



BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2009-10

**Computing the Accuracy of Complex
Non-Random Sampling Methods:
The Case of the Bank of Canada's
*Business Outlook Survey***

by Daniel de Munnik, David Dupuis, and Mark Illing

Bank of Canada Working Paper 2009-10

March 2009

**Computing the Accuracy of Complex
Non-Random Sampling Methods:
The Case of the Bank of Canada's
*Business Outlook Survey***

by

Daniel de Munnik, David Dupuis, and Mark Illing

Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
dmunnik@bankofcanada.ca

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

Acknowledgements

We would like to thank Tara Ainsworth for her assistance in compiling historical data on the Business Outlook Survey. In addition, we thank Brigid Brady, Paul Fenton, Sharon Kozicki, Lise Pichette, Lori Rennison, Raphael Solomon, Greg Tkacz, Michael Yake and the other members of the Regional Analysis Division and participants at the Bank of Canada seminar, CEA conference session (Vancouver - June 2008), and the Computational Economics Society Meeting (Paris - July 2008) for providing helpful comments. The authors take full responsibility for any errors and/or omissions.

Abstract

A number of central banks publish their own business conditions surveys based on complex non-random sampling methods. The results of these surveys influence monetary policy decisions and thus affect expectations in financial markets. To date, however, no one has computed the accuracy of these surveys because their respective non-random sampling method renders this assessment non-trivial. This paper describes a methodology for modeling complex non-random sampling behaviour, and computing relevant measures of statistical confidence, based on a given survey's historical selection practice. We apply this framework to the Bank of Canada's Business Outlook Survey by describing the sampling method in terms of rules-based criteria, historical practices, and Bayesian probabilities. This allows us to replicate the firm selection process using Monte Carlo simulations on a comprehensive micro-dataset of Canadian firms. We find, under certain assumptions, no evidence that the Bank's firm selection process results in biased estimates and/or wider confidence intervals.

JEL classification: C42, C81, C90

Bank classification: Econometric and statistical methods; Central bank research; Regional economic developments

Résumé

Nombre de banques centrales publient leur propre enquête de conjoncture, qu'elles réalisent auprès d'entreprises sélectionnées de façon non aléatoire. Les résultats de leurs coups de sonde influencent les décisions de politique monétaire et, de ce fait, les attentes des marchés financiers. Jusqu'à présent, personne n'a mesuré la précision statistique de ces enquêtes, car il est difficile de l'évaluer en raison du mode d'échantillonnage non aléatoire. Les auteurs décrivent une méthodologie qui permet de modéliser des processus complexes d'échantillonnage non aléatoire et de calculer des indicateurs de confiance statistique pertinents sur la base du mode de sélection des entreprises utilisé pour une enquête donnée. Ils appliquent leur cadre méthodologique à l'enquête de la Banque du Canada sur les perspectives des entreprises en se fondant sur les pratiques suivies par le passé et les probabilités bayésiennes pour caractériser le mode d'échantillonnage. Les auteurs peuvent ainsi reproduire le processus de sélection des entreprises à l'aide d'une simulation de Monte-Carlo menée sur un riche ensemble de microdonnées constitué de firmes canadiennes. Sous certaines conditions, constatent-ils, le processus d'échantillonnage retenu par la Banque n'entraîne ni estimations biaisées ni intervalles de confiance élargis.

Classification JEL : C42, C81, C90

Classification de la Banque : Méthodes économétriques et statistiques; Recherches menées par les banques centrales; Évolution économique régionale

(1) Introduction

Numerous central banks rely on in-house non-random surveys to assess economic conditions through business sentiment. The results from these surveys, including among others the Federal Reserve's *Beige Book*, the Bank of England's *Agents' Scores*, and the Bank of Canada's *Business Outlook Survey*, are valued and timely inputs into the monetary policy decision-making process,¹ and as such are closely monitored by financial market participants. In Canada, since its inception over a decade ago, the Business Outlook Survey (BOS) has become an important piece of information in the Bank of Canada's monetary policy decision-making process. Its influence on economic and financial markets has also gained significance since the document's first publication in the spring of 2004, as can be judged in part by the ongoing press coverage. Given the survey's small sample size of 100 firms and its non-random quota sampling² approach, the accuracy of the results is of particular interest. Although some effort has been deployed to investigate the information content of the BOS and other business surveys (see for example Martin and Papile 2004 and Trebing 1998), prior research has not investigated the statistical accuracy of the survey results. This is because the respective non-random sampling methods render this assessment non-trivial.

This paper aims to fill this void in the literature by presenting a new approach to modeling complex non-random sampling behaviour. Using bootstrapping techniques, our approach allows us to assess the precision of, and confidence intervals around, the BOS quantitative results. Our results show that under certain assumptions, the Bank of Canada Business Outlook Survey's firm selection process do not results in wider confidence intervals over an above the random selection case, and do not appear to produce any particular bias in the estimates.

The rest of the paper will be organized as follows. Section 2 provides a short description of the Bank of Canada's Business Outlook Survey and its quota sampling methodology. Section 3 describes the Monte Carlo simulation and discusses some of the relevant literature. Section 4 presents the data used in our model, and Section 5 explains our selection model. Analyses of the results and a sensitivity analysis are reported in Section 6 and 7. Some concluding remarks follow in Section 8.

(2) A Description of the Bank of Canada's BOS

The Business Outlook Survey (BOS) consultation allows Bank of Canada economists to engage in two-way conversations with Canadian businesses about developments in the Canadian economy.³ The BOS is conducted on a quarterly basis by the Bank's regional offices, located in Halifax, Montreal, Toronto, Calgary,

¹ See Macklem (2002) for a discussion of the monetary policy discussion-making process in Canada.

² A glossary of relevant survey terminology is included at the end of this document.

³ For a full review of the Business Outlook Survey procedures, the reader should refer to Martin and Papile (2004).

and Vancouver. For each round of consultations, 100 private sector company head-offices are targeted to obtain a representative profile of the Canadian economy, the regional and industrial mix of companies approximating their representation in business sector GDP.⁴ Efforts are also made so that businesses selected by each region reflect the composition of that region's GDP. In order to get a good representation of business sentiment, a cross-section of small, medium, and large companies are consulted.⁵ The survey is conducted on a voluntary basis. If a selected firm is unavailable, another suitable firm is substituted to maintain (as much as possible) the sample size as well as the regional and industrial pre-determined quotas. Meetings are usually scheduled over a three- to four-week period during each quarter. Such a set-up insures that the cost of acquiring intelligence is kept relatively low.

In general, the BOS is designed to acquire information on four broad themes: (1) the company's past business conditions; (2) its outlook for various aspects of business activity – including sales, investment, and employment; (3) an evaluation of the pressures on the firm's production capacity; and (4) its outlook for wages, prices, and inflation.

The survey questions typically use two formats to summarize information. The most common is a three-part scale for measuring qualitative responses: positive/higher, no change/the same, or negative/lower. Then, balances of opinion are constructed to summarize the information collected by subtracting the proportion of negative responses from the proportion of positive responses. As such, balance of opinion values can range from -100 to +100.⁶ The second format, appropriate for questions on firms' ability to meet demand and labour shortages, reports the proportions of respondents (i.e. 'yes' or 'no' questions) experiencing constraints. Hence, sample proportions for these two questions are bounded between 0 and 100.

The value of the BOS comes in part from the information provided through confidential discussions with business representatives. These discussions offer valuable qualitative knowledge of business conditions and expectations. In addition, the BOS's quantitative results, as reported through balances of opinions and population proportions, also appear to provide informative measures of current business conditions and expected future activity as was reported by Martin and Papile's (2004) correlation analysis. As they argued, "*[t]he method of sample selection ensures a good cross-section of opinion. Nevertheless, the statistical reliability of the survey is limited, given the small sample size.*" Hence, the goal of this paper is to assess the statistical reliability of the survey by providing an

⁴ The Business Outlook Survey regional, industrial and size mix is presented in Appendix B to this document.

⁵ Each quarter, the aim is to balance the survey sample with more or less one-third of small, medium, and large firms. For the purpose of the survey, firm size is defined by the number of employee: small (1 to 100), medium (101 to 500), or large (more than 500). A comparison of between the BOS quotas, GDP weights, and the Dunn and Bradstreet Database of firms is covered in Appendix A.

⁶ Another way to report this three-part scale data is offer by the ISM diffusion indexes. These are calculated by taking the percentage of respondents that report that the activity has increased and adding it to one-half of the percentage that report the activity has not changed and adding the two percentages. Using half of the "Same" percentage effectively measures the bias toward a positive (above 50 percent) or negative index. Both the balance of opinion and the diffusion index offer the same information, as they are in fact monotonic transformations of another.

understanding of variance estimators and confidence intervals around the unknown population parameters.

(3) Building Our Selection Model: Data and Method

Constructing the Population

Theoretically speaking, the population of firms that can be invited to participate in the BOS is composed of all private for-profit Canadian head offices and all Canadian subsidiaries of foreign companies. However, it is extremely difficult to know at any given point in time the exact number of firms that would fit the statement, given for example, the high entrance and exit rate of companies. Therefore, to create a pseudo-population of firms we use the Dunn & Bradstreet database (D&B), a private business registry. Based on an August 2007 download, it provided information on 14 452 Canadian firms (based on the number of head offices).⁷ This dataset provided five pieces of information required to build our sampling scheme:

1. firm name (used to identify the sampling unit)
2. head office city/town (required in for the Rotation/Cluster sampling)
3. head office region (Atlantic, Quebec, Ontario, Prairies, British Columbia)
4. firm-wide Canadian employee count (used to define firm size)⁸
5. major industry grouping (six broad categories based on 3-digits NAICS codes)⁹

How each of these variables is used in the construction of the selection model is highlighted in the next section.

Prior to the sample selection, an “opinion” is randomly assigned to each firm present in the database for the two types of BOS questions. For the balance of opinion type question, a positive, neutral or negative view is assigned to each firm. For the population proportion questions, which are essentially “yes or no” questions, we pre-determine in which of the two camps a firm will lie. The probability with which a firm is attributed a response is conditioned on the firm’s key characteristics (i.e.: region, industry, and size), and these probabilities were calibrated in accordance with the BOS recorded historical responses for the past sales (balance of opinion) and capacity constraints (population proportion) questions.¹⁰ The fact that the Canadian economy appears to have gone through a

⁷ This dataset includes head quarters only and excludes the religious, not-for-profit, health care, education, and public administration sectors. Generally, only firms that borrow external funds are included (although efforts are made to include others), which may result in an undercounting of smaller firms or foreign firms with a Canadian head office that undertakes these operations.

⁸ Where small firms have less than 100 employees, medium-sized firms have between 100 and 499 employees, and large firms have more than 500 employees.

⁹ The six broad categories are: 1- the primary sector (NAICS 100 to 219); 2- the manufacturing sector (NAICS 300 to 340); 3- the trade sector (NAICS 410 to 479); 4- the construction, information, transportation and utilities sector (NAICS 220 to 239 and 480 to 519); 5- the finance, insurance and real estate sector (NAICS 520 to 539); and 6- the community, personal, and business services sector (NAICS 540 to 569 and 810 to 819).

¹⁰ The calibration of this exercise can be done an infinite number of ways. For example, one could choose to calibrate the different sub-samples in such a way as to assume that the Central Canadian manufacturing sector is in recession, while the primary sector is expanding at a rapid pace. As a first step,

full business cycle over the survey period comforts us in this assumption. The probability distributions for each of the ninety possible sub-samples for the two types of questions are presented in Appendix B.¹¹ From this information set, the reader will be able to appreciate, for example, that a small Atlantic Canadian firm operating in the primary sector is, on average over recorded BOS history, less optimistic than a large Ontarian manufacturer when reporting its opinion about its own past sales momentum.

A general framework for firm selection simulations

The difficulty when working with several other non-random sampling elements (such as quota sampling) is that the probability of selection of its sampling units is unknown. For this reason, there is no conventional closed form statistical equation for calculating the variance of an estimate derived from this type of non-random selection process. However, using Monte Carlo techniques, various re-sampling schemes (i.e. a model-based approach) can be used to analyze complex probability surveys (see for example, Rao and Wu 1988, Sitter 1992, Shao 2003, and Davidson and Mackinnon 2004). While we take from this literature the general approach, we uniquely incorporate a number of non-random selection elements into our model, (so-called) ‘build-up’ from the random case.

Before presenting the model mimicking the BOS selection process, we present a general framework for the random case as an introduction to our methodology. In the general framework, the first step is to model the data generating process in order to create a finite ‘pseudo-population’.¹² In order to generate the pseudo-population with survey responses ex-ante, it is divided into N firms in L non-overlapping strata of N_1, N_2, \dots, N_L units such that the total population size is, using Sitter’s (1992) notation:

$$\sum_{i=1}^L N_i = N .$$

The sample sizes within each stratum are denoted by n_1, n_2, \dots, n_L , such that the total sample size is:

$$\sum_{i=1}^L n_i = n .$$

Next a vector of K characteristics has to be specified, represented by

$y_{ij} = (y_{1ij}, y_{2ij}, \dots, y_{Kij})^T$, where i refers to the stratum, and j refers to the j th unit of the i th stratum. For our purposes, we assign a survey question response to each of the firms based on their K characteristics (such as region, industry and firm size). The probability of firm y_{ij} responding positively, neutral or negatively to a question is based on the historical BOS responses.¹³

we decided to abstract from any potential business cycle effect and take the average BOS historical responses for each sub-samples as a simplifying assumption. Further research will address the sensitivity of this assumption.

¹¹ The 90 subsamples ($5 \times 6 \times 3 = 90$) are determined by firm’s head-office location (5), broad industry grouping (6), and size (3).

¹² This ‘pseudo-population’ can then be calibrated using the moments of the original population.

¹³ A sensitivity analysis to this approach is discussed in Section 7.

Once the population has been identified and assigned a response, we can obtain an estimate of the population parameter $\theta = \theta(S)$, where $S = \{y_{ij} : i = 1, 2, \dots, L; j = 1, 2, \dots, n\}$ one carries out the following steps:

1. Draw a random sample $\{y_{ij}^* \}_{j=1}^{n_i}$, such that $n_i^* \geq 1$, with replacement from $\{y_{ij}\}_{j=1}^{n_i}$. (The * denotes that the variable or parameter estimate is obtained from a random draw).
2. Repeat step 1 a large number of times, B , to obtain multiple estimates of the population parameter denoted $\hat{\theta}_1^*, \hat{\theta}_2^*, \dots, \hat{\theta}_B^*$.
3. Calculate the variance of each $\hat{\theta}^*$ relative to the 'pseudo-population' parameter θ :

$$v_r = E(\theta - E \hat{\theta}^*)^2,$$

or its Monte Carlo approximation:

$$v_r = \frac{1}{B-1} \sum_{j=1}^B (\theta - \hat{\theta}_j^*)^2$$

When the frequency distribution is symmetric, Davidson and MacKinnon (2004) suggest using the percentile-method, for which Monte Carlo confidence intervals have the form:

$$[\theta - \hat{\theta}_j^* |_{(1-\alpha/2)(B+1)}, \theta - \hat{\theta}_j^* |_{(\alpha/2)(B+1)}]$$

where α is typically set to 0.05 for a 95 per cent confidence level. For a simulation of $B = 1000$, this is equivalent to taking the 25th smallest and the 975th largest estimated parameter values to be the confidence interval.

The Federal Reserve uses a similar procedure to construct confidence intervals for their *National Survey of Small Business Finances* (NSSBF).¹⁴ The NSSBF is similar to the BOS in that it stratifies targeted businesses by size and region. "The sample design is sufficiently complex that it is not possible to apply standard methods to estimate sampling variance without a number of assumptions. For this survey, we use a bootstrap replication procedure to capture the important dimensions of variation in the original sample selection and the adjustments made at the weighting stage." (Board of Governors, 1996). The authors create a pseudo-population based on 4,637 firms, and then run 1000 with-replacement simulations. The 4,637 firms come from the two NSSBF surveys of 1987 and 1993 (the NSSBF was subsequently conducted in 1998 and 2003).

Another example of the application of this approach is Ho (1993). The author considers whether a reduction in the sample size of the IRS' Taxpayer Compliance Measurement Program survey would yield any sizable increase in the confidence interval of the estimate of the parameter of interest. The population from which Ho (1993) conducts random draws is in fact the sample of firms that the IRS had

¹⁴ Nevertheless, not for the *Beige Book*, perhaps because of its more qualitative nature.

audited. From Ho's perspective, the parameters of the true population (i.e., the universe of tax paying firms) are irrelevant. Ho's approach merely seeks to compare the parameter estimates of smaller-sized draws relative to the known parameters of the large sample of firms.

Aparicio-Pérez and Lorca (2005) is a more recent empirical application that resembles the BOS predicament. The authors use with-replacement simulations to estimate confidence intervals for imputed (missing) values in Spain's stratified Structural Industrial Business Survey. In their case, the data generating process is the imputation method, whereas in our case it is various aspects of the BOS sampling process.

The key challenge in assessing the confidence intervals around the BOS results stem from the fact that the survey uses a complex non-random sampling methodology, and as such, the probabilities of selection for a given sample unit are unknown. The random sampling scheme uncovered above will therefore be conditioned to make use of the prior knowledge of the systematic non-random elements in the BOS sampling process. In other words, the random draws in step 1 above will be subjected to non-random constraints characteristic to the BOS sampling process. The random cases, as uncovered above, will serve as our basis for comparison.

(4) The Case of the BOS: Accounting for non-random selection

There are several systematic non-random elements in the BOS sampling process that we can model within the Monte Carlo framework to account for its non-probabilistic nature.¹⁵

1. *Rotation/Cluster sampling* (firms from specific urban clusters are surveyed in each of the four rounds per year)
2. *Quota selection* (region, industry and size targets fixed in advanced)
3. *Familiarity sampling* (certain firms are sampled on a regular basis)
4. *Non-response bias* (based on the analysis of non-response firm characteristics by Ainsworth and Pichette 2006).

We believe that these four factors are the most important non-random elements in the BOS sampling process. Evidently, there are hosts of other behaviours inherent to opinion surveys that we cannot model. They include, but are not limited to: 1) survey design, 2) interviewer bias, 3) timing effects, and 4) survivor bias. Obviously, these intricacies affect all surveys and there is no reason to believe that they would be systematically greater for the Bank of Canada's BOS. Evidently, idiosyncratic bias proper to the BOS sampling process could remain. We have not been able to

¹⁵ Complete definitions are available in Appendix B.

identify any beyond the four behaviours listed earlier that might be systematically important.

The selection model presented in the following section is designed to mimic the BOS sampling procedure. We aim to incorporate the above-listed non-random selection elements into a re-sampling scheme in order to estimate the additional variance and potential bias that these non-random elements induce into the results compared to a simple random selection model. Evidently, the general logic can be applied to the quantitative results (diffusion index, balance of opinions, population proportion, etc.) of any non-random complex sampling process. Once a specific non-random behaviour is identified, one can redesign a sampling procedure and study its effect on sample variance and bias.

(5) The Selection Model

Simulating the rotation/cluster constraint

The first step in the BOS sampling process involves selecting a cluster of cities and towns that interviewers will visit. Modeling this process will account for any potential clustering bias in the BOS sampling procedure. In practice, the choice of locations depends on the relative size of the various urban economies in each region, travel times between urban areas, and rotational practices at the regional offices.

The 14,452 Canadian firms in the Dunn & Bradstreet database are headquartered in 1,206 various cities, suburbs, and towns across Canada. Bank of Canada interviewers periodically visit 38 metropolitan clusters encompassing 437 of these separately named locations. A total of 10,497 firms from the Dunn & Bradstreet database are located in these clusters. When we impose the rotation/clustering constraint to our model, only these firms have a chance of selection. Hence, firms in cities never visited by the Bank's regional economists, have a zero probability of being selected.

Secondly, most major urban areas, including the cities in which each regional office is located, are visited every survey round while others are visited only on alternating rounds, twice per year, or sometimes even less frequently. In order to account for these relative sampling frequencies we introduce the concept of a 'rotation'. Specifically, we consider four rotations representing the four survey periods in a year. Each city is placed into one or more rotation depending on how often it is visited. The relative frequencies are based on historical practice at the regional offices.

Simulation of the two steps described above is graphically depicted below in Figure 1. First, every new draw, the rotation is determined within the simulation based on probabilities $\rho_1 = \rho_2 = \rho_3 = \rho_4 = \{0.25\}$, where the subscripts 1 through 4 refer to the four rotations. Next, the locations to be sampled are determined with probability $\lambda = \{0, 1\}$. Note that λ is conditional on ρ . This selection constraint has

the potential to bias the estimated population parameter if the response distribution of the different clusters displays a high degree of variability. For example, suppose that the second largest urban cluster in a particular region is affected by idiosyncratic shocks that nevertheless would affect the true population parameters. Now, what if this region is visited only once a year (i.e. it has a much lower probability of being selected than a firm located the largest regional city) then the estimated sample parameters could be biased. In other words, the purpose of this constraint is to capture the fact that some firms have a higher or lower probability of selection based on their geographical location. Based on the regional offices historical practice and the BOS response distribution, 60 per cent of firms are available for sampling in each of the four rounds, as most of the head-offices are located in the home city.

Simulating the quota constraints

Next, we model the BOS quota constraints. The sampling quotas are designed to mirror private sector GDP shares based on a firm's region of activity, industry grouping, and firm size.¹⁶

First, the simulation randomly draws a firm from one of the pre-determined rotational clusters. Then the characteristics of the selected firm (region, industry, and size) are compared to the respective BOS quotas. If the firm meets all three conditions (i.e. the quotas are not full), the simulation proceeds to the next step.

This process of verifying the quotas is illustrated below (Figure 2). In this example, a randomly selected firm from Atlantic Canada (region α_1), operating in the trade sector (industry β_3), employing between 100 and 499 people (firm size σ_2) is mapped against the reselected quotas. If the quotas for α , β , and σ are not yet filled, the firm passes to the next step. If one of these quotas is already filled, the firm is then rejected.¹⁷

Simulating the familiarity constraint

Once it has been determined that a firm meets the quota constraints, we model the familiarity behaviour inherent in the BOS sampling process. We do this by verifying whether the firm is on the regional offices' list of existing contacts. Existing contacts are firms that have participated in one or more past surveys. Typically, such firms are surveyed at most once every two years. The model incorporates this familiarity constraint by cross tabulating the regional office contact list with the Dunn &

¹⁶ The regional quota is determined by location of head office, which is an imperfect but necessary proxy for the firm's region of activity. Firm size quotas are devised as per footnote 9. Industry grouping are as follows: Primary (NAICS 2002 > 100 to <220); Manufacturing (NAICS 2002 >310 to <340); Trade (NAICS 2002 >409 to <480); Construction, information and cultural industries, transportation and utilities (NAICS 2002 >219 to <240 and >479 to <520); Finance, insurance, real estate and leasing (NAICS 2002 >519 to <540), and lastly, Commercial, personal and business services (NAICS 2002 >539).

¹⁷ Obviously, in real-life survey setting once a selection quota is filled, the surveyors will stop contacting firms displaying the relevant quota characteristics, however, this non-random behaviour is hard to code in our model setting. Note that these rejections are not to be confused with the non-responses discussed further in the text.

Bradstreet database of company names (including abbreviations, common names, and holding company names).

The regional office contact list contains 2,628 individual companies, of which 1,201 can also be found in the Dunn & Bradstreet database. The latter are coded as 'existing' contacts in the simulation model. Historically, in any given survey round existing contacts make up about one-third of the BOS sample. Accounting for the variation around this number, the selection model is programmed to fill at least 20 per cent and no more than 40 per cent of the sample with existing contacts. The familiarity constraint's impact should be limited since we have not assumed that familiar firms have systematically different opinions than non-familiar firms. However, this constraint does restrict the population of firms available for sampling in each rotation, as only 8% of firms present in the database are considered 'familiar' firms. Hence, over repeated sampling these firms are bound to be over-represented and may nevertheless generate a potentially important source of bias.

Modeling non-responses

The next step in the model is to account for the non-responses. Even existing contacts will for a variety of reasons be unable or unwilling to participate in a given survey round. Ainsworth and Pichette (2006) document the historical non-response rates for the BOS based on firm characteristics (region, industry, size). We use the industry specific response rates to determine whether a selected firm enters the sample within our simulations.¹⁸ We use a random number generator to determine whether firm j in industry i provides a response. Specifically, the firm will provide a response if $\pi_j < \pi_i$, where π_i is the industry specific historical response rate.¹⁹ If a firm provides a response, it is captured as part of the simulated survey. The survey quotas for the firm's region, industry, and size are reduced by one respectively. If a firm does not provide a response, then its response is placed into a counterfactual survey.

All of the above steps are repeated until the simulated survey contains the required 100 observations. This marks the completion of one simulated Business Outlook Survey. The 100 responses are averaged to produce a balance of opinion or compiled to obtain sample proportions depending on the BOS question under investigation. Likewise, the counterfactual responses are collected. The simulation is re-run 1000 times yielding vectors of balances of opinions or sample proportions, as well as their counterfactual counterparts.

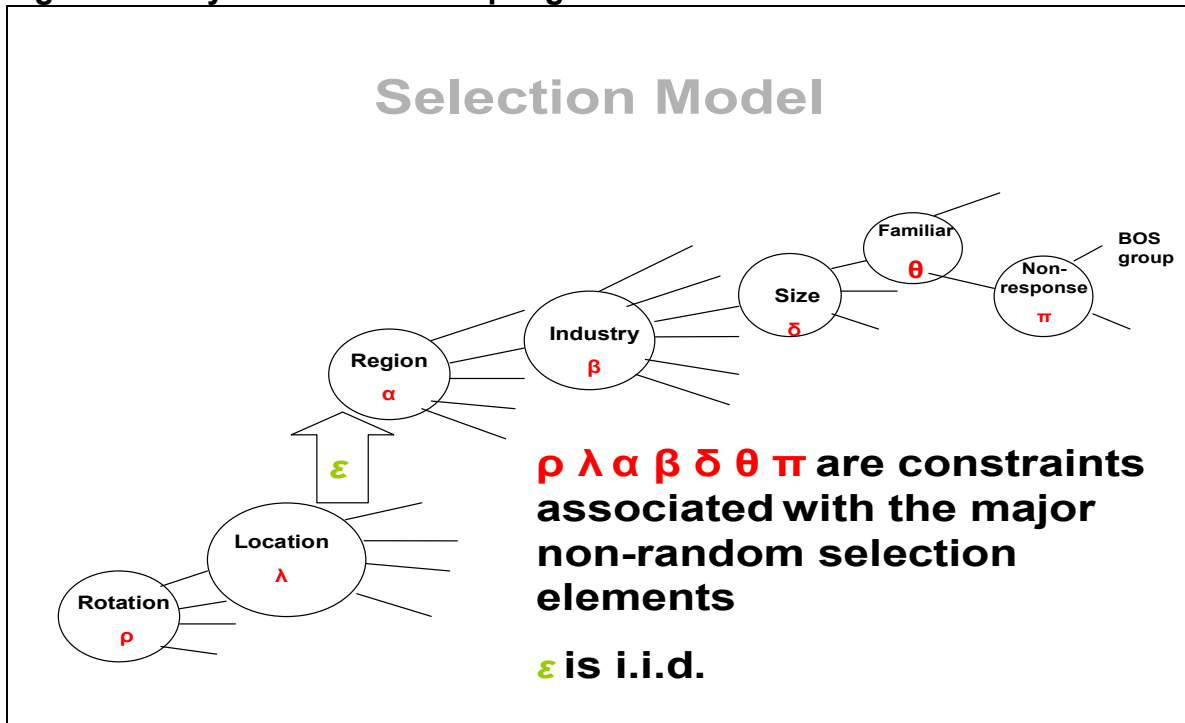
The diagram below summarizes the fully constrained sampling model. As discussed above, we first select a specific rotation period (p) during which specific locations are visited with probability λ . Once our rotation cluster has been predetermined we select a firm at random. This operation is represented on the diagram by the random number generator function $\varepsilon \sim N(0,1)$. The selected firm is then matched to the region, industry, and size quotas. If one of the quotas is

¹⁸ Because the joint region, industry, firm size response rates are difficult to calculate, added the regional and firm size rates into the model is left to future research.

¹⁹ The reader should see that π is nested between 0 and 1.

already filled, the firm is cast aside and another firm is drawn at random from the rotation cluster. If all three quotas are as yet unfilled, the model verifies whether the firm is an existing contact (familiar) and whether the familiarity constraint has been met. No more than 20 to 40 per cent of firms in the sample can be familiar firms. Once a firm has past all of the above constraints a random number generator determines whether the firm responds to the survey or not (based on its industry specific historical response rate). If a firm 'agrees' to respond, responses are gathered as part of a simulated BOS sample. Otherwise, it is simply discarded.

Figure 1: Fully constrained sampling model



(6) Selected Simulation Results

Based on bootstrapping techniques we set out to model the survey's complex non-random sampling behaviour that will allow us to assess the variance estimation and confidence intervals around BOS results. The framework presented above will also be useful in determining the impact of the BOS's individual selection behaviours on the survey's precision and help identify potential biases. In order to make a valid assessment of the BOS sampling behaviour, we will compare our results to the simple random sample case for which 1000 simulations were also run.

General results are presented in the form of estimated kernel density functions. To evaluate the confidence intervals we use the two methods identified above: 1) the Monte Carlo approximation and 2) the percentile method first introduced by Efron

and Tibshirani (1993).²⁰ The same two methods will be used to assess any potential bias: 1) the Monte Carlo approximation which compares the simulated mean to the known 'pseudo-population' mean, and 2) the percentile method which will compare the median response (i.e.: the 500th observation) to the true 'pseudo-population' mean. By choosing to compare simulation result to the true population mean, we implicitly assume in our opinion survey setting that all firms' opinions are created equal.

Eight different sampling schemes will be employed in our analysis, including the random sampling methods against which all other methods will be tested. These sampling schemes are:

1. The Random Sample Model (i.e. no constraints imposed)
2. The Regional Quota Model (i.e. only the regional quota constraints are imposed)
3. The Industry Quota Model (i.e. only the industry constraints are imposed)
4. The Firm-size Quota Model (i.e. only the size constraints are imposed)
5. The Rotation/Cluster Constraint (i.e. restricting firm selection based on location)
6. The Familiarity Constraint (i.e. restricting selection based on our visit history)
7. The Non-response Constraint (i.e. the probability a firm will respond to the survey by industry classification)
8. The Fully Constrained Sampling Model (i.e. all selection constraints imposed).

Each of these sampling schemes will be tested to establish whether they are significantly different from the random model using the Shapiro-Francia test for normality and the Kolmogorov-Smirnov equality test.

Section 6.1: Selection Model Results relevant for the Balance of Opinion

The results for the estimated balance of opinion are reported in Table 1. Before considering the fully constrained model, it is interesting to see how each constraint individually affects the estimate of the population parameter compared to the simple random sample case. The first row shows results for the random sample model that was generated by randomly selecting 100 firms from a universe of 14,452 firms. This model provides a virtually unbiased estimate of the population parameter with a 95% confidence interval of +/- 16.6 percentage points.

The results from the three quota constraints (region, industry, and size), when controlled individually are presented in rows 2-4 of Table 1. These results suggest that the regional quota positively biases the population parameter estimate by 2 percentage points compared to the mean of the known pseudo population. In contrast, the industry and firm size models negatively bias the estimator by about 2 and 3 percentage points, respectively. When investigated individually, only the

²⁰ We will use the 25th and 975th observation to estimate the 95% confidence interval, the 170th and 830th observations to estimate the 66% confidence interval.

industry and the firm size quotas are found to widen the confidence interval bands when compared to the random sample model, at +/- 17.5 and +/- 17.7, respectively.

Various Selection Models	Bias Percentile method^a	Bias vs. 'pseudo-population'^b	Confidence Band - 95%^c (66%)	Shapiro-Francia Normality Test (p<Z-vaule)	Kolmogorov-Smirnov Two-sample Equal dispersion Test (p-value)^d
1- Random Sample	0	0.06	16.6 (8.2)	0.467	(tested against)
2- Regional Quota	2	2.00	16.6 (8.2)	1.000	0.000
3- Industry Quota	-2	-2.07	17.5 (8.6)	0.990	0.000
4- Firm Size Quota	-3	-2.78	17.7 (8.7)	0.358	0.000
5- Rotation Constraint	0	0.17	16.7 (8.2)	0.515	0.888
6- Familiarity Constraint	0	-0.23	17.0 (8.4)	0.463	0.794
7- Non-response Constraint	0	-0.10	16.7 (8.4)	0.602	0.500
8- Fully Constrained Model	-1	-0.23	16.8 (8.3)	0.648	0.859

^a These are based on the Davidson and MacKinnon suggested comparison of the 500th simulation and the known 'pseudo-population' mean.

^b The bias is estimated as the simulated mean deviation from the known 'pseudo-population' mean.

^c The 66% and 95% confidence bands are calculated from the known 'pseudo-population' mean.

^d This test compares the distribution of the kernel density function estimates of each constrained model to that of the Random Sample Model results (evaluated at the same 63 points – covering the range of outcomes - for each model results; Epanechikov Kernel was used).

Why do the quota constraints bias the balance of opinion? Because the BOS selection method is constrained by the quota targets, the degree of over or under representation of specific quotas relative to the D&B population of firms could introduce biases in the results. As reported in Appendix B, we can easily see that the lower than average balance of opinion for Ontario (2.6% vs. 16.1% for the 'pseudo-population'), in combination with Ontario being under represented (25 firms out of 100 sampled compared to 42% of the D&B population) could likely explain the positive regional quota bias. Similarly, the lower balance of opinion for the over represented manufacturing sector (28 firms out of 100 sampled compared to 15% of the D&B population) may largely explain the negative bias generated by the industry constraint. Lastly, the negative bias generated by the firm size constraint could be driven by the fact that both medium and large sized firms have a lower balance than smaller firms, which represent 85% of the D&B database.

The results from the rotation/clustering, familiarity, and non-response constraints are reported in rows 5-7 of Table 1. None of these BOS sampling elements individually biases the parameter estimate. The homogeneity of firms' balance of opinion within the four rotations' pseudo-population, between familiar and non-familiar firms, as well as between respondents and non-respondents, suggests that the results are unlikely to differ from the random sample case (Table 2). Sensitivity

testing on these assumptions is left for future research.²¹ Nevertheless, the fact that we model these clusters now will induce some degree of variability that could, compounded with other constraints, affect the final results.

	Total Sample	Rotation 1	Rotation 2	Rotation 3	Rotation 4
% of firms available for sampling from pseudo-population	100%	61.27%	59.50%	61.52%	61.78%
population proportion	16.09%	16.35%	16.26%	16.23%	16.08%

The final selection model, the fully constrained simulation, provides us with an estimate of the confidence band and potential bias for the BOS by constraining each firm selection on all major non-random elements of the BOS selection process. As reported in Table 1, our best estimate of the impact of these factors on the BOS confidence interval is that it is close to the random normal case. In other words, although the quota constraints results in biases on the parameter estimate when controlled individually, these biases remain small and appears to be largely offsetting when the model is calibrated using average historical responses.

Lastly, the Shapiro-Francia test for normality suggests that the results for all eight-simulation models are normally distributed. This suggests little variation in the distribution of errors (no skewness or kurtosis). The Kolmogorov-Smirnov equality of distributions tests confirms that the three quota constraints produce a biased estimate for balance of opinion type questions, while the other constraints in isolation and the fully constrained model do not when the model is calibrated on the average historical response by region, by sector and by firm size.

Section 6.2: Selection Model Results relevant for the Population Proportions

The labour shortage and capacity constraint questions are different from the other published BOS questions because they seek to determine a single population proportion (i.e. they ask a yes or no question). In order to re-estimate the selection model and obtain result for this type of question, we made use of the capacity constraints²² historical data to calibrate a population proportion response vector (Appendix B, Table ?). Hence, the firms' responses were calibrated this time around a population proportion type question that only allows a yes (+1) or no (0) answer. The results presented below were estimated using the methodology presented in Section 5. As for section 6.1, Table 3 compares the random-selection model results to seven constrained models. Although the fully constrained model is

²¹ In real-life survey settings, regional issues can very well become important, especially in geographical clusters experiencing idiosyncratic shocks, such as mono-industrial regions for example.

²² The survey questions reads as follows: "How would you rate your firm's current ability to meet an unexpected increase in demand or sales?" The percentage of firms that report having some or significant difficulty meeting an unexpected increase in demand is our population proportion.

the focus of our investigation, it is interesting to consider each of the constraints in isolation to provide some insight to how they affect population proportion estimate.

Table 3: Comparing Population Proportion Simulation Results

Various Selection Models	Bias Percentile method ^a	Bias vs. 'pseudo-population' ^b	Confidence Band - 95% ^c (66%)	Shapiro-Francia Normality Test (p<Z-vaule)	Kolmogorov-Smirnov Two-sample Equal dispersion Test (p-value) ^d
1- Random Sample	0	0.06	9.49 (4.67)	0.877	(tested against)
2- Regional Quota	2	2.00	9.33 (4.60)	0.944	0.024
3- Industry Quota	-2	-2.07	12.25 (6.03)	0.748	0.308
4- Firm Size Quota	-3	-2.78	9.81 (4.82)	0.012	0.096
5- Rotation Constraint	0	0.17	10.13 (4.99)	0.632	0.067
6- Familiarity Constraint	0	-0.23	9.62 (4.74)	0.815	0.036
7- Non-response Constraint	0	-0.10	9.91 (4.80)	0.160	0.054
8- Fully Constrained Model	-1	-0.23	10.01 (4.93)	0.197	0.152

^a These are based on the Davidson and MacKinnon suggested comparison of the 500th simulation and the known 'pseudo-population' mean.

^b The bias is estimated as the simulated mean deviation from the known 'pseudo-population' mean.

^c The 66% and 95% confidence bands are calculated from the known 'pseudo-population' mean.

^d This test compares the distribution of the kernel density function estimates of each constrained model to that of the Random Sample Model results (evaluated at the same 40 points – covering the range of outcomes – for each model results; Epanechikov Kernel was used).

As expected, the random selection model (row one, Table 3) produces the estimated population proportion that is almost exacting the true (known) pseudo-population proportion. The three quota constraints, modeled individually, are presented in rows 2-4. While the regional constraint does not significantly influence the proportion estimate or the confidence bands, both the industry and the firm size constraints bias the results by +4% and -1% per cent, respectively. Most interestingly, the confidence bands for the industry constraint are about +/- 3% wider than the random case at the 95% confidence level. This result largely reflects the fact that the trade sector, which represents a large number head offices, has a significantly lower average population proportion than the sample mean (Appendix B). Despite these differences, however, the industry quota constraint model results distribution is normally distributed and not significantly different from the random case according to the Kolmogorov-Smirnov test (last column of Table 3). With respect to the firm size constraint's negative bias, the lower average proportion for large firms could account for the difference. This constraint produces the only non-normally distributed set of results (Table 3).

Rows 5-7 of Table 3 present the results from the rotation/clustering, familiarity, and the non-response constraint models. While the familiarity and non-response simulations produced results close to the random selection model, as they did for the balance of opinion estimation, the rotation constraints does produce a slight

positive bias and wider confidence intervals. Further testing suggests that the distribution is not significantly different from the normal distribution, however. Interestingly, and unlike the results presented in Section 6.1 for the balance of opinion simulations, the rotation constraint produces a small positive bias. This is reflective of the fact that the four rotations have a similar sized bias owing to the fact that firms located in cities that are not visited by the Bank’s regional economists have a slightly lower average incidence of capacity constraints because of their firm size and industry mix (see Table 4 below). These differences, however, are a function of the calibration of the response vector, which is based on the BOS average historical responses. Overall, only the regional constraint and familiarity constraint, however, produce a distribution of results that are significantly different from the random model (Table 3).

	Total Sample	Rotation 1	Rotation 2	Rotation 3	Rotation 4
% of firms available for sampling from pseudo-population	100%	61.27%	59.50%	61.52%	61.78%
population proportion	38.13%	38.77%	38.75%	38.70%	38.69%

Finally, row 8 of Table 3 presents the results from the fully constrained model. Once all seven constraints are imposed, the model estimates the population proportion is slightly positively biased and the confidence intervals are slightly wider. Although it is difficult to pinpoint the exact source of the bias, the industry constraint could probably account for it. Significance testing suggests that the results are normally distributed and not significantly different from the random case.

The key points to retain from section 6.1 and 6.2 is that, when response vectors are calibrated using the BOS average historical responses, the induced bias on the fully constrained model is small and confidence intervals around responses are generally close to what would be obtained if the survey was done in a purely random fashion. Results could prove to be different if the underlying response vector was calibrated differently, as it is likely to happen throughout the different regional or industrial business cycles.

Section 7: Sensitivity Analysis

Hence, how sensitive are our results to the assumptions underlying the calibration of the response vectors between the 90 identifiable strata? In the first set of results (Section 6), we relied on the historical past sales question distribution of positive, negative and neutral responses for each of the 90 possible strata shown in appendix B for balance of opinion and population proportion questions. Although we believe this assumption is reasonable, we propose a sensitivity test based on a

single distribution of responses for all firm (i.e.: all firms have the same probability of responding positive, negative or neutral to the balance of opinion question). This calibration will allow us to isolate the BOS sampling procedure and determine whether it is a source of bias, and eliminates our dependence on the historical data by allowing the variation between firm responses to be purely random.²³

Once the response vector is calibrated, we run the model according to the methodology in Section 5. Once again, the only difference between the results presented in Section 6 and the results below is that the responses assigned to the 14,452 firms in the pseudo-population were assigned differently.²⁴ The results, presented in Table 5 suggest that: (i) the distribution of balances of opinion generated by the model is not significantly different from normal; and (ii) the distribution of results from the fully constrained model is not significantly different from the random model. Interestingly, however, the confidence intervals are about half a percentage point wider and the estimated balance of opinion is biased by about +3.5%, which is more than the estimate produced using the historical data (+1%). However, despite the fact that response vectors were generate using the same probability function, since 75 of the 90 strata post a sub-sample size of less than 300 firms, there will still be a high degree of variation in the posted balance of opinion between strata which will play a major role in the appearance of the bias. The main point here remains that despite all the constraint imposed on the selection model, the bias (when present) remains small and the confidence intervals are not significantly different from what could be obtained, given the nature of the balance of opinion, from a purely random sampling strategy.

Various Selection Models	Bias Percentile method^a	Bias vs 'pseudo-population'^b	Confidence Band - 95%^c	Confidence Band - 66%^c
1- Random Sample	0	-0.18	16.48	8.12
2- Regional Quota	-1	-0.97	16.45	8.10
3- Industry Quota	1	0.58	17.12	8.43
4- Firm Size Quota	-2	-2.25	17.73	8.73
5- Rotation Constraint	1	0.48	16.75	8.25
6- Familiarity Constraint	0	-0.07	16.71	8.23
7- Non-response Constraint	0	0.19	16.68	8.21
8- Fully Constrained Model	-3	-3.55	17.44	8.59

^a These are based on the Davidson and MacKinnon suggested comparison of the 500th simulation and the known 'pseudo-population' mean (which we round to 16% (from 16.19%) since our balance is out of 100 firms).

^b The bias is estimated as the simulated mean deviation from the known 'pseudo-population' mean.

^c The 66% and 95% confidence bands are calculated from the known 'pseudo-population' mean.

²³ There is a number of different sensitivity test that we could conduct. What other possible sensitivity tests are we considering? First with respect to the non-response and the clustering constraint, one alternative is to give non-responders and various clusters a systematical different response from an alternative generating process. Second, and perhaps more interestingly, we could consider large regional and industrial disparities that reflect macroeconomic cycles to generate our response distribution. These are left to future research.

²⁴ More specifically, the calibration is set so that 46% of the firms would respond positively (+1) and 30 per cent would response negatively (-1), while all others would response neutral (0). However, since 75 of the 90 strata post a sub-sample size of less than 300 firms, there will still be a high degree of variation in the posted balance of opinion between strata despite the fact that these were generate using the same probability function.

(7) Conclusions

The Bank of Canada's Business Outlook Survey uses a complex non-random sampling methodology. It is complex because surveyors employ various sampling rules and practices to insure a good representation of the Canadian economy and the quality of the data obtained. Since statistical inference is based on random probability theory, past research has urged strong caution in the use of survey data by non-random methods such as quota sampling. However, we can describe the Bank of Canada's BOS sampling process in terms of Bayesian probabilities and historical practices in order to account for the major (but not all possible) non-random elements. After modeling the three quota constraints, the clustering (sample rotation) constraint, the familiarity, and the non-response constraint, our simulations suggest that there is little evidence that the BOS questions are meaningfully biased or have larger confidence intervals.

As many Central Banks around the world build in-house surveys and have to execute them with limited budgets, assessing the accuracy of these surveys is of great important to policymakers. In addition to answering key questions about the accuracy of the BOS, this work provides a unique contribution to the literature on survey methodology by describing a model based approach that can be used to assess similar potentially systematically biased surveys. Future research should aim to determine how sensitive our estimates of confidence intervals and biases are to changes in the calibration of the different subjective probabilities of responses between the regions, industries, firm sizes, and other constraints imposed on the BOS sample selection process.

References

- Ainsworth, T. and L. Pichette (2006). "The BOS Nonresponses Rate." RN-06-086.
- Aparicio-Pérez, F. and D. Lorca (2005). "Performance of resampling variance estimation techniques with imputed-survey data." Instituto Nacional de Estadística, submitted to work session on 'tatistical Data Editing'in Ottawa, Canada, May 2005.
- Board of Governors (1996). "National Survey of Small Business Finances Methodology Report." Board of Governors of the Federal Reserve System, *Occasional Staff Studies*.
- Davidson, R. and J. MacKinnon (1993). "Estimation and Inference in Econometrics." New York, Oxford University Press.
- Davidson, R. and J. MacKinnon (2004). "Bootstrap Methods in Econometrics." Working paper.
- Dobrescu, S. (2007). "The Business Outlook Survey and Other Canadian Surveys." mimeo.
- Dorofeev, S. and P. Grant (2006). *Statistics for Real-Life Sample Surveys*. Cambridge University Press. Cambridge, United Kingdom.
- Efron, B. (1994). "Missing data, imputation, and the bootstrap." *Journal of the American Statistical Association* 89, 463-479.
- Ho, C.-C. (1993). "A bootstrap strategy for enhancing the confidence interval estimation of small samples." *American Statistical Association 1993 Proceedings of the Section on Survey Research Methods*, 872—876.
- Lohr, S.L. (1999). *Sampling: Design and Analysis*. Duxbury Press. Pacific Grove, CA.
- Macklem, T. (2002). "Information and Analysis for Monetary Policy: Coming to a Decision." *Bank of Canada Review* (Summer): 11–18.
- Shao, J. (2003). "Impact of the bootstrap on sample surveys." *Statistical Science* 18, 191—198.
- Sitter, R. R. (1992). "A resampling procedure for complex survey data." *Journal of the American Statistical Association* 87, 755—765.
- Trebing, M. (1998). "What's Happening in Manufacturing: 'Survey Says...'" *Philadelphia Federal Reserve Business Review*. September/October.

Appendix A: Glossary of Relevant Survey Terms

The following terms are presented thematically, not in alphabetical order.

Probability Sampling requires that surveyors select survey units (firms in the case of the BOS) with a known probability of selection. It is not required that each firm has an equal probability, but that probability is known and that that selection is therefore random at least within a given strata or segment of the survey populations. In order to select a firm with a known probability, the population (or the *sampling frame*) must be well-defined.

Simple random sampling is the most well-known method of probability sampling. Each unit within the sample frame is selected independently with identical probability.

Stratified sampling involves breaking up the population into smaller mutually exclusive groups before using random sampling.

Non-probability sampling is a survey method where the probability of a unit's selection cannot be quantified and therefore the results cannot be replicated. This may take many forms, such as 'conveyance sampling' (asking questions on a street corner) or more complex survey designs such as quota sampling.

Quota sampling is a type of non-probability sampling commonly used to generate representative samples. Like stratified sampling, the quotas, or target proportions, are set in advance and in such a way as to ensure that the sample is similarly proportioned to the population. The selection of units, however, is non-random.

Parallel quotas allow quotas to overlap with little or no selection preference between them. For example, the BOS is designed to be representative by region, industry, and firm-size; however, *within* the industry quota there is no target for firm size.²⁵

Interlaced quotas are when two or more quotas have a predetermined relationship. For the BOS, the industry composition of a region is known prior to the survey's execution (such as the number of manufactures to be surveyed in Ontario). Each survey, the industrial composition of each region is rotated somewhat, although some region-specific concentrations are accounted for when the targets are set (for example, the concentration of auto manufacturing in Ontario). The total number of visits within each region are also predetermined and do not change quarter-to-quarter. Once the industry composition is set, regional offices begin selecting firms to match the industry and size quotas.

Population refers to the universe of units from which samples can be drawn. For the BOS, firms are selected from a combination of in-house databases and more formal lists such as the Dun and Bradstreet Selectory. While these lists are large

²⁵ There are a few exceptions: Because there are a number of large national firms in retail, auto manufacturing, financial services, etc., regional economists do target (informally) these large nationals in each survey.

and cover a broad range of private sector firms in Canada, they are not exhaustive, and therefore are only approximations of the total population.

Clustering bias occurs when units are selected from concentrated locations and responses are correlated with location. Cluster surveying is used in the BOS to reduce the time and cost of the survey. For example, the Ontario office will find firms which match their industrial and size quotas in the Greater Toronto Area (GTA), Barrie and Windsor one round, and the GTA, Niagara and Thunder Bay the next round.

Familiarity bias occurs when units are intentionally selected in subsequent surveys, typically because they are amenable to being surveyed, but their responses differ systematically from other units. In the case of the BOS, approximately two-third of the sample is made up of new firms while one-thirds has been previously surveyed.²⁶

Non-response bias occurs when a unit is selected but does not respond to all or part of a survey. In the BOS, a firm with similar quota characteristics is selected to replace non-responders. The reasons for not responding vary greatly and may be correlated with the responses (rather than the quota characteristic), thus introducing bias. Lohr (1999) argues that the best way to deal with non-response bias is to avoid it. Although the regional offices have undertaken a concerted effort to understand the nature for non-response and to limit its frequency (Ainsworth and Pichette 2006), no 'silver bullet' has been found. Regions have created a series of informal best practises in calling companies for industry visits and are currently recording the firms' characterises and reasons for non-response to track our progress. The current non-response rate is near 40 per cent.

²⁶ Firms are surveyed at most once every 18 months. Exceptions to this rule include a few firms which operate in highly concentrated sectors.

Appendix B: Tables

Table B1 reports a number of interesting features of the BOS dataset. The distribution of firms by region in the BOS sample weights the Atlantic and B.C. regions more heavily than their actual GDP shares. Similarly, the manufacturing sector is over-weighted in the BOS sample relative to its true GDP weight.

Region	BOS Targets	GDP Share (2000-2006)	D&B Database (Aug 07)	# of obs. (% of sample)	Historical Balance of Opinion	Historical Population Proportion
Atlantic	15 firms	6%	6%	403 (15.2%)	11.7%**	44.2%
Quebec	20 firms	20%	19%	541 (20.5%)	12.4%**	33.9%**
Ontario	25 firms	42%	42%	665 (25.1%)	2.6%**	34.6%**
Prairies	20 firms	18%	20%	516 (19.5%)	20.7%**	52.8%**
BC	20 firms	12%	13%	520 (19.7%)	21.9%**	47.6%
Industry*	BOS Targets	GDP Share (2000-2006)	D&B Database (Aug 07)	# of obs. (% of sample)	Historical Balance of Opinion	Historical Population Proportion
Primary	9 firms	8%	3%	236 (8.9%)	11.0%**	61.4%**
Manufacturing	28 firms	21%	15%	737 (27.9%)	6.0%**	42.0%
CITU	20 firms	23%	12%	522 (19.7%)	18.6%**	50.6%**
FIRE	13 firms	15%	11%	333 (12.6%)	24.0%**	37.7%
Trade	15 firms	17%	35%	392 (14.8%)	15.3%**	27.7%**
CPBS	15 firms	16%	24%	425 (16.1%)	10.6%**	37.5%
Firm Size	BOS Targets	GDP Share (2000-2006)	D&B Database (Aug 07)	# of obs. (% of sample)	Historical Balance of Opinion	Historical Population Proportion
Small	33	42.9%	85%	777 (29.4%)	17.8%**	42.5%
Medium	33	17.1%	13%	831 (31.5%)	8.8%**	43.7%
Large	34	40.0%	2%	1037 (39.2%)	13.6%**	40.3%
Total	100	100%	100%	2645	13.3%	42.0%

** Statistically different from the overall balance of opinion at the 99% confidence level or greater.

^a Industry targets have changed over time to match the detailed 3-digit NAICS codes composition of the economy. These are approximate numbers.

* Industry aggregates are defined by NAICS as follows: Primary (100 to 219); Manufacturing (300 to 339); Trade (410 to 479); Construction, information and cultural industries, transportation and utilities (CITU 220 to 239 and 480 to 519); Finance, insurance, real estate, and leasing (FIRE 520 to 539); Commercial, personal, and business services (CPBS >540).

The definition of firm size depends on the metric used (e.g., employment, assets, capital, etc.) and the choice of thresholds. The BOS categories are based on employment, with medium-sized firms having between 100 and 499 employees. By this metric, the BOS over-weights medium-sized firms relative to what the Survey of Employment, Payroll and Hours (SEPH) would suggest.

The fourth column in Table B1 report the average balances of opinion for past sales over the sample period. Here we treat the data set as a cross-section (ignoring, for the moment the time-varying nature of the data²⁷). What this illustrates is that there is variation between the regional, industry, and firm-size responses over the sample horizon. All of the balances of opinion are statistically different from one another at the 99 per cent confidence level (indicated by **), except for two: those for future sales for the Trade sector and medium-sized firms. *A priori*, we would expect variation between regional and industry aggregate balances owing to regional- and industrial-specific shocks. An explanation for variation between firm sizes, however, is less obvious.

The variation in Table B2 suggests that imposing quota constraints at the regional, industry and firm-size level may improve accuracy relative to a random selection methodology. We note that this is just a starting point. In order to test the robustness of our results, we would also want to consider pseudo populations with non-typical distributions (e.g., a quarter in which there was a large regional or sector shock). We propose to do this as an extension to the initial results.

²⁷ At any rate, the quotas change little through time, so there should be no inter-group bias based on time-varying factors.

Table B2: Response Distribution for the Balance of Opinion

Atlantic	Small			Medium			Large		
	+1	0	-1	+1	0	-1	+1	0	-1
Primary	0.32	0.26	0.42	0.63	0.00	0.38	0.33	0.56	0.11
Manufacturing	0.36	0.36	0.28	0.50	0.18	0.32	0.35	0.45	0.19
CITU	0.39	0.22	0.39	0.24	0.18	0.59	0.38	0.50	0.12
FIRE	0.38	0.29	0.33	0.31	0.19	0.50	0.50	0.38	0.13
Trade	0.57	0.14	0.29	0.53	0.21	0.26	0.40	0.60	0.00
CBPS	0.48	0.26	0.26	0.36	0.36	0.29	0.36	0.18	0.45
Quebec	Small			Medium			Large		
	+1	0	-1	+1	0	-1	+1	0	-1
Primary	0.33	0.50	0.17	0.50	0.00	0.50	0.36	0.43	0.21
Manufacturing	0.31	0.34	0.34	0.35	0.27	0.38	0.37	0.27	0.36
CITU	0.35	0.35	0.30	0.52	0.22	0.26	0.62	0.13	0.26
FIRE	0.63	0.17	0.21	0.40	0.40	0.20	0.58	0.26	0.16
Trade	0.50	0.36	0.14	0.36	0.18	0.45	0.49	0.23	0.28
CBPS	0.39	0.36	0.24	0.60	0.10	0.30	0.38	0.28	0.34
Ontario	Small			Medium			Large		
	+1	0	-1	+1	0	-1	+1	0	-1
Primary	0.43	0.29	0.29	0.50	0.50	0.00	0.44	0.11	0.44
Manufacturing	0.27	0.29	0.45	0.49	0.17	0.34	0.43	0.20	0.37
CITU	0.41	0.22	0.37	0.47	0.27	0.25	0.40	0.16	0.44
FIRE	0.30	0.40	0.30	0.39	0.30	0.30	0.45	0.20	0.35
Trade	0.48	0.24	0.29	0.32	0.18	0.50	0.41	0.13	0.46
CBPS	0.45	0.18	0.36	0.38	0.08	0.55	0.36	0.28	0.36
Prairies	Small			Medium			Large		
	+1	0	-1	+1	0	-1	+1	0	-1
Primary	0.56	0.00	0.44	0.50	0.10	0.40	0.40	0.28	0.32
Manufacturing	0.48	0.40	0.12	0.48	0.14	0.38	0.48	0.09	0.43
CITU	0.60	0.20	0.20	0.52	0.15	0.33	0.52	0.28	0.21
FIRE	0.47	0.33	0.20	0.60	0.20	0.20	0.50	0.20	0.30
Trade	0.48	0.19	0.33	0.43	0.05	0.52	0.69	0.14	0.17
CBPS	0.52	0.26	0.22	0.35	0.25	0.40	0.59	0.29	0.12
British Columbia	Small			Medium			Large		
	+1	0	-1	+1	0	-1	+1	0	-1
Primary	0.45	0.32	0.23	0.60	0.15	0.25	0.40	0.17	0.43
Manufacturing	0.59	0.09	0.31	0.46	0.14	0.41	0.42	0.15	0.42
CITU	0.60	0.19	0.21	0.44	0.17	0.39	0.49	0.24	0.27
FIRE	0.69	0.14	0.17	0.57	0.29	0.14	0.61	0.17	0.22
Trade	0.55	0.24	0.21	0.48	0.24	0.29	0.45	0.27	0.27
CBPS	0.51	0.15	0.33	0.39	0.15	0.45	0.72	0.06	0.22

Table B3: Response Distribution for Population Proportion

Atlantic	Small		Medium		Large	
	+1	0	+1	0	+1	0
Primary	0.45	0.55	0.62	0.38	0.55	0.45
Manufacturing	0.36	0.64	0.41	0.59	0.61	0.39
CITU	0.26	0.74	0.53	0.47	0.65	0.35
FIRE	0.52	0.48	0.63	0.37	0.50	0.50
Trade	0.24	0.76	0.32	0.68	0.00	1.00
CBPS	0.55	0.45	0.46	0.54	0.18	0.82
Quebec	Small		Medium		Large	
	+1	0	+1	0	+1	0
Primary	0.33	0.67	0.67	0.33	0.43	0.37
Manufacturing	0.39	0.61	0.38	0.62	0.34	0.66
CITU	0.32	0.68	0.37	0.63	0.44	0.56
FIRE	0.22	0.78	0.40	0.60	0.05	0.95
Trade	0.43	0.57	0.30	0.70	0.21	0.79
CBPS	0.38	0.62	0.48	0.52	0.28	0.72
Ontario	Small		Medium		Large	
	+1	0	+1	0	+1	0
Primary	0.57	0.63	1.00	0.00	0.44	0.56
Manufacturing	0.34	0.66	0.47	0.53	0.35	0.65
CITU	0.59	0.41	0.59	0.61	0.33	0.67
FIRE	0.20	0.80	0.30	0.70	0.23	0.77
Trade	0.29	0.71	0.14	0.86	0.24	0.76
CBPS	0.48	0.52	0.23	0.77	0.21	0.79
Prairies	Small		Medium		Large	
	+1	0	+1	0	+1	0
Primary	0.67	0.33	0.65	0.35	0.75	0.25
Manufacturing	0.36	0.64	0.53	0.47	0.52	0.48
CITU	0.68	0.32	0.52	0.48	0.58	0.42
FIRE	0.73	0.27	0.40	0.60	0.67	0.33
Trade	0.14	0.86	0.33	0.67	0.37	0.63
CBPS	0.35	0.65	0.55	0.45	0.35	0.65
British Columbia	Small		Medium		Large	
	+1	0	+1	0	+1	0
Primary	0.55	0.45	0.89	0.11	0.47	0.53
Manufacturing	0.59	0.41	0.53	0.47	0.46	0.54
CITU	0.57	0.43	0.39	0.61	0.62	0.38
FIRE	0.43	0.57	0.29	0.71	0.29	0.71
Trade	0.33	0.67	0.38	0.62	0.27	0.73
CBPS	0.46	0.54	0.42	0.58	0.44	0.56

Table B4: Regional, Industry and Firm Size Weights					
Atlantic					
	Small	Medium	Large	TOTAL	
Primary	19	3	0	22	
Manufacturing	75	20	5	100	
CITU	106	17	4	127	
FIRE	98	8	1	107	
Trade	317	13	1	331	
CBPS	205	16	1	222	Regional Total: 909
Quebec					
	Small	Medium	Large	TOTAL	
Primary	26	6	2	34	
Manufacturing	323	138	19	480	
CITU	204	55	11	270	
FIRE	242	31	11	284	
Trade	851	85	4	940	
CBPS	618	83	11	712	Regional Total: 2720
Ontario					
	Small	Medium	Large	TOTAL	
Primary	62	5	1	68	
Manufacturing	681	336	59	1076	
CITU	525	137	35	697	
FIRE	583	92	32	707	
Trade	1883	168	20	2071	
CBPS	1298	165	42	1505	Regional Total: 6124
Prairies					
	Small	Medium	Large	TOTAL	
Primary	182	36	13	231	
Manufacturing	220	69	15	304	
CITU	325	69	13	407	
FIRE	282	27	5	314	
Trade	902	49	2	953	
CBPS	539	59	11	609	Regional Total: 2818
British Columbia					
	Small	Medium	Large	TOTAL	
Primary	58	6	0	64	
Manufacturing	191	58	5	254	
CITU	224	33	5	262	
FIRE	178	21	3	202	
Trade	661	27	3	691	
CBPS	372	31	5	408	Regional Total: 1881
					National Total: 14452

Table B5: Non-Response Frequencies (Based on Regional Office records)			
Region	% Average	% Min	% Max
Atlantic	30.4	0.0	52.9
Quebec	47.8	25.0	70.8
Ontario	45.2	24.2	64.8
Prairies	41.8	23.1	67.7
BC	22.3	4.8	39.4
Industry	% Average	% Min	% Max
Primary	29.5	0.0	61.5
Manufacturing	46.3	34.1	59.1
CITU	41.5	18.5	62.3
FIRE	41.8	18.8	61.9
Trade	31.7	7.7	50.0
CPBS	37.0	20.0	67.3
Firm Size	% Average	% Min	% Max
Small	43.7	25.0	56.1
Medium	47.4	39.1	55.4
Large	32.2	17.3	50.0