association, Applications Section, September 1984, Volume 79, Number 387.

Roy L. and Laroche D. "The Internet Response Method: Impact on the Canadian Census of Population Data", Proceedings of the Survey Research Methods Section, American Statistical Association, 2006.

Yiptong J. "Best Practices in Mixed-Mode Surveys, Summary of Literature Research and a Proposal for Future Work", Internal document, Statistics Canada, May 2007.