Cohort patterns in Canadian earnings: assessing the role of skill premia in inequality trends

Paul Beaudry and David A. Green *Department of Economics, University of British Columbia*

Abstract. In this paper we document the pattern of change in age-earnings profiles across cohorts and evaluate its implications. Using synthetic cohorts from the Survey of Consumer Finances over the period 1971 to 1993, we show that the age-earnings profiles of Canadian men have been deteriorating for more recent cohorts in comparison with older cohorts. We find this pattern for both high school and university educated workers. In no case do we find evidence that the return to gaining experience has been increasing over time, nor do we find increased within-cohort dispersion of earnings. We view these findings as conflicting with the hypothesis that increased skill premia largely explain the observed increase in dispersion of male weekly earnings. JEL Classification: J31, O33

Effets de cohorte dans les gains au Canada: évaluation du rôle des primes d'habilités dans l'explication des tendances vers l'inégalité. Ce mémoire documente le pattern de changements dans les relations âges-gains d'une cohorte à l'autre, et en analyse les implications. Utilisant des cohortes synthétiques pour la période 1971–93, on montre que les profils âges-gains des hommes au Canada se sont détériorés dans les cohortes les plus récentes. On ne trouve pas de résultats qui montreraient que le rendement sur l'expérience accumulée s'est accru dans le temps ou qu'il y a dispersion accrue des gains à l'intérieur des cohortes. Les auteurs suggèrent que ces résultats contredisent l'hypothèse que des primes accrues pour les habiletés expliqueraient l'accroissement qu'on a observé dans la dispersion des gains des hommes au Canada.

1. Introduction

Changes in the distribution of earnings have become the focus of considerable interest. It is now well established that, over the last fifteen years, the dispersion of male earnings in Canada has widened considerably (see Richardson 1997; Moris-

The authors wish to thank Alan Stark and Maud Rivard for excellent research assistance. We also thank seminar participants at the UBC labour workshop and the CIAR 1996 fall meeting for their comments. Email: green@econ.ubc.ca

Canadian Journal of Economics / Revue canadienne d'Economique, Vol. 33, No. 4
November / novembre 2000. Printed in Canada / Imprimé au Canada

0008-4085 / 00 / 907–936 / © Canadian Economics Association

sette, Myles, and Picot 1996). This observation holds for both annual and weekly earnings. One of the salient features of the increase in dispersion in Canada is its relationship to the age of workers. For example, between 1981 and 1993 the difference between the average earnings of male workers aged 45 to 55 and those aged 25 to 35 increased by 18 percentage points. There are two interpretations of such an observation. On the one hand, such a pattern may indicate that experience is more valuable today, and therefore current young workers should expect higher wage growth as they age than previous labour market entrants experienced. Alternatively, the increased age differential may imply that recent labour market entrants are and will be paid less than their older counterparts throughout their lifetime. Under this interpretation, the increased age differential reflects a deterioration in lifetime earnings opportunities for newer labour market entrants. To this point, the existing literature has assumed almost universally that the increased age differential reflects higher returns to experience.[1] As such, the increased differential is viewed as part of a general increase in returns to skills of all types. These increased skill premia are viewed, in turn, as being a reflection of an economy-wide increase in demand for skills. In parts of the U.S. literature on inequality, the notion that a skill-biased shift in demand is the prime force affecting wage patterns is taken as a given, and effort is devoted to uncovering the source of the demand shift (e.g., Borjas and Ramey 1996; Berman, Bound, and Griliches 1994).

Our objectives in this paper are twofold. First, we want to document how the age-earnings relationship has changed across birth cohorts. In order to provide a comprehensive picture of this process, we report results for different educational categories and for both the mean and the 90-10 decile differences. This evidence allows us to examine whether the increased age differential for males is more likely a reflection of increased returns to experience or of successively worse labour market performance for younger workers. Answering this question is of interest, in part, as an aid in setting policies, such as training, that have different impacts across generations. Our second objective is to use the cohort-based earnings patterns as a tool for re-evaluating the relevance of the skill-biased demand shift hypothesis for Canada.

In pursuing these objectives, and in contrast to much of the previous literature on skill-biased demand shifts, we present results for both male and female workers. In recent work it has been shown, however, that females did not experience the same increases in inequality as males over the 1980s in Canada (Beach and Slotsve 1996).[2] In any attempt to assess theories involving broad demand shifts, such as shifts

1 The only exceptions are Morissette and Berube (1997) and Burbidge, Magee, and Robb (1997), both of which take a cohort-based approach. Our study differs from Morissette and Berube (1997) because they use tax data, which does not permit an educational breakdown, and from Burbidge, Magee, and Robb (1997) because we use the results to examine alternative theories of the causes of increased inequality.

2 This is true in an examination of all workers and all earnings sources, which is essentially what we use in this paper. Burbidge, Magee, and Robb (1996) show that female and male inequality trends are very similar for full-year, full-time, non-self-employed workers.

generated by skill-biased technical change, it is important to examine wage patterns for all workers.

Our approach to the first objective is simple. Using data from the Surveys of Consumer Finances between 1971 and 1993, we follow the age-earnings path of different birth cohorts in order to examine whether returns to added years of experience are changing over time. We do so after removing business cycle effects in order to highlight long-term trends.

In order to meet our second objective, we illustrate the effects of a skill-biased demand shift on cohort-specific age-earnings profiles in the context of a simple labour demand model. We then compare the predicted effects with the observed patterns. Overall, our reading of the data is that increased skill premia contribute surprisingly little to an understanding of the changes in wage inequality in Canada (when skill is interpreted along the dimensions of education, experience, and unobserved quality). Instead, we find that deteriorating labour market outcomes for younger male workers provide a more succinct characterization of the prime trends in earnings patterns. We do not provide a single explanation for this deterioration as an alternative to the skill-biased demand shift hypothesis. We argue that, with attention shifted away from a single demand side explanation, several interesting alternative hypotheses present themselves. These alternatives deserve further investigation.

The paper is structured as follows. In section 2 we describe the data. In section 3 we document how age-earnings profiles have been changing across cohorts. We begin by documenting the behaviour of the mean of cohort-specific age-earnings profiles, followed by a description of the 90-10 percentile difference. In section 4 we discuss the implications of our observations in terms of hypotheses that focus on increased skill premia due to biased labour demand shifts. Finally, in section 5 we offer some concluding comments.

2. Data

Our empirical work centres on following the earnings of cohorts of workers through time. We do this using data from the Surveys of Consumer Finance for the years 1971, 1973, 1975, 1977, 1979, 1981, 1982, 1984, 1986, 1988, 1990, 1992, and 1993. The data for the years before 1981 are drawn from Census Family Files while those from 1981 on are taken from the Individual Files. In order to create a consistent series over time, we restrict the samples from the Individual Files to include only individuals who are heads or spouses of census households.[3]

We define an entry cohort as a group of individuals who were age 25 or 26 in an even-numbered year. Thus, the 1972 entry cohort consists of individuals who turned

3 The individual files do not include a specific variable indicating that an individual is the head or spouse of a census family. Our method for assigning head or spouse status is detailed in Beaudry and Green (1997). In earlier work (Beaudry and Green 1996) we followed a smaller set of cohorts through the 1980s using only the Individual Files. The results from that exercise are very similar to those we obtain here for the 1980s following census family heads and spouses.

25 or 26 in 1972. We use two-year cohorts as a trade-off between single-year cohorts that would allow considerable definition in establishing trends and broader cohorts that would provide smoother estimates because of larger numbers of observations within each cohort. We choose age 25 as the entry age to the mature labour market.[4] This is not meant to imply that some individuals do not work steadily before age 25. Instead, it is intended to focus attention on a period after individuals have largely ended their major educational investments and during which they become more permanently attached to the labour market. We choose this focus, in part, to ease the measurement exercise that follows, since the greater flexibility exhibited by younger individuals entails more substantial selection issues. We also believe that this is an interesting place to focus attention, however, since this is the age of a crucial transition to stable work patterns, accelerating careers and family formation. If changes in the economy damage earnings patterns from this age on, this is a point of major policy concern.

The specific set of entry cohorts we follow are those entering in 1962 and after. We do not examine all cohorts in all years to avoid trying to make inferences too far out of our sample. More specifically, we believe that the older cohorts observed at the start of our sample period experienced a very different labour market relative to those who entered the labour market in the 1960s and afterward. If that is true, then attempts to extrapolate backward from our data to predictions about the earliest parts of the age-earnings profiles of these workers in order to compare their outcomes with more recent cohorts is simply misguided. We restrict our observations to individuals who are under age 56 in any given sample year in order to avoid fluctuations revolving around early retirement.

In our analysis, we divide our entry cohorts into subgroups defined by education level and examine two education groups: (1) those with some or completed high-school education; and (2) those with a university degree or more. The first group includes individuals who have some post-secondary education but have not obtained any post-secondary certificate or degree. A preliminary investigation of the data indicated that these post-secondary non-completers have wage and employment rate patterns very similar to those with a high-school education. Note that we do not present results relating either to individuals with less than a high-school education or to individuals who have completed a post-secondary certificate or degree other than a university degree. In table 7 the cell sizes for each of the gender/education/cohort groups in 1992 are provided. The sizes generally are large enough to permit confident use of cell means for weekly earnings. Further, one can see steady rises in cohort sizes until the early 1980s, associated with the entry of the baby boom into our age window and the almost continuous rise in the proportion of females with a university education. These trends are important to keep in mind when alternative hypotheses concerning the sources of the inequality trends are evaluated.

4 Card and Lemieux (1996) provide evidence on a range of outcomes for younger workers over a similar sample period. Morissette (1997) also examines earnings and employment outcomes for youths using longitudinal tax files.

One major concern in using this set of data arises from changes in the definitions of educational categories in data for 1990 and after. The main change of concern for our purposes is the switch in assignment of post-secondary education that does not require a high-school diploma from the high-school to the post-secondary education categories. To the extent that individuals after 1990 enter but do not complete these programs, they will continue to be grouped with high-school educated individuals under our assignment system. Individuals who do earn a certificate from a program not requiring a high-school diploma, however, will be categorized as having a high-school education before 1990 and a post-secondary certificate afterward. This will not affect our university degree group but will alter the composition of our high-school educated group. Given that this involves a removal of more educated individuals from the high-school group, it can be anticipated that the change will cause a fall in average earnings for the low-educated group after 1990. To assess the impact of the educational change, we reran all of our statistical specifications, including a post-1990 dummy variable. For males, the inclusion of this variable did not change any of the conclusions discussed below. For females, including the post-1990 variable changed some results, and we describe those changes at the appropriate points in the paper.

After we define cohorts, our analysis consists of examining earnings measures for each cohort that is present in each of our sample years. For example, we examine the earnings of all individuals in the 1968 entry cohort in 1971, 1973, 1975, and so forth. In principle, this provides us with a picture of the earnings path followed by this group of individuals over time. Since the SCFs do not form a true panel, we are not actually following the same group of individuals over time in this exercise. As long as the composition of the group being followed does not change over time, however, earnings measures, such as average earnings in each year for this 'synthetic' cohort, will provide an accurate picture of the average experience for individuals in the cohort. It is possible that the composition of the cohort groups changes over time as individuals acquire more education. We believe that problems of this sort are minimized by the fact that we examine only individuals over age 25.[5]

The earnings measure we examine is real weekly earnings in 1980 dollars. The conversion to real dollars is made using the Consumer Price Index for Canada for the relevant year. Weekly earnings are created for each individual in a given sample year by dividing total earnings by the number of weeks worked in the year.[6] We

5 Immigration is a further source of potential compositional changes, since new members may be added to a cohort from one data year to the next. To examine the implications of immigration, we recreated all the figures and tables presented below using samples of non-immigrant males. Results using non-immigrant data for high-school educated males are virtually identical to those presented for the whole sample in the remainder of the paper. Results for the university educated are somewhat less similar to those presented here, but basic data patterns, and conclusions based upon them, do not change when immigrants are removed.

6 Although our measure of total earnings represents mainly wages and salaries, it does include self-employment earnings and rental property earnings. We use total earnings, as opposed to only wages and salaries, because we cannot identify the self-employed or self-employment earnings in early sample years.

omit individuals with zero weeks worked and/or non-positive earnings.[7] Once these earnings are calculated for each individual in a cohort, we calculate summary measures such as the percentiles of the earnings distribution and the average weekly earnings for individuals in a given cohort in a given year.

3. Changes in age-earnings profiles

Our objective in this section is to document how age-earnings profiles have been changing for successive cohorts of labour market entrants. Since age-earnings profiles are likely to be different for different segments of the labour market, we examine changes in age-earnings profiles for four labour markets groups: high-school educated males, high-school educated females, university educated males, and university educated females. As mentioned in section 2, we focus on weekly earnings, and we restrict attention to workers between 25 and 56 years of age.

3.1. A first look at the data

3.1.1. Males

In figures 1a and 1b male real average weekly earnings for several different cohorts as they age are plotted. Figure 1a relates to the high-school educated sample, while figure 1b relates to the university educated sample. We refer to any particular cohort by the year in which it enters our age window; that is, the 1964 cohort is composed of workers who turned 25 or 26 in 1964. Note that these figures are mainly for illustrative purposes, since we do not graph the full sample of cohorts available in our data. Raw plots of the data are erratic and difficult to read both because of sampling variability and the effects of business cycles, which affect different cohorts at different ages.[8] Thus, to help to isolate trends, we plot smoothed cohort-specific profiles. We generate the smoothed profiles by estimating a cubic age-earnings profile for each cohort while simultaneously controlling for business cycle conditions. The business cycle indicator used is the quadratically detrended unemployment rate of males 45 to 54 years of age.[9] The coefficient on this detrended

7 Of course, movements in average weekly wages may partly reflect movements in hours worked per week, which is not our primary focus. We recreated all the results in the paper using only full-year (50 or more weeks worked in the sample year), full-time (30 or more usual hours per week) (FYFT) workers. The results in terms of cohort-related patterns are very similar to those reported here and none of our conclusions is altered. The only notable difference is that cyclical effects are diminished with FYFT workers. Partly because we are interested in describing overall distributional changes and partly to avoid selection effects related to changes in the proportion of workers who are FYFT, we chose to present the results for all workers.

8 Raw data plots are presented in Beaudry and Green (1997).

9 In these regressions and all the regressions that follow, we instrument for the unemployment rate variable using the detrended U.S. unemployment rate and a dummy variable capturing the period after 1982 when the Canadian and U.S. rates moved apart. This is done to allow for the possibility that underlying factors driving the wage patterns we are studying also affect the unemployment rate. The U.S. detrended unemployment rate enters as a highly significant regressor in the regression of the Canadian detrended unemployment rate on the U.S. detrended rate and the post-1982 dummy variable.

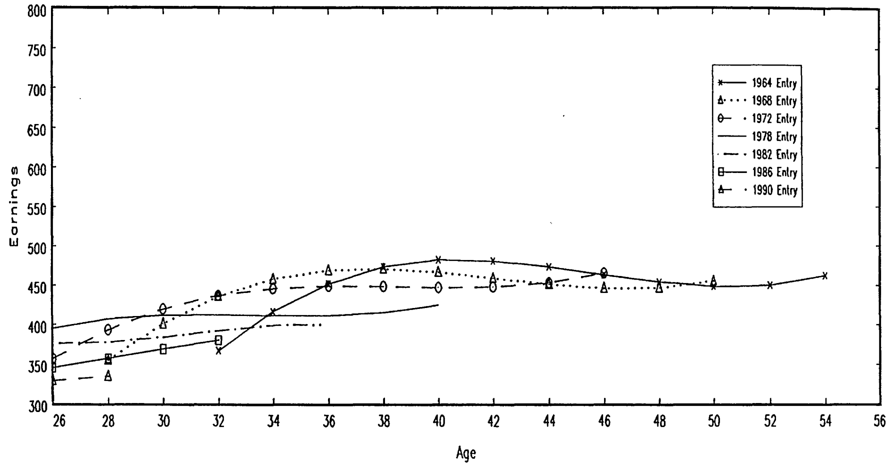

FIGURE 1a Real weekly earnings by cohort, smoothed and cycle effects removed: males, some or completed high-school education

unemployment variable is restricted to be the same across cohorts.[10] We experimented with other indicators for business cycle conditions, all of which gave very similar results in terms of patterns across and within cohorts.

Two facts are immediately apparent from figures 1a and 1b. First, different cohorts experience substantially different earnings profiles as they age. Second, the university educated sample experiences much higher wage growth, on average, as it ages. In terms of long-term trends, for the university educated sample, the most noticeable pattern is the appearance of a downward shift in cohort-specific profiles. In particular, at almost all ages in our window, the different cohorts keep a strict ordering: the older cohorts earn more and the later cohorts earn less. Moreover, the differences in magnitude across cohorts are substantial. For example, at age 32, the 1986 cohort earns approximately 20 percent less than the 1964 cohort.

Another important observation that emerges from figure 1b relates to the slopes of the earnings profiles. With the possible exception of the 1972 cohort, there does not appear to be any systematic pattern of younger cohorts experiencing greater wage growth as they age than older cohorts. Therefore, increased earnings differentials by experience for university educated men appear to arise from worsening outcomes for more recent cohorts rather than increasing returns to experience. We will examine this claim more systematically (using all the cohorts available in our data set) in subsection 3.2.

10 We also investigated specifications in which the measure of the cycle was interacted with age. Since these added interactions were neither economically substantial nor statistically significant and did not alter any of the results presented here, we focus on the simpler specification.

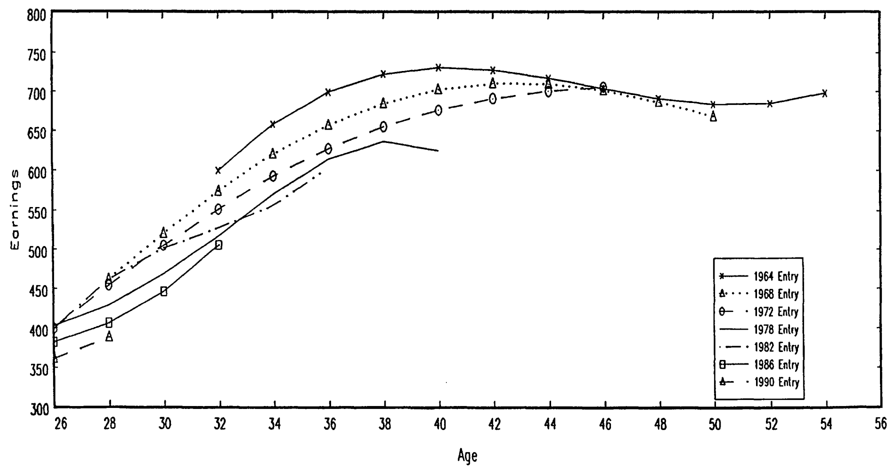

FIGURE 1b Real weekly earnings by cohort, smoothed and cycle effects removed: males, university education

The pattern of change affecting the age-earnings profiles of high-school educated males, depicted in figure 1a, is more complicated than that for university educated males. Nevertheless, a pattern still can be discerned. Part of this pattern is a rotation in cohort-specific age-profiles for the 1964 cohort through to the 1978 cohort. More precisely, entry-level wages successively improve for these cohorts, while wage growth as they age declines. This pattern appears to stop with the 1978 cohort. Then, from the 1978 cohort on, the cohort profiles successively shift down, with no discernable pattern of change in the slope. The size of the latter decline is quite substantial, being of the order of 20 percent in just twelve years. The observed pattern is consistent with improving labour market conditions for the high-school educated in the 1970s, creating both increased wage growth for older workers in those years and higher entry-level wages for newer workers.

3.1.2. Females

Figures 2a and 2b contain the smoothed age-earnings profiles for high-school and university educated women, respectively.[11] For males, one can see higher earnings growth with age for university educated versus high-school educated workers. For females, however, university educated workers also have entry-level earnings that are considerably higher than those for the high-school educated. When we compare figures 1 and 2, a key difference between males and females is the much lower weekly earnings levels for females with comparable education and age (note that

11 The smoothing and removal of cycle effects are performed in the same way as for the male data discussed in section 3.1.1.

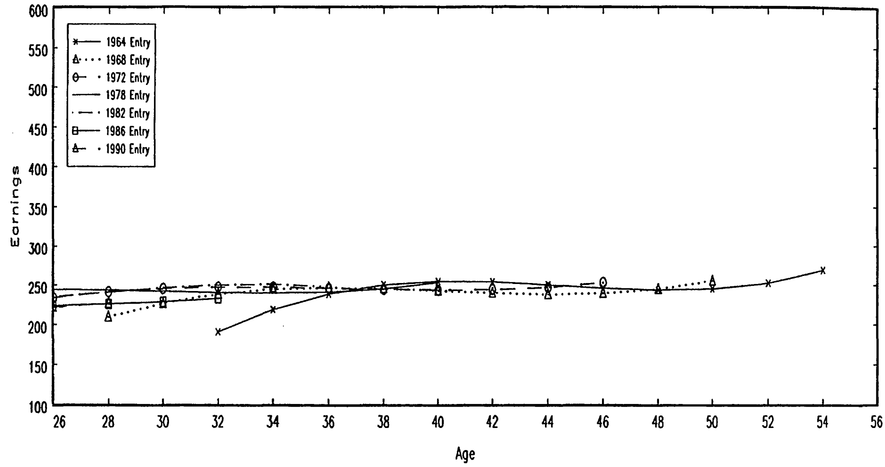

FIGURE 2a Real weekly earnings by cohort, smoothed and cycle effects removed: females, some or completed high-school education

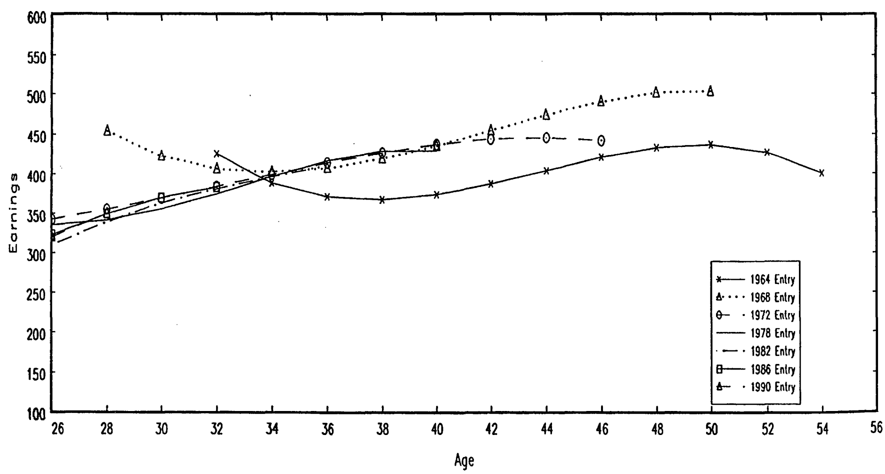

FIGURE 2b Real weekly earnings by cohort, smoothed and cycle effects removed: females, university education

the vertical axes in the male and female figures cover the same earnings difference, but the female axis is set at a lower level). Also, the slopes of the earnings profiles are much greater for men than for women. This is particularly true for the university educated and for the high-school educated in cohorts that entered before 1978.

Perhaps the most important difference between the male and female plots for our purposes is the lack of a strong cohort pattern in the female data.[12] For high-school educated females, all the cohort profiles are essentially flat and lie in approximately a $50 range. For university educated females, the profiles have a positive slope, and, with the exception of the 1964 entry cohort, the profiles for all the different cohorts lie very close to one another. It is important to keep in mind that the period covered by this data is one of large movements of females into the labour market. Thus, the patterns in these figures likely reflect significant selection effects along with whatever influences affect male average earnings levels. We return to this point in section 4.

3.2. Statistical analysis for average weekly earnings

We now turn to a statistical and more systematic examination of our data. The objective remains to characterize the major trends affecting cohort-specific age-earnings profiles.

3.2.1. Males

In tables 1 and 2 we report regression results associated with parameterizing changes in age-earnings profiles for high-school and university educated males, respectively. In each case, the dependent variable is the logarithm of the average wage for a given cohort-education group in a given year. The different columns in the tables correspond to different regression specifications. The specification in the first column corresponds to the estimates associated with regressing the log average real weekly earnings of an age-education-year cell on a quadratic in the cohort entry year, a cubic in age, an interaction of the linear age and cohort terms, and the detrended unemployment rate for males age 45 to 54.[13] We again include the latter variable to control for variation over the business cycle and instrument using the detrended U.S. unemployment rate. The standard errors are corrected for the form of heteroscedasticity associated with observations that themselves are averages of other data.[14]

When we focus on the cohort variables, we see that column 1 of table 1 indicates a pattern very similar to that suggested by the visual inspection of figure 1a. In particular, the cohort and cohort-squared variables indicate that the age-earnings profiles moved up for the early cohorts but eventually shift back down. The cohort-age interaction variable indicates that age-earnings profiles have been getting

12 Burbidge, Magee, and Robb (1997) present results showing that this lack of a cohort pattern for females is true primarily for the post-1960 entry cohorts. For pre-1960 cohorts, successive cohorts experienced substantial upward shifts in their age-earnings profiles.

13 The cohort variable is defined such that the first cohort in our sample, the 1962 cohort, is given a value of 1 and successive cohorts are counted up incrementally.

14 In particular, we assume that the variance-covariance matrix of the disturbance terms equals a scalar, sigma, times a diagonal matrix with elements equal to the number of individual observations used to construct average earnings for each cohort-age-sex-education cell. The latter matrix is normalized so that its trace equals 148 (the number of observations used in our regressions). In forming the standard errors, we use the squared standard error of the relevant regression as an estimate of sigma and multiply this by the normalized, diagonal matrix just described.

TABLE 1
High-school educated men

	1	2	3
Cohort	0.040	0.040	.
	(0.008)	(0.008)	
Cohort Sq.	−0.0028	−0.0028	.
	(0.0004)	(0.004)	
1978-Cohort	.	.	0.157
			(0.045)
1992-Cohort	.	.	−0.011
			(0.061)
Cohort * Age	−0.0021	.	−0.0020
	(0.0004)		(0.0004)
Age * Year	.	−0.0011	.
		(0.0002)	
Age	0.040	0.032	0.042
	(0.005)	(0.004)	(0.005)
Age square	−0.0015	−0.0005	−0.0016
	(0.0003)	(0.0003)	(0.0003)
Age cube	0.00002	0.00002	0.00002
	(0.000008)	(0.000008)	(0.000007)
U. Rate	−1.42	−1.42	−1.69
	(0.55)	(0.55)	(0.41)
Intercept	5.81	5.81	5.79
	(0.038)	(0.038)	(0.043)
Adj-R2	0.807	0.790	0.795

NOTES: Standard errors are in parentheses. Dependent variable: log of real avg. weekly earnings. 148 observations

flatter – not steeper – for more recent cohorts. Moreover, the specification in the second column, in which the age-cohort interaction is replaced with an age-year interaction, shows that age profiles, in general, have been getting flatter in recent years. Note that the latter effect is highly significant and therefore places in doubt the view that the increasing age-earnings differentials observed in cross-sectional data for males in Canada reflect increased returns to experience.[15]

The combination of the shifting and flattening of cohort-profiles can best be seen by plotting fitted cohort-specific age-earnings profiles. In figure 3a we plot fitted age-earnings profiles calculated from the estimates presented in column 1 of table 1.[16] Parts of these fitted profiles, particularly the parts at older ages for the newer cohorts, are projections well out of sample. We present them because they help to amplify the trends we are trying to point out, but the specific out-of-sample values should be treated with some caution.

15 Morissette and Berube's (1996) finding, with true longitudinal data, that more recent cohorts of labour market entrants face longer spells of low earnings than earlier cohorts fits with our results on returns to experience.

16 For these fitted profiles, the detrended unemployment rate is held constant at zero. Plots obtained using results in the other columns of table 1 are virtually identical. We chose to base our plots on the simplest specification.

TABLE 2
University educated men

	1	2	3	4
Cohort	−0.012	−0.012	−0.014	.
	(0.01)	(0.01)	(0.003)	.
Cohort Sq.	−0.0001	−0.0001	.	.
	(0.0006)	(0.0006)		
1978-Cohort	.	.	.	−0.117
				(0.056)
1992-Cohort	.	.	.	−0.215
				(0.076)
Cohort * Age	−0.0004	.	−0.0003	−0.0003
	(0.0006)		(0.0003)	(0.0005)
Age * Year	.	−0.0002	.	.
		(0.0003)		
Age	0.065	0.063	0.064	0.063
	(0.007)	(0.005)	(0.005)	(0.006)
Age square	−0.0027	−0.0025	−0.0030	−0.0026
	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Age cube	0.00003	0.00003	0.00003	0.00003
	(0.000009)	(0.00001)	(0.000009)	(0.000009)
U. Rate	−1.66	−1.66	−1.74	−2.46
	(0.686)	(0.686)	(0.682)	(0.665)
Intercept	6.14	6.14	6.15	6.14
	(0.056)	(0.056)	(0.028)	(0.059)
Adj-R2	0.924	0.924	0.924	0.923

NOTES: Standard errors are in parentheses. Dependent variable: log of real avg. weekly earnings. 148 observations

In figure 3a we provide a striking depiction of the flattening of the age-earnings profiles. Moreover, there is a clear pattern of shifts up in the cohort-specific profiles until the 1978 entry cohort, followed by accelerating shifts down for more recent cohorts as the cohort squared term in the specification begins to dominate. The solid line in the figure shows the cross-sectional age-earnings profile for the year 1990. Note that this cross-sectional profile is much steeper than any of the cohort-specific age-earning profiles. Thus, for this sample period, the cross-sectional profile does not reflect the actual wage progression with age for any cohort of individuals and so does not provide a useful basis for predicting future outcomes for recent labour market entrants. Closer examination of the figures reveals that for an earlier year, when the cohort profiles are not as spread out, the cross-sectional profile would be flatter. This, combined with the flattening of the cohort specific profiles, reinforces the point that the observed increase in experience differentials across successive cross-sectional datasets is due to changes in earnings outcomes across cohorts rather than a steepening of the age-earnings profile within cohorts.

Column 3 of table 1 contains results from a regression in which the quadratic specification for cohort is replaced by a fully flexible profile with a cohort dummy for each cohort. The omitted category is the 1962 cohort. Instead of reporting the

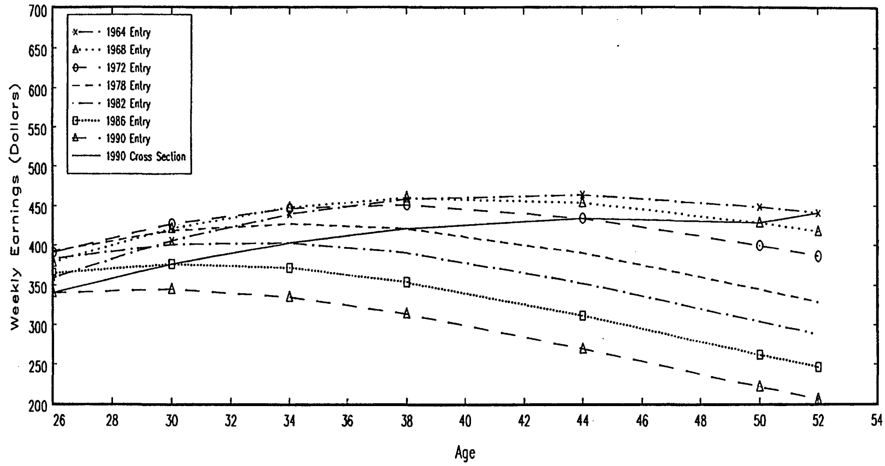

FIGURE 3a Age-earnings profiles allowing differing slopes by cohort: males, some or completed high-school education

values of all fifteen cohort dummies, we report only the values for the 1978 and 1992 cohorts, since they summarize well the overall pattern. As was suggested by the quadratic specification, the cohort-dummy variable specification indicates improving cohort performance until 1978 and then deteriorating performance afterwards. Moreover, the cohort-age interaction term again is significantly negative, indicating that more recent cohorts experience less wage growth as they age than previous cohorts experienced.

Overall, in table 1 and figures 1a and 3a a consistent picture of the changes in age-earning profiles for high-school educated men over the last two decades is offered. The evidence clearly suggests that labour market entrants since 1978 have been performing poorly in comparison with previous cohorts. In particular, recent cohorts appear to start at lower wages and experience slower wage growth as they age than earlier cohorts did. A simple projection of these trends suggests that more recent cohorts are unlikely to catch up to the wages of older cohorts unless there is a major reversal in the underlying forces driving labour market trends.

For university educated men, the estimated cohort terms in column 1 of table 2 indicate that starting wages for successive cohorts have been falling continuously over the entire 1971 to 1993 period. Moreover, the cohort-age interaction variable does not indicate that age-earnings profiles have been steepening across cohorts. The age-year interaction term in the second column also shows no appreciable change in the age slope over time. Again, there is no evidence of increasing returns to experience. The quantitative implications of these estimates are seen in the fitted cohort-specific age-earnings profiles plotted in figure 3b. The pattern observed in figure 3b is simpler than that of figure 3a, but the overall message is similar: there

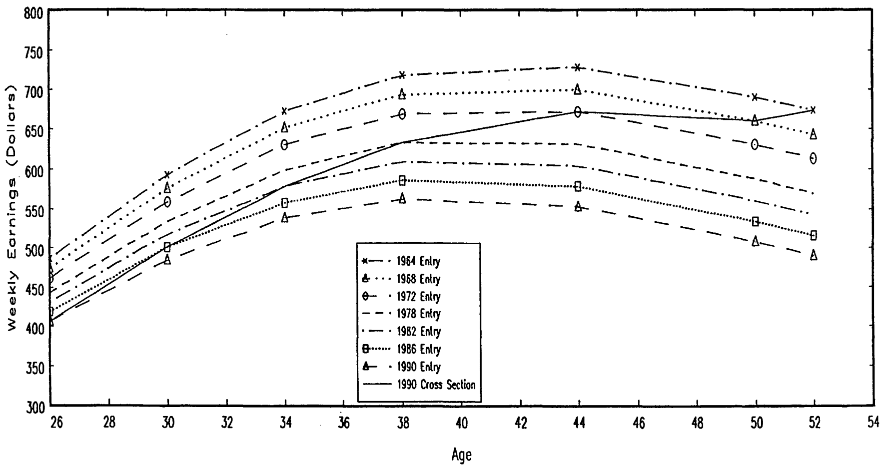

FIGURE 3b Age-earnings profiles allowing differing slopes by cohort: males, university education

is no indication that recent cohorts, who are starting at lower entry-level wages, should be expected to experience higher wage growth as they age than that experienced by older workers, and thus, there is no basis for predicting that they will attain similar wage levels. Once again, the 1990 cross-sectional profile is substantially steeper than the age-earnings progression actually experienced by any cohort. Given the linear downward trend in cohort earnings experiences, however, there is no suggestion that the cross-sectional profiles are getting steeper with time. This, in fact, is what is found in the cross-sectional literature, which indicates that experience differentials increased substantially in the 1980s and early 1990s for less educated workers but showed at most mild increases for university educated workers.

In columns 3 and 4 in table 2 we explore the robustness of the pattern reported in column 1. Since the estimates on cohort and cohort squared in column 1 are not individually significantly different from zero (although they are jointly significant with a p-value less than 0.001), we drop the cohort squared term in column 3 in order to verify the robustness of the observation that cohorts are successively starting at lower wages and are not experiencing higher wage growth. This is, indeed, supported by the results in column 3.

Column 4 of table 2 corresponds to a specification in which the quadratic cohort term is replaced by a full set of cohort-dummies. As in table 1, here we report only two of the fifteen cohort dummies, since that is all that is needed to represent the overall pattern. Both the 1978 and 1992 cohort dummies reported in the table indicate once again the general pattern of deteriorating labour market performance for successive waves of labour market entrants. The point estimates indicate that the 1992 cohort earns approximately 20 percent less than their counterpart did twenty years earlier.

A final point of interest relates to individual choice of education level. Consider a simple signalling model in which education is used to signal ability, and wages equal the marginal product of the worker. Assume that the relation of levels of education to ability changes, but the overall average wage, once signalling accurately reveals worker productivity, does not change. If the level of education needed to signal high ability is raised, the result will be a shift of high-ability high-school workers into the university educated category. Since the shifting workers will be of low ability relative to others in the university category, the average wage for both high-school and university workers will fall. Thus, one might observe the cohort patterns detailed here without falling average wages for the cohorts as a whole. To investigate this possibility, we re-estimated the specification in column 1 of the tables using overall average wages for all workers with either a high school or university education in a given cohort. The result of this estimation is very similar to that reported in column 1 for high-school educated men. Thus, a simple compositional shift story cannot by itself explain the observed wage patterns.

3.2.2. Females

Table 3 contains regression results for high-school educated female workers. The specification and econometric approach are identical to those described for males. At first glance, the cohort variable results in column 1 appear to indicate that less educated females are experiencing similar cross-cohort patterns to less educated males. The plots of the cohort-specific age profiles derived from the column 1 estimates presented in figure 4, however, indicate very little similarity. In contrast to their male counterparts, less-educated females show little in the way of differences across cohorts in our sample period. This occurs because the cohort and cohort squared terms presented in table 3 offset each other. The table and the figure also indicate that there is little variation in the age profile of any cohort with age and that the age profiles have not been steepening in recent years.

For the university female results in table 4, one again sees evidence of cohort patterns that are similar to those presented for males. As with the more educated males, the specification that excludes the cohort squared term indicates statistically significant falls across cohorts. The plots in figure 6b, however, indicate that the drops in the intercepts of the age profiles across cohorts do not at all rival the falls for comparable males. Moreover, in contrast to the males, the declines in the intercept are offset by increases in the slopes of the age profiles over time. University educated females are the one group for whom one could argue there has been an increase in the return to experience in recent years. This shows up in the figure in the form of very large age slopes for the most recent cohorts. Since the high earnings levels predicted for the most recent cohorts later in life correspond to age ranges that are well out of sample, such extrapolation should be interpreted with caution.

Finally, as discussed earlier, attempts to correct for effects of the 1990 education category changes do not affect estimates using male data but do have an impact on female results. In specifications including a post-1990 dummy variable to control

TABLE 3
High-school educated women

	1	2	3
Cohort	0.032	0.032	.
	(0.012)	(0.012)	
Cohort Sq.	−0.0019	−0.0019	.
	(0.0007)	(0.0007)	
1978-Cohort	.	.	0.173
			(0.067)
1992-Cohort	.	.	0.081
			(0.083)
Cohort * Age	−0.0008	.	−0.0009
	(0.0006)		(0.0006)
Age * Year	.	−0.0004	.
		(0.0003)	
Age	0.010	0.007	0.012
	(0.007)	(0.005)	(0.007)
Age square	−0.00008	0.0003	−0.0002
	(0.0004)	(0.0004)	(0.0004)
Age cube	−0.000002	−0.000003	−0.000001
	(0.00001)	(0.00001)	(0.00001)
U. Rate	−2.18	−2.18	−2.23
	(0.75)	(0.75)	(0.77)
Intercept	5.33	5.33	5.31
	(0.057)	(0.057)	(0.043)
Adj-R2	0.179	0.179	0.144

NOTES: Standard errors are in parentheses. Dependent variable: log of real avg. weekly earnings. 148 observations

for level shifts created by the education category changes, the female results are more similar to those for males. In particular, for high-school educated females, the cohort and cohort squared coefficients are larger in absolute value and the cohort-age interaction term becomes negative and significant. For university educated females, the cohort-age interaction term ceases to be statistically signficantly different from zero. Thus, conclusions based on the male data are reinforced.

3.3. Changes in distribution

In the previous two sections we documented changes in the age-profiles for mean earnings across cohorts. While tracking cohort means may provide insight into overall earnings trends, it yields only a partial picture for evaluating changes in inequality. In particular, in many cross-section-based studies of male earnings inequality, changes in inequality for workers in the same observable skill group have been found to contribute substantially to the overall inequality growth. Since we want to examine this issue after controlling for differences across cohorts, the appropriate concept in our context is to follow earnings inequality over time for specific cohort-education groups. We choose to follow changes in the shape of the earnings distri-

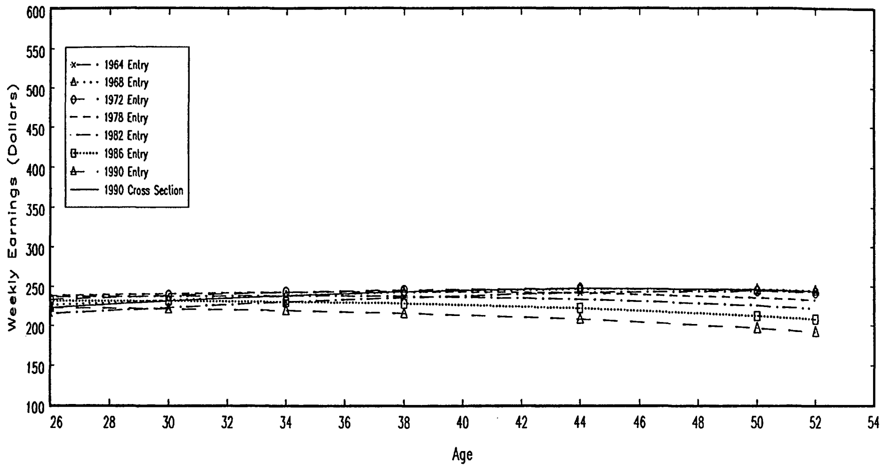

FIGURE 4a Age-earnings profiles allowing differing slopes by cohort: females, some or completed high-school education

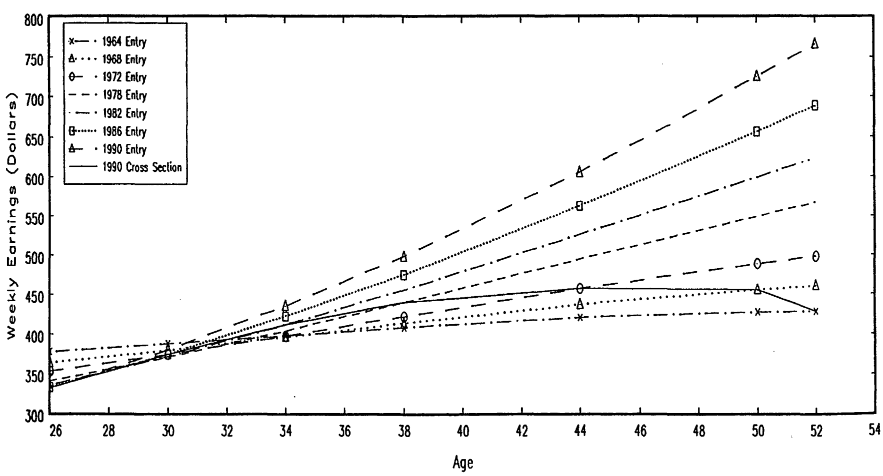

FIGURE 4b Age-earnings profiles allowing differing slopes by cohort: females, university education

bution for specific cohorts using the 90-10 percentile differential: the difference between the weekly earnings observation at the 90th percentile of the distribution minus the observation at the 10th percentile. We calculate this statistic for each of our age-education-year cells.

TABLE 4
University educated women

	1	2	3	4
Cohort	−0.024	−0.024	−0.009	.
	(0.019)	(0.020)	(0.004)	.
Cohort Sq.	0.0008	0.0008	.	.
	(0.0009)	(0.0009)		
1978-Cohort	.	.	.	−0.102
				(0.100)
1992-Cohort	.	.	.	−0.145
				(0.105)
Cohort * Age	0.0021	.	0.0015	0.0018
	(0.0008)		(0.0003)	(0.0008)
Age * Year	.	0.0011	.	.
		(0.0004)		
Age	0.004	0.014	0.064	0.007
	(0.009)	(0.006)	(0.005)	(0.009)
Age square	0.00009	−0.0009	−0.0001	−0.00009
	(0.0005)	(0.0005)	(0.0004)	(0.0005)
Age cube	−0.000005	−0.000005	−0.000003	−0.000001
	(0.00001)	(0.00001)	(0.00001)	(0.00001)
U. Rate	−1.237	−1.237	−1.179	−1.214
	(0.895)	(0.895)	(0.890)	(0.873)
Intercept	5.95	5.95	5.89	5.92
	(0.092)	(0.092)	(0.039)	(0.097)
Adj-R2	0.619	0.619	0.618	0.644

NOTES: Standard errors are in parentheses. Dependent variable: log of real avg. weekly earnings. 148 observations

3.3.1. Males

Figures 5a and 5b plot the smoothed age-profile of the 90-10 differential for several cohorts of males.[17] The smoothed profiles are obtained as in the average earnings figures by regressing the 90-10 differentials on the detrended unemployment rate and a cohort-specific cubic in age. It is clear from figures 5a and 5b that the 90-10 differential tends to increase with age. This a well-known fact that has been studied extensively in the search-matching literature. It also fits with Mincer's (1974) derivations from the human capital model.[18]

The somewhat surprising observation from figures 5a and 5b is the rather stable pattern of the smoothed 90-10 profiles. Close inspection of the figures nevertheless reveals a pattern in which the 90-10 differential increased rapidly with age at young ages for the earliest cohorts. The age-earnings differential slope becomes much flatter across successive cohorts up to the 1978 entry cohort and remains

17 The raw profile plots are available upon request.
18 Dooley and Gottschalk (1984) examine the empirical implications of Mincer's statements using cohort data for the United States.

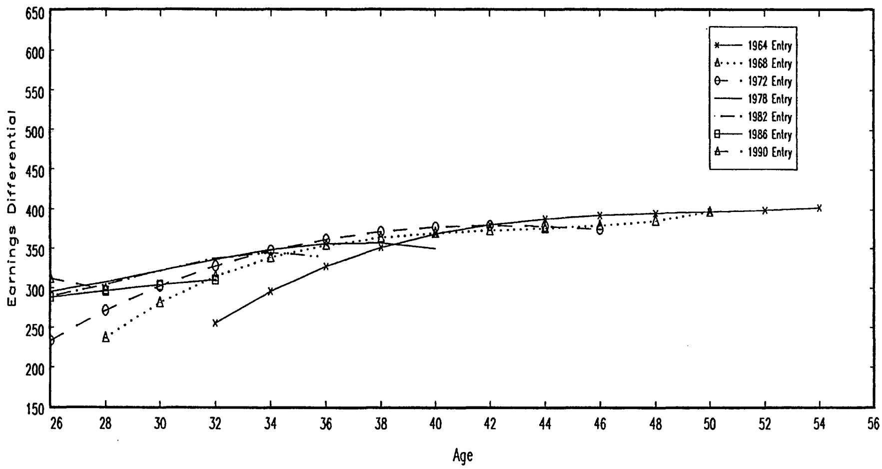

FIGURE 5a 90-10 differential for real weekly earnings by cohort, smoothed and cycle removed: males, some or completed high-school education

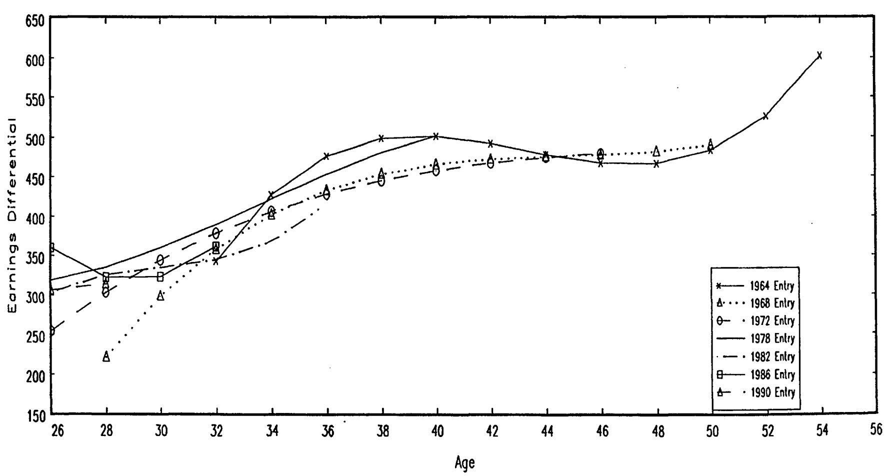

FIGURE 5b 90-10 differential for real weekly earnings by cohort, smoothed and cycle removed: males, university education

remarkably stable for successive cohorts. This pattern is particularly evident for the high-school educated workers. Inspection of figures for the 10th and 90th percentiles separately (not presented here) show that the main difference between earlier and later cohorts is found in the behaviour of the 90th percentile with age.

The 90th and 10th percentiles were closer at age 25 for earlier cohorts, and the 90th percentile grew faster with age for those cohorts. For later cohorts, the 90th percentile is relatively larger at age 25, and the gap does not increase substantially with age.[19]

The pattern in figures 5a and 5b is important because it indicates that increases in male inequality in the 1980s in Canada did not occur because of increases in inequality within cohort groups. Put in other terms, this result indicates that the declining outcomes of successive cohorts reflected in the mean earnings results were shared by all members of the cohorts: the earnings distributions shift down but maintain a constant spread.

In order to describe the movements in the 90-10 differentials more systematically, in table 5 we report regression results associated with parameterizing those movements. The first two columns of table 5 correspond to results for the high-school educated group, while the last two columns correspond to results for the university educated group. In all cases the 90-10 differential is regressed on a full set of cohort dummies, a cubic in age, a linear-cohort-age interaction term, and detrended unemployment.[20] The dependent variable in columns 1 and 3 is the 90-10 differential measured in levels while in columns 2 and 4 the difference in the logs of the 90th and 10th percentiles are used. The latter results are included to provide an easier match to the existing cross-sectional literature, in most of which percentile differences in log earnings as a measure are used.

As the inspection of figures 5a and 5b suggested, columns 1 and 3 of table 5 indicate that within-cohort inequality increases substantially across successive cohorts up to the 1978 cohort and then virtually stops. For the high-school educated workers, the results in logs in column 2 indicate that the rise in inequality continued after 1978. The difference between the pattern suggested by columns 1 and 2 is easily reconciled once it is recalled that the mean earnings for newer cohorts of high-school educated men fell rapidly after 1979. The constant dispersion in levels combined with a falling mean results in increasing relative dispersion, which is what is captured by the difference in logs. Therefore, the continual, increased widening of the within-cohort distribution suggested by column 2 can be viewed as driven entirely by the fact that the mean cohort-specific earnings were falling, not by an expanding distribution around the mean.[21] A similar, though somewhat weaker, pattern is evident for the university educated. It should also be noted that, even for the 90-10 differential in log earnings, the increase after the 1978 cohort occurs at a much slower rate than that observed for cohorts prior to 1978.

19 Riddell and Sweetman (1997) show that later cohorts of workers officially classified as high-school educated are more likely to have acquired extra educational certificates. These extra signals could make it easier for employers to differentiate among high-school educated workers at time of hiring. This, in turn, could account for larger differentials at younger ages for more recent cohorts.

20 We again instrument for the detrended unemployment rate using the detrended U.S. unemployment rate and a post-1982 dummy variable. The standard errors are corrected as described in footnote 8.

21 MaCurdy and Mroz (1995) make a similar observation for the United States.

TABLE 5
90-10 percentile differences: males

	HS Level 1	HS Log 2	UNIV Level 3	UNIV Log 4
1964-Cohort	34.2	0.08	−70.0	0.01
	(8.7)	(0.04)	(24.3)	(0.08)
1968-Cohort	80.4	0.21	3.8	0.09
	(11.1)	(0.05)	(32.3)	(0.09)
1972-Cohort	112.4	0.29	43.1	0.26
	(13.5)	(0.06)	(40.5)	(0.13)
1976-Cohort	136.1	0.40	78.5	0.41
	(15.2)	(0.07)	(46.3)	(0.14)
1980-Cohort	151.7	0.47	98.4	0.49
	(16.1)	(0.07)	(50.1)	(0.16)
1984-Cohort	156.9	0.55	97.2	0.56
	(17.0)	(0.07)	(53.3)	(0.17)
1988-Cohort	137.8	0.52	107.5	0.60
	(17.6)	(0.08)	(53.9)	(0.17)
1992-Cohort	138.4	0.66	84.9	0.54
	(21.1)	(0.09)	(59.5)	(0.19)
Cohort * Age	−1.27	−0.002	−1.11	−0.003
	(0.15)	(0.0006)	(0.43)	(0.0013)
Age	21.3	0.038	31.7	0.039
	(1.8)	(0.008)	(5.1)	(0.016)
Age square	−0.54	−0.0003	−1.07	−0.0013
	(0.12)	(0.0005)	(0.30)	(0.001)
Age cube	0.004	0.0000001	0.016	0.00004
	(0.003)	(0.00001)	(0.007)	(0.00002)
U. Rate	252.9	1.42	269.3	0.23
	(197.5)	(0.87)	(521.5)	(1.62)
Intercept	150.5	0.59	226.3	0.54
	(14.7)	(0.06)	(46.4)	(0.14)
Adj-R2	0.842	0.718	0.783	0.435

NOTES: Standard errors are in parentheses. 148 observations

3.3.2. Females

The smoothed 90-10 differential profiles for females are presented in figures 6a and 6b. For high-school educated females, the profiles are noticeably lower and flatter than for their male counterparts. For university educated females, levels of inequality are similar at young ages to those for males but do not rise as quickly with age. As in the case of males, there is little clear evidence of changes in the level of inequality at any age across cohorts.

In table 6 we provide a more systematic evaluation of changes in the 90-10 differential for females. As in the case of high-school educated males, the cohort effects for less-educated females in column 1 contain a rising trend until 1978 and thereafter are flat or declining. The size of the inequality changes, however, are much smaller than for the comparable male group. The results for university educated females in column 3 show no discernable trends after the 1970 entry cohort

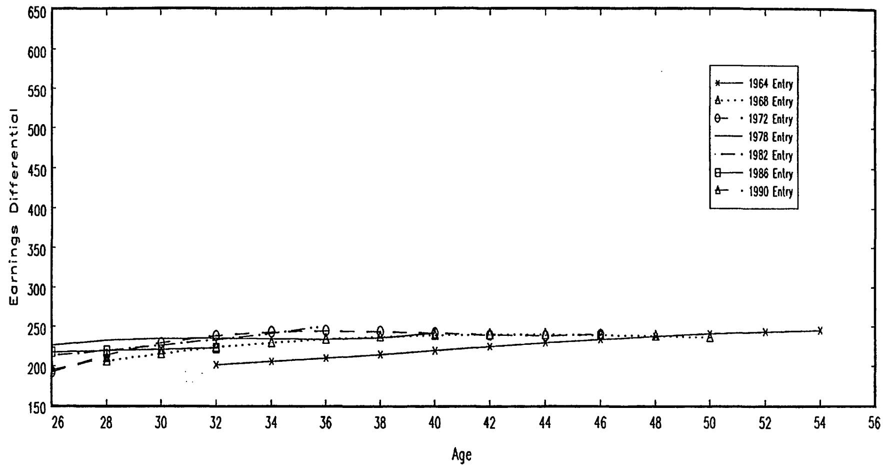

FIGURE 6a 90-10 differential for real weekly earnings by cohort, smoothed and cycle removed: females, some or completed high-school education

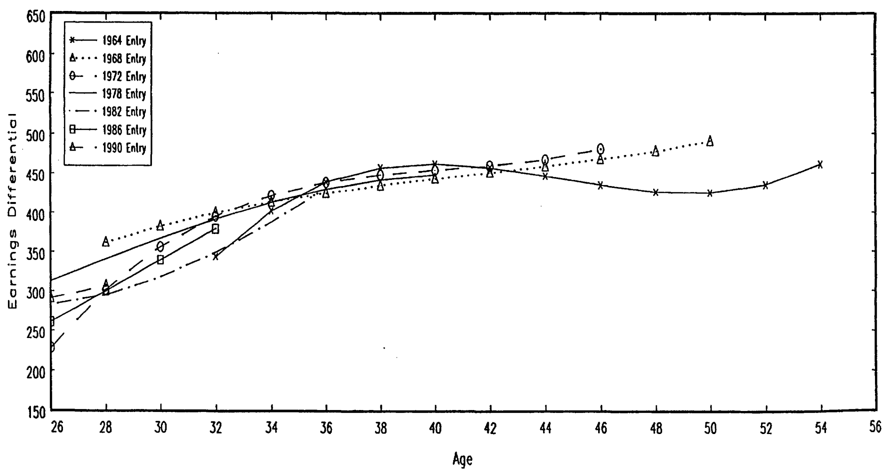

FIGURE 6b 90-10 differential for real weekly earnings by cohort, smoothed and cycle removed: females, university education

and no cohort differences that are statistically significant at conventional levels. In contrast to the situation for males, the log differential results in columns 2 and 4 do not display sharply different patterns relative to the results based on differentials in levels. This is because there is no strong trend up or down in average earnings across cohorts for females.

TABLE 6
90-10 percentile differences: females

	HS Level 1	HS Log 2	UNIV Level 3	UNIV Log 4.
1964-Cohort	7.9	0.04	7.8	0.20
	(6.9)	(0.06)	(24.0)	(0.16)
1968-Cohort	30.8	0.26	35.3	0.17
	(9.0)	(0.08)	(32.7)	(0.21)
1972-Cohort	45.1	0.32	32.7	0.40
	(11.3)	(0.10)	(42.0)	(0.28)
1976-Cohort	52.1	0.37	47.4	0.52
	(12.7)	(0.11)	(47.7)	(0.31)
1980-Cohort	53.1	0.41	45.6	0.60
	(13.5)	(0.12)	(51.1)	(0.34)
1984-Cohort	58.0	0.49	33.7	0.65
	(14.1)	(0.13)	(53.1)	(0.35)
1988-Cohort	38.5	0.38	13.5	0.37
	(14.4)	(0.13)	(52.6)	(0.34)
1992-Cohort	41.1	0.52	48.5	0.60
	(16.3)	(0.15)	(51.8)	(0.34)
Cohort * Age	−0.39	−0.005	−0.32	−0.005
	(0.12)	(0.0011)	(0.41)	(0.0027)
Age	8.1	0.089	27.2	0.18
	(1.4)	(0.013)	(4.5)	(0.030)
Age square	−0.34	−0.0005	−1.32	−0.011
	(0.09)	(0.0008)	(0.24)	(0.002)
Age cube	0.005	0.00008	0.021	0.0002
	(0.002)	(0.00002)	(0.006)	(0.00004)
U. Rate	−96.5	−0.65	−149.9	−2.54
	(151.5)	(1.4)	(431.5)	(2.83)
Intercept	164.4	1.05	262.0	0.73
	(12.5)	(0.11)	(48.4)	(0.32)
Adj-R2	0.392	0.260	0.729	0.383

NOTES: Standard errors are in parentheses. 147 observations

3.4. Summary

For high-school educated men, an examination of the weekly earnings outcomes of labour market entry cohorts over the period 1971 to 1993 reveals that successive cohorts up to the 1978 cohort experienced higher entry wages, more within-cohort earnings dispersion, and flattening age-earnings profiles. For more recent cohorts, the pattern has changed substantially. Successive waves of new entrants since 1978 have received lower entry wages, age-earnings profiles that are either flattening further or not changing, and no discernable change in the distribution of earnings around the cohort-specific mean at a given age. For university educated men, there has been a gradual but continuous deterioration in the labour market performance of newer cohorts. This deterioration in performance was accompanied by increasing within-cohort inequality up to approximately the 1978 cohort. Thereafter, within-cohort dispersion has been relatively constant across cohorts for any given age.

Moreover, for both groups of workers, there is no evidence to suggest that the observed increases in earnings differentials across age groups in cross-sectional data reflect an increased return to gaining experience, and therefore this increased differential should not be taken as indication of likely faster wage growth in the future for currently young workers. Together, these results raise grave concerns about intergenerational equity among males in the labour market. Recent male labour market entrants, regardless of their education, are earning dramatically less than their predecessors earned at the same age, and there is little reason, based on current trends, to expect them to catch up.

For female workers, the patterns are very different. In contrast to the males, neither high-school or university educated females experienced substantial differences in average weekly earnings across cohorts. As in the situation of males, there is no evidence of increased dispersion in earnings across cohorts at any age. The university educated females are the only group who experience increasing age profiles of average earnings. For this group, it may therefore be the case that there has been an increase in the returns to experience. It must be noted, however, that the relationship between age and work experience has been changing substantially as women have become more attached to the labour market.

4. Interpretation and implications

4.1. Changes in age-earnings profiles and skilled-biased demand shift hypotheses

As discussed in the introduction, the source of increased male earnings inequality over time in both Canada and other countries has been the focus of considerable study. There appear to be at least three main contenders for explaining the trends: (1) technical change favouring skilled workers at the expense of less skilled workers; (2) increased trade with lower-skilled countries, effectively increasing competition for low-skilled workers; and (3) an erosion of labour market institutions, such as minimum wages and unions, that support the lower tail of the earnings distribution. Evidence for the skill-biased technical change hypothesis has been built mostly on U.S. data in a series of articles, including Juhn, Murphy, and Pierce (1993) and Bound and Johnson (1992). The former argues that increased educational and experience differentials over time in the United States point to a skill-biased increase in demand. Further, increases in relative dispersion within skill groups is interpreted as reflecting increased demand for unobservable skills over time. Together these results are taken as evidence that changes in the male earnings structure over the 1980s were generated by a general increase in skill premia, as induced by a skill-biased demand shift.

The skill-biased demand shift hypothesis has also been advanced to explain increased male inequality in Canada. Not all the Canadian evidence, however, lines up as neatly as that presented for the United States by Juhn, Murphy, and Pierce (1993). For example, the educational differential did not increase substantially in Canada in the 1980s (see Bar-Or et al. 1995). Nonetheless, proponents of the skill-biased demand shift hypothesis point to increased experience-age differentials over time and increased within-skill group dispersion as evidence in favour of their position.[22] Our objective in this section is to interpret the results of the previous section in the light of the debate on the importance of increased skill premia in explaining movements in relative wages.

4.2. Structural interpretations and cohort, time, and age decompositions

In the analysis of cohort data, it is common in the literature to provide a breakdown in terms of time, cohort, and age effects (with no interaction between these terms).[23] It is well known that providing such a decomposition involves solving a basic identification problem (see Heckman and Robb 1985). The identification problem arises because, if we know any two of a person's age in the data year, the data year, and the year of cohort entry into the labour force, we can exactly determine the third. Thus, separating out the three effects is not possible without incorporating extra information in the form of restrictions.

It is important to note that the results presented in the previous sections are not meant to offer a decomposition in terms of time, age, and cohort effects. Instead, the empirical evidence presented in section 3 should be viewed as offering a reduced-form representation of trend movements in the cohort-earnings data. Our approach, in both our regression specifications and our smoothed data plots, has been only to remove common temporary (cyclical) effects and smooth out the trend effects. This procedure does not imply, however, that we we have eliminated all time effects or separated them out from cohort and age effects.[24] For example, the deterioration in earnings across cohorts documented in section 3 could equally well be generated by a downward trend affecting members of all cohorts (i.e., a time effect) that is disguised for older cohorts by offsetting age effects or by real (structurally meaningful) cohort effects.[25] At this point, we remain agnostic as to the split of these observed

22 The lack of an increase in educational differentials is explained as arising from offsetting increases in supply.

23 Age effects are defined as changes in earnings as individuals age that are common across entry cohorts and independent of aggregate shifts in the economy. The effects of aggregate shifts, assumed to affect workers of all cohorts and all ages in the same way, are termed 'time effects.' Finally, cohort effects are defined as permanent differences in earnings differentials across entry cohorts that exist regardless of the age at which we observe the cohort or the size of the common time effect.

24 Although our cycle correction is not intended to remove all time effects, the reader may be concerned that our results are sensitive to the way we remove cyclicality. This is not the case. Smoothed plots without removal of cyclical effects are very similar to the smoothed plots seen in section 3 (these results are available upon request). We also tried different means of removing cyclical effects, including regressions containing year effects that were constrained to be equal at cyclically similar points. These provided virtually identical results to those in the paper.

25 A structurally meaningful cohort effect is a permanent effect that is determined before a specific cohort enters the labour market and determines earnings levels regardless of the economic times the cohort lives through. Differences in school quality across cohorts could generate this type of cohort effect.

trends into age, time, and cohort effects. Instead, we argue that the reduced-form evidence can be used directly to evaluate a particular hypothesis regarding the cause of increased dispersion in male weekly earnings, without the need to identify separate time, age, and cohort effects.[26]

In order to make the discussion transparent, it is helpful to derive explicitly the implications of a general skill-biased labour demand shift for cohort-profiles in the context of a simple model. In this model, we allow there to be two skill classes and we allow productivity to depend on age (experience). Aggregate production, Y_t, is given by

$$Y_t = F(L_{1,t}, L_{2,t}, t) \tag{1}$$

where $L_{i,t}$ is the quantity of effective labour of skill class i in period t. The explicit dependence of the production function on time t captures the possibility of technological change and/or changes in the relative prices of domestic- versus foreign-produced goods. The quantity of effective labour for a particular skill class is given, in turn, by

$$L_{i,t} = \sum_{j \in J} g^i(a_{j,t}, t) l_{i,j,t} \tag{2}$$

where j indexes cohorts, $a_{j,t}$ is the age of cohort j at time t, $l_{i,j,t}$ is the level of employment in skill class i of cohort j at t and $g^i(a,t)$ is the effective units of labour of a worker of age a in skill class i at time t. The explicit dependence of $g(\cdot)$ on time is meant to capture the possibility that the productivity value of experience may change over time.

Under the assumption that workers are paid their marginal products, and denoting the derivative of $F(\cdot)$ with respect to its ith argument by $F^i(\cdot)$, we can see that the log wage at t of a worker of skill class i and cohort j is

$$\ln w_{i,j,t} = \ln(F^i(L_{1,t}, L_{2,t}, t)) + \ln(g_i(a_{j,t}, t)). \tag{3}$$

The above specification implies that log wages of skill class i can be decomposed in a time effect $(\ln(F^i(L_{1,t}, L_{2,t}, t)))$ and a time-varying age effect $(\ln(g_i(a_{j,t}, t)))$, where each of these effects can have a pattern that differs across skill groups.

Within the above framework, if $i = 1$ represents the low-skill group and $i = 2$ represents the high-skill group, the skill-biased demand shift hypothesis implies that over time (1) $F^1(\cdot)$ will decrease and $F^2(\cdot)$ will increase, and (2) $g_i^a(\cdot)$ will increase. In other words, over time a skill-biased demand shift will simultaneously increase the value of marginal product of skilled workers, decrease that of unskilled

26 This reduced-form evidence also can be seen as a useful basis for short-term prediction. For example, a projection of current trends does not lead one to expect that recent male labour force entrants should expect to attain earnings comparable to those of older cohorts of males.

workers, and increase the return to age(experience). In order to see what such a change implies in terms of changes in cohort-specific age-earning profiles, it is best to consider a simple parameterization of the induced time effects and the time-varying age effects derived above. The following parameterization captures these two effects in its simplest form:

$$\ln w_{i,j,t} = \alpha_{i,0} + \alpha_{i,1} * t + \alpha_{i,2} a_{j,t} + \alpha_{i,3} a_{j,t}^2 + \alpha_{i,4} a_{j,t} * t. \tag{4}$$

In the above parameterization, the skilled-biased demand shift hypothesis can be interpreted as implying[27] that $\alpha_{1,1} < 0$ (decrease in the value of low skill), $\alpha_{2,1} > 0$ (increase in the value of high skill), and $\alpha_{i,4} > 0$ (increase in the value of experience).

Noting that since calendar time t can be expressed as the sum of cohort time and age, that is, $t = (j + a_{j,t})$, the above expression can rewritten to emphasize its implications for cohort-specific age-earning profiles as follows:

$$\ln w_{i,j,t} = \alpha_{i,0} + \alpha_{i,1} * j + (\alpha_{i,2} + \alpha_{i,1}) a_{j,t} + (\alpha_{i,3} + \alpha_{i,4}) a_{j,t}^2 + \alpha_{i,4} a_{j,t} * j. \tag{5}$$

The above equation captures the main implications of the skill-biased demand shift hypothesis for cohort-specific age-earning profiles. In particular, it suggests that for the low-skill group, the age-earnings profiles should be rotating upward, since the intercept is falling across cohorts while the slope of the age profile is increasing. In contrast, for the more skilled group, the age-earnings profiles should be increasing over time and fanning out, since the intercept and the slope are increasing. The evidence on male earnings presented in section 3 is almost exactly the opposite of this implied pattern if we interpret the high-school educated group as the low-skill group and the university educated group as the high-skill group. In this sense, the hypothesis that a skill-biased demand shift is the main driving force determining wage patterns appears to be at odds with the Canadian experience. Note that an extension of the skill-biased demand shift model to include supply shifts will not resurrect the hypothesis, because there are no increases in supply for high-school educated men after the early 1980s (see table 7).[28]

Our observations regarding within-cohort dispersion also appear to us to be inconsistent with the skill-biased demand shift hypothesis. If increased relative demand for skills of all kinds were occurring, one would expect to see evidence of it in the form of the wages of the workers at the top of the earnings distribution (those with the most unobserved skills) being bid up over time. Instead, we observe the wages of all workers falling together. It is difficult to conceive of a demand shift

27 We are disregarding for now possible supply effects that could offset the effects of the skill-biased demand shift.

28 It is possible that we do not observe skill-biased demand type patterns in the reduced form because of measurement problems. In particular, continual and increasing overstatement of the CPI, due to insufficient accounting for changing of the quality of products with technological change, is in the right direction. However, the implied errors would have to be huge: recall that real wages at a given age fell by 20 per cent for university men over our sample period, while the demand shift model would have predicted real wage increases. Further, one would have to explain differing apparent CPI effects for men versus women.

TABLE 7
Data cell sizes

	Males		Females	
Cohort	High Sch.	Univ.	High Sch.	Univ.
1962	602	120	612	51
1964	543	152	532	89
1966	574	204	614	88
1968	564	198	694	148
1970	582	223	721	139
1972	692	246	808	178
1974	713	309	799	214
1976	726	334	779	255
1978	726	281	891	248
1980	705	293	853	266
1982	801	303	853	235
1984	744	276	808	246
1986	846	278	909	264
1988	761	264	866	274
1990	705	252	713	253
1992	516	182	660	212

NOTE: See the text for a description of sample definitions.

that does not generate any gainers, even at the top of the skilled-wage distribution. Earlier claims of increasing within-group dispersion arise only because of the use of relative measures of dispersion. Thus, it appears to us that emphasis should be placed on explaining the declining average performance of successive cohorts of males rather than on explaining increased within-group dispersion and returns to experience.

In summary, the previous literature has argued in favour of a skill-biased demand shift as the prime explanation of the changing earnings patterns in Canada. Such a shift is hypothesized to cause increased educational differentials, returns to experience, and earnings dispersion within skill groups. The first of these three effects, increased educational differentials, is not very strong in Canada, as has been shown by other researchers. We have shown that there has not been any systematic increase in returns to experience or increases in absolute earnings dispersion within education-cohort groups over time in Canada. Thus, none of the three main facts in support of this hypothesis appears to exist in Canada.

It is worth emphasizing that these results do not imply a complete absence of a role for shifts in relative labour demand. The fact that the education differential has changed little in the face of substantial increases in the relative supply of university educated workers over the last fifteen years in Canada might suggest that there have been some offsetting relative demand shifts for workers of different education levels at the same time. The results do indicate that focusing on a single shift in relative demand towards workers with skills of all types is potentially misleading.

5. Conclusion

In this paper, we present new evidence on earnings inequality trends in Canada by following cohorts of labour market entrants across successive datasets. This contrasts with the previous literature on the topic, which has concentrated on examinations of earnings differentials within cross-sectional datasets at two or more points in time. The earlier literature established that, for males, earnings differentials by experience level have increased over the last fifteen years. This is often interpreted as reflecting an increase in returns to experience over time. We find, instead, that the increased differential stems from much poorer earnings outcomes for more recent male labour market entrants relative to those who entered the labour market in the 1960s and 1970s. More specifically, we find that cohort-specific age-earnings profiles have shifted down for successive cohorts since the 1978 entry cohort for high-school educated males and from an even earlier point for university educated males. Further, there is no evidence that these downward shifts have been offset by increased returns to gaining experience for recent cohorts. This raises grave concerns about intergenerational equity among males in the labour market. Finally, we find that the absolute dispersion of earnings within cohorts has not increased in recent years. For women, we find no substantial changes in earnings experiences across cohorts, either in the mean weekly earnings level at a given age or in the dispersion about that mean. We argue that these observations viewed together place in doubt the relevance of the hypothesis that increased skill premia (induced by a skilled-biased demand shift) largely explain the movements in earnings inequality observed in Canada. Other explanations for movements in inequality, including institutional changes (e.g., changes in unionization among the young), supply effects (caused by dramatic increases in post-secondary school enrolment and female labour force participation), and selection effects, should hence be explored.

References

Baker, Michael, and Dwayne Benjamin (1994) 'The performance of immigrants in the Canadian labor market,' *Journal of Labor Economics* 12, 369–405

Bar-Or, Yoram, John Burbidge, Lonnee Magee, and A. Leslie Robb (1995) 'The wage premium to a university education in Canada, 1971–1991,' *Journal of Labor Economics*, vol 13, no. 4: 762–794

Beach, Charles M., and George A. Slotsve (1996) *Are We Becoming Two Societies: Income Polarization and the Myth of the Declining Middle Class in Canada*, Social Policy Challenge Series, vol. 12 (Toronto: C.D. Howe Institute)

Beaudry, Paul, and David A. Green (1996) 'Individual Responses,' study prepared for Industry Canada, September

— (1997) 'Employment outcomes in Canada: a cohort analysis,' paper presented at the IRPP Conference on Adapting Public Policy to a Labour Market in Transition, Montreal

— (1997) 'Cohort patterns in Canadian earnings: assessing the role of skill premia in inequality trends,' NBER Working Paper 6132

Berman, Eli, John Bound, and Zvi Griliches (1994) 'Changes in the demand for skilled labor within U.S. manufacturing: evidence from the Annual Survey of Manufactures,' *Quarterly Journal of Economics* 109, 367–98
Borjas, George J., and Valerie A. Ramey (1995) 'Foreign competition, market power, and wage inequality,' *Quarterly Journal of Economics* 110, 1075–110
Bound, John, and George Johnson (1992) 'Changes in the structure of wages in the 1980's: an evaluation of alternative explanations,' *American Economic Review* 82, 371–92
Burbidge, John B., Lonnee Magee, and A. Leslie Robb (1996) 'Canadian wage inequality over the last two decades,' Department of Economics, McMaster University
— (1997) 'Cohort, year and age effects in Canadian wage data,' Department of Economics, McMaster University
Card, David, and Thomas Lemieux (1996) 'Multiple modes of adjustment: a comparative study of youth in the US and Canada, 1970–1995,' University of Montreal
Dooley, Martin D., and Peter Gottschalk (1984) 'Earnings inequality among males in the United States: trends and the effect of labor force growth,' *Journal of Political Economy* 92, 59–89
Heckman, James J., and Robert Robb (1985) 'Using longitudinal data to estimate age, period and cohort effects in earnings equations,' in *Analyzing Longitudinal Data for Age, Period and Cohort Effects*, ed. S. Feinberg and W. Mason (New York: Academic Press)
Juhn, Chinhui, Kevin Murphy, and Brooks Pierce (1993) 'Wage inequality and the rise in returns to skill,' *Journal of Political Economy* 101, 410–42
MaCurdy, Thomas, and Thomas Mroz (1995) 'Measuring macroeconomic shifts in wages from cohort specifications,' Stanford University
Mincer, Jacob A. (1974) *Schooling, Experience and Earnings* (New York: Columbia University Press)
Morissette, René (1995) 'Has inequality in weekly earnings increased in Canada?' Statistics Canada Analytical Studies Branch Research Paper no. 80
— (1997) 'The declining labour market status of young men,' Business and Labour Market Analysis Division, Statistics Canada
Morissette, René, and C. Berube (1996) 'Longitudinal aspects of earnings inequality in Canada,' Statistics Canada Analytical Studies Branch Research Paper no. 94
Morissette, René, John Myles, and Garnett Picot (1996) 'Earnings polarization in Canada: 1969–1991,' in *Labour Market Polarization and Social Policy Reform*, ed. Keith Banting and Charles Beach (Kingston: School of Policy Studies)
Richardson, David (1997) 'Changes in the distribution of wages in Canada: 1981–1992,' *Canadian Journal of Economics* 30, 662–43
Riddell, W. Craig, and Arthur Sweetman (1997) 'Educational attainment in Canada,' paper presented at the IRPP Conference on Adapting Public Policy to a Labour Market in Transition, Montreal