Catalogue no. 11F0019MIE — No. 235
ISSN: 1205-9153
ISBN: 0-662-38674-4

Research Paper

Analytical Studies Branch Research Paper Series

The Evolution of the Gender Earnings Gap Amongst Canadian University Graduates

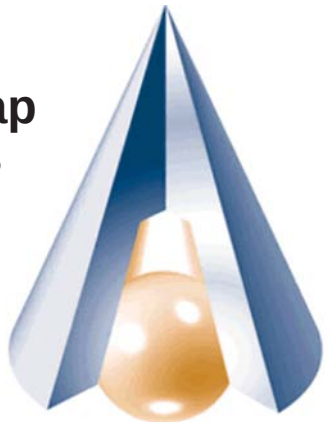

By Ross Finnie and Ted Wannell

Business and Labour Market Analysis Division
24-F floor, R.H. Coats Building, Ottawa, K1A 0T6

Telephone: 1 800 263-1136

This paper represents the views of the authors and does not necessarily reflect the opinions of Statistics Canada.

Statistics Canada Statistique Canada

Canada

The Evolution of the Gender Earnings Gap Amongst Canadian University Graduates

by Ross Finnie and Ted Wannell

11F0019MIE No. 235
ISSN: 1205-9153
ISBN: 0-662-38674-4

Business and Labour Market Analysis Division
24-F, R.H. Coats Building, Ottawa, K1A 0T6
Statistics Canada

How to obtain more information :
National inquiries line: 1 800 263-1136
E-Mail inquiries: infostats@statcan.ca

The paper is available on Internet: (www.statcan.ca)

November 2004

This research was financially supported by the Human Capital and Education Studies Division in the Applied Research Branch of Human Resources Development Canada. A Social Sciences and Humanities Research Council grant awarded to Ross Finnie provided assistance for an earlier phase of the work. Helpful comments were received from Richard Chaykowsi, Doug Giddings, Philip Jennings, Lisa Powell, and an anonymous referee. Excellent research assistance was provided by Marc Frenette.

Published by authority of the Minister responsible for Statistics Canada

Aussi disponible en français

Table of Contents

Abstract

This paper reports the results of an empirical analysis of the gender earnings gap amongst recent Canadian Bachelor's level university graduates. The overall gap, as of two years leaving university, narrowed significantly across successive cohorts of graduates, but widened significantly from two to five years after graduation for all groups. Differences in the exogenous variables "explain" from about 40 percent to essentially the entire gap across the different periods, this portion rising from two to five years out and across cohorts. By the final group, all of the gap is thus "explained" at the two-year point in time, and most of it is explained at the five-year mark, with labour market returns (measured in this manner) largely gender-neutral for the last group of graduates. Hours of work are the single most important influence, while past work experience, job characteristics, family status, province of residence, and language have smaller and more mixed effects.

Keywords: gender earnings gap, post-secondary graduates, school-to-work transition

I. Introduction

Due to a relatively large and still growing literature, we now know a good deal about the gender earnings gap in Canada (as elsewhere) and how it has been shifting over time.[1] Virtually all previous studies have, however, been based on cross-sectional databases and covered workers of all ages (and usually many different types), meaning that while we have a broad understanding of the general structure of the gender earnings gap and how it varies across workers of different ages (and types), we know much less regarding precisely how the gap evolves over the life cycle for given cohorts of workers or how these dynamics have been shifting over time.

The contribution of this paper is to report the findings of an empirical analysis of the gender earnings gap amongst Canadian Bachelor's level university graduates over the first five years following graduation and to compare these dynamics for three separate cohorts of recent graduates. The work is based on three waves of the recently available National Graduates Surveys (NGS) databases, which comprise large, representative samples of individuals who successfully completed their programmes at Canadian universities in 1982, 1986, and 1990, with the data gathered during interviews conducted two and five years after graduation for each group of graduates (1984/87, 1988/91, 1992/95). The longitudinal, cross-cohort structure of these databases—along with their size, representative nature, and the range of variables available—make these Canadian data rather uniquely well-suited to such a focused and dynamic analysis of the gender earnings patterns amongst this important group of workers.

The analysis highlights the significant increases which have occurred in the "starting" earnings levels of female graduates versus declines for men, thus resulting in a significant narrowing of the gender earnings gap across successive cohorts of graduates—and this over a period of just eight years (from 1984 through 1988 to 1992). There have, however, not been commensurate changes in the patterns of earnings growth from two to five years following graduation, so that the gender gap has increased significantly from the second to fifth year following graduation for each class, and as much (or more) for the most recent cohort as the earlier ones. This is an extremely interesting and important finding with particularly significant implications for longer-term gender earnings patterns, especially in light of evidence that the major portion of real lifetime earnings growth occurs during the first few years of young peoples' post-schooling careers. In short, women might not be catching up as much in a fuller life cycle context as immediate post-graduation earnings levels suggest.

The paper first offers some descriptive statistics which track the earnings patterns of male and female graduates in the years following graduation, and then reports the findings of a standard regression-based decomposition analysis which breaks the overall male-female difference in mean earnings into the part due to differences in the mean values of the explanatory variables and the part due to the differences in the associated earnings model coefficient estimates. The explanatory variables considered in the analysis include field of study, post-graduation labour force experience, hours of work, temporary work or being

1. Gunderson (1985, 1989), Gunderson and Riddell (1991) in References provide good overviews of the gender earnings gap in Canada in terms of the relevant estimation issues, the empirical evidence, and the related policy issues, while more recent work includes Baker et al. (1995), Christofides and Swidinsky (1994), Doiron and Riddell (1994), Miller (1987), and Reilly and Wirjanto (1999a, b).

self-employed, marital status and the presence of children, and province of residence and language spoken. We are thus able to assess the contribution of these factors to the structure of the earnings gap at each point in time, its evolution from two to five years following graduation, and its shifts across cohorts.[2]

II. The NGS data and the construction of the working samples

II.1 The National Graduates Surveys

The National Graduates Surveys (and Follow-Up) databases, developed by Statistics Canada in conjunction with Human Resources Development Canada, are well suited to this analysis for a number of reasons. First, the NGS files are representative of the underlying national population of university graduates in the given years and include large numbers of observations, including in the neighbourhood of 10,000 at the Bachelor's level, thus providing abundant samples for the analysis.[3]

Second, the NGS databases have a longitudinal element which derives from the two interviews carried out for each cohort, two and five years after graduation. This allows for a dynamic analysis of the school-to-work transition covering the seminal years of graduates' working careers, with the view precisely situated in terms of the post-graduation period corresponding to the two interview dates.

Third, the availability of data for three separate cohorts of graduates—representing those who successfully completed their university programmes in 1982, 1986, and 1990—permits the comparison of outcomes across successive "generations" over a period typically thought to have been characterised by important changes in labour markets, as well as a time during which women's earnings have generally continued to catch up to men's, while also bringing the record up to the fairly recent past.

Finally, the NGS databases possess various arrays of variables covering graduates' educational experiences, job characteristics, and basic demographic information (as listed below) which are important to any analysis of earnings patterns.

In summary, the three NGS databases provide for a detailed, dynamic analysis of early labour market outcomes amongst Canadian university graduates in the critical early years following graduation from the early 1980s into the mid-1990s. The NGS data are, furthermore, interesting and unique not only in a Canadian context, but are perhaps unequalled in the world in terms of offering large, consistent,

2. Finnie and Wannell (1999) analyses various other early career outcomes by gender.

3. The NGS databases are based on a stratified sampling scheme, with stratifications by degree level, discipline and province. All results reported here are based on the appropriate weights. The sample framework is established through the use of institutions' administration files on graduates, with those records also providing educational information, such as programme and discipline of study. Response rates were generally around 80 percent for each of the first interviews, and in the neighbourhood of 90 percent of these respondents were successfully contacted a second time. Checks of certain outcomes for those who were successfully interviewed just once versus those who completed both interviews suggest that any sample bias appears to be quite small (Finnie (2000)).

representative, partly-longitudinal surveys of post-secondary graduates covering various elements of the school-to-work transition over the last decade and a half.[4]

II.2 Selection of the working samples

The analysis excludes graduates who had accumulated five or more years of full-time work experience by the time of graduation or were 35 years of age or older upon completing their studies, thus focusing on "fresh" graduates who had followed more-or-less conventional career profiles with respect to school and work. These deletions exclude, in particular, women returning to school after having spent time raising children, as well as both men and women who undertook major re-tooling of their human capital in terms of education. While such individuals certainly comprise interesting groups of graduates, their study is best left to a separate project.

Second, graduates who obtained an additional degree by one of the two interview dates were deleted from the analysis at that point. This was done, first of all, because many such graduates no longer belonged to the original education group (*e.g.*, a Bachelor's graduate might have become a Master's graduate and perhaps changed disciplines) and had, in any event, been mixing school and work in a way likely to affect the labour market outcomes upon which this analysis is focused. A second, related reason for this selection rule is that including "additional degree graduates" would have thrown off the precise post-graduation time frame corresponding to the two interview dates (*i.e.*, two and five years after graduation) which holds for the non-continuing group. Finally, it is impossible to identify the specific field of study (an important variable) in which any new degree was obtained as of the 1984 survey for the 1982 graduates.

Third, and in a similar vein, part-time workers who cited school as the reason for their only partial involvement in the labour market were excluded on the grounds that such individuals were—by definition—still principally students and had not yet entered the school-to-work transition in earnest. Other part-time workers were, on the other hand, generally included in the analysis, thus lending it a broad sample base.

Fourth, the small number (less than one-half of one percent in each year) of non-regular workers (*i.e.*, not paid or self-employed) were also eliminated, since employment status, earnings levels, and other job outcomes of these family employees, volunteers, and other such workers differ from those of others (as verified empirically). A similarly small number of workers deemed to have unreasonably low earnings (full-time workers with an annual rate of earnings of less than $5,000), were also dropped.

4. Dolton, O'Neill and Sweetman (1996) is closest to this study conceptually and methodologically in that they analyze the evolution of the earnings gender gap amongst university graduates in the UK using samples of men and women who completed their studies in 1960, 1970, and 1980, interviewed in 1967, 1977, and 1986. They also employ a decomposition approach similar to the one adopted here (although they explicitly decompose the *change* in the gender wage differential across periods rather than decomposing its structure in each year as we do, thus allowing for direct comparisons from two to five years out as well as from one cohort to another), but also carry out a detailed breakdown of the "residual" component which represents the portion of the gap which is not accounted for by either differences in the coefficient estimates or the levels of the explanatory variables associated with the regressors included in the models.

Finally, observations were deleted if the required information was missing for any of the variables used in the analysis, resulting in a few additional deletions.

III. Raw earnings patterns

Table 1 reports the earnings levels of graduates in 1995 constant dollars. These are presented in a number of ways: the mean earnings of all graduates who were working as of each of the interview dates (and otherwise met the relevant sample inclusion criteria); the mean earnings of full-time workers only; median earnings for the same groups (all workers, full-time only); and mean earnings by quintile for the two sets of workers. We focus, however, on the first set of figures—the mean earnings of all workers taken together, as these provide the best context for the decomposition analysis presented below. The key elements of the patterns discussed here are, however, generally similar to what is found in the medians, the more detailed breakdowns, and the full-time groups.

Male graduates' mean earnings were, not surprisingly, higher than female graduates' in every survey year. More interesting, however, are the magnitudes of those gaps and their movements over time. The cross-cohort trends clearly favour female graduates: male graduates' earnings generally declined over time, with the mean earnings of the third group of graduates laying 6.3 percent below those of the first cohort as of the first interview, and 8.1 percent lower as of the second interview (these occurring at roughly comparable points in the economic cycle), whereas the mean earnings of the female graduates of the third cohort were 6.7 and 2.6 percent *above* those of the first cohort as of the first and second interviews respectively. The gender earnings gap was, therefore, narrower for each subsequent cohort relative to the preceding one as of each of the interview dates.

At the same time, the NGS data allow us to see rather precisely how graduates' earnings rose over the early years in the labour market and—of particular relevance to this paper—how male graduates' growth rates outstripped those of females in every case: 31 versus 21 percent for the first cohort, 22 versus 18 percent in the second, and 28 as opposed to 16 percent for the final group of graduates (see the relevant columns in Table 1). The gender gap thus generally widened from the first interview to the second for each set of graduates.

More precisely, female mean earnings were 82 and 76 percent as high as male's earnings as of the first and second interviews, respectively, for the first cohort; 84 and 81 as high for the second group; and 93 and 84 percent as high for the final set of graduates: that is, progressions of 82, 84, and 93 (first interview), and of 76, 81, and 84 (second interview). The gender earnings gap as of two years following graduation was, therefore, reduced by 61 percent over the relevant eight year period, but by a considerably smaller 33 percent as of five years following graduation.[5]

5. These differences compare to earnings gaps ranging from .10 to .40 found for other groups of workers, with the gap generally being greater for broader groups of workers, smaller for more specific groups, and narrowing over time (Baker et al. (1995), Gunderson (1985)).

The gender earnings gap thus narrowed steadily across cohorts, but the *growth* in the gap from two to five years following graduation was actually greatest of all for the most recent group of graduates. In other words, although the gender earnings gap started off smaller for each successive cohort, it continued to widen in the post-graduation years as much for the most recent group as—or even more so—for the earlier ones.

So while female graduates' earnings profiles appear to have been shifting up towards males' with each succeeding cohort in terms of starting *levels*, the relative *slopes* of those profiles might not have changed commensurately. These findings have particularly interesting implications for the longer-term earnings profiles of graduates, as they suggest that longer-run ("permanent") reductions in the earnings gap amongst recent cohorts of graduates might not be nearly as great as the immediate post-graduation record suggests. The decomposition analysis will now identify the factors underlying these patterns.

IV. The decomposition analysis

IV.1 The decomposition approach

The decomposition analysis reported here essentially follows standard conventions as adapted to the specific groups of workers and data at hand.[6] First, separate earnings models were estimated for male and female graduates for each year of data available: 1984/87 for the 1982 graduates, 1988/91 for the 1986 graduates, and 1992/95 for the 1990 graduates. The gender differences in mean earnings were then decomposed into the part due to differences in the mean values of the explanatory variables (field of study, job experience, province of residence, individual characteristics, etc.), and the part due to the differences in the associated coefficient estimates.

The effects of the differences in the explanatory variables (the first component of the decomposition) were evaluated using the male coefficient estimates, while the effects of the different coefficient estimates were evaluated using the female mean levels. As in any such exercise, the decomposition algebra could have been reversed, but this approach allows us to view the earnings gap in terms of how much less women earn due to the "discriminatory" rates of return applied to their given sets of characteristics and how much less they earn due to their different characteristics evaluated at "fair" market rates of return.

IV.2 The underlying earnings models

The regression models underlying the decomposition analysis represent conventional human capital earnings equations as adapted to the post-graduation period covered and information available in the NGS data, with earnings taken to be a function of various sets of variables representing individuals' human capital and other factors that affect earnings (see below).

6. The method is generally associated with Oaxaca (1973). See Cain (1986) for a general exposition of the decomposition method employed here, Gunderson (1989) for another general review in the context of the gender earnings gap *per se*, and Gunderson (1985) and Gunderson and Riddell (1991) for a focus on Canadian studies of the gender earnings gap.

The earnings variable available on the NGS databases is based on the question: "Working your usual number of hours, approximately what would be your annual earnings before taxes and deductions at that job?" The variable thus represents what the person would earn on an annual basis were the job to last the full year, regardless of the actual job status—effectively a normalised annual rate of pay. This is a somewhat unconventional measure, but is well-defined, analytically interesting, and presumably well reported (being a figure individuals should either know or be able to calculate rather easily). All earnings values have been converted to constant 1995 dollars, rounded to the nearest thousand, and capped at the $99,000 upper limit which characterizes the 1984 data (the lowest bound in the six databases), or $143,035 in constant 1995 dollars.

The models are specified in log-linear form, meaning that the coefficient estimates can be interpreted as the proportional effect on earnings of a one unit change in the given independent variable (for small changes), with the decompositions framed in a similar context. Furthermore, since earnings levels are generally fairly similar across cohorts (as seen above), changes in the estimated proportional effects from one group of graduates to another reflect approximately similar changes in the absolute (real) dollar effects. Such a straightforward relationship between the proportional and actual dollar hold a little less closely, however, for the first versus second interview results, since mean earnings rise substantially over this interval.

The regressors included in the earnings models are generally restricted to factors which could more easily be presumed to be exogenous in order to obviate the need to deal with issues relating to the any potential endogeneity. In particular, industry and occupation are omitted (although models with these included have been estimated and are available from the authors upon request).

Hours of work are, on the other hand, included (either a part-time indicator or a series of dummy variables indicating various ranges of hours worked, the latter fitting the data better than a continuous hours measure). While there could potentially be an issue of endogeneity here, it should be much less of a concern in a context where the dependent variable is annual earnings rather than the hourly wage.[7] Furthermore, hours have a significant direct influence on annual earnings and—it turns out—play an important role in the gender earnings gap, meaning that it is important to include these measures in the analysis. Finally, tests indicated that hours were in fact not endogenous and could thus be included in a straightforward manner in the models (in the form specified).

Apart from this, the models include regressors representing field of study, already holding a higher degree, post-graduation work experience, temporary work status or self-employment, marital status and the presence of children, province of residence, and language, all of which are defined in a consistent fashion across cohorts. Detailed definitions of the variables included in the models are provided in the appendix.

7. That is, while the number of hours usually worked in a week are generally likely to be a function of the hourly wage (standard labour supply theory), it is much less obvious as to why hours worked would be a function of annual earnings.

V. The decomposition results

V.1 Models including the part-time work dummy (all cohorts)

Table 2 reports the first set of decomposition results, where hours of work are measured with a simple part-time indicator (as opposed to the reference full-time group). The effects associated with each of the variables included in the regressions are grouped as follows: the intercept; educational characteristics (predominantly field of study but also including whether the person already held a degree beyond the Bachelor's); work experience since graduation (dummy variables to capture the differences in earnings between those in full-time versus part-time work or not in the labour force as of selected dates following graduation); part-time employment status in the current job; other current job characteristics (self-employed, temporary position); marital status and the presence of children; province and language.

The overall beta and explanatory effects

Following convention, results are presented in terms of the effects of the differences in the coefficient estimates ("Beta") and the effects of the explanatory variables ("X"). These are given in 'Total' (the last row of Table 2), as well as for each group of variables. These sum (approximately) to the overall gender earnings gaps, expressed in the percentage terms which conform to the log earnings model specification used in the regressions.[8]

Looking first at the overall effects, for the first cohort, male-female differences in the explanatory variables accounted for 41 percent of the overall mean earnings gap of 20 percent as of the first interview (two years following graduation), then rose to a moderately higher 52 percent share of the larger 27 percent gap which held as of the second interview (five years out). The effects fell in absolute terms but the relative shares of the explanatory variables rose moderately for the smaller overall gaps (16 and 23 percent) which held for the second cohort: 50 percent and 55 percent as of the two interview dates. Perhaps most interestingly, the Beta effects were effectively zero for the 1990 graduates as of the first interview, suggesting a fully "nondiscriminatory" earnings-generating process, with practically the full share (93 percent) of the much smaller gap (8 percent) "explained" by difference in field of study, the rate of part-time work, and the other job characteristics captured in the models. The explanatory share could not but drop for the larger second interview gap (17 percent), but remained high (74 percent), and increased by slightly more than the coefficient effects in absolute terms.

There were, therefore, two principal time trends. First, the increases in the gender earnings gap from two to five years following graduation were driven more by increased differences in observable characteristics than by differences in the returns to those factors, the explanatory effects accounting for a little more than half to almost three-quarters of the overall gap at the later point for each cohort. Second,

8. The overall gaps are not exactly the same as those implied by the mean earnings tabulations shown above due to the slightly different samples used in the regressions (observations with missing values of the explanatory variables were deleted) and due to the approximation which is implicit in the log-linear regression specification (the indicated "percentage effects" hold precisely only for small changes in the independent variables in the neighbourhood of the sample means).

differences in observable characteristics were associated with slightly smaller gender differences in earnings in each later cohort in absolute terms, whereas the differences due to male-female differences in the returns to those factors declined in a steady and fairly dramatic fashion. In short, there was an increase in the relative and absolute importance of the explanatory variables over time for a given cohort, and a pronounced decrease in the unexplained differences in male-female earnings in each later cohort so that for the latest cohort the earnings-determination process appeared to be wholly—or nearly wholly—gender neutral.

The variable-specific effects

Taken at face value, the Beta effects associated with the intercept terms represent the differences in the mean earnings of men and women which are unrelated to any of the explanatory variables included in the models—"general" differences which apply to all graduates. More strictly speaking, however, the intercept captures the earnings gap effects with all the explanatory (dummy) variables set to zero. (In our case, this represents a social science graduate with no previously held higher degree, who worked full-time at each of the dates following graduation at which labour force activity was ascertained, and who currently held a full-time permanent job, who was single and childless, and who first spoke English and was currently living in Ontario.) The Beta effects of the other groups of variables then represent the effects of the male-female differences in the associated coefficient estimates relative to the omitted group (e.g., female graduates doing relatively worse than males in fields other than the omitted social sciences category).

Choosing a different set of omitted variables for the categorical variables would, therefore, generally result in different "intercept" or "general" effects as well as different influences for each of the particular groups of variables.[9] This is not, however, the case for the explanatory portion of the gap, where such variable-by-variable contributions are completely unambiguous (note that the intercept has zero effect in this regard). In the analysis which follows, we therefore offer comments about the overall intercept effects and the contribution of each group of variables to the overall gender gap in the appropriate cautionary manner, with the comparisons benefiting from the consistent structure of the models—and decompositions—over time.[10]

With those caveats in mind, the general effects associated with the intercept and the omitted categories of the categorical variables are largest for the first cohort and then decline over time, being associated with gender earnings differences of 18.7 and 22.8 percent as of two and five years following graduation for the 1982 cohort, differences of 12.9 and 11.0 as of the same points in time percent for the class of 1986, but much lower differences of just 0.3 and 4.0 percent for the last set of graduates. Thus, as a

9. For example, if women did better in the omitted educational category (the social sciences in our case) relative to other fields than was the case for men, that would tend to drive the contribution of the intercept effect down and the education effects up relative to what would hold were one of the relatively lower-earnings fields of study used as the omitted education variable in the models.

10. See Jones (1983) Oaxaca and Ransom (1999) on this point. The authors are grateful for the comments of an anonymous referee in this regard.

large proportion of the relatively large gender earnings gap which existed for the earlier groups of graduates appears to have been quite generalised, so was much of the catching up to male graduates which female graduates accomplished across cohorts.

The set of educational characteristic variables, predominantly the field of study indicators, have mixed effects on the overall earnings gap. The associated X effects indicate that women have—perhaps as anticipated—tended to be over-represented in disciplines which have generally low earnings, with the impact of these differences growing from two to five years following graduation, especially for the earlier cohorts. On the other hand, the negative Beta effects for the first two cohorts—representing influences which tended to diminish the earnings gap—indicate that the earnings patterns between the baseline social science group and graduates in other disciplines generally favoured women; that is, earnings have, on average, been relatively higher for women in the other disciplines than was the case for men. This was, however, no longer the case for the 1990 cohort, where the educational Beta effects are positive.

Taking the intercept and field effects together, in every period at least one-half of the gender earnings gap was related to the intercept term, the other omitted variable categories, and the field of study effects. In other words, a large portion of the gender earnings gap amongst recent graduates has been associated with a general tendency for female graduates of a given field of study to have lower earnings than males regardless of the specific nature of their current job characteristics, post-graduation work experience, or personal attributes.

Furthermore, this "non-specific component" grew significantly from two to five years following graduation, meaning that much of the increases in the overall gender earnings gap over this interval have been widespread, independent of specific post-graduation career profiles or personal characteristics. On the other hand, the narrowing of the gender earnings gap across cohorts has been equally widespread—seen in the declines of the intercept and educational characteristic effects from the earlier to later groups of graduates.[11]

One consistently important specific factor, however, has been part-time work, with women's typically shorter work week hours driving their overall earnings levels from 2.8 to 7.0 percent lower (depending on the particular period) than those of men. Furthermore, the part-time effects are uniformly larger for the second interview decompositions than the earlier ones, meaning that gender divergences in part-time work patterns have been a significant factor in the widening of the gender earnings gap from two to five years following graduation. The part-time Beta effects are, on the other hand, more mixed and generally smaller (and sometimes negative), indicating that the penalties for working part-time in terms of reduced annual earnings have sometimes been greater for men, but in other cases greater for women.

11. Regarding the dependency of the Beta effects on the precise specification of the model (discussed above), the small Beta effects associated with the other groups of variables imply that an alternative formulation would not likely change this result.

The effects of the other variables included in the models are generally smaller and less consistent across the different periods.[12] The job characteristic effects are, for example, quite mixed as of the first interview, and although somewhat more consistent, still fairly small contributors to the gender earnings gap in the later period in each case. The results reflect both a greater incidence of permanent job holding and self-employment amongst men and greater returns to self-employment for males than females (specific effects not shown).

The influences of marital status and the presence of children—after controlling for past work experience, current part-time work status, and the other influences captured by the variables included in the models—account for moderate portions of the earnings gap for the first cohort (both points in time), and for the later year for the second cohort, but have effectively no (direct) effect for the third cohort. The work experience variables—which capture the effects of having had periods of working only part-time work or of being completely out of the labour force between graduation and the current interview—also show varying effects across periods, with the negative Beta effects in 1987 and 1991 suggesting that female graduates' earnings have not suffered as much as men's from past breaks from full-time work.[13] The province and language effects are mostly quite small, indicating that the distribution of graduates across the provinces and the related earnings effects have been quite similar for male and female graduates.

V.2 Models including usual hours worked (1986 and 1990 cohorts only)

We now turn to the models which replace the single part- versus full-time indicator used above with a series of dummy variables representing hours of work. These results are shown in Table 3. The overall earnings gap changes at each point in time with the addition of the hours dummies, but this is only the result of dropping observations where the usual hours of work information was missing and the differences are not great.

The beta and explanatory effects

The proportion of the total gender earnings gap accounted for by the "X" effects is now greater but greater only where the part- versus full-time differentiation was permitted, in every case these effects growing from the first interview to the second, as average hours by gender diverged significantly over the early years in the labour market.

The 5.3 percent effect of the male-female difference in hours worked thus accounts for 30 percent of the total gender earnings gap amongst 1986 graduates two years after graduation, while the 8.6 percent effect accounts for 38 percent of the considerably larger gap which held five years after graduation.

12. This also means that the "general effects" interpretation of the intercept term—or the intercept in combination with the more important explanatory variables (education in particular)—are less likely to be affected by changing the omitted categories used in the underlying earnings models.
13. The relatively small effects of the experience variables may be at least partly due to the proxy aspect of the measures employed, necessitated by the fact that the NGS databases lack direct measures of work experience following graduation.

These effects were even greater for the class of 1990: the two year effect of 5.9 percent represents 61 percent of the total earnings gap in 1992, and the five year effect of 9.3 percent represents 53 percent of the 1995 gap. Male-female differences in hours were, therefore, central factors in the overall gender earnings gap and its increase over the early years in the labour market for the 1986 and 1990 graduates, for which this information is available—an important finding.
The result is also interesting for showing how a significantly greater portion of the gender earnings gap can be explained when a better measure is available. The implication is that some portion of the remaining "unexplained" portion of the gap could be similarly reduced with improvements in other measures or the addition of other variables to the analysis.

Adding the hours dummies affects the estimated effects of some of the other variables, but the differences are generally not very large and the earlier discussions generally hold. One might have thought that the family status variables, in particular, would have changed significantly as the better measures of hours were included in the models (on the assumption that the hours effects were at least partially captured by those other variables in the absence of the more precise measure), but the changes are relatively small—largely because hours of worked already diverged quite significantly for unattached men and women (see the following).

Male-female differences in hours

Pursuing the importance of hours more deeply, at all time points, Table 4a shows that full-time men averaged at least two hours more work per week than female graduates, with the gap growing over time within each cohort. Amongst 1986 graduates, average hours for men increased from 44.3 in 1988 to 45.1 in 1991, while female hours increased from 41.8 to 42.4, with the hours gap thus widening from 2.5 to 2.7. For the 1990 graduates average male hours went from 44.7 in 1992 to 45.3 in 1995, the average female work week remained steady at 41.8 hours in both years, and the hours gap grew commensurately from 2.9 to 3.5. Differences in hours worked were, therefore, greater and increased more over time for the later cohort—thus running counter to the convergence in earning levels themselves across subsequent cohorts of graduates.

The major reason for the hours gap, in the accounting sense, is that many more male than female graduates worked very long hours (i.e., more than 50 hours per week), with more than one-quarter of full-time employed men working greater than 50 hours per week in each cohort, as compared to just 17.1 percent of the female graduates in 1988 and 19.3 percent in 1992 (Table 4b). Furthermore, this gap grew in the subsequent three years amongst both sets of graduates: for men, the incidence of long hours increased to over 30 percent in both cohorts, while for women the rate grew from 17 to 20 percent for the earlier class and inched up just half a point for the later group. Thus, within 5 years of graduation, one and a half times as many men were working very long hours in comparison to women.

The hours gap is, not surprisingly, greatest amongst married graduates with children. Married mothers in full-time work averaged at least four hours less work per week than their male counterparts as of two years following graduation, with the difference increasing to 5.1 hours by the second interview for the 1986 cohort and to 6.4 hours for the 1990 group. There is, however, also a sizeable gap amongst

single, childless graduates. For the class of 1986, unattached males averaged 1.8 hours per week more than their female counterparts in both 1988 and 1991, and the gap was even higher in the class of 1990: 2.8 hours per week in 1992 and 2.4 hours per week in 1995.

VI. Conclusion

This paper has reported the results of an empirical study which exploits the rather unique strengths of the National Graduates Survey databases to analyse the evolution of the gender earnings gap amongst recent Canadian Bachelor's level university graduates over the last decade and a half. The major findings may be sumarized as follows:

- There was a substantial narrowing of the overall gender earnings gap across cohorts—the result of increases in female graduates' earnings and decreases in males' earnings—but the narrowing was much greater two years following graduation than five years out, as men's earnings grew considerably more strongly than women's over this interval, even for the later groups of graduates. Thus, while female graduates started out on much more equal footing in the 1990s, their earnings quickly began to trail behind those of their male classmates as much as in earlier years.
- A large part of the gender earnings gap at each interview date, and much of the increase in the overall earnings gap from two to five years following graduation, appears to have been of a generalised nature, unrelated to specific job characteristics, experience profiles, or individual attributes. At the same time, much of the narrowing of the gap across cohorts has been of a similarly general nature.
- There was an increase in the relative and absolute importance of the explained portion of the gender gap over time for a given cohort, and a pronounced decrease in the unexplained differences (Beta effects) in each later cohort, being zero or near-zero for the latest group of graduates, thus indicating gender neutrality (or near to it) in the returns to various factors in the labour market.
- As for more specific influences, hours of work was an important determinant of the earnings gap at each point in time, as well as of its increase in the years following graduation.
- Other factors, such as past work experience, specific job characteristics, family status, and province of residence and language spoken, have played only smaller and generally more mixed roles in the gender gap amongst these Bachelor's graduates.

Probably the most interesting and important result is, then, the extent to which the gender earnings gap has increased so significantly from two to five years following graduation for each set of graduates, with this dynamic continuing even as the initial post-graduation gap has narrowed significantly from one cohort to the next. The obvious question is "Why?" In particular, if certain forces – be they on the demand side or the supply side – have been causing such a significant shift in starting earnings levels, why have these same forces not affected the male-female patterns of earnings growth as well?

One clue to these dynamics is perhaps found in the relatively important role that hours of work appears to have played in the initial earnings gap and its widening in the subsequent post-graduation years. One straightforward interpretation of these results is that men and women have been making—and, it would

appear, continue to make—different labour supply decisions which have had important direct effects on their earnings levels. Put most simply, women have been working fewer hours than men and have been receiving commensurately lower earnings levels.
At the same time, the hours effects explain only a portion of the observed earnings patterns, and thus we need to continue to search for the other underlying causes of these dynamics. Furthermore, the hours patterns, as well as any other "explanatory" factors, might themselves be at least partly determined by a larger set of processes which could include various types of discrimination—directly in the labour market itself, in other institutions (including schools and universities), or on a broader social level. Identifying—and quantifying—such influences is one of the great challenges in modern social science research but represents a task which lays beyond the scope of this paper.

On the other hand, if female working hours are consistently below male hours throughout their careers, then the job experience of women relative to men is likely to be mis-measured—biased upwards—with implications for other analyses. For example, women could appear to receive a lower rate of return to experience when this only reflects the different "quality" of experience accumulated over time.

Regarding directions for further work, it would be interesting to see if the dynamic view of the evolution of the gender earnings gap amongst these particular groups of workers provided here holds up for other workers—in Canada or elsewhere—particularly as these results point to a different evolution of future gender earnings patterns (over the life cycle) than the more common static views might suggest. In short, with female graduates' earnings profiles shifting upwards and closer to men's immediately after entering the labour market but remaining much flatter in the following years, the gender earnings gap is likely to narrow much less in the longer run than any early record might suggest. Any further understanding of the structure of this dynamic, such as identifying the significant role played by hours of work identified here, would clearly help our understanding of human capital investments, labour supply decisions, and the operation of labour markets in a gender context.

APPENDIX: DEFINITIONS OF THE VARIABLES INCLUDED IN THE MODELS

Educational characteristics:

- ***Field of study***: 15 discipline groups, reflecting fields of roughly similar natures with respect to subject material, approach, etc., and comparable mean earnings patterns as determined through a preliminary analysis focused on this issue. (*"Other" (general) social sciences* is the omitted group in the regressions.)
- ***A previously obtained higher degree***: a dummy variable indicating that the individual held a higher degree (Master's or Ph.D.) before completing the Bachelor's degree which is the basis for inclusion in these samples.

Post-graduation work experience:

- Proxied with the part-time and full-time employment status at various points in time since graduation—two dates between graduation and the first interview in the case of the first interview regressions (October in the year following graduation (1983/87/91) and June of the year after that (1984/88/92)) and the addition of the first interview employment status to the second interview regressions. (*Working full-time* is the omitted category in each case.)

Part-time employment status and hours of work:

- A dummy variable indicating that the individual normally worked less than thirty hours per week (standard definition).
- For 1986 and 1990 graduates, a separate set of models usually includes a series of dummy variables representing the usual hours of work. (This information was not collected for the 1982 graduates.) Hours of work dummies rather than the number of hours are used for two reasons: such a specification allows for non-linear and/or threshold effects, while the inclusion of a linear hours variable tends to introduce different scaling effects (essentially shifting the intercept) which makes comparisons across years somewhat less straightforward. The hours of work dummies used are: <20 hours, 20-29 hours, 30-34 hours, 35-39 hours, 40 hours (the omitted group), 41-44 hours, 45-49 hours, and 50 hours or more.

Job characteristics:

- ***Temporary job status***: A dummy variable indicating that the individual held a temporary (as opposed to permanent) job.
- ***Self-employed***: A dummy variable indicating that the individual was self-employed (as opposed to being a paid worker).

Marital status/presence of children:

The specific variables included allow for the effects of children to vary by marital status:

- ***Single (never married), with/without children***.
- ***Married, with/without children***.
- ***Widowed/Separated/Divorced, with/without children***.

(*Single, no children* is the omitted category.)

Province/language:

The choice of variables allows for "majority" and "minority" language effects along with provincial/regional differences, with the language variable representing the one first spoken by the individual (the most clearly defined and consistent language measure in the NGS data):

- ***Atlantic Canada, Quebec, Ontario, Manitoba/Saskatchewan, Alberta, British Columbia and the Territories***: a series of indicator variables reflecting the graduate's current province of residence (*Ontario* is the omitted category.)
- ***Quebec-English***: captures the effect of being an anglophone in Quebec, leaving the Quebec variable alone (above) to represent (primarily) francophone Québecois.
- ***Minority French***: captures the effect of being a francophone outside of Quebec, leaving the provincial/regional indicators noted above to represent anglophones in those areas (assuming a common effect across provinces).
- ***Other language***: captures non-French/English speakers (again assuming a common effect across provinces—including Quebec).

Table 1: Mean earnings (1995 Constant Dollars)

	1982 Cohort			1986 Cohort			1990 Cohort		
	1984 $	1987 $	Change %	1988 $	1991 $	Change %	1992 $	1995 $	Change %
Males									
Mean earnings									
All workers	35,000	45,800	31	35,200	43,100	22	32,800	42,100	28
Full-time workers	35,800	47,100	32	35,700	44,800	25	33,700	43,500	29
Median earnings									
All workers	33,200	40,900	23	32,000	39,100	22	31,300	39,000	25
Full-time workers	33,200	40,900	23	32,000	40,200	26	31,300	40,000	28
Mean earnings by quintile									
All workers									
Top	55,400	76,600	38	55,400	67,700	22	50,100	66,800	33
4th	39,200	48,100	23	37,200	45,000	21	36,500	46,000	26
3rd	32,600	40,500	24	31,700	39,000	23	30,800	38,500	25
2nd	25,400	33,600	32	25,800	32,200	25	22,300	30,900	39
Bottom	15,600	21,300	37	15,800	20,400	29	10,200	18,200	78
Full-time workers									
Top	55,800	78,500	41	55,600	69,300	25	50,700	67,800	34
4th	39,500	49,100	24	37,600	45,800	22	37,000	46,600	26
3rd	33,300	41,500	25	32,200	39,900	24	31,400	39,700	26
2nd	26,400	35,000	33	26,700	33,700	26	24,100	32,500	35
Bottom	17,700	24,000	36	17,400	23,000	32	12,900	20,100	56
Females									
Mean earnings									
All workers	28,600	34,600	21	29,700	35,100	18	30,500	35,500	16
Full-time workers	30,000	36,500	22	30,700	36,900	20	31,600	36,900	17
Median earnings									
All workers	28,900	32,000	11	29,500	33,800	15	28,100	33,500	19
Full-time workers	28,900	34,500	19	29,500	34,900	18	29,200	35,000	20
Mean earnings by quintile									
All workers									
Top	43,200	53,000	23	46,200	54,500	18	47,200	53,800	14
4th	34,100	38,400	13	33,100	38,900	18	33,700	39,400	17
3rd	28,400	33,000	16	28,700	33,400	16	28,400	33,500	18
2nd	21,100	26,100	24	22,200	25,800	16	19,100	26,000	36
Bottom	12,100	14,500	20	12,800	11,800	-8	7,600	13,800	82
Full-time workers									
Top	43,600	55,100	26	46,300	58,000	25	47,500	55,300	16
4th	34,800	40,000	15	33,900	40,200	19	34,100	40,400	18
3rd	29,600	34,400	16	29,600	34,400	16	29,300	35,100	20
2nd	23,200	28,700	24	23,900	28,000	17	22,000	28,300	29
Bottom	15,400	18,000	17	15,700	15,100	-4	11,600	16,600	43

Source: National Graduates Survey databases.

Table 2: Decomposition results

	1st Interview			2nd Interview		
	Beta %	X %	Total %	Beta %	X %	Total %
1982 Cohort						
Intercept	18.7	0.0	18.7	22.8	0.0	22.8
Educational characteristics	-6.3	3.8	-2.5	-8.0	6.0	-2.0
Work experience	0.1	-0.2	-0.1	-5.7	2.0	-3.7
Part-time work	0.5	2.8	3.4	1.7	3.6	5.3
Job characteristics	-2.3	1.3	-1.0	0.5	2.0	2.5
Marital status/children	1.9	0.3	2.2	1.9	0.7	2.6
Province/language	-0.8	0.3	-0.5	-0.1	-0.1	-0.2
Total	11.7	8.3	20.1	13.1	14.2	27.3
1986 Cohort						
Intercept	12.9	0.0	12.9	11.0	0.0	11.0
Educational characteristics	-3.3	2.4	-0.9	-0.3	3.2	2.9
Work experience	1.3	0.2	1.5	-4.0	1.8	-2.2
Part-time work	-2.0	4.2	2.2	1.1	7.0	8.1
Job characteristics	-1.4	1.5	0.1	1.0	1.3	2.3
Marital status/children	0.7	-0.2	0.5	4.3	-0.5	3.8
Province/language	1.2	0.1	1.3	-2.9	-0.3	-3.2
Total	8.1	8.0	16.1	10.1	12.5	22.6
1990 Cohort						
Intercept	0.3	0.0	0.3	4.0	0.0	4.0
Educational characteristics	2.6	3.3	5.9	2.7	3.7	6.5
Work experience	1.0	-0.4	0.7	-0.7	0.9	0.2
Part-time work	-1.0	3.7	2.7	-1.9	6.7	4.8
Job characteristics	-1.9	1.3	-0.6	0.4	1.6	2.0
Marital status/children	-0.5	-0.3	-0.8	0.0	-0.5	-0.5
Province/language	1.2	0.1	1.4	-0.1	0.6	0.5
Total	0.6	7.6	8.2	4.5	12.5	17.0

Note: The Beta effects associated with each group of variables depend on the precise specification of the model and, therefore, are not unique. See the text for further discussion.

Source: Decomposition analysis based upon the National Graduates Survey data.

Table 3: Decomposition results with hours of work included

	1st Interview			2nd Interview		
	Beta %	X %	Total %	Beta %	X %	Total %
1986 Cohort						
Intercept	10.9	0.0	10.9	11.0	0.0	11.0
Educational characteristics	-2.8	2.3	-0.5	-0.2	2.8	2.6
Work experience	1.5	0.2	1.6	-3.4	1.5	-1.9
Hours of work	-1.1	5.3	4.2	1.0	8.6	9.6
Job characteristics	-1.2	1.4	0.2	0.9	1.1	2.0
Marital status/children	0.3	-0.2	0.1	3.1	-0.5	2.6
Province/language	0.7	0.1	0.9	-2.8	-0.3	-3.1
Total	8.3	9.1	17.4	9.6	13.3	22.9
1990 Cohort						
Intercept	-3.9	0.0	-3.9	5.0	0.0	5.0
Educational characteristics	3.0	3.1	6.1	1.6	4.0	5.6
Work experience	0.7	-0.4	0.3	-0.9	0.7	-0.2
Hours of work	0.9	5.9	6.8	-1.9	9.3	7.4
Job characteristics	-1.7	1.3	-0.5	-0.4	1.4	1.0
Marital status/children	-0.5	-0.3	-0.8	-1.0	-0.4	-1.4
Province/language	1.7	0.1	1.8	-0.2	0.5	0.3
Total	0.0	9.7	9.7	2.3	15.4	17.7

Source: Decomposition analysis based upon the National Graduates Survey data.

Table 4a: Average hours of work by family status, full-time only

		Two Years After		Five Years After	
		1988	1992	1991	1995
Men	Single no children	43.9	44.4	44.9	45.0
	Married no children	44.9	45.6	44.9	45.2
	Married with children	46.3	45.5	45.7	45.9
	All men	44.3	44.7	45.1	45.3
Women	Single no children	42.1	41.6	43.1	42.6
	Married no children	41.3	42.3	42.4	42.4
	Married with children	41.8	41.5	40.6	39.5
	All women	41.8	41.8	42.4	41.8
Difference	Single no children	1.8	2.8	1.8	2.4
	Married no children	3.6	3.3	2.5	2.8
	Married with children	4.5	4.0	5.1	6.4
	All	2.5	2.9	2.7	3.5

Source: National Graduates Survey.

Table 4b: Long work weeks (> 50 Hours), Full-time only

% of full-time working more than 50 hours per week				
	Two years after		Five years after	
	1988	1992	1991	1995
	%	%	%	%
Men	26	31	28	33
Women	17	20	19	20

Source: National Graduates Survey.

References

Baker, Michael, Dwayne Benjamin, Andrée Desaulniers and Mary Grant. 1995. "The Distribution of the Male/Female Earnings Differential, 1970-1990." *Canadian Journal of Economics*, 28(3): 479-501.

Cain, Glen. 1986. "The Economic Analysis of Labor Market Discrimination: A Survey." Chapter 13 (pp. 693-785) in Orley Ashenfelter and Richard Layard (eds.), *Handbook of Labor Economics*, Vol. 1. Oxford: Elsevier/North-Holland.

Chrisotfides, Louis N. and Robert Swidinsky. 1994. "Wage Determination by Gender and Visible Minority Status: Evidence from the LMAS." *Canadian Public Policy*, 20(1): 34-51.

Doiron, Denise and Craig Riddell. 1994. "The Impact of Unionization on Male-Female Earnings Differences in Canada." *The Journal of Human Resources*, 29(2): 504-534.

Dolton, Peter, Donal O'Neill and Olive Sweetman. 1996. "Gender Differences in the Changing Labor Market: The Role of Legislation and Inequality in Changing the Wage Gap for Qualified Workers in Britain." *The Journal of Human Resources*, 31(3): 549-565.

Finnie, Ross. 2000. "From School to Work: The Evolution of Early Labour Market Outcomes of Canadian Post-Secondary Graduates." *Canadian Public Policy*, 26(2): 197-224.

Finnie, Ross and Ted Wannell. 1999. "The Gender Earnings Gap Amongst Canadian Bachelor's Level University Graduates: A Cross-Cohort, Longitudinal Analysis," in Richard Chaykowski and Lisa Powell (eds.), *Women and Work*. John Deutsch Institute, Queen's University.
Montréal: McGill-Queen's University Press.

Gunderson, Morley. 1989. "Male-Female Wage Differentials and Policy Responses." *Journal of Economic Literature*, 27(1): 46-72.

Gunderson, Morley. 1985. "Discrimination, Equal Pay, and Equal Opportunities in the Labour Market," in Craig Riddell (ed.), *Work and Pay: The Canadian Labour Market*. Toronto: University of Toronto Press.

Gunderson, Morley and Riddell, Craig. 1991. "Economics of Women's Wages in Canada," in Steven L. Willborn (ed.), *International Review of Comparative Public Policy*, V. 3: *Women's Wages - Stability and Change in Six Industrialized Countries*.

Jones, F. L. 1983. "On Decomposing the Wage Gap: A Critical Comment on Blinder's Method." *The Journal of Human Resources*, 18(1): 126-130.

Kidd, Michael P. and Michael Shannon. 1994. "An Update and Extension of the Canadian Evidence on Gender Wage Differentials." *Canadian Journal of Economics*, 27(4): 918-938.

_____ 1996. "The Gender Wage Gap: A Comparison of Australia and Canada." *Industrial and Labour Relations Review*, 49(4): 729-746.

Miller, Paul. 1987. "Gender Differences in Observed and Offered Wages in Canada, 1980." *Canadian Journal of Economics*, 20(2): 225-244.

Oaxaca, Ronald L. 1973. "Male-Female Wage Differentials in Urban Labor Markets." *International Economic Review*, 14(3): 693-709.

Oaxaca, Ronald L. and Michael R. Ransom. 1999. "Identification in Detailed Wage Decompositions." *The Review of Economics and Statistics*, 81(1): 154-157.

Reilly, Kevin T. and Tony S. Wirjanto. 1999a. "Does More Mean Less? The Male/Female Wage Gap and the Proportion of Females at the Establishment Level." *Canadian Journal of Economics*, 32(4): 906-929.

_____1999b. "Does More Mean Less? The Proportion of Females in the Establishment: Discrimination, Preferences and Technology." *Canadian Public Policy*, 25(Supplement 1): S73-S94.

Stata Corporation.1997. *Stata Reference Manual*, Release 5, Volume 2: 329-334. Texas: Stata Press.