

# Consumption, Income, and Wealth Inequality in Canada\*

Matthew Brzozowski  
*York University*  
[brzozows@yorku.ca](mailto:brzozows@yorku.ca)

Martin Gervais  
*University of Southampton*  
and *IFS*  
[gervais@soton.ac.uk](mailto:gervais@soton.ac.uk)

Paul Klein  
*University of Western Ontario*  
[pklein2@uwo.ca](mailto:pklein2@uwo.ca)

Michio Suzuki  
*Bank of Japan*  
[michio.suzuki@boj.or.jp](mailto:michio.suzuki@boj.or.jp)

June 9, 2009

## Abstract

In this paper, we document some features of the distribution of income, consumption and wealth in Canada using survey data from many different sources. We find that wage and income inequality have increased substantially over the last 30 years, but that much of this rise was offset by the tax and transfer system. As a result, the rise in consumption inequality has been relatively mild. We also document that wealth inequality has remained fairly stable since 1999. Using both confidential data and publicly available data, we are able to gauge the extent to which the publicly available data conceals aspects of inequality that confidential data reveals.

*Journal of Economic Literature* Classification Numbers: D12; D31

*Keywords:* Income Inequality; Consumption Inequality; Wealth Inequality

---

\*We would like to thank Jim Davies and David Green for their many comments and suggestions on this project, as well as participants at the Philadelphia conference on Heterogeneity in Macroeconomics and the Bank of Japan for useful comments. Gervais and Klein gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada. The views expressed in this paper are those of the authors and do not necessarily reflect the official views of the Bank of Japan.

# 1 Introduction

In this paper, we document some salient facts concerning the distribution of wages, hours worked, income, consumption and wealth in Canada. We use a variety of data sources to do so, as described in section 2. In general, these datasets have two access levels, one publicly available (Public-Use Files, or PUF) and one only available through Research Data Centres (RDC) administered by Statistics Canada. While our main conclusions about means (section 3) and inequality (section 4) are drawn from RDC data, we discuss in section 5 the extent to which noise introduced by Statistics Canada to protect individuals' identity distorts various measures of inequality.

Income inequality over the last 30 years or so has risen quite substantially in Canada. Wage inequality, as measured by the variance of log wages, doubled from 1977 to 2005. The level of wage inequality as well as its trend, however measured, are remarkably similar to the U.S. over the same period of time (see [Heathcote et al. \(2009\)](#) in this issue). In Canada, most of this rise occurred within skill groups, as the skill premium remained fairly stable until the mid-1990's, in contrast to the U.S. where the skill premium has been rising consistently since 1980.<sup>1</sup> As we move towards more inclusive measures of income, such as family earnings or total income before taxes and transfers, U.S. and Canadian inequality remain noticeably similar.

Perhaps the most striking finding of this study is the remarkable role played by the tax and transfer system both to compress inequality and to absorb changes in before-tax inequality. Disposable income inequality, as measured by the variance of the log, was essentially flat from 1977 until 1990, a period over which the variance of log pre-government income increased by more than 10 points. Although disposable income inequality has increased since 1990, its rise pales in comparison to that of pre-tax income inequality.<sup>2</sup> While transfer payments—the main ones being, in order of importance in mitigating income inequality: social assistance, unemployment benefits, and various child benefit programs—are mainly responsible for absorbing changes in pre-government income inequality, both taxes and transfers play a signifi-

---

<sup>1</sup>[Boudarbat et al. \(2006\)](#) point out, however, that the rise in the education premium in Canada starts much earlier once experience is controlled for.

<sup>2</sup>It should be noted that the tax and transfer system appears to compress inequality especially at the bottom of the distribution, as evidenced by much smaller movements in the variance of log disposable income than the Gini coefficient.

cant role in compressing income inequality. Given the evolution of disposable income inequality, it should come as no surprise that consumption inequality also rose only moderately over our sample period.

An interesting feature of income inequality in Canada is that its evolution revolves around recessions. From 1976 to 2005, the Canadian economy experienced only two recessions: one at the beginning of the 1980's and one at the beginning of the 1990's. Unlike the U.S., Canada did not experience a recession at the turn of the century. During each recession, wage and income inequality rose substantially, and the declines that followed were not sufficient to offset the rise, resulting in more inequality over time. While a similar pattern can be detected for disposable income inequality, the movements are much milder and, at least in the 1980's, the rise during the recession was fully offset thereafter.

The previous findings also bear on patterns of inequality over the life-cycle, in the sense that the age-profile of inequality is much flatter for disposable income (and consumption) than for pre-government measures of income. Interestingly, wage inequality increases almost linearly over the life-cycle, suggesting the presence of highly persistent wage shocks (see [Storesletten et al. \(2004\)](#)). However, the rise for earnings inequality tends to start later in life, and fails to show a clear monotonic (let alone linear) pattern. This lack of linearity leads us to question the validity of a unit root process as the main driving force for earnings, although this specification seems reasonable for wages.

Notwithstanding the caveats stated in the previous paragraph, we estimated wage and earnings (unit root) processes from our Canadian data. Our main finding is that a very high fraction of the overall cross-sectional variance and also a high fraction of the risk faced by households is accounted for by the “permanent” component as opposed to the transitory component of the process. However, this result is sensitive to the specification of the statistical model. In the Conclusion, we discuss an approach that would yield different implications in this regard. Irrespective of the many caveats, it is interesting to note that in line with our main results discussed above, the tax and transfer system substantially reduces both permanent and transitory earnings risk.

A highly desirable property of distributions is log normality. With that property, a distribution can be fully characterized by its first two moments, and two distributions

are unambiguously comparable with respect to the degree of inequality. In our data, the cross-sectional distribution of consumption is much closer to log-normal than that of income, much like [Battistin et al. \(2007\)](#) find in U.S. and U.K. data. Interestingly, but perhaps not surprisingly given our discussion above, disposable income is also found to be more log-normal than pre-government income.

The main results of this study are closely related to those found in [Frenette et al. \(2007\)](#). Using income data from the Census, they find that the 1980's were characterized by a strong rise in before-tax income inequality, but that most of that rise was absorbed by the Canadian tax and transfer system, with the result that after-tax (and transfer) income inequality remained constant. They further report that while before-tax inequality also increased in the 1990's, but that the tax and transfer system did not offset this rise. As a result, after-tax income inequality also rose in that decade, albeit not to the same extent. While our findings are similar, we choose to stress how small the increase in disposable is relative to the increase in before-tax income inequality, as opposed to their emphasis on the fact that disposable income inequality did rise in the 1990's.

Now the reason why [Frenette et al. \(2007\)](#) use Census data rather than survey data is mainly because they doubt the validity of results obtained through survey data. Indeed, [Frenette et al. \(2006\)](#) show that income inequality trends in survey data are inconsistent with both Census and tax data. However, perhaps because of these inconsistencies, Statistics Canada implemented a revision to the weights in survey data, mainly in order for the surveys to be consistent with information (from the Canadian Revenue Agency) on wages and salaries.<sup>3</sup> Although such a revision almost necessarily entails distorting other aspects of the data (such as employment to population ratios) and brings about a break in the series because the revision was only applied retroactively to 1990, it seems to have brought income inequality results closer to those that emerge from Census data. Indeed, a look at income inequality from Public-Use Files, which for the most part still feature the pre-revision weights, confirms that the choice of [Frenette et al. \(2007\)](#) to use Census data was warranted. Our results from Public-Use Files look much like theirs from survey data, and when we compare those results to the results emerging from the revised weights, we are

---

<sup>3</sup>The details of this revision and its impact on survey income data can be found in [Lathe \(2005\)](#).

inclined to support their view that the old survey weights led to a distorted depiction of the evolution of income inequality in Canada.<sup>4</sup>

While using Census data may have advantages over survey data, it also raises major concerns, of which [Frenette et al. \(2007\)](#) are of course well aware. The most serious drawback of using Census data is that it contains no information on taxes. To remedy this problem, [Frenette et al. \(2007\)](#) impute taxes using administrative tax data. Essentially, they regress taxes paid on income and characteristics in tax data, and use it to predict taxes for each Census family. Another problem with Census data is that these data are only collected every 5 years, making it difficult to capture events that occur at higher frequencies, such as recessions, which we find to be central to the evolution of income inequality over the last 30 years. For these reasons, we prefer to use the survey data with the revised weights rather than the Census data.

## 2 Sources of Data

### 2.1 Datasets

As is typically the case, Canada does not have a single survey from which information on consumption, income and wealth is available. [Table 1](#) lists the data sources with the main variables used and the sample period covered by each data set, which we briefly describe below (see [Appendix A](#) for details).

We take consumption data from the *Survey of Family Expenditures* (FAMEX) for 1969, 1978, 1982, 1986, 1992, and 1996 and the *Survey of Households Spending* (SHS) for 1997–2005.<sup>5</sup> In addition to consumption, both surveys contain data on income and standard characteristics (other than education from 1997 to 2003) at the *household* level, or something close to a household (similar to consumer units in the Consumer Expenditure Survey for the U.S.). See [Appendix A.1](#) for details.

For earnings and hours worked, we use data from the *Survey of Consumer Finances* (SCF) for 1977 and 1979–1997 and the *Survey of Labour and Income Dynam-*

---

<sup>4</sup>Our results are not directly comparable, however, as they use a very different sample and a different way to equalize earnings.

<sup>5</sup>FAMEX data are also available for 1974, 1984, and 1990. However, since the survey covers only urban areas and have a smaller sample for those years, we do not include the data in our analysis.

Table 1: Data Sources

Data Source	Variables	Sample Period
FAMEX	income, expenditure	1969, 1978, 1982, 1986, 1992, 1996
SHS	income, expenditure	1997–2005
SCF	income, hours worked	1977, 1979–1997
SLID	income, hours worked	1996–2005
SFS	income, wealth	1999, 2005

*ics* (SLID) for 1996–2005.<sup>6</sup> While the SCF provides only cross-sectional data, the SLID has a panel dimension. Specifically, the SLID consists of overlapping samples interviewed in six consecutive years, with new waves introduced in 1993, 1996, 1999, 2002, and 2005. In any given year since 1996, then, there are two waves of data available. Because the sample size of the SLID is relatively small from 1993 to 1995, the SCF is deemed to contain better information for these years. Accordingly, we do not use SLID data prior to 1996. See Appendix A.2 for details.

Wealth data is hard to come by in Canada. The SCF had a wealth supplement in 1977 and 1984, called the ADSCF, and the *Survey of Financial Securities* (SFS) was introduced in 1999 and ran again in 2005. Because of inconsistencies between the two Surveys, we only use the SFS in this paper.<sup>7</sup> The unit of analysis in the SFS is the *economic family*. This survey offers standard balance sheet data from which net worth can be constructed, as well as income and some characteristics of the economic family and the head of the family. See Appendix A.3 for details.

## 2.2 RDC Data versus Public-Use Files

Some of the data sets described in Section 2.1 have two access levels: one publicly available and one only available within the confines of Research Data Centres (RDC) located at various institutions throughout the country and administered by Statistics Canada. While the data that we made available emanate from Public-Use Files, most results contained in the paper emanate from proprietary or RDC data. The

<sup>6</sup>We do not use SCF data for 1978 because data on usual hours worked are missing for that year. SCF data was also collected in the early 1970’s. However, we chose not to use years prior to 1976 for consistency reasons. Appendix A.2 documents the inconsistencies in detail.

<sup>7</sup>Morissette et al. (2002) outline ways to make the ADSCF compatible with the SFS.

main advantage of using RDC data is that these data are not modified in any way or, if they are, a flag identifying observations that have been modified is supplied in RDC data.

The severity of distortions introduced in public-use data varies across surveys. At one extreme, public-use files for consumption data (FAMEX/SHS) are for our purposes identical to RDC data. At the other extreme the panel dimension of SLID data is only available through RDC.<sup>8</sup> But for income and wealth data, Public-Use Files have some characteristics removed (SCF/SLID) and data have been shuffled (SLID) to preserve the anonymity of some of the respondents, mainly households at the top of the distribution. Finally, wealth data (SFS) from our Public-Use Files suffers from the typical top-coding problem that work with many survey data must confront. In Section 5 we investigate the extent to which using Public-Use Files (PUF) can replicate results obtained from the original, RDC data.

## 2.3 Survey Weights

The sampling frames of all the surveys used in this study are based on the Labour Force Survey (LFS), a monthly survey equivalent of the Current Population Survey (CPS) in the U.S.. Initially, each individual in the LFS sample is given a weight equal to the inverse of the probability of selection. These weights are then adjusted to account for non-response and to match census estimates (or projections) for various age-sex groups by province and major sub-provincial areas.

Our income surveys are typically conducted as a supplement to the April (SCF) or January (SLID) LFS, much like the March Supplement to the CPS. Since both SCF and SLID consist of sub-samples of the LFS, the weights for these surveys are initially based on LFS weights. These weights are first adjusted for non-response in SCF/SLID (but not for LFS non-response, as this has already been taken care of in the LFS weights), that is, the weight of non-respondents are distributed to respondents. The non-response adjusted weights are then calibrated to ensure that estimates on relevant characteristics of the population (age, sex, province, family size, and household size) match up with census data (or projections from a recent census).

---

<sup>8</sup>The Cross-National Equivalent Files (CNEF) prepared by the Department of Policy Analysis and Management at Cornell University contain a version of the panel aspect of SLID.

Because of contradicting evidence on the degree and evolution of income inequality coming from different sources of data (tax data, census data, and survey data), as evidenced for instance by the work of [Frenette et al. \(2006\)](#), Statistics Canada recently revised their strategy to calibrate the survey weights. This new strategy adds wages and salary to the set of targets to calibrate survey weights. The targets themselves come from T4 slips (equivalent to W4 slips in the U.S.), which are employer remittances to Canada Revenue Agency (formerly Revenue Canada) for payroll tax purposes. This new strategy to impute weights is currently used for all surveys that we use in this study, and has been implemented retroactively to 1990 for SCF/SLID, and to 1996 for FAMEX/SHS. All SFS data uses this new methodology as well (see [Lathe \(2005\)](#) for more details).

The result of this new calibration strategy, documented in [Rea and Greenlee \(2005\)](#), is that more weight is given to observations with low/zero wages and salary as well as to observations with high wages and salary. The reason is that a disproportionate fraction of respondents are from the middle of the wages and salary distribution. Practically, as far as SCF data is concerned, this revision implies a higher level of inequality in wages and salary, which carries over to pre-tax income inequality and to a lesser extent to after-tax or disposable income inequality at the individual level. Ultimately, this revision had a similar impact on family/household level income inequality. This revision and its impact should therefore be kept in mind when interpreting the evolution of inequality over time in the rest of the paper.

While the new strategy is also used for consumption (FAMEX/SHS) and wealth (SHS) data, no impact can be detected for these data either because all years of data have been revised (SFS) or because the survey was not conducted sufficiently frequently (FAMEX) to produce a noticeable break. It should be noted, however, that while the consumption surveys are also based on the LFS, 10 to 15 percent of SFS samples is drawn from geographic areas in which a large proportion of households had what was defined as “high-income”, while the rest of the sample is based on LFS sampling. In other words, like the Survey of Consumer Finances in the U.S., the SFS attempts to over-sample the rich in order to improve wealth distribution estimates.



## 2.4 Sample Selection and Statistics

The details of our samples are contained in [Appendix B](#). Each sample used in this study is labeled in this appendix. We refer to these labels throughout the paper. Briefly, we restrict our samples to households whose head is at least 25 years or age and no older than 60.<sup>9</sup> We also exclude observations with missing characteristics (gender, education, family type, marital status, and province of residence). For income data, we exclude any observation whose income was missing or imputed, as well as observations with a wage rate less than half the minimum wage. We also exclude observations with positive earnings and zero hours worked. For consumption data, we exclude part-year households and households with annual food expenditure less than 400 Canadian dollars.<sup>10</sup> In addition, since the territories (Northwest Territories, Nunavut, and Yukon) were not covered by FAMEX surveys, we exclude households from the territories in SHS as well. Finally, for wealth data we exclude outliers and families with missing wealth information. After these restrictions have been imposed, we only exclude observations as necessary to compute the statistics of a particular figure or table—typically excluding non-positive values of a particular variable.

Table 2 compares our samples along broad dimensions for three different years.<sup>11</sup> Except for particularly low fractions of people with a university degree in FAMEX data, and a particularly high fractions of married couples in FAMEX/SHS data, our data sets are fairly consistent with one another.

---

<sup>9</sup>The head of household is defined as follows. 1. If the household has just one member, that is the head. 2. If the household has several members, take the set of eldest males between 25 and 60. If this set is a singleton, the only member is the head. 3. If the set is empty, take the set of eldest females between 25 and 60. If that set is a singleton then the only member is the head. 4. If at this stage we have a set with more than one member (either a set of eldest males or a set of eldest females), choose the member with the highest sum of labour earnings plus the absolute value of business income. If that still doesn't break the tie, appoint as head the member with the lowest personal id number.

<sup>10</sup>A full-year household is defined as having at least one full-year member, i.e. a member who was in the household for the entire year; a part-year household is composed entirely of part-year members.

<sup>11</sup>The samples used for this Table consist of Sample A from [Appendix B](#) for each respective dataset.

Table 2: Summary Statistics

	1986		1996			2005		
	FAMEX	SCF	FAMEX	SCF	SLID	SHS	SLID	SFS
Age	40.76	40.25	41.92	41.57	40.81	43.27	42.79	43.01
Family Size	3.11	2.82	2.87	2.69	2.64	2.85	2.42	2.64
% University	0.14	0.18	0.19	0.22	0.20	0.30	0.26	0.26
With Spouse	0.71	0.65	0.67	0.60	0.59	0.67	0.54	0.62
% Immigrant	0.19	0.19	0.21	0.18	0.17	n/a	0.16	0.20
Observations	7142	14094	7196	16134	13779	9028	10814	<3868

Notes. The 2005 SHS only provides categorized data on the age of the head of the household. We assign each household head the middle value of the age range to which the person belongs, and report the average over the assigned ages.

### 3 Means

The purpose of this section is to compare the evolution of average income and consumption from micro data to their counterparts from national statistics. A similar exercise is performed on employment rates. Before presenting these comparisons, we provide a brief overview of the Canadian economy.

Table 3 displays aggregate statistics for the Canadian economy over the 1976 to 2005 period, the years covered by most of our micro data. Over those 30 years, the population of Canada grew at an average annual rate of 1.1 percent, from 23 to over 32 million. Meanwhile GDP grew at an average annual rate of 2.8 percent, so GDP per head grew at 1.7 percent per year. Both figures are slightly lower than in the U.S., where GDP and GDP per head grew at average rates of 3.0 and 2.0 percent per year, respectively. Similarly, GDP per hour only grew at an average annual rate of 1.2 percent over that period in Canada, compared to 1.5 percent in the U.S.

Figure 1 displays real disposable income per capita as reported in the System of National Accounts (SNA) as well as its implied counterparts constructed from survey data.<sup>12</sup> While the gap between SNA and survey data increased over the years covered by the SCF, it is reassuring that it has been shrinking under SLID. In the last survey

<sup>12</sup>All real series in this section were deflated by the CPI using base year 2003. For all survey data, the full sample was used, that is, no sample restrictions were imposed.

Table 3: Aggregate Statistics: 1976–2005

	from	1976	1976	1980	1985	1990	1995	2000
	to	2005	1979	1984	1989	1994	1999	2005
Population		1.10	1.05	1.09	1.35	1.15	0.92	1.02
Real GDP		2.83	3.81	2.08	3.51	1.52	3.81	2.48
Real GDP Per head		1.73	2.76	0.99	2.15	0.37	2.89	1.46
Real GDP per hour		1.20	0.46	1.77	0.10	1.80	1.64	1.20
Employment rate		0.67	0.63	0.64	0.67	0.66	0.67	0.71

Notes. The first four rows display average annual growth rates of the variable. Employment rates are annual average employment rates for the 20–69 age group over the relevant period.

Sources. Population: Statistics Canada, Table 051-0001; GDP: Statistics Canada, Table 380-0002; Hours worked and employment rates: Labour Force Survey


year available, as shown in the first panel of Table 4, average disposable income was over 90% of the SNA figure, either from SHS or SLID data.

Figure 2 shows that a similar situation prevails for consumption, at least over the SHS years (since 1997), with a gap sitting between 80 and 82 percent (see second panel of Table 4). Furthermore, the ratio of non-durable consumption to income from SHS is consistently only around 5 percentage points lower than the same ratio from SNA data, as shown in the third panel of Table 4 for 2005. Thus, the situation in Canada for consumption data is not as dire as in the U.S., where the gap has been rising consistently over the last 25 years.<sup>13</sup> For example, Slesnick (2001) reports that the ratio of per capita consumption in the Consumer Expenditures Survey (CEX) to the NIPA estimate of Personal Consumption Expenditures was 0.68 in 1980, 0.61 in 1990, and 0.56 in 1995. Similarly, Battistin (2003) finds that the ratio of non-durable expenditures in CEX to the NIPA estimate fell from around 80 percent in 1985 to around 65 percent in the late 1990’s.

Figure 3 displays employment to population ratios at the aggregate level and from survey data over the sample period. It should be noted that the aggregate employment series also comes from a survey—the monthly Labour Force Survey (LFS). Also note

<sup>13</sup>See Crossley and Pendakur (2006) for a discussion of the quality of expenditure survey data in Canada.

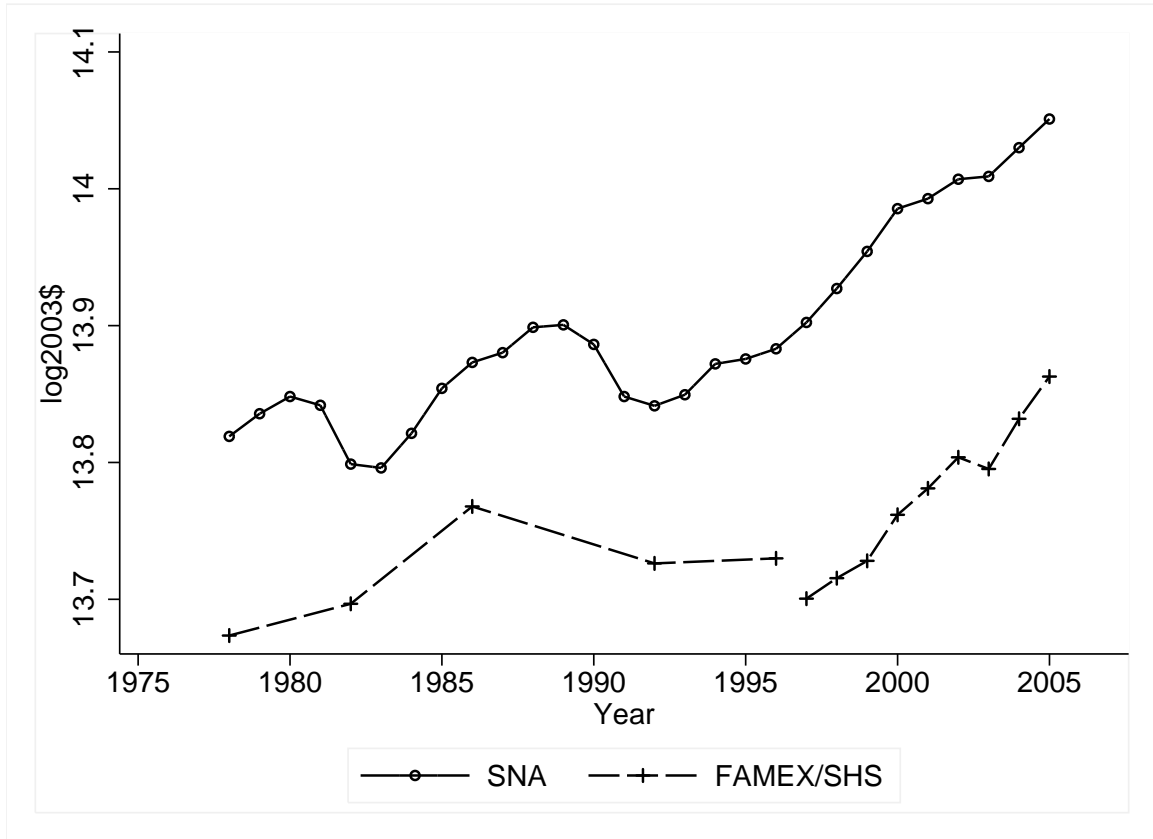
Figure 1: Disposable Income Per-Capita



that while most questions in our income surveys pertain to the year prior to the survey year, employment questions in SCF data pertain to the current month (typically April). By contrast, employment information in SLID is available for every month of the previous year. Furthermore, as mentioned in Section 2.3, both SCF and SLID surveys are conducted as a supplement to the LFS, that is, they consist of subsamples of the LFS.<sup>14</sup> It is therefore surprising to see significant discrepancies between the two series. The discrepancies emanate from two distinct sources. First, the SCF had unusually small samples in 1977, 1979, 1981 and 1984, which explains the

<sup>14</sup>Since the SCF is typically conducted as a supplement to the April LFS, both series in Figure 3 refer to the month of April. For 1976 and 1983, SCF data pertained to the month of May, and so we chose that month for the aggregate series as well for these two years. SLID offers monthly labour force data, so we also used the month of April for the years covered by SLID.

Figure 2: Non-Durable Consumption Per-Capita



anomalies prior to 1990.<sup>15</sup> The post 1990 discrepancies can be explained by the 2003 historical revision of the weights discussed in Section 2.3. As a result of this revision, which was applied retroactively to 1990, higher weights were assigned to low/no earnings individuals as well as high earnings individuals and had the effect of increasing substantially the fraction of the population not in the labour force. This is reflected in Figure 3 as a low employment to population ratio post-1990 in SCF data.<sup>16</sup>

<sup>15</sup>Survey years 1977 and 1979, which contain data pertaining to 1976 and 1978, have other problems—in particular hours worked is missing—so we dropped these years for the rest of the analysis.

<sup>16</sup>This is confirmed by a similar figure (not shown) constructed using weights available prior to the revision, which shows essentially no discrepancy between the SCF and LFS numbers over the 1990's.

Table 4: Survey to Aggregate Ratios (%)

	1978	1986	1996	2005
<u>Mean Disposable Income</u>				
FAMEX/SHS	93.1	90.3	95.7	91.8
SCF	88.3	84.9	83.4	n/a
SLID	n/a	n/a	84.1	91.0
<u>Mean Non-Durable Expenditures</u>				
FAMEX/SHS	86.5	90.0	85.8	82.8
<u>Ratio of Non-Durable Expenditures to Disposable Income</u>				
FAMEX/SHS	49.9	51.6	46.5	48.8
SNA	53.7	51.8	51.9	54.0

## 4 Dimensions of Inequality


### 4.1 Individual Level Inequality

Figure 4 displays four measures of wage inequality.<sup>17</sup> The wage is defined as average hourly earnings, and earnings include wages and salaries as well as the labour share (62%) of income from self-employment.<sup>18</sup> The top-left panel shows that wage inequality, as measured by the variance of log wages, increased substantially in the late 1970's and early 1980's, and somewhat more modestly in the early 1990's and around 2000. The top-right panel of that figure shows that a similar pattern emerges when wage inequality is measured by the Gini coefficient. Over the entire sample, the variance of log wages doubled and the Gini coefficient increased by 9 points, from 0.28 to 0.37. It is interesting to note that while high wage earners have been consistently gaining on the median over the sample period (lower right panel), low wage earners

<sup>17</sup>The sample used in this section consists of Sample B from Tables 12 and 14.

<sup>18</sup>The labour share of income is taken from the SNA; an average is computed over the period 1961:1 to 2002:1. Specifically, the labour share is defined as the ratio of wages, salaries & supplementary labour income plus taxes less subsidies on factors of production and gross domestic product at market prices less accrued net income of farm operators from production less net income of non-farm unincorporated business including rent. The precise data source is Matrix 6520 from CANSIM 1.

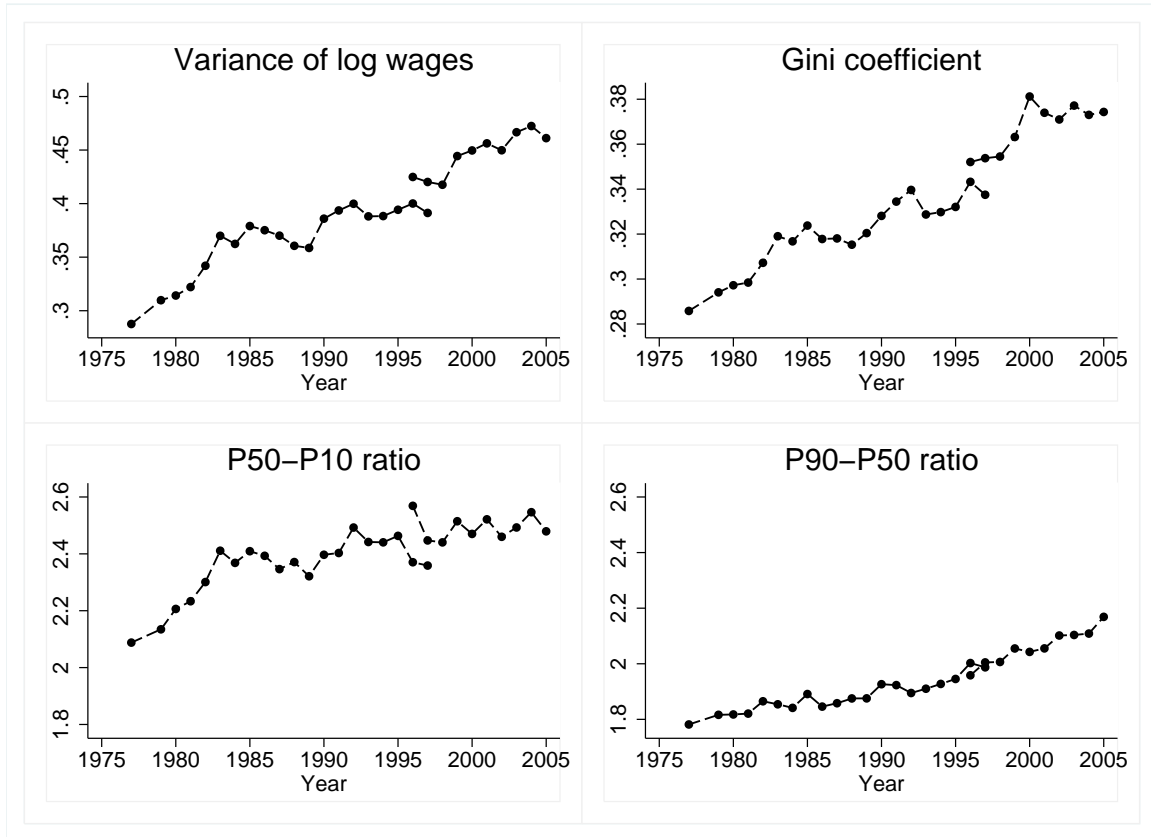
Figure 3: Employment to Population Ratio



falling behind the median is the main reason behind the increase in wage inequality in the late 1970's and early 1980's (lower left panel). Notice that while inequality tends to increase during recessions, its tendency to go back down in between recessions is very mild.

The bottom-right panel of Figure 5 shows that the above pattern of wage inequality is very similar to the pattern of residual wage inequality—inequality unexplained by observables—suggesting that most of the increase in wage inequality cannot be explained by observables. The top-left panel of Figure 5 shows the education premium, defined as the the average wage of males with at least a university certificate, a diploma, or a bachelor's degree, relative to the average wage of males without any such degrees. Interestingly, the education premium remained fairly constant from the

Figure 4: Basic Inequality in Hourly Wages



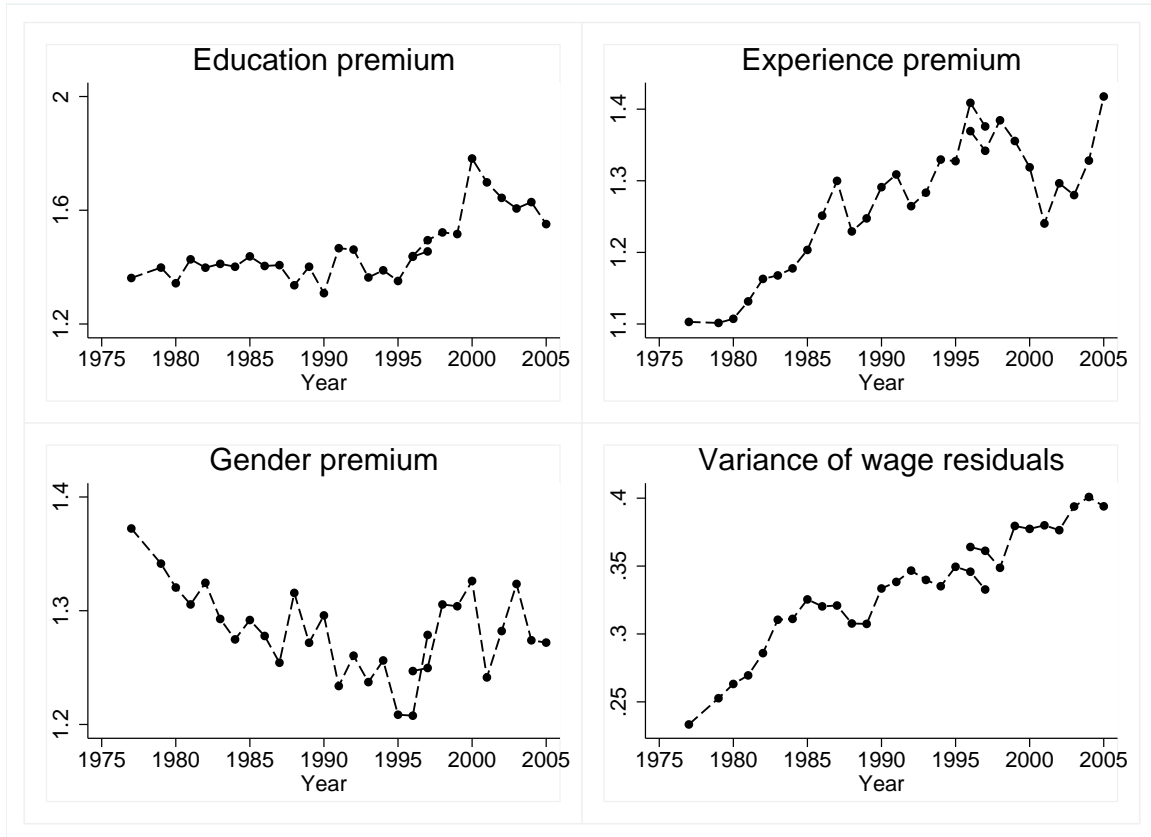
beginning of the sample until the mid 1990's.<sup>19</sup> Since then, the education premium rose from an average of around 40 percent from the beginning of the sample until 1995 to just below 60 percent in 2005. Meanwhile, the fraction of the population with a university degree increased monotonically from 14 percent in 1977 to 28 percent in 2005. These patterns are quite different from those observed in the U.S., where the college premium *and* the supply of college graduates have been rising simultaneously since the late 1970's (Heathcote et al. (2008)). These observations suggest that as opposed to the U.S., the rise in wage inequality in Canada has mainly occurred within skill groups.

The bottom-left panel of Figure 5 displays the gender premium, that is, the average

<sup>19</sup>Boudarbat et al. (2006) point out, however, that the rise in the education premium starts much earlier once experience is controlled for.

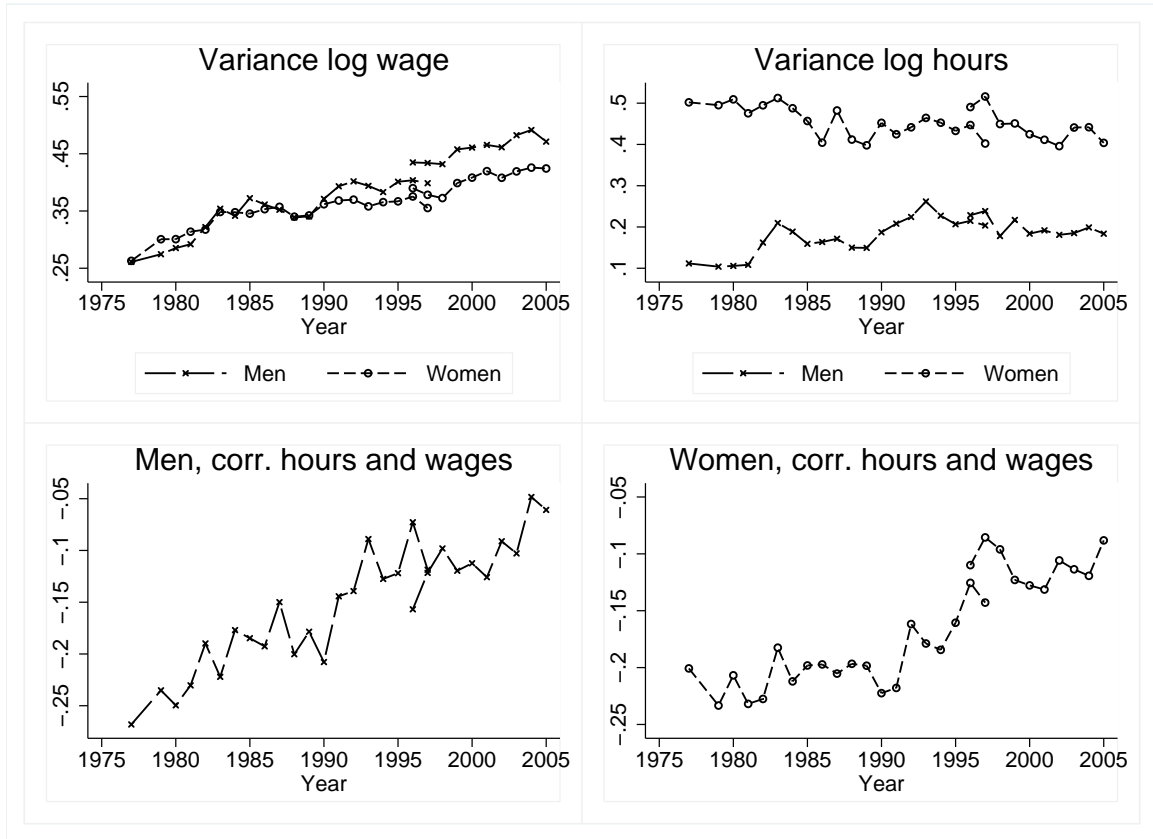


Figure 5: Wage Premia and Residual Wage Inequality



wage of males relative to that of females. The gender premium narrowed considerably until the mid 1990's, from 37 percent in 1977 to just over 20 percent in 1995, but it has failed to decline in more recent years, averaging around 27 percent over the SLID years. As a result, the gender gap in Canada is now close to that in the U.S., but the drop in the gender premium since the late 1970's was more spectacular in the U.S., where the wage premium was at a much higher level (65 percent) than in Canada (37 percent) in the late 1970's ([Heathcote et al. \(2008\)](#)). The top-right panel suggests that the reverse has happened to the experience premium, measured as the average wage of males aged 45–55 relative to those between 25 and 35 years of age. This premium increased until the mid 1990's and declined sharply during the dot-com years of the late 1990's. However, the experience premium came back up over the last 5 years of the sample.

Figure 6: Inequality of labour Supply



The top-left panel of Figure 6 shows that, qualitatively, the evolution of wage inequality for both men and women mimics that of the entire population (top-left panel of Figure 4). The gap between male and female wage inequality, which opens up in the 1990's, may be an artifact of the retroactive revision of survey weights.<sup>20</sup> The last 10 years of SLID data suggest that the variance of log wages for men is around 5 log points higher than that of women.

Inequality of hours worked, as shown in the top-right panel of Figure 6, paints a drastically different picture for men and women. First, as one would expect, the level of hours inequality is much higher for women than for men throughout the sample. Second, inequality of hours worked for men is much more responsive to business cycle

<sup>20</sup>Recall that as a result of the weight revision, more weight was given to observations with low or no income, which may have affected male and female wages differently.

fluctuations than that of women: from 1981 to 1983 the variance of log hours for men increased by 10 points (from 0.11 to 0.21) while it only increased by 3 points for women (from 0.48 to 0.51); similarly, the variance of log hours from 1990 to 1993 increased by 7 points for men (from 0.19 to 0.26) and by 1 point for women (from 0.45 to 0.46).<sup>21</sup> It is also interesting to note that over the entire sample period, the variance of hours worked for women trended down, while that of men has been essentially stable since the recession of the early 1980's. Finally, the bottom panels of Figure 6 shows that the correlation between hours worked and wages for men and women alike has been rising throughout the sample period, approaching zero from below. It is interesting to note that the same pattern is observed in the U.S. (see [Heathcote et al. \(2009\)](#).)

## 4.2 Family Level Inequality

This section pertains to inequality at the *economic family* level, where a family is defined as a group of individuals living at the same address and related by blood, marriage or adoption.<sup>22</sup>

The next two figures display various measures of family earnings inequality, where family earnings is the sum of the family members' wages and salaries plus the labour part (0.62) of self-employment income.<sup>23</sup> Figure 7 shows that observables (family type, education, age and region of residence) do not explain much of the trend in earnings inequality, despite the fact that the variance explained by education has generally trended up over the sample period. On average, observables explain around 23 percent of the variance of equivalized earnings. All measures of inequality in Figure 8 show a large rise over the two recessions: the variance of log equivalized earnings increased by 14 points (from 0.59 to 0.65) from 1981 to 1983 and by 20 points (from 0.59 to 0.79) from 1990 to 1993. Both episodes were marked by the median gaining on the poor and the rich gaining on the median, although the magnitude of

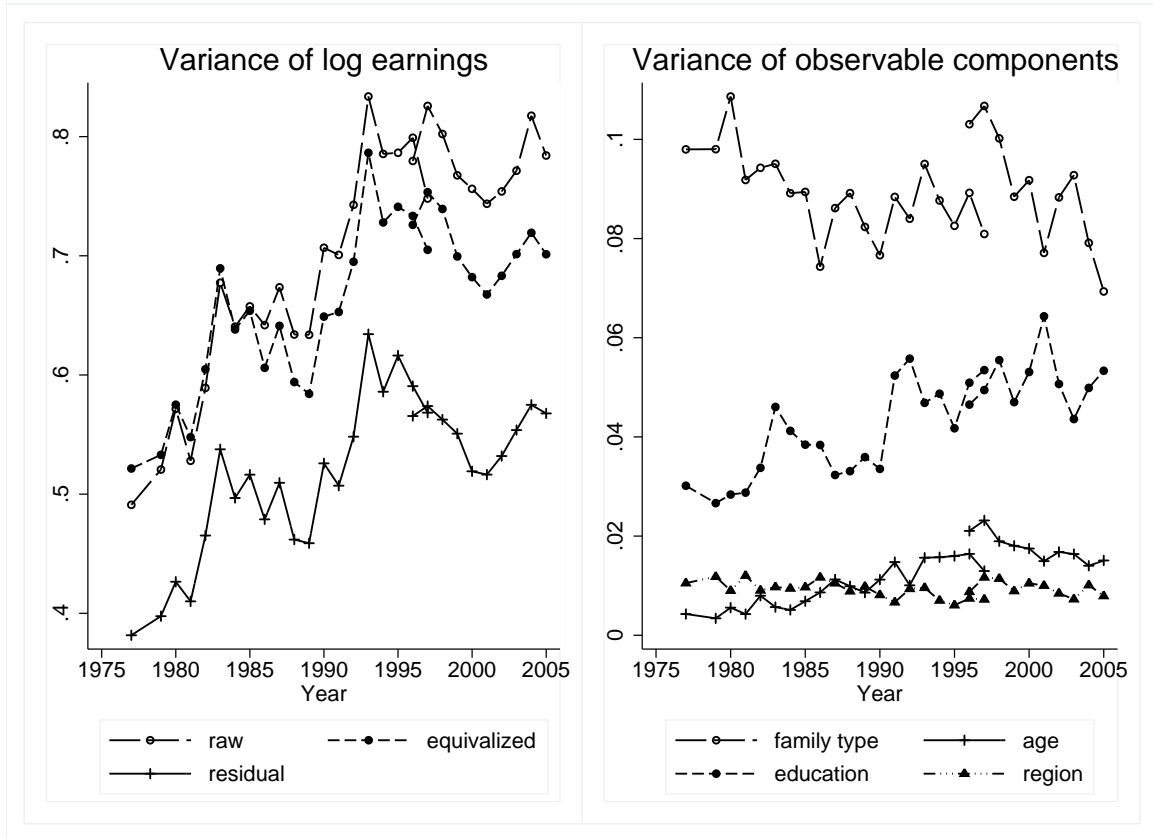
---

<sup>21</sup>We similarly found the following variables to fluctuate more for males than for females over business cycles: the employment to population ratio, unemployment rate, and participation rate for the population aged 15 and above in Canada.

<sup>22</sup>Economic families may differ from households, defined as a person or a group of persons living at the same address.

<sup>23</sup>The sample used for these two Figures consists of Sample D from Tables 13 and 15.

Figure 7: Earnings Inequality and its Decomposition




the former is much more pronounced than the latter. Notice that the pattern around recessions is slightly different from wage inequality (see Figure 4), as family earnings inequality measured by the variance of the log tends to decrease after recessions. Note, however, that the level of the Gini coefficient kept increasing over the last 10 years of the sample. Over the entire sample, the variance of log earnings increased by 18 points (from 0.52 to 0.70) and the Gini coefficient by 7 points (from 0.31 to 0.38).

Figure 9 shows the variance of log income for many measures of income as well as earnings of household heads.<sup>24</sup> The first thing to note here is that the distance between pre-government earnings of the head ( $y^H$ ) and of the family ( $y^L$ ), which is typically around 10 to 12 points. This relatively large distance between these two

<sup>24</sup>The sample used for this Figure consists of Sample C from Tables 13 and 15.


Figure 8: Basic Inequality of Equivalized Earnings



measures of income inequality suggests that there is substantial family insurance. But the most startling feature of this Figure stands out when comparing the variance of log disposable income ( $y^D$ ) with that of various measures of pre-government earnings or income ( $y^L$ ,  $y^{L+}$  and  $y$ ). Not only does disposable income exhibit much less inequality than pre-government earnings/income measures, but the degree of inequality is also much less variable than that of pre-government earnings/income. This suggests both that Canadian policy has been and remains redistributive, and that it smooths cyclical shocks to pre-tax income inequality.

To get some perspective on the redistributive effects of public policy, it is useful to look at the effects of different aspects of policy separately, as illustrated in Fig-

Figure 9: From Wages to Disposable Income




Definitions.  $y^H$ : pre-government earnings of head;  $y^L$ : pre-government family earnings;  $y^{L+}$ : pre-government non-financial income of family;  $y$ : pre-government family income;  $y^D$ : total family disposable income.

ure 10.<sup>25</sup> The most obvious distinction is that between tax policy on the one hand and transfers on the other hand, noting that this distinction is to some extent a matter of arbitrary fiat. The most important government transfer programs in Canada come under three headings: unemployment benefits, social assistance and various child benefit programs. Together these programs are more significant in bringing down inequality than the tax system is. Of the three transfer programs, social assistance

<sup>25</sup>The sample used for this Figure consists of Sample B from Tables 13 and 15 with the bottom 0.5% of pre-government income trimmed. We chose to use a broader sample for this Figure as it is important to include observations with zero earnings in examining the role of government policies such as the unemployment insurance and social assistance. However, because the sample selection is very inclusive, the statistics became sensitive to very few observations, which explains our trimming strategy.

Figure 10: From Pre-Government Income to Disposable Income



Definitions.  $y$ : pre-government family income;  $y + trans$ : pre-government family income plus total transfers;  $y - tax$ : pre-government family income minus total taxes paid;  $y + ui$ : pre-government family income plus unemployment benefits;  $y + socass$ : pre-government family income plus total social assistance transfers;  $y + chldben$ : pre-government family income plus child benefits transfers;  $y^D$ : total family disposable income.

has the largest impact, followed by unemployment benefit. Note that the tax system nevertheless plays an important role in compressing income inequality, as the distance between the bottom two lines on Figure 10 is around 7 to 8 points.

This Figure also shows some noticeable changes over time. The first is that the correlation between pre-government inequality and post-government inequality seems to have disappeared. Post-government inequality exhibits very small cyclical fluctuations after 1996 and those movements that do occur are uncorrelated with the corresponding movements in pre-government earnings and income. Finally, while after-tax inequality remained fairly stable over the first half of our sample period, it increased

quite substantially (8 points) in the second half, from 0.24 in 1989 to 0.32 in 2005.<sup>26</sup>

Table 5 presents more evidence of our last result.<sup>27</sup> This table is also meant to compare our results to those of Frenette et al. (2007), who use income data from the Census to characterize the evolution of income inequality over the 1980–2000 period. They conclude that most of the strong rise in before-tax income inequality over the first decade was absorbed by the tax and transfer system, leaving after-tax (and transfer) inequality unchanged. However, the equally strong increase in before-tax inequality in the 1990’s was accompanied by an increase in after-tax (and transfer) inequality, albeit not of the same size. Our results in Table 5 reinforce that conclusion by reporting increases of much larger magnitude than their estimates from Census data. As we will see in Section 5, the weight revision was instrumental in obtaining these results.

### 4.3 Consumption Inequality

We now turn our attention to consumption inequality, where consumption refers to non-durable expenditures.

Figure 11 shows that consumption inequality exhibits a rising trend.<sup>28</sup> In the “raw” data, where no adjustment has been made for family size, the variance of the log has almost doubled in 30 years. However, this is largely because of the rise in single-member households; the fraction of households that had a single member went up from 8 percent in 1969 to 21 percent in 2005. Once the data have been divided by the number of adult-equivalents (equivalized), the rise in inequality is less pronounced and only amounts to about 3 points, from 0.2086 in 1969 to 0.2446 in 2005.

Figure 12 juxtaposes the changes in consumption inequality as measured in four different ways with the corresponding changes in the inequality of disposable income.<sup>29</sup> The main thing to take away from that figure is that consumption inequality is consistently lower than disposable income inequality. For the variance of the log,

---

<sup>26</sup>This suggests that public policy has become gradually less redistributive in its effects, though not necessarily in its intent, since the beginning of the 1990’s.

<sup>27</sup>The same sample as Figure 10 is used for this Table.

<sup>28</sup>The sample used for this Figure consists of Sample A from Table 11.

<sup>29</sup>The sample used for this Figure consists of Sample B from Table 11.



Table 5: Pre- and After-Tax Income Inequality

year	Pre-Tax Income				After-Tax Income			
	Varlog	50/10	90/50	Gini	Varlog	50/10	90/50	Gini
1980	0.4930	2.4427	1.9946	0.3217	0.2669	1.9509	1.8399	0.2741
1985	0.5858	2.6175	1.9971	0.3396	0.2785	1.9675	1.8025	0.2810
1990	0.5990	2.7109	2.0016	0.3411	0.2622	1.9169	1.8055	0.2708
1995	0.7217	3.0276	2.0719	0.3661	0.2884	1.9554	1.8433	0.2847
2000	0.6846	2.8973	2.0508	0.3703	0.2853	1.9647	1.8466	0.2926
2005	0.6639	3.0021	2.1509	0.3805	0.3188	2.0464	1.9132	0.3058
<u>growth (%)</u>								
1980–2000	44.40	18.58	5.66	15.88	8.00	-2.91	2.04	6.56
1980–1990	21.51	10.98	0.35	6.03	-1.75	-1.74	-1.87	-1.21
1990–2000	18.14	5.99	5.29	9.33	9.92	-1.13	3.96	7.84
1980–1985	18.83	7.16	0.13	5.57	4.37	0.86	-2.03	2.51
1985–1990	2.25	3.57	0.23	0.44	-5.86	-2.58	0.17	-3.64
1990–1995	20.49	11.68	3.51	7.33	9.98	2.01	2.09	5.14
1995–2000	-2.61	-6.00	1.72	1.92	-0.03	-3.14	1.85	2.68
2000–2005	-3.02	3.62	4.89	2.76	11.75	4.16	3.61	4.53


Notes. All changes that involve years over which SCF and SLID overlap are computed as the sum of the change in SCF until 1996 and the change in SLID from 1996.

consumption inequality is typically 10 to 15 points lower than disposable income inequality. Also, at least for the variance of the logs, the two tend to move up and down together, although consumption inequality is less volatile. Note, however, that these conclusions are based on relatively few years of observation.

#### 4.4 Inequality over the Life-Cycle

The evidence on life-cycle profiles of the variance of hours, wages, earnings and consumption is reported in Figure 13, where we control for time effects and Figure 14,


Figure 11: Consumption Inequality and its Decomposition



where we control for cohort effects.<sup>30</sup> Interestingly, both figures are qualitatively similar, although controlling for cohort effects systematically leads to steeper inequality profiles—albeit not to the same extent as with U.S. data (see [Heathcote et al. \(2009\)](#)). For wages, the picture is clear: variance increases monotonically over the life cycle. For raw earnings, the increase is more or less monotonic, but the profile is convex, providing some support for the “heterogenous income profile” (see [Güvenen \(2009\)](#)). For equivalized earnings, there is no monotonic increase whether we control for time or cohort effects, the pattern is instead J-shaped. A reasonable conjecture is that earnings inherit this J-shape as a combination of linear log-wage variance profile and

<sup>30</sup>The samples used in these Figures consists of Sample B of Tables 12 and 14 for individual level income data, and Sample D of Tables 13 and 15 for family level income data. For consumption data, Sample B from Table 11 is used.

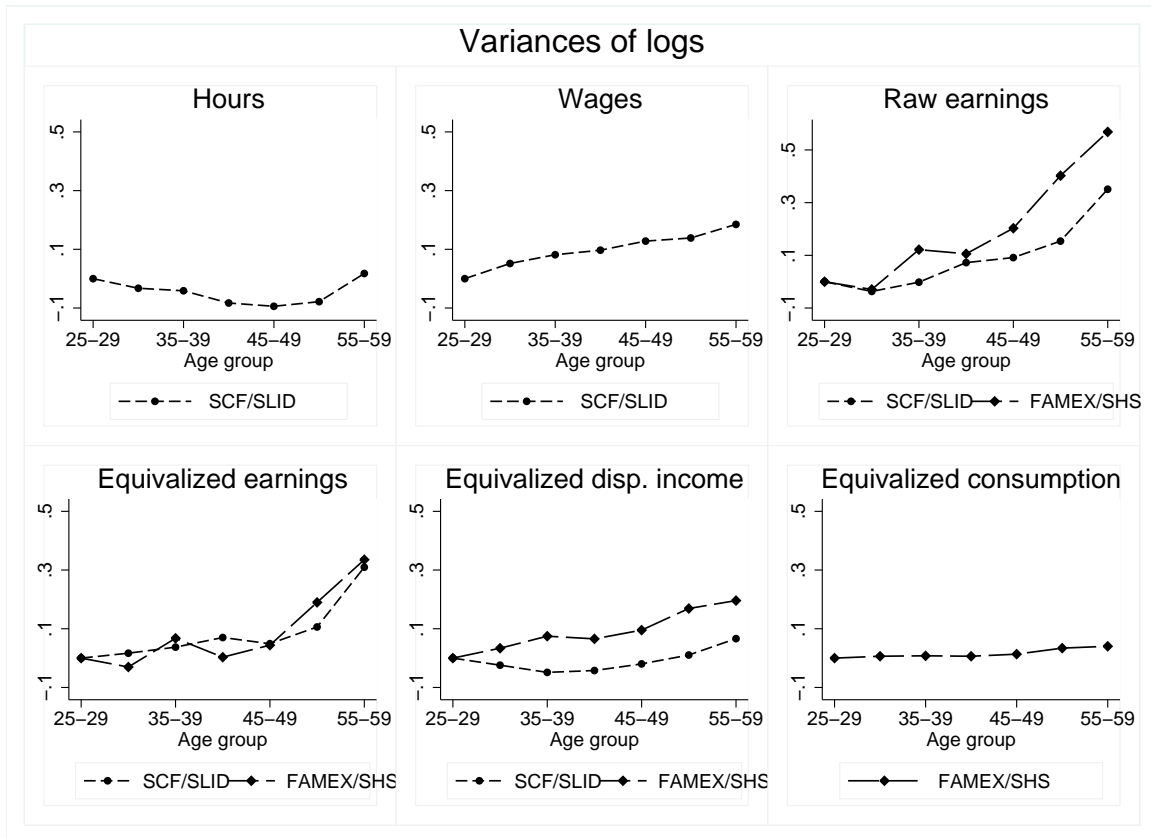
Figure 12: From Disposable Income to Consumption



the U-shape profile for the variance of hours. Finally, the bottom panels of Figures 13 and 14 show that the variance of both disposable income and consumption exhibit very little changes in inequality over the life-cycle. This is perhaps a reflection of the role played by government policy in compressing inequality which we emphasized previously.

The behavior of earnings inequality over the life cycle is particularly interesting. The lack of a clear monotonic (let alone linear) pattern is in sharp contrast to the results constructed and presented in Storesletten et al. (2004) for the U.S.. This matters because it is precisely the linearity of the cross-sectional variance reported there that has motivated the profession to write down a unit root process as the main driving force for earnings; indeed we follow that approach in this paper as well. The

Figure 13: Inequality over the Life-Cycle (Time Effects)

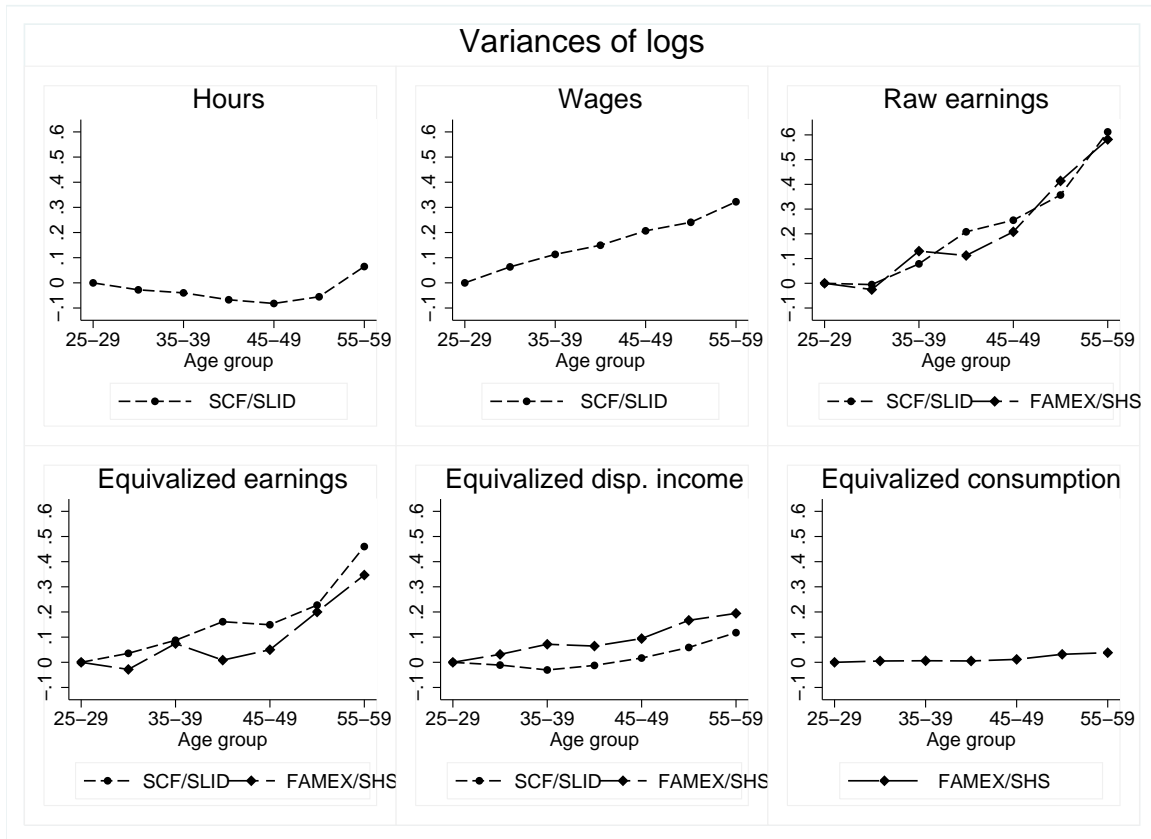


facts reported in Figures 13 and 14 cast some doubt on the appropriateness of that approach, at least when applied to Canadian earnings, though it does seem to hold for wages.

## 4.5 Wealth Inequality

Morissette et al. (2002) document that wealth inequality went up in Canada from 1984 to 1999. The Gini coefficient for equivalized net wealth went up from 0.678 to 0.723, and this difference is statistically significant. We extend this analysis by considering the year 2005 as well. Note that our numbers are not directly comparable to theirs, chiefly because we confine our attention to working-age-headed households.

Figure 14: Inequality over the Life-Cycle (Cohort Effects)



We look at three measures of wealth which we call net financial wealth, net total wealth and net worth, respectively. Net total wealth is defined as net financial wealth plus net residential real estate and business equity. Net worth is simply total assets minus total liabilities.

Table 6 reports wealth to after-tax (disposable) income ratios and Gini coefficients for our income and wealth measures.<sup>31</sup> For all three wealth measures, the wealth to income ratio increased between 1999 and 2005. The Gini coefficients of the three wealth measures are about twice as large as those of disposable income. For all income and wealth measures, the Gini coefficient remained stable between 1999 and 2005. Table 7 reports correlation between income and wealth and the share of the top

<sup>31</sup>The sample used for this Table consists of Sample A from Table 16.

1 and 5 percentile in the income and wealth distributions. For all wealth measures, the correlation with disposable income has increased between 1999 and 2005. The share of the top 1 percentile in the wealth distribution is about three times as high as that for disposable income.

Relative to the U.S., wealth is much less concentrated in Canada. For example, the level of the Gini coefficient, around 0.66, is low relative to the 0.80 reported by [Budría Rodríguez et al. \(2002\)](#) for the U.S. in 1998. More spectacular is the comparison of the fraction of the wealth held by the top of the wealth distribution. While the top 5% in Canada hold about 35% of the wealth, the top 1% in the U.S. hold about about the same fraction of the wealth. The top 1% in Canada only hold about 16% of the wealth. In addition, it appears that income and wealth are more correlated in the U.S. than in Canada: in 1998, the correlation between income and wealth in the U.S. was 0.6, whereas that between disposable income and wealth in Canada was only 0.35 in 1999.<sup>32</sup>

Table 6: Disposable Income and Net Wealth

	1999	2005
Net financial wealth/disposable income	2.8841	3.2328
Net total wealth/disposable income	5.6318	6.8045
Net worth/disposable income	6.3066	7.4307
Gini index of disposable income	0.3355	0.3359
Gini index of net financial wealth	0.7054	0.6992
Gini index of net total wealth	0.6664	0.6636
Gini index of net worth	0.6603	0.6590

Notes. Net total wealth is net financial wealth plus net residential real estate and business equity. Net worth is net total wealth plus other net wealth including vehicles. (Source: SFS-RDC)

<sup>32</sup>Note that for the U.S. the correlation is with respect to pre-government income, as the Survey of Consumer Finances in the U.S. does not have any tax information.

Table 7: Disposable Income and Net Wealth

	1999	2005
<u>Correlation coefficient</u>		
(Net financial wealth, disposable income)	0.3655	0.4377
(Net total wealth, disposable income)	0.3551	0.4146
(Net worth, disposable income)	0.3470	0.4255
<u>Share of the top 1%</u>		
Disposable income	0.0541	0.0583
Net financial wealth	0.1635	0.1449
Net total wealth	0.1665	0.1749
Net worth	0.1568	0.1668
<u>Share of the top 5%</u>		
Disposable income	0.1634	0.1719
Net financial wealth	0.4143	0.3926
Net total wealth	0.3702	0.3817
Net worth	0.3514	0.3639

Notes. Net total wealth is net financial wealth plus net residential real estate and business equity. Net worth is net total wealth plus other net wealth including vehicles. (Source: SFS-RDC)

## 4.6 Wage/Income Processes

In this section, we examine (exogenous) idiosyncratic income risk that individuals and families face by estimating a stochastic process representing the dynamics of individual wages, equivalized family earnings and equivalized family disposable income.<sup>33</sup> For this estimation, we use data from SLID exclusively as only this survey provides panel data on income variables as well as individual and family characteristics. In order to extract the idiosyncratic component of a given variable, we first regress the variable of interest on observables and take residuals from the regression.<sup>34</sup> We control for age, gender, marital status, education, province of residence, immigration status and mother tongue in individual wages, while we control for head's and spouse's (if

<sup>33</sup>The samples used in this section consist of Sample B from Tables 12 and 14 for individual level data, and Sample C from tables 13 and 15 for family level data.

<sup>34</sup>We run the cross-sectional regressions year by year, allowing coefficients to vary over time.

present) characteristics listed above as well as family type and the number of earners for family earnings and disposable income. We model the dynamics of the residuals by the following stochastic process:

$$\begin{aligned} u_{it} &= \alpha_{it} + \varepsilon_{it}, \text{ with } \varepsilon_{it} \sim N(0, \sigma_{\varepsilon,t}^2) \\ \alpha_{it} &= \alpha_{it-1} + \eta_{it}, \text{ with } \eta_{it} \sim N(0, \sigma_{\eta,t}^2), \end{aligned}$$

where  $u_{it}$  represents the residual earnings/income of person/family  $i$  at time  $t$ ,  $\alpha_{it}$  represents the permanent component of the residual earnings/income,  $\varepsilon_{it}$  represents the transitory component, and  $\eta_{it}$  represents innovation to the permanent component. We assume that  $\eta_{it}$  and  $\varepsilon_{it}$  are i.i.d., that  $E(\varepsilon_{is}\eta_{it}) = 0$  for all  $s, t$  and that  $E(\varepsilon_{is}\varepsilon_{it}) = E(\eta_{is}\eta_{it}) = 0$  for all  $s \neq t$ .

In Table 8 we present results from estimating processes for individual wages, equivalized family earnings and equivalized family disposable income. The estimation process is constructed so as to match the autocovariance function of log differences.<sup>35</sup> The result is that a very high fraction of the overall cross-sectional variance and also a high fraction of the risk faced by households is accounted for by the permanent component as opposed to the transitory component.<sup>36</sup> Inspecting the moment conditions, we see that the magnitude of the  $\sigma_{\varepsilon}^2$  estimates depends on the covariance between subsequent log differences. If subsequent log differences are not highly correlated, which they are not, we are led to conclude that transitory shocks account for a small fraction of the overall variance and risk. Indeed, given our estimates on wages, that fraction is below 3 percent for 60-year-olds; for family earnings and family disposable income, the corresponding number is below one percent. Of course, much higher numbers obtain for the young.<sup>37</sup>

In Table 8, we report estimates for  $\sigma_{\eta}^2$  (permanent risk) and  $\sigma_{\varepsilon}^2$  (transitory risk) for post-government disposable family income as well as those for family earnings.

---

<sup>35</sup>See [Heathcote et al. \(2009\)](#) in this volume for details of the estimation procedure.

<sup>36</sup>The estimated variance of the permanent component for wages, at around 0.055, implies much higher inequality over the life-cycle than is reported in Figures 13 and 14.

<sup>37</sup>These fractions are likely to be much lower than we would get from an estimation designed either to match the age profile of cross-sectional variance (as in [Storesletten et al. \(2004\)](#)) or if we sought to match the autocovariance function of levels (as opposed to differences). Thus the conclusion that transitory shocks are unimportant should be taken with a very large grain of salt. Indeed, preliminary evidence suggests that transitory shocks become more important in accounting for the cross-sectional variance of log wages than permanent shocks when the estimation tries to match moments in levels.



Table 8: Parameter estimates for the earnings process

	Individual wage		Equivalent family earnings		Equiv. family disposable income	
	$\sigma_\varepsilon^2$	$\sigma_\eta^2$	$\sigma_\varepsilon^2$	$\sigma_\eta^2$	$\sigma_\varepsilon^2$	$\sigma_\eta^2$
1994	0.0598	0.0588	0.0296	0.1039	0.0096	0.0300
1995	0.0535	0.0588	0.0306	0.0980	0.0084	0.0326
1996	0.0578	0.0592	0.0227	0.0966	0.0079	0.0335
1997	0.0632	0.0523	0.0269	0.0903	0.0089	0.0316
1998	0.0572	0.0538	0.0265	0.0858	0.0080	0.0306
1999	0.0733	0.0480	0.0182	0.0954	0.0091	0.0341
2000	0.0721	0.0485	0.0192	0.0809	0.0105	0.0292
2001	0.0639	0.0536	0.0185	0.0857	0.0088	0.0297
2002	0.0596	0.0566	0.0245	0.0731	0.0155	0.0302
2003	0.0579	0.0536	0.0240	0.0685	0.0122	0.0224
2004	0.0513	0.0583	0.0294	0.0673	0.0092	0.0269
2005	0.0513	0.0691	0.0294	0.0852	0.0092	0.0298
Mean	0.0609	0.0547	0.0250	0.0859	0.0098	0.0301

Notes.  $\sigma_\varepsilon^2$  and  $\sigma_\eta^2$  are the variance of the transitory shock and that of the permanent shock, respectively.

Government policies such as progressive taxation, social welfare programs, and employment insurance are designed to mitigate earnings risk that families face. In order to examine the extent to which government policies reduce earnings risk, we compare estimates for equalized family earnings and disposable income. Much like our previous findings, our results suggest that the tax and transfer system substantially reduces both permanent and transitory earnings risk, with the permanent and transitory risk reduced on average by 5.58 point (from 0.0859 to 0.0301) and 1.52 point (from 0.0250 to 0.0098), respectively.

## 4.7 Log-Normality Tests

A highly desirable property of distributions is log normality. With that property, a distribution can be fully characterized by its first two moments, and two distributions

are unambiguously comparable with respect to the degree of inequality. In particular, if earnings are log-normally distributed, it can never happen that the Gini coefficient goes up but the variance of log earnings goes down.

We investigate log-normality using two approaches. The first is formal tests, including the Kolmogorov-Smirnov and the Skewness-Kurtosis test.<sup>38</sup> These tests tell us whether the deviations from normality that we observe are statistically significant. A further question, though, is whether these deviations are sufficiently large to be important. To address that issue, we use [Epanechnikov \(1969\)](#)'s kernel density estimation in order to compare non-parameterically estimated densities with the corresponding normal densities.

The results from formal tests are on the whole negative. For earnings, we can always reject log-normality, no matter what the test, no matter what the year, no matter what the cohort and no matter what the (reasonable) significance level. The situation with consumption is rather different. As pointed out in [Battistin et al. \(2007\)](#), consumption is more log-normal than earnings. In fact, for observations in 1982, 1986 and 1992 and using the Kolmogorov-Smirnov test, we cannot reject, at the 5 percent level, log-normality of the cross-sectional distribution of consumption among two cohort groups: those born in the 1940's and those born in the 1950's. For other tests and other samples, however, we can reject normality.

Turning our attention to the second approach, consider [Figures 15–16](#). These figures plot non-parameterically estimated probability density functions along with their normal distribution counterparts for the 1940's birth cohort.<sup>39</sup> It is very clear from these figures that log consumption is closer to normal than log earnings. Indeed, an inspection of the skewness and kurtosis of log consumption reveals that the distribution is very nearly symmetric and exhibits only a little more kurtosis than the normal distribution. The difference is barely visible to the naked eye. Thus log-normality is an excellent approximation of the distribution of consumption.

For earnings, things are quite different. The distributions are visibly skewed to the left (the mean is to the left of the mode) and strongly leptokurtic (the density has a sharper peak and fatter tails than the normal distribution). Log-normality is

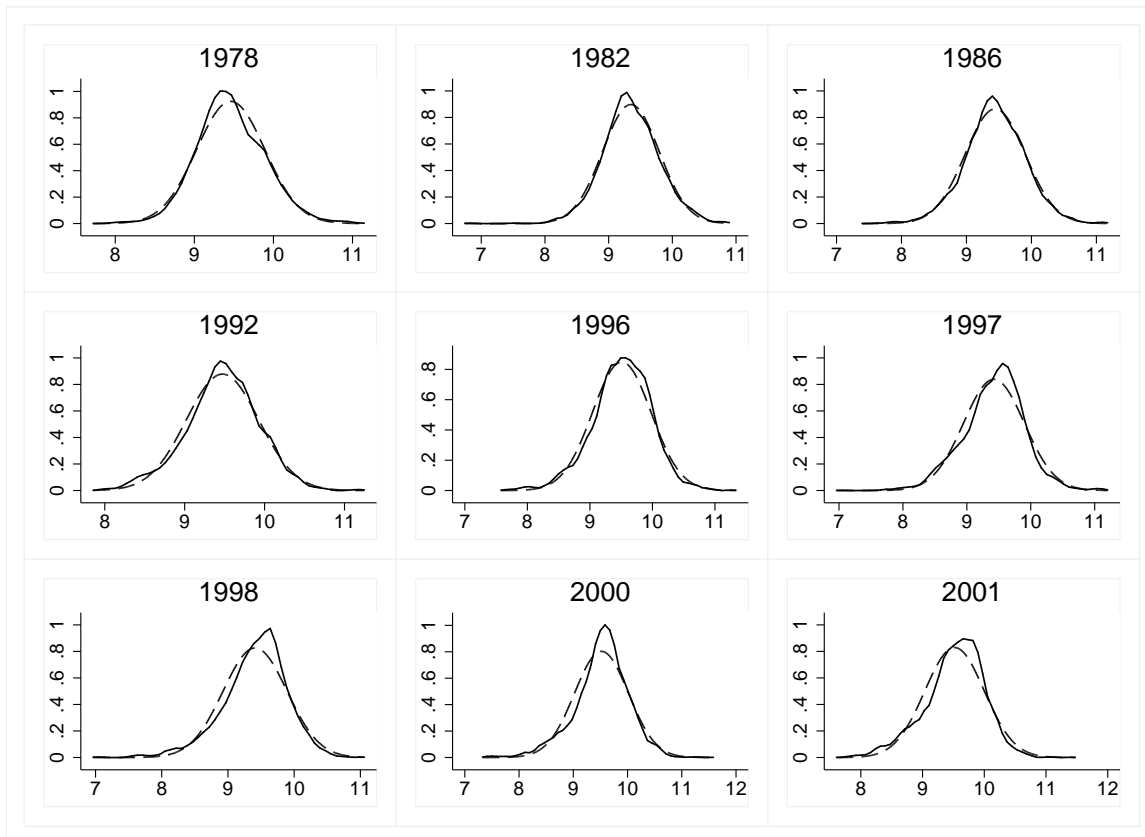
---

<sup>38</sup>All the results in this section are based on data from FAMEX/SHS using sample B from [Table 11](#).

<sup>39</sup>Similar pictures emerge from other cohorts.

not a particularly good approximation of the distribution of earnings.<sup>40</sup>

Figure 15: How log normal is consumption (1940's birth cohort)



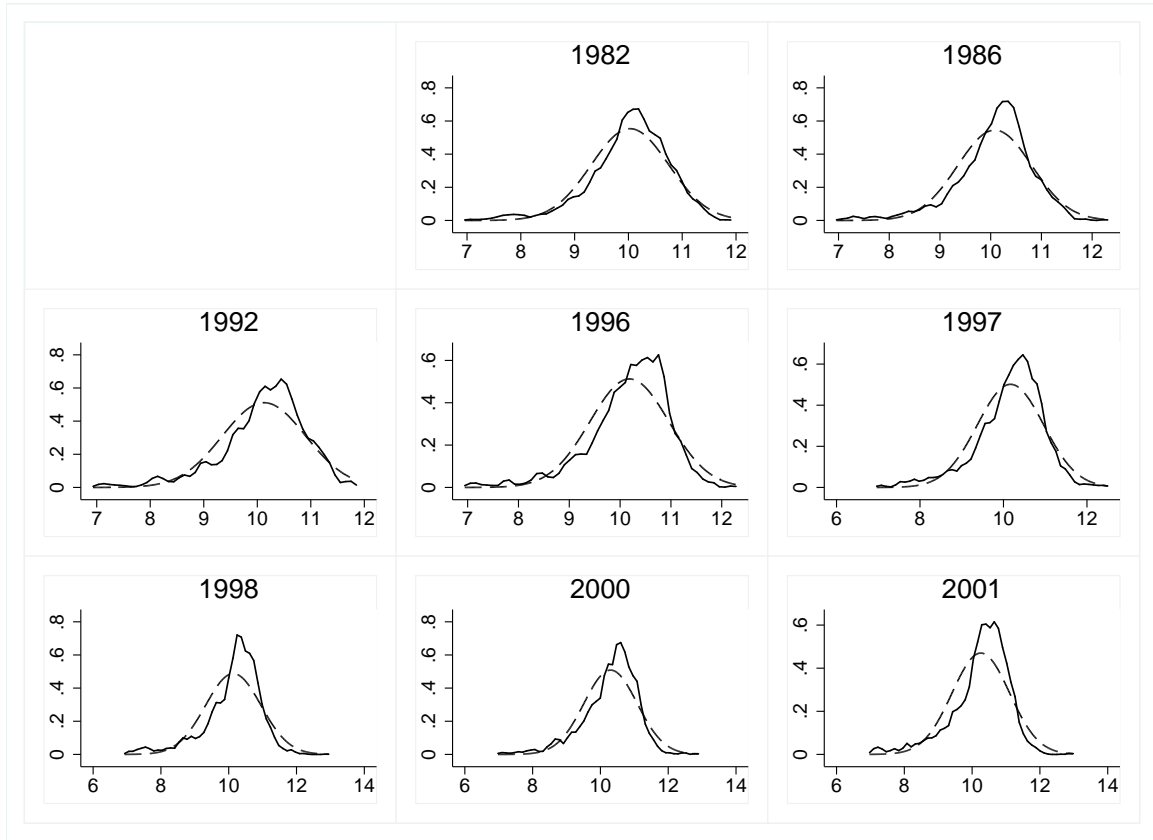
## 5 A View of Inequality from Public-Use Files

It may be of some interest to compare the conclusions about trends in inequality we draw in section 4 from the ones that would be drawn from Public-Use Files (PUF). PUF data differ from RDC data in two important respects. First, Statistics Canada introduces noise in PUF data in order to preserve the identity of the respondents.

---

<sup>40</sup>We also estimated densities for wages from SCF/SLID data, obtaining similar results. Interestingly, log-normality is not as bad an approximation for disposable income. Indeed, there is an instance where we could not reject log-normality of the cross-sectional distribution of disposable income. This was for the 1950's cohort for observations in 1982.

Figure 16: How log normal are earnings (1940's birth cohort)



Second, the revised weights are not currently available in PUF for the SCF sample. Below we briefly review how the main message from RDC data is altered when using public data for income and wealth—recall that for our purpose, there is no distinction between RDC and PUF for consumption data.

## 5.1 Income Inequality

To see what Public-Use Files tell us about trends in income inequality, consider Table 9 and contrast it with the corresponding results from RDC data, displayed in Table 5. One striking result is that the level of inequality, however measured, is understated in Public-Use Files. In terms of changes, the difference is even more

striking. For example, the Public-Use Files suggest that the level of inequality has barely changed from 1986 to 2000. In particular, it completely misses the increase in inequality in the 1990's, which led [Frenette et al. \(2007\)](#) to abandon survey data in favor of Census data.<sup>41</sup>

Table 9: Pre- and After-Tax Income Inequality from Public-Use Files

year	Pre-Tax Income				After-Tax Income			
	varlog	50/10	90/50	Gini	varlog	50/10	90/50	Gini
1980	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
1986	0.5639	2.6657	1.9936	0.3284	0.2670	1.9610	1.8254	0.2709
1990	0.5610	2.6621	1.9552	0.3273	0.2498	1.9210	1.7682	0.2611
1995	0.6359	2.8613	2.0036	0.3402	0.2622	1.9373	1.7919	0.2665
2000	0.6080	2.8182	2.0267	0.3438	0.2724	1.9639	1.8179	0.2740
2005	0.6751	3.0673	2.1370	0.3676	0.3103	2.0677	1.8801	0.2920
<u>growth (%)</u>								
1986–2000	0.45	-3.16	1.81	0.46	-1.95	-2.99	0.26	-1.62
1986–1990	-0.51	-0.13	-1.93	-0.33	-6.45	-2.04	-3.13	-3.60
1990–2000	1.03	-3.02	3.81	0.81	4.92	-0.94	3.45	2.09
1980–1985	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
1986–1990	-0.51	-0.13	-1.93	-0.33	-6.45	-2.04	-3.13	-3.60
1990–1995	13.34	7.48	2.48	3.95	4.98	0.84	1.34	2.04
1995–2000	-12.29	-10.55	1.30	-3.18	-0.13	-1.78	2.10	0.03
2000–2005	11.03	8.84	5.44	6.91	13.89	5.29	3.42	6.57

Notes. All changes that involve years over which SCF and SLID overlap are computed as the sum of the change in SCF until 1996 and the change in SLID from 1996.

## 5.2 Wealth Inequality

One aspect that PUF data typically have problems with is the very top of the distribution (because of top coding), and that is of course particularly important for

<sup>41</sup>Some of the changes between 2000 and 2005 are due to the revision of survey weights as SLID data for 2000 feature old weights based on Census 1996, while the SLID data for 2005 feature the revised weights based on Census 2001.

the wealth distribution. Table 10 sheds some light on that. The Public-Use Files for the SFS suffer from top-coding problems where each observation above a certain threshold is given the value of that threshold. Since we have access to RDC data, in which no observation is top-coded, we can compare the accuracy of different statistical methods to impute top-coded observations in terms of various measures of inequality.

What we do is the following, taking net financial wealth as an example. First we consider the components of net financial wealth, like mutual fund holdings, mortgage debt, etc. For each such component, we consider the top decile with the idea that that top decile has been drawn from a Pareto distribution. Excluding the top-coded observations, we then use the remaining observations in the top decile to estimate the single (remaining) unknown parameter of the Pareto distribution. This can be done either by least squares or maximum likelihood and we report the results from both approaches in Table 10, labeled Public 3 and Public 4 respectively. Having estimated the parameter, one can then mechanically compute the mean of the distribution and replace the top-coded values with that mean. As Table 10 reveals, both methods work pretty well. For the Gini coefficient of any of the wealth categories, the estimate is at most 1 percent away from the RDC value. However, the estimates are not as precise for the top 1%.

## 6 Concluding Remarks

This paper characterizes various aspects of inequality in Canada over the last 30 years or so. All measures of income inequality have increased over these years. Most notable is the rise in wage and earnings inequality, which were to a large extent offset by the tax and transfer system to produce a moderate rise in after-tax income. As a result, the rise in consumption inequality has also been fairly mild, though still noticeable. While we emphasize the role of government policy in compressing inequality, much work remains to be done to identify the precise role of specific components of the tax and transfer system in mitigating the trend as well as movements in income

Table 10: Comparison of public-use and confidential SFS 1999.

	Public 1	Public 2	Public 3	Public 4	RDC Data
<u>Share of top 1%</u>					
Disposable income	0.0481	0.0481	0.0482	0.0483	0.0541
Net financial wealth	0.1246	0.1431	0.1841	0.1966	0.1635
Net total wealth	0.1161	0.1424	0.1623	0.1831	0.1665
Net worth	0.1066	0.1323	0.1509	0.1683	0.1568
<u>Gini index</u>					
Disposable income	0.3252	0.3303	0.3319	0.3320	0.3355
Net financial wealth	0.6798	0.6926	0.7014	0.7032	0.7054
Net total wealth	0.6348	0.6536	0.6616	0.6685	0.6664
Net worth	0.6274	0.6456	0.6532	0.6591	0.6603

Notes. The column designated Public 1 reports estimates for the sample excluding observations with top-coded values. Public 2 reports the estimates for the sample including observations with top-coded values. Public 3 corrects for top-coding by OLS estimation suggested by David Domeij. Public 4 corrects for top-coding by maximum likelihood estimation. The estimation is conducted for each income and wealth component, since top-coding is done at the component level. For the estimation, the sample of families with the given variable above the 90th percentile and below the top-coding threshold is used.

inequality.<sup>42</sup>

In light of the recent work of [Frenette et al. \(2006\)](#) and [Frenette et al. \(2007\)](#), our results suggest that the revision of the weights that Statistics Canada has recently implemented was valuable, in the sense that these revisions brought measures of inequality in Survey data much closer to the evidence emanating from Census data. In particular, our results tend to support the broad conclusion on the evolution of income inequality reached by [Frenette et al. \(2007\)](#).

In our opinion, future work should focus on finding a better specification for earnings/income and wage processes. [Domeij and Flodén \(2009\)](#) (in this volume) criticize the specification and also the estimation strategy (based on first differences) used here; in these concluding remarks we provide a distinct, but equally critical, perspec-

<sup>42</sup>The most important reforms are probably the replacement in 1991 of the manufacturers sales tax (MST) with the Goods and Services tax (GST), which is a federal value added tax, as well as the income tax reform of 1987.

tive.<sup>43</sup> The point we want to make here is that the estimated process described in Section 4.6 has counterfactual implications for the autocovariance function of (residual log) earnings (and wages). This autocovariance function has some very distinctive and very robust features, and these features are common across datasets from several countries, including the U.S. PSID and GSOEP from Germany.

To describe these features, fix an age  $s$  and a calendar date  $t$ ; it doesn't matter that much what  $s$  is provided it is not too high and for simplicity we suppress the calendar date for the sake of this argument. Consider now the covariance between (residual) earnings at age  $s$  and earnings at age  $s+k$ , denoted by  $\Gamma_{s,s+k}$ . In the data, the difference between  $\Gamma_{s,s}$  and  $\Gamma_{s,s+1}$  is much larger than the difference between  $\Gamma_{s,s+1}$  and  $\Gamma_{s,s+2}$ . In fact, the ratio  $\Gamma_{s,s}/\Gamma_{s,s+1}$  is about two in the data used in this paper and the ratios  $\Gamma_{s,s+k}/\Gamma_{s,s+k+1}$  for  $k \geq 1$  are not too far from one. This is not consistent with the estimates reported in Section 4.6. To be more precise, we can make assumptions about the variance of some "initial" value for the permanent component of earnings to make the model consistent with the variances and covariances for the value of  $s$  corresponding to that initial age. But then it will be inconsistent with the autocovariance functions for the other values of  $s$ . Technically, this is because  $\sigma_\varepsilon^2$  is too small relative to  $\sigma_\eta^2$ .

The fact that  $\Gamma_{s,s}/\Gamma_{s,s+1}$  is so much larger than  $\Gamma_{s,s}/\Gamma_{s,s+1}$  is prima facie evidence that the transitory component of earnings is relatively much more important than what the results in Section 4.6 suggest, and that the estimation strategy there tends to underestimate the importance of that component.

An alternative specification that can easily capture both the autocovariance functions and the age-variance profile decomposes (residual log) earnings into three components: a transitory, a persistent and a fixed component. Further work is needed to examine what further features might be needed to adequately capture the key features of the data.

---

<sup>43</sup>Another specification, applied to Canada, can be found in [Baker and Solon \(2003\)](#). For U.S. data, an important alternative perspective is given by [Guvenen \(2009\)](#).



# Appendix A : Detailed Sources of Data

## A.1 Consumption Data

The *Survey of Family Expenditures* (FAMEX) and the *Survey of Household Spending* (SHS) provide cross-sectional data on individual and household characteristics, income, and expenditure. For both surveys, income and expenditure data refer to a calendar year (called reference year). Data are collected through personal interviews conducted in the first quarter following a given reference year. FAMEX data are available for 1969, 1974, 1978, 1982, 1984, 1986, 1990, 1992, and 1996. But the surveys for 1974, 1984, and 1990 cover only urban areas and have a smaller sample (around 7,000 households instead of 14,000 or so for the years with national coverage). The SHS replaced the FAMEX in 1997 and has been conducted annually thereafter. Currently, the latest data are for 2005.

For 1969–1986, expenditure data are collected for *spending units*. Spending unit is defined as a group of persons dependent on a common or pooled income for the major items of expense and living in the same dwelling or one financially independent individual living alone. In 1990, the unit of analysis was replaced by *households*, which is defined as a person or a group of persons occupying one dwelling unit.

Disposable income after taxes and transfers ( $y^D$ ) can be constructed for the whole sample period. However, pre-tax household earnings ( $y^L$ ), which is defined as the sum of wages and salaries and the labour part of self-employment income, cannot be constructed in a consistent manner over time, since wages/salaries and self-employment income are not separately available over the 1997–2005 period. We use nondurable goods expenditure as our measure of consumption. Nondurable goods expenditure consists of the expenditure on food, alcohol and tobacco, personal care supplies and services, water, fuel, and electricity, household operation, public transportation, operation of automobiles and trucks, apparel, recreation, reading, education, traveler accommodation, health care, and miscellaneous nondurable expenditures. We deflate income by region-specific CPI for all items. We deflate each expenditure component by the corresponding region-specific CPI.

Unfortunately, no information on education is available from 1997 to 2003. Thus,

we cannot conduct analysis involving education level for those years.

## A.2 Income Data

The *Survey of Consumer Finance* (SCF) and the *Survey of Labour and Income Dynamics* (SLID) provide cross-sectional data on individual and household characteristics as well as income components. Both surveys share the same sample design, as their sample consists of rotation groups from the *Labour Force Survey* (LFS) and includes roughly 15,000 households.<sup>44</sup> For both surveys, income data refer to a calendar year (called reference year). SCF data are collected through personal interviews conducted in April (so that individuals have their tax return in mind) following a given reference year. As for SLID, characteristics are collected in January following a given reference year, and income data is collected in May following a given reference year. SLID respondents are given the option to skip the May interview by giving Statistics Canada permission to access their T1 (income) tax information. Over 80% of respondents give their consent to the use of administrative records. The change in collection method may have had an impact on the comparability of data across the two surveys.<sup>45</sup>

Unlike SCF, SLID also has a panel dimension. Waves are introduced every three years, and each wave of respondents is interviewed for 6 years, with the first wave starting in 1993. The panels, which initially constitute a representative cross-sectional sample of the population, are thus as follows: Panel 1 from 1993 to 1998; Panel 2 from 1996 to 2001; Panel 3 from 1999 to 2004; Panel 4 from 2002 to 2005; Panel 5 from 2005. For each sampled household in SLID, up to 12 interviews are conducted

---

<sup>44</sup>The LFS is like the CPS, and the SCF is like the March supplement to the CPS. The SLID is also like a March supplement, only each sample is interviewed over the next 6 years.

<sup>45</sup>Statistics Canada reports that “[C]omparisons of figures produced from the SCF with other sources of data (Census of Population, Longitudinal Administrative Data, National Economic and Financial Accounts) reveal that certain income components, such as investment, self-employment earnings, social assistance payments and EI benefits, are under-reported in the SCF.” Furthermore, “SLID’s estimates of the number of income recipients, aggregate individual income and average family income are higher than the corresponding estimates from the SCF data.” It should also be pointed out that the response rates, while relatively high in both the SCF and SLID, are higher in SLID than in SCF data: in the SCF it ranged from 78.1% in 1989 to 82.1% in 1995, while SLID’s cross-sectional rate of response was 87.1% in 1996. The higher SLID response rate is primarily due to the use of administrative data from the tax files.

over a six-year period.

SCF data are available every other year from 1965 to 1971, and yearly thereafter until 1997. For all years, SCF data are available at the individual, census family and economic family level. In addition, household-level data are available for 1976-1997. We have SLID data from 1993 to 2005 at all levels. In this paper, we use individual- and economic family-level data.

Though SCF data are available for odd years between 1965 and 1975, we do not use the data prior to 1976 for the following two reasons. First, some of the data necessary for our analysis are not available over these years. For example, data on labour force status and usual hours worked are not available for these years.<sup>46</sup> Second, SCF started adjusting non-responses by imputation in 1977, which is necessary to construct integrated survey weights at household, economic family, census family, and individual levels. (Non-responses were adjusted by re-weighting before 1977.) Though Statistics Canada dropped households with non-respondent members from the 1976 sample in order to apply integrated weights based on Census 1996, they have not applied the same adjustment to the data for 1965–1975. Without such adjustment, the data for 1965–1975 cannot receive revised weights.

With the detailed income data available in SCF and SLID, we are able to compute individual and family earnings ( $y^L$ ), pre-government family non-financial income ( $y^{L+}$ ), net financial income ( $y^A$ ), pre-government family income ( $y$ ), and post-government family income ( $y^D$ ). In SCF, we compute annual hours worked by multiplying usual hours worked per week and weeks worked during the reference year. SLID provides data on total hours paid at all jobs during the reference year. We compute hourly wage ( $w$ ) by dividing individual earnings by the person's annual hours worked. We deflate the income data by province-specific CPI.

SCF adopted a new set of questions on education, introduced by the Labour Force Survey in the 1989 reference year. The following three points are relevant to our analysis: 1. In 1976–1988, post-secondary education requires high-school graduation. In 1989–1997, any education that could be counted towards a degree, certificate, or diploma from an educational institution is taken as post-secondary education.

---

<sup>46</sup>Note that LFS employment and unemployment estimates before and after 1976 are not directly comparable since the LFS questionnaire underwent significant changes in 1976.

For example, trades programs offered through apprenticeship, vocational schools or private trade schools do not always require high school graduation. Such education is considered as post-secondary in 1989–1997 while only primary or secondary would have been recognized prior to 1989. Because of this change, one cannot construct a post-secondary education category that is consistent over time. 2. One cannot identify individuals whose educational attainment is exactly high school graduation in SCF data prior to 1989, while it is possible for SCF 1989–1997 and for all SLID years. 3. Over the 1976–1988 period, one cannot distinguish between university (or college) certificate/diploma below bachelor’s degree and university degree (bachelor’s degree or above). For SCF 1989–1997 and for all SLID years, those two education levels can be identified separately. Due to the lack of detailed information before 1989, we use the following broad education categories in our analysis for consistency: 1. high school dropouts or graduates; 2. some secondary education below university; 3. university certificate/diploma/degree.

### A.3 Wealth Data

The *Survey of Financial Securities* (SFS) provides various details of individual and family characteristics as well as income and wealth data. The data are available for 1999 and 2005. The SFS sample is drawn from two sources. The main sample is selected from the Labour Force Survey (LFS) sampling frame. The second portion of the sample, accounting for approximately 10 to 15 percent of the total sample, is drawn from geographic areas in which a large proportion of households have what is defined as “high-income”. In other words, much like the Survey of Consumer Finances in the U.S., the SFS over-samples the rich in order to improve wealth estimates.

Demographic data are available at the individual level, while the questions pertaining to wealth are asked with reference to the economic family. With detailed data on income, assets and debts, we can construct a measure of disposable income ( $y^D$ ), net financial wealth ( $a$ ), net total wealth ( $a^+$ ), and net worth for 1999 and 2005. The disposable income and the three wealth measures are deflated by province-specific CPI.

## Appendix B: Sample Selection

The following tables document the various samples used in the text. Table 11 displays our FAMEX/SHS sample selection. Tables 12 and 13 present our samples for individual-level and family level data for the SCF, while Tables 14 and 15 are for SLID data. Finally, our sample selection for the SFS is displayed in Table 16.

Table 11: Sample Selection in FAMEX-SHS

	Observations deleted	Remaining observations
Original data set		225360
Observed in 1974, 1984, and 1990	15991	209369
Part-year household (SHS)	4603	204766
Aged less than 25 or more than 60	63474	141292
Missing main characteristics	1630	139662
Yukon, North-west Territories or Nunavut	3868	135794
Food consumption less than 400	114	135680
<b>Sample A</b>		<b>135680</b>
non-positive disposable income	172	135508
<b>Sample B</b>		<b>135508</b>

Notes. Main characteristics include education, region of residence, and family type. Yukon, North-west Territories and Nunavut are only covered in the SHS survey, not in FAMEX.

Table 12: Sample Selection in Individual-level SCF-RDC

	Observations deleted	Remaining observations
Original data set (1976-1997)		1965517
Observed in 1976 or 1978	79407	1886110
Aged less than 25 or more than 60	994857	891253
Income imputed	164839	726414
Missing main characteristics	0	726414
Missing earnings or hours worked	2880	723534
Wage less than half min wage	26821	696713
Positive earnings and zero hours worked	91117	605596
<b>Sample A</b>		<b>605596</b>
Non-positive earnings	150761	454835
Wage greater than 100 and annual hours worked less than 100	775	454060
<b>Sample B</b>		<b>454060</b>

Notes. Main characteristics include gender, education, marital status, and province of residence.

Table 13: Sample Selection in Family-level SCF-RDC

	Observations deleted	Remaining observations
Original data set (1976-1997)		750396
Observed in 1976 or 1978	27077	723319
Aged less than 25 or more than 60	187911	535408
Income imputed	124512	410896
Missing main characteristics	0	410896
Missing income or head's hours worked	2274	408622
Head's wage less than half min wage	13823	394799
Positive earnings and zero hours worked	89548	305251
<b>Sample A</b>		<b>305251</b>
Initial number of observations (Sample A)		305251
Non-positive disposable income	1098	304153
Non-positive pre-government income	21439	282714
<b>Sample B</b>		<b>282714</b>
Initial number of observations (Sample B)		282714
Non-positive pre-gov't non-financial income	3210	279504
Non-positive household earnings	9357	270147
Non-positive head's earnings	12848	257299
<b>Sample C</b>		<b>257299</b>
Initial number of observations (Sample A)		305251
Non-positive household earnings	35074	270177
<b>Sample D</b>		<b>270177</b>

Notes. Main characteristics include head's education, province of residence, and family type. Note that these samples are not necessarily nested.

Table 14: Sample Selection in Individual-level SLID-RDC

	Observations deleted	Remaining observations
Original data set (1993-2005)		875597
Aged less than 25 or more than 60	441891	433706
Income imputed	91145	342561
Missing main characteristics	16438	326123
Missing earnings or hours worked	17497	308626
Wage less than half min wage	16480	292146
Positive earnings and zero hours worked	9504	282642
<b>Sample A</b>		<b>282642</b>
Non-positive earnings	37178	245464
Wage greater than 100 and annual hours worked less than 100	684	244780
<b>Sample B</b>		<b>244780</b>

Notes. Main characteristics include gender, education, marital status, and province of residence.



Table 15: Sample Selection in Family-level SLID-RDC

	Observations deleted	Remaining observations
Original data set (1993-2005)		319934
Aged less than 25 or more than 60	57278	262656
Income imputed	90882	171774
Missing main characteristics	7983	163791
Missing income or head's hours worked	7561	156230
Head's wage less than half min wage	7934	148296
Positive earnings and zero hours worked	11986	136310
<b>Sample A</b>		<b>136310</b>
Initial number of observations (Sample A)		136310
Non-positive disposable income	53	136257
Non-positive pre-government income	6562	129695
<b>Sample B</b>		<b>129695</b>
Initial number of observations (Sample B)		129695
Non-positive pre-gov't non-financial income	793	128902
Non-positive household earnings	3865	125037
Non-positive head's earnings	4023	121014
<b>Sample C</b>		<b>121014</b>
Initial number of observations (Sample A)		136310
Non-positive household earnings	11239	125071
<b>Sample D</b>		<b>125071</b>

Notes. Main characteristics include head's education, province of residence, and family type. Note that these samples are not necessarily nested.

Table 16: Sample Selection in Family-level SFS-RDC

	Observations deleted	Remaining observations
Original data set (1999, 2005)		21215
Aged less than 25 or more than 60	5504	15711
Missing main characteristics	0	15711
Missing income or wealth	0	15711
Outliers in wealth	$x$	$15711 - x$
<b>Sample A</b>		<b><math>15711 - x</math></b>

Notes. We dropped outliers with respect to wealth. However, because the number of observations dropped violates the minimum threshold set by Statistics Canada for public release, we cannot report the exact number of observations affected by this criterion. Main characteristics include head's education, province of residence, and family type.

## References

- Baker, M. and G. Solon (2003). Earnings dynamics and inequality among Canadian men, 1976–1992: Evidence from longitudinal income tax records. *Journal of Labor Economics* 21(2), 289–321.
- Battistin, E. (2003). Errors in survey reports of consumption expenditures. Institute for Fiscal Studies Working Paper W03/07.
- Battistin, E., R. Blundell, and A. Lewbel (2007). Why is consumption more log normal than income? Gibrat’s Law revisited. Institute for Fiscal Studies Working Paper 08/07.
- Boudarbat, B., T. Lemieux, and W. C. Riddell (2006). Recent trends in wage inequality and the wage structure in Canada. In D. A. Green and J. R. Kesselman (Eds.), *Dimensions of Inequality in Canada*. Vancouver, Canada: University of British Columbia Press.
- Budría Rodríguez, S., J. Díaz-Giménez, V. Quadrini, and J.-V. Ríos-Rull (2002). Updated facts on the U.S. distributions of earnings, income, and wealth. *Federal Reserve Bank of Minneapolis Quarterly Review* 26(3), 2–35.
- Crossley, T. and K. Pendakur (2006). Consumption inequality in Canada. In D. A. Green and J. R. Kesselman (Eds.), *Dimensions of Inequality in Canada*. Vancouver, Canada: University of British Columbia Press.
- Domeij, D. and M. Flodén (2009). Inequality trends in Sweden 1978–2004. Unpublished manuscript.
- Epanechnikov, V. A. (1969). Nonparametric estimation of a multivariate probability density. *Theory of Probability and Its Applications* 14, 153–158.
- Frenette, M., D. A. Green, and K. Milligan (2007). The tale of the tails: Canadian income inequality in the 1980s and 1990s. *Canadian Journal of Economics* 40(3), 734–764.
- Frenette, M., D. A. Green, and G. Picot (2006). Rising income inequality in the 1990s: An exploration of three data sources. In D. A. Green and J. R. Kesselman

- (Eds.), *Dimensions of Inequality in Canada*. Vancouver, Canada: University of British Columbia Press.
- Guvenen, F. (2009). An empirical investigation of labor income processes. *Review of Economic Dynamics* 12(1), 58–79.
- Heathcote, J., F. Perri, and G. L. Violante (2009). Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006. Unpublished manuscript.
- Heathcote, J., K. Storesletten, and G. L. Violante (2008). The macroeconomic implications of rising wage inequality in the United States. Unpublished manuscript.
- Lathe, H. (2005). Survey of labour and income dynamics: 2003 historical revision. Statistics Canada, Income Statistics Division, Income Research Paper Series.
- Morissette, R., X. Zhang, and M. Drolet (2002). The evolution of wealth inequality in Canada, 1984-1999. Statistics Canada Working Paper No. 187.
- Rea, W. and E. Greenlee (2005). The effects of the revised estimation methodology on estimates from household expenditure surveys. Statistics Canada, Income Statistics Division, Household Expenditure Research Paper Series.
- Slesnick, D. T. (2001). *Consumption and Social Welfare. Living Standards and Their Distribution in the United States*. Cambridge: Cambridge University Press.
- Storesletten, K., C. I. Telmer, and A. Yaron (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics* 51(3), 609–633.