Evolution of the Female Labour Force Participation Rate in Canada, 1976-1994: a Cohort Analysis

*Paul Beaudry and Thomas Lemieux**

One striking feature of the Canadian labour market during the 1970s and 1980s was the sustained growth in the numbers of working women. For women aged 25-64 the participation rate in the labour force increased from less than 50 per cent in the mid 1970s to 70 per cent in the late 1980s (panel a of Chart 1). An equally striking development is the recent stagnation in the female participation rate, which has remained around 70 per cent throughout the 1990s. While the participation rate of men of the same age group has declined during this latter period, Chart 1a shows that the 1990s represented a much less dramatic departure from previous trends — indicated by a dotted line in the charts — for men than for women.[1]

Why has female participation in the labour force stagnated? Is it a temporary phenomenon tied to the poor growth performance of the economy during the 1990s, or does it represent a permanent change in the behaviour of women? For example, it could indicate that the process of integration of women into the work force is almost complete. A cohort analysis is performed using data from the Survey of Consumer Finances (1976 to 1994) to examine this issue.

Our methodology involves isolating the effect of three separate factors on the participation rate of women — we follow cohorts of women over time, that is, we track the participation rate of representative groups of women who entered the work force at a given point in time (e.g. those who were 25 years old in 1976). We then "decompose" a cohort's participation rate into three effects. The first is a macroeconomic effect that by definition is common across cohorts. Recession and structural phenomena such as the generosity of the employment insurance system are some of the factors that may cause a macroeconomic effect. The sec-

Chart 1 Participation and employment rates: Men and women age 25 to 64

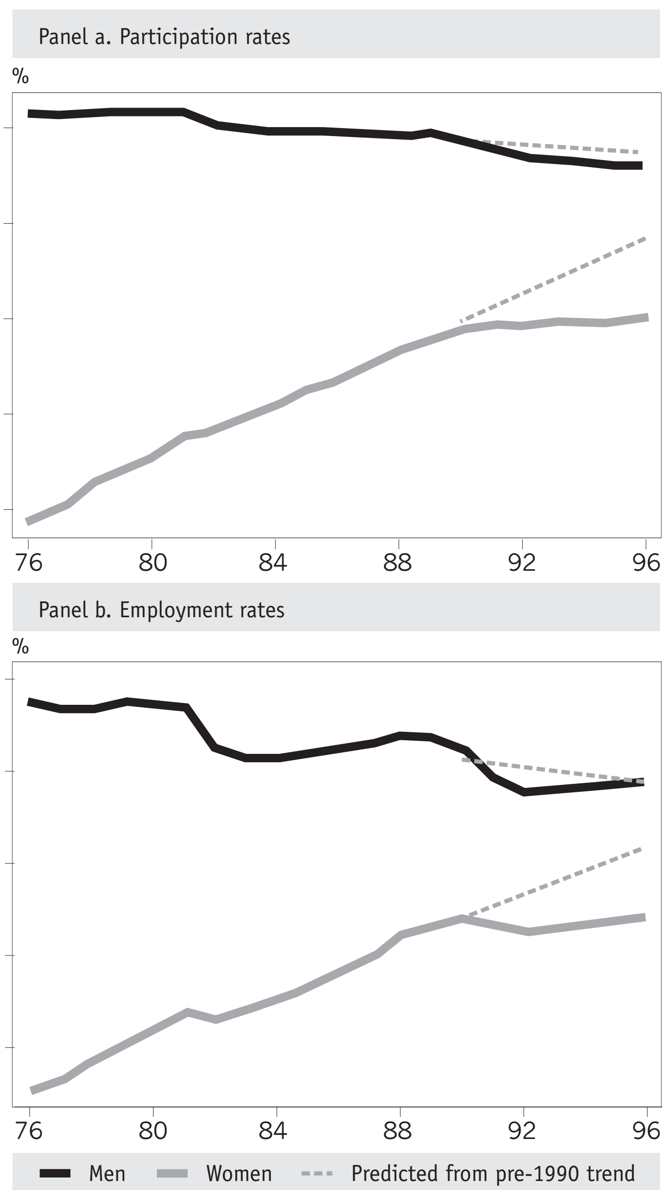

Table 1 Age structure and education level for women 25-64, 1976-94

Year	Age distribution in per cent				Per cent with high school or less	Number of obs.
	25-34	35-44	45-54	55-64		
	(1)	(2)	(3)	(4)	(5)	(6)
1976	40.9	24.0	19.8	15.3	73.6	12269
1978	39.1	24.9	20.5	15.5	75.0	17372
1980	39.9	26.1	19.1	14.9	74.7	18212
1982	39.5	26.2	19.1	15.3	73.5	18881
1983	39.2	27.5	18.9	14.4	72.7	19775
1985	38.7	29.3	17.7	14.3	71.0	19664
1987	37.9	29.9	18.0	14.3	70.4	17949
1989	37.1	31.0	18.3	13.6	68.3	21117
1991	34.9	32.3	19.3	13.6	59.7	26033
1993	31.4	32.8	21.6	14.2	56.2	22592
1994	32.8	32.7	21.2	13.3	54.3	22420
Average	37.0	29.2	19.4	14.3	67.2	216284

ond factor is the age or life cycle effect, which shows how the cohort's participation rate changes as its members age. The third is the cohort specific effect, which shows differences between cohorts for a given age and macroeconomic effect. For example, if the cohort that entered the labour force in 1976 has a participation rate 10 per cent higher than that of the cohort that entered the labour force in 1966 at the same age and under similar macroeconomic conditions, the 1976 cohort is said to exhibit a 10 per cent cohort effect relative to the 1966 cohort.

Our results indicate that cohort effects are likely the dominant factor in explaining the recent stagnation in female labour force participation rates. The same result was obtained when the labour force participation rate was replaced by the employment rate. Cohort effects help explain both the large increase in participation and employment rates during the 1970s and 1980s, as well as their stagnation in the 1990s. The 1989-94 recession merely amplified the stagnation phenomenon; it also explains the observed decline in the participation rate for some demographic groups. These results show, however, that stagnation would have occurred, albeit later in the 1990s, even if more favourable macroeconomic conditions had prevailed.

Data and Descriptive Statistics

Our date were obtained from the Survey of Consumer Finances (SCF) for 1976, 1978, 1980, 1982, 1983, 1985, 1987, 1989, 1991, 1993 and 1994 (survey years).[2] They were chosen because: from 1976-82 the survey was taken only every other year; since then it has been taken every year, except 1984; and 1994 was the last year available when we began this study. The years represented (about every other year) provide a fairly coherent sample across time. One important advantage of the SCF over other existing data sets is that the age of respondents is available in the public use samples provided by Statistics Canada.[3]

Individuals were grouped into two-year cohorts according to their date of permanent entry into the labour force. This was defined, somewhat arbitrarily, as the even-numbered year in which the woman was 25 or 26 years old (e.g. a women born in 1941 is in the "entering" cohort for 1966). Thus for each even year, all women aged 25-64 are divided into 20 cohorts (25-26, 27-28, ..., 63-64).[4] In total, 29 cohorts entered the labour force between 1936 and 1992.

Note that Statistics Canada's public use files of the SCF from before 1982 provide data only for heads of households and spouses; we therefore confined our analysis to this sub-sample for the entire 1976-94 period.[5] Labour force activity (employment, unemployment or non-participation) is determined based on individual responses to the usual Labour Force Survey questions (for April in the SCF). Chart 2 shows the evolution of the labour force participation rate, represented by a solid line, and the employment rate, represented by a dotted line, for the cohorts that entered the labour force in 1940, 1950, 1960, 1970, 1980 and 1990. These six cohorts are a relatively representative sample of the full 29 cohorts that are not shown to simplify the chart. Those entering the labour force first (e.g. the 1940 cohort) appear only in their last years in the labour market, while those that entered last (e.g. the 1990 cohort) appear only in their earlier years. The chart shows that the evolution of the labour force participation and employment rates is similar for all cohorts. Both these rates tend to increase from the age of 25 to 45-50 years, then decrease rapidly until age 65. Participation and employment rates are obviously higher for those cohorts that entered the labour force most recently.

Other descriptive statistics are shown in Table 1, which illustrates age composition and education levels (percentage of women with a high school education or less) for each year studied. The table shows quite a young population in 1976-94. About 65 per cent of women in the 25-64 age group during these years were 44 or younger. The impact of the baby boom/baby bust on the population's age composition is also clearly visible. This helps explain the increase in the proportion of women aged 35-44 since the beginning of the 1980s; the first wave of boomers born in 1946 reached the age of 35 in 1981. The same phenomenon occurred at the beginning of the 1990s as the first of the boomers reached 45. It is now the baby bust generation, those women born after 1965, who make up the 25-34 segment.

The statistics in Table 1 also show a steady increase in level of education: the percentage of women with a high school education or less

Chart 2 Female participation rate and employment rate by cohort

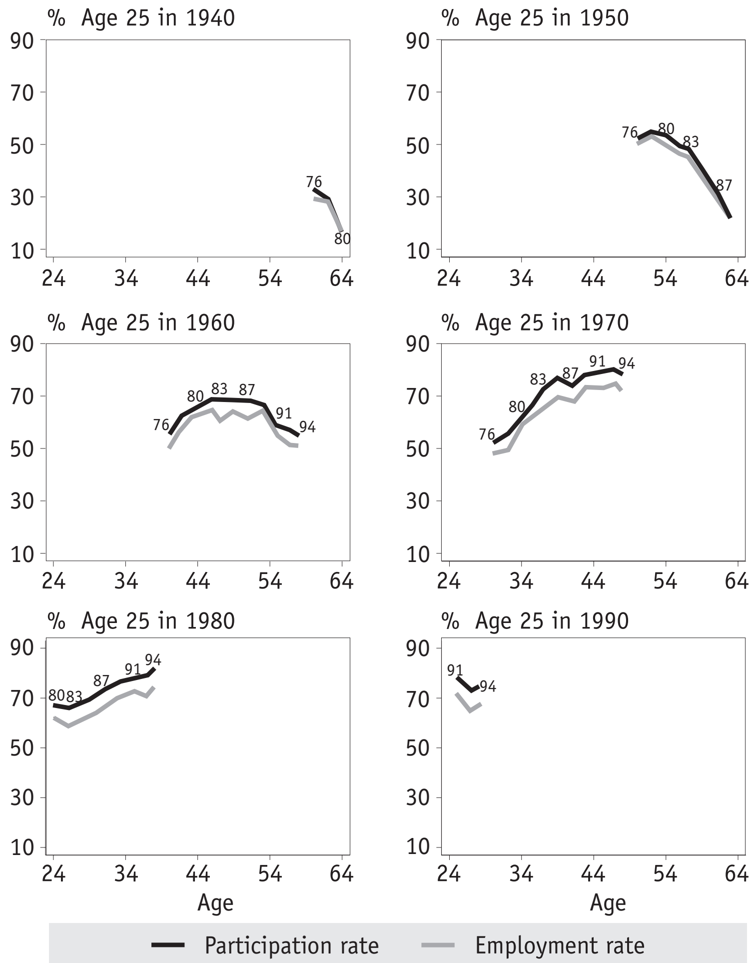

dropped from 73.6 per cent in 1976 to 54.3 per cent in 1994. This trend, however, is slightly exaggerated by the revamping of the questions on education in the LFS in 1990.[6]

Cohort Analysis: Econometric Model

Our cohort analysis uses an econometric model to examine the separate roles played by the macroeconomic, cohort and age effects on labour force participation and employment rates. The dependent variable used in the regressions is the participation (or employment) rate p_{jt} for cohort j at time t expressed in "log-odds" form $\ln(p_{jt}/(1-p_{jt}))$.[7] For example, $p_{74,84}$ represents the labour force participation rate for the cohort that entered the labour force in 1974 ($j = 74$) during the year 1984 ($t = 84$). This functional form is used to account for the special nature of variable p_{jt}, whose

Chart 3 Participation rate and age profiles

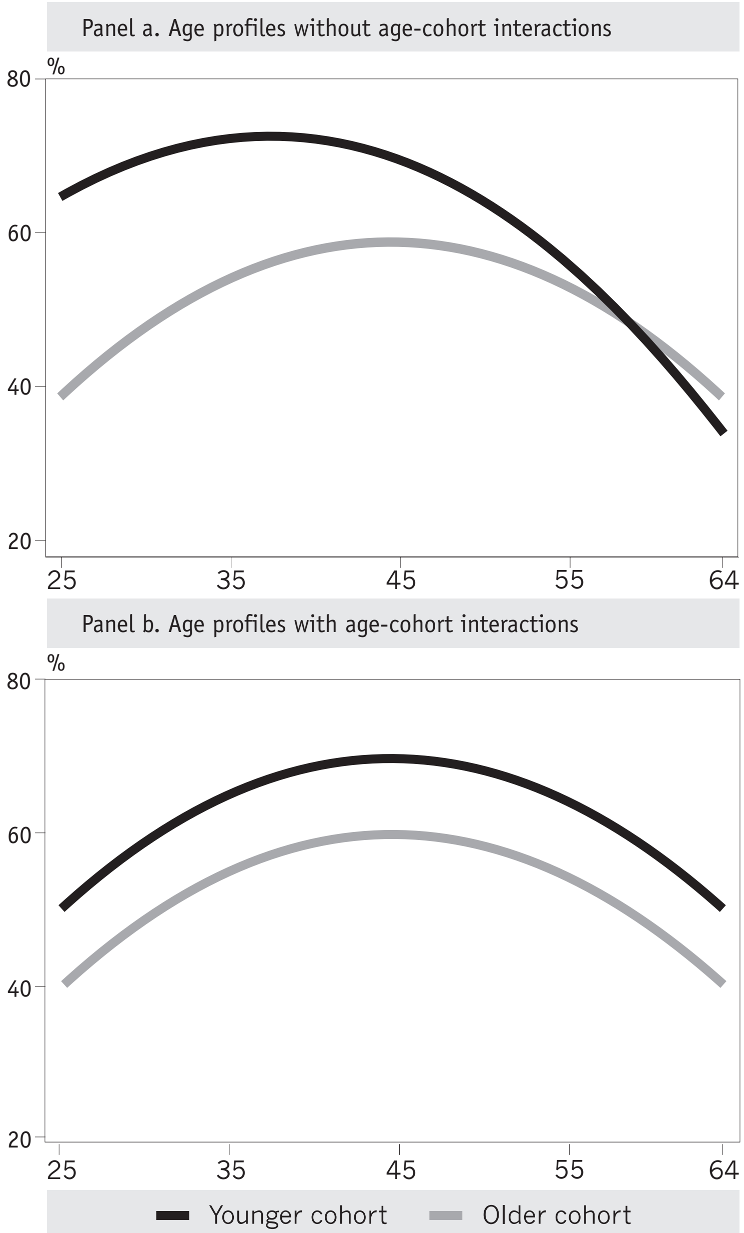

value is always between 0 and 1. It ensures the predicted value will always be between 0 and 1, which would not be the case if a standard linear specification were used instead.[8]

In most of the estimated models, only one macroeconomic variable is used, the unemployment rate among men aged 25-44. Although certain long-term trends in this rate may be determined by structural factors, it is clear its short-term fluctuations mainly reflect the evolution of the economic climate. Other variables such as the output gap may be used in as well as the unemployment rate, but we prefer to concentrate on the latter, because of its simplicity; however, the results must be interpreted with caution.[9] The scope of the macroeconomic effect will, however, be broadened by adding other variables in our discussion on robustness analysis later in this article.

The model is completed by adding a flexible specification for age, or life cycle, effects (fourth degree polynomial) and for cohort effects (third degree polynomial). This yields the equation:

$$(1) \quad \ln (p_{jt}/(l-p_{jt})) = \alpha + \delta ur_t + \beta_1 j + \beta_2 j^2 + \beta_3 j^3 + \gamma_1 a_{jt} + \gamma_2 a_{jt}^2 + \gamma_3 a_{jt}^3 + \gamma_4 a_{jt}^4$$

where δur_t represents the unemployment rate among men 25-44, j represents the cohort, and a_{jt} represents the age of women from cohort j in year t.

One characteristic of equation (1) is that the age profile for each cohort, i.e. the evolution of the labour force participation rate over the life cycle, is similar for each cohort; they differ only in terms of the intercept. In other words, the model allows a vertical displacement of the life cycle profile from one cohort to another while forcing the shape of the profile, and thus the slope, to be identical for each cohort. A more general model is produced by introducing age-cohort interaction terms to allow the age effect to vary from one cohort to another. This was done with the following model, which incorporates an age-cohort (a_{jt}j) and an age-cohort squared interaction term:[10]

$$(2) \quad \ln (p_{jt}/(1-p_{jt})) = \alpha + \delta ur_t + \beta_1 j + \beta_2 j^2 + \gamma_1 a_{jt} + \gamma_2 a_{jt}^2 + \gamma_3 a_{jt}^3 + \gamma_4 a_{jt}^4 + \theta_1 a_{jt} j + \theta_2 a_{jt} j^2.$$

If second or higher order polynomial terms are omitted, equation (2) shows that the effect of age on $\ln(p_{jt}/(1-p_{jt}))$ is equal to $\gamma_1 + \theta_1$j. If θ_1 is positive, the age effect will be greater for those cohorts that most recently entered the labour force (highest j) than for the others, and vice versa. Coefficient θ_1 thus allows the life cycle profile to vary from one cohort to another.

Chart 3 illustrates the advantages of equation (2) over equation (1), which does not include the age-cohort interaction term. Without such interaction terms, the intercept is the only difference between different cohorts' age profiles (panel a). The same increase in participation at career outset and the same decrease in participation at career end is shown for every cohort. The age profile is clearly more flexible in panel b where interaction

Table 2 Detailed results of regressions

(standard deviation in brackets)

	Employment rate			Participation rate		
	1976-1994	1976-1994	1976-1989	1976-1994	1976-1994	1976-1989
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.792	-3.102	-2.912	0.997	-3.398	-2.746
	(0.143)	(0.291)	(0.344)	(0.158)	(0.245)	(0.177)
Unemployment rate.	-0.019	-0.029	-0.031	-0.004	-0.016	-0.014
	(0.011)	(0.004)	(0.003)	(0.011)	(0.003)	(0.002)
Cohort effect: a						
co	-1.830	3.804	3.506	-2.227	4.162	3.267
	(0.154)	(0.383)	(0.465)	(0.151)	(0.319)	(0.215)
co^2	1.609	-0.792	-0.681	1.892	-0.851	-0.561
	(0.131)	(0.118)	(0.150)	(0.102)	(0.100)	(0.062)
co^3	-0.289	—-	—-	-0.336	—-	—-
	(0.025)			(0.019)		
Age effect:						
Age/10	0.389	2.706	2.780	0.405	3.004	2.749
	(0.053)	(0.175)	(0.216)	(0.051)	(0.124)	(0.126)
(Age/10)^2	-0.445	-1.742	-0.676	-0.471	-0.760	-0.701
	(0.024)	(0.101)	(0.027)	(0.021)	(0.032)	(0.019)
(Age/10)^3	-0.130	-0.151	-0.161	-0.139	-0.162	-0.167
	(0.007)	(0.007)	(0.009)	(0.006)	(0.006)	(0.009)
(Age/10)^4	0.031	0.032	0.024	0.032	0.034	0.027
	(0.006)	(0.006)	(0.004)	(0.005)	(0.005)	(0.005)
Interactions:						
Age/10	—-	-1.742	-1.902	—-	-1.963	-1.886
*co		(0.101)	(0.140)		(0.067)	(0.091)
(Age/10)	—-	0.226	0.298	—-	0.258	0.305
*co^2		(0.021)	(0.021)		(0.016)	(0.021)
P-value of						
pred. test:	0.001	0.096	—-	0.000	0.036	—-
R squared:	0.941	0.954	0.951	0.953	0.964	0.962
Number of observations	224	224	164	224	224	164

a co=(year of entry into the labour force — 1932)/20

b p-value of a (Chow) test that the 1990-94 data is correctly predicted by the model estimated from the 1976-89 data.

Chart 4 Effect of the entry of new cohorts on the participation rate

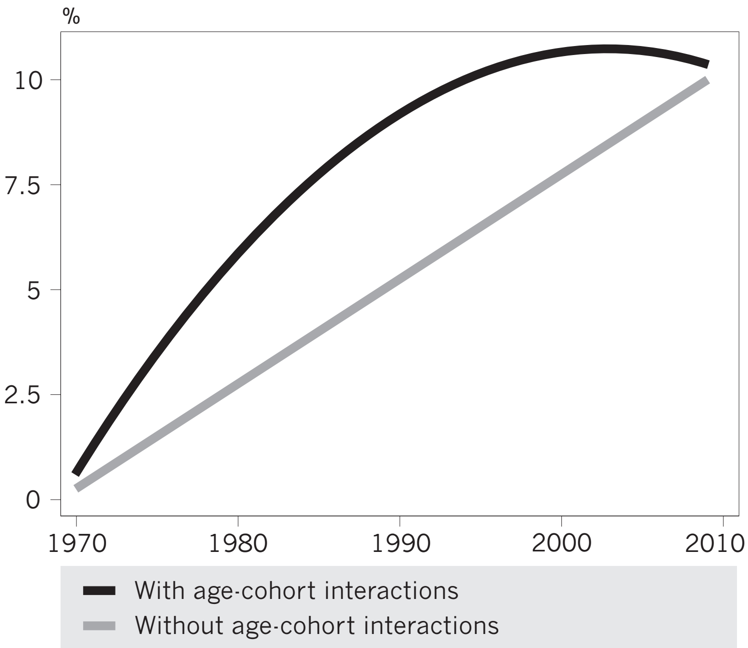

terms are introduced. Here, the "newer" cohort has both a higher ordinate value at the origin and a shallower slope. This results in a higher and more stable age profile at career outset than in the previous cohorts (the "older" cohorts). This profile is also more similar to that for men, whose participation rates are fairly high and stable until the age of about 50. The situation shown in panel b is therefore more consistent with the idea of a convergence between men's and women's labour force participation rates, or increasing participation of women in the labour force, than that shown in panel a.

In panel b, the cohort effect is concentrated at career outset, participation rates before the age of 40 for the newer cohort being much higher than those for the older cohort, while the rates are reasonably comparable after age 50. The impact of the entry of the newer cohort on the aggregate labour participation rate would thus be felt most strongly during the first 10 or 20 years after its arrival, while in panel a its influence is shown as continuing throughout the life cycle. In other words, the entry of newer cohorts in panel b should result in a rapid increase in the aggregate participation rate, followed by a period of stagnation. Panel a, on the other hand, implies a constant increase in the aggregate participation rate.

Chart 4 shows the impact of the arrival of newer cohorts in the two cases discussed above, with those cohorts entering the labour force after 1970 considered "newer" and those entering before 1970 considered "older."[11] The chart clearly indicates that only the presence of an age-cohort interaction effect can explain the stagnation phenomenon.

Though the case presented in Charts 3 and 4 is only an illustrative example, it provides insights into the importance of the age-cohort interaction in explaining stagnation of participation rates in the 1990s. The specification we actually estimate (equation 2) is a flexible version of the dichotomous case considered in the two charts. It allows for smooth changes in age profiles across cohorts.[12]

An often-mentioned problem with cohort analyses is the impossibility of separately identifying cohort effects, year effects (macroeconomic effects), and age effects because of the linear dependence between them. In fact, since $a_{jt} = 25 + t - j$, the three variables (a_{jt}, j and t) are perfectly collinear. Our work assumes that variable δur_t captures any systematic macroeconomic effect and that there is no other temporal trend in this effect. If there were a time trend in the female participation rate, for example, because of the feminist movement, we would be attributing this trend to cohort effects. Note, however, that as long as the time trend is linear, it cannot account in itself for the stagnation in participation rates (a break in trend) in the 1990s. So although there may be some issues regarding the interpretation of our results, we do not think our analysis of the sources of the stagnation in female participation rates will be affected by these issues.

These considerations aside, econometric models (such as (1) and (2)) can never explain all the variations in the data (R squared is less than 1). As a rule, a residual macroeconomic effect is obtained, representing the macroeconomic variation in the data that cannot be explained by other variables in the model. If during a period, say the 1990s, we were to find a large residual we would interpret this as indicating that participation in this period has experienced a macroeconomic effect not captured by its standard co-movement with the unemployment rate.

Results

Equation (1) and (2) were estimated using weighted least squares, with cohort size j at time t

used as the weights. The results are shown in Table 2 for employment rates (columns 1-3) and participation rates (columns 4-6). For model (1), note that all the coefficients are significant, except for the rate of unemployment among men 25-44 (columns 1 and 4). That effect becomes significant, however, when age-cohort interaction terms from model (2) are introduced (columns 2 and 5). Also note that interaction term coefficients are highly significant, and that R-squared for model (2) is higher than for model (1).

We also present the results of the regressions when the sample is limited to the 1976-89 period (columns 3 and 6). The purpose of this exercise is to assess whether the levelling off of participation and employment rates in the 1990s was predictable from the behaviour of these rates before 1990. The results indicate that the estimated parameters for 1976-89 are relatively similar to those for the period as a whole. In fact, the p-value of the Chow-test statistic reported at the bottom of column 2 indicates we cannot reject the null hypothesis that the estimated employment rate models are the same for 1976-94 (column 2) and 1976-89 (column 3).[13] The estimated models for 1976-94 and 1976-89 are only marginally different (p-value of 0.036 in column 5) in the case of the participation rate. We shall return to the question of the stagnation of the participation and employment rates.

To facilitate presentation of the results, it is simpler to use a graphical approach than to examine the numbers presented in Table 2 in detail. For each rate (participation and employment) and each model (1 and 2), Chart 5 presents four panels. Panel a shows the cohort effect at age 44, i.e. the variations in the participation and employment rates attributable to the cohort effect at a precise point in the life cycle. Panel b shows the age effect throughout the life cycle for a typical cohort (the one that entered the labour force in 1964). Panel c presents a similar result for the six representative cohorts of Chart 2 (entering the labour force in 1940, 1950, 1960, 1970, 1980 and 1990) to illustrate cohort differences over the entire age profile. It should be noted, however, that panel b shows a predicted age profile for the entire life cycle, while panel c shows the profile only for the ages at which the cohort in question is observed in the data (1976-94). Finally, panel d indicates the degree to which the estimated model for

Chart 5 Participation rate model without age-cohort interactions

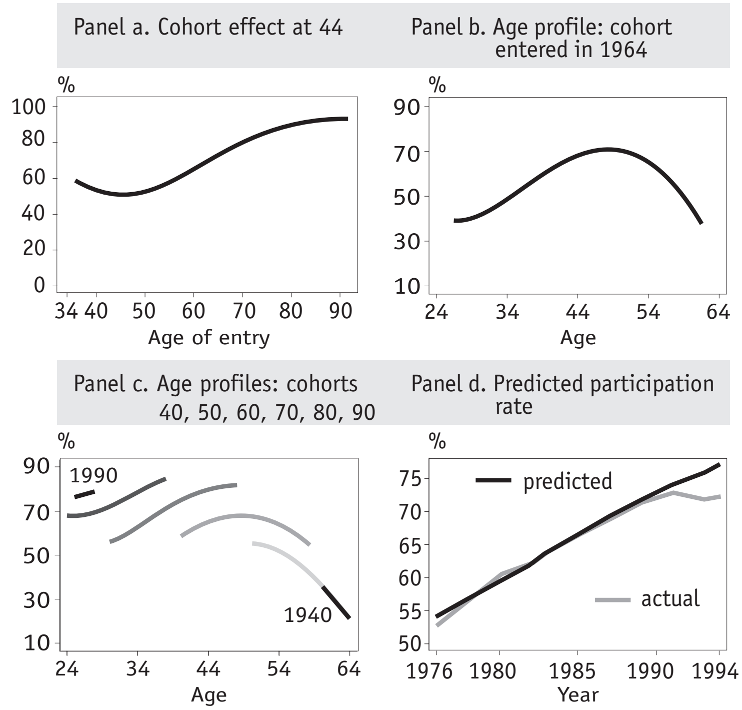

Chart 6 Participation rate model with age-cohort interactions

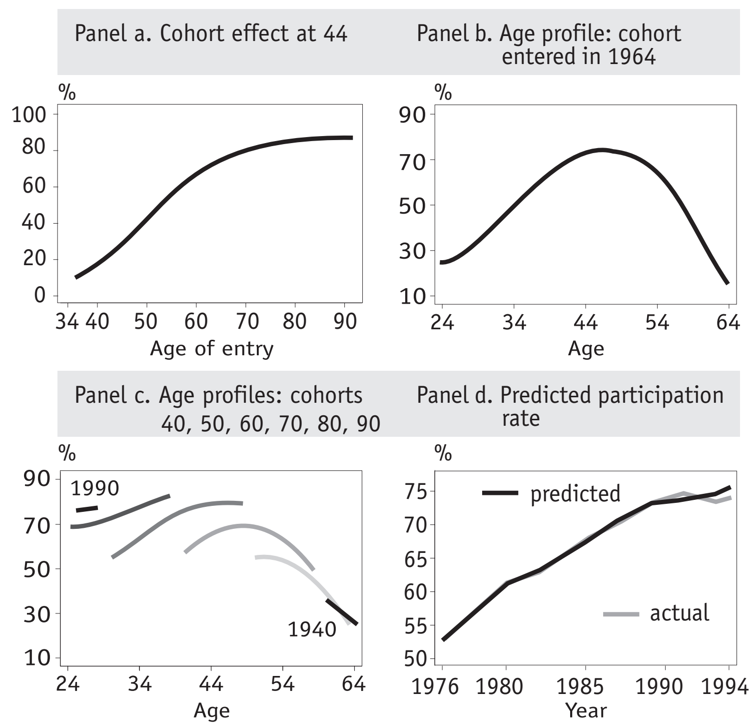

Chart 7 Employment rate model without age-cohort interactions

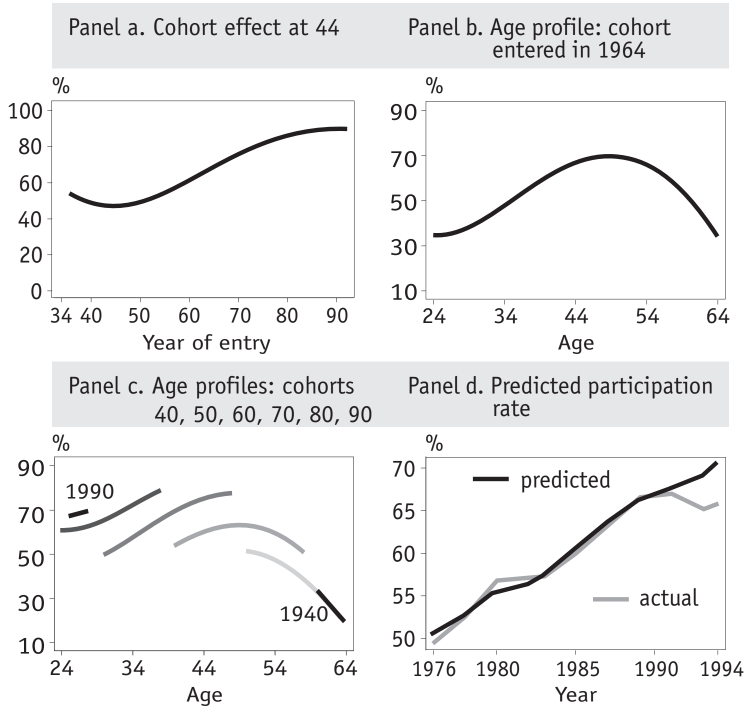

Chart 8 Employment rate model with age-cohort interactions

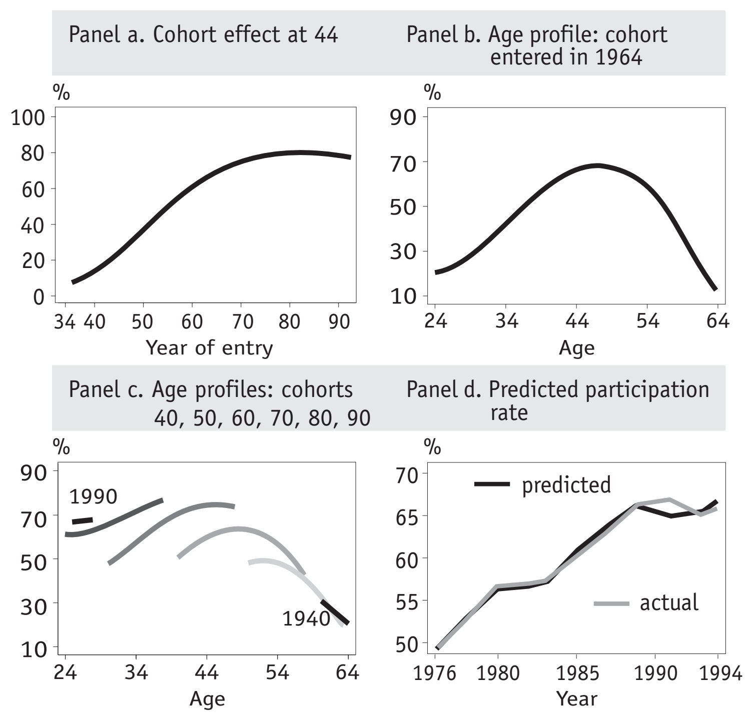

1976-89 can be used to predict participation and employment rates for the entire 1976-94 period.

Rather than discussing each of the graphs in detail, we will confine ourselves to noting a few highlights:

- All the estimated models indicate a levelling off in cohort effects (panel a of Charts 5-8); the participation and employment rates for the 1992 cohort are comparable to the ones for the cohorts that entered the labour force in the 1980s (or at least, the cohorts will all be comparable once they have reached age 44).
- The participation and employment rates peak around age 50 (panel b of Charts 5-8).
- The younger cohorts have flatter age profiles (shallower slopes) early in their careers. This pattern is particularly pronounced for model 2, which includes age-cohort interaction (panels c of Chart 6 and 8).[14]
- The model without interactions (model 1) provides no explanation whatsoever of the stagnant participation rate (panel d in Chart 5) and falling employment rate (panel d in Chart 7) observed in the 1990s. On the other hand, these phenomena can be predicted more accurately from the model with interactions (model 2) estimated for 1976-89 (panel d in Charts 6 and 8).

This last finding is the most interesting. The charts suggest model 2 does a much better job predicting post-1990 patterns from pre-1990 data than model 1. The prediction tests reported at the bottom of Table 2 is a formal way of evaluating this hypothesis. The p-value for the specifications corresponding to model 1 (columns 1 and 4) are well below the critical value of 0.05, which suggests that the models fit to the 1976-89 period do not predict accurately the post-1990 data. By contrast, the specifications corresponding to model 2 predict more accurately the post-1990 data, especially for the employment rate model (p-value of 0.096). Though the p-value is higher in column 5 (model 2) than in column 4 (model 1), we nevertheless reject the null hypothesis of an accurate prediction for the participation rate.

The reason why the prediction test fails for the participation rate can be seen in panel d of Chart 6. This shows that the model predicts some growth in the participation rate in the early 1990s while the actual data show a clear stagnation.[15] Relative to the pre-1990 period for which the fit of the model is almost perfect (predicted and actual values look almost identical), the discrepancy between the predicted and actual values in the 1990s seems sizable.

All in all, however, model 2 provides a much more accurate prediction of what happened in the 1990s than model 1. Looking at the average for 1991, 1992, and 1994 (three years of observation for the 1990s), model 1 overpredicts the employment and participation rates by three percentage points and 3.2 percentage points, respectively. By contrast, model 2 underpredicts the employment rate by 0.4 percentage points and overpredicts the participation rate by 0.5 percentage points.[16]

Our results seem to indicate that, in addition to unfavourable macroeconomic conditions, the levelling off of cohort effects also contributed to the trend observed in the 1990s. This hypothesis will be examined in detail in our discussion of decomposition. Note also that the stable cohort effects for cohorts entering after 1970 yields some support for the assumption there is no time trend in participation rates (Section 2a). If there were a positive time trend in participation rates, the cohort effects for cohorts entering after 1970 would have to follow a negative trend (youngest cohorts less active than those that entered in 1970), which is a surprising result.[17]

In our more formal analysis of the role of different factors in the recent evolution of aggregate participation and employment rates for women aged 25-64, we break down this evolution into four components: the macroeconomic effect related to the unemployment rate among men aged 25-44 (the economic cycle), the residual macroeconomic effect, the age effect and cohort effects. In terms of equation (2), it is relatively easy to identify the first two factors, which correspond to the term δur_t and to the residuals of this equation.[18]

More precisely, we first calculate the participation (or employment) rate for each year, taking the weighted average of p_{jt} values for each t. The observed rate (p_{jt}) is then replaced by the predicted rate ${}_{jt}$ from the estimated model.[19] The average of the ${}_{jt}$ values for each t therefore represents the aggregate rate predicted by the model. The difference between the observed aggregate rate and the predicted rate represents the residual macroeconomic effect.

Chart 9 Decomposition of the participation rate

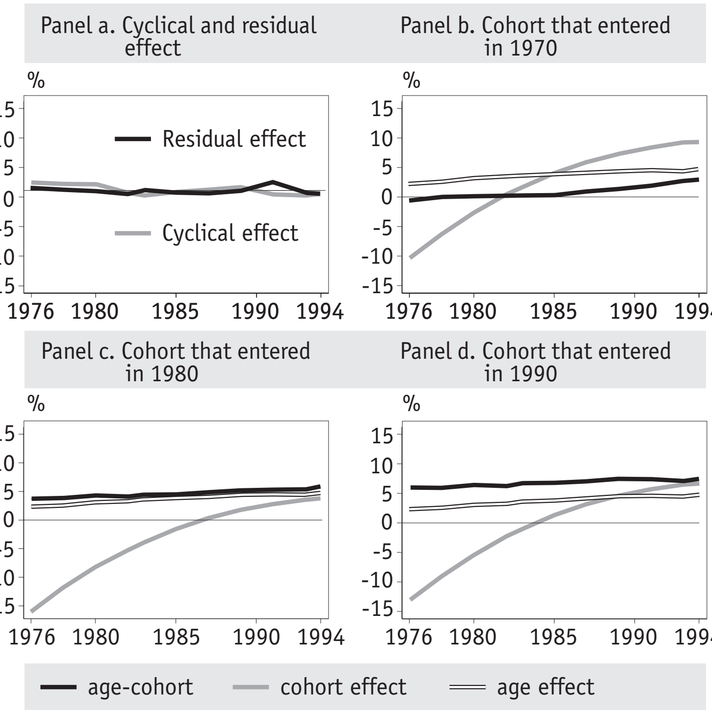

We then recalculate the prediction by replacing the observed unemployment rate by the average of the unemployment rates over the entire sample (8.2 per cent). The difference between this new prediction and the preceding prediction represents the macroeconomic effect related to the male unemployment rate, which we also call the cyclical effect.

The cyclical and residual effects obtained in this manner are presented in panel a of Chart 9 for the participation rate. Note that the charts showing the decomposition results for the employment rate are qualitatively similar to those for the participation rate and are not reported here.[20] We nevertheless discuss the main results both for the participation and the employment rate in the text.

During the 1990s, the cyclical effect is about -1 per cent for the participation rate and -2 per cent for the employment rate. In other words, the female participation rate would have been 1 per cent higher in the first part of the 1990s if the male unemployment rate had held steady at 8.2 per cent.

Age and cohort effects are somewhat more complicated to understand because of the interaction terms in equation (2). It should be noted, first of all, the age effect comes into play only to the

Chart 10 Predicted and actual cohort effects participation rate

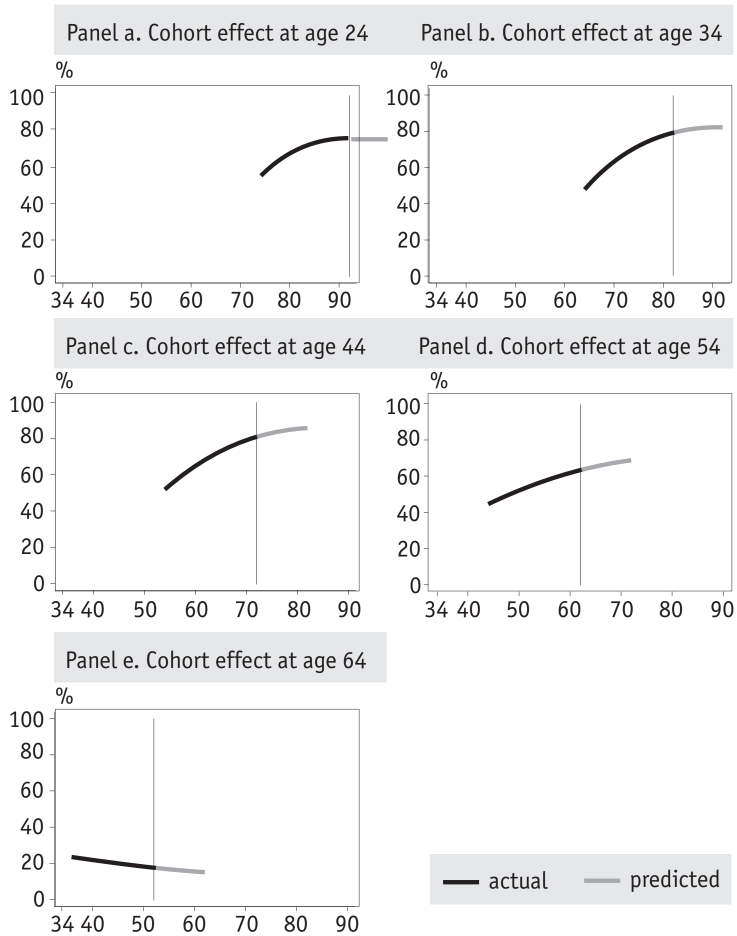

extent that the population's age composition changes over time. For example, the arrival of the baby boomers in the labour force in the early 1970s considerably rejuvenated the 25-64 population as a whole. As these young women had below-average participation rates, it should have been expected this change in composition would have had a negative effect on the aggregate participation rate and vice versa.

It might therefore be supposed that to identify the age effect, it is enough to recalculate the predicted rate using a uniform age composition (5 per cent of the population aged 25-64 in each two-year age group) instead of the observed age composition. The problem with this procedure is that it depends on the cohorts present in the labour force in each year, since the age profile is dependent on the cohorts through the interaction terms. This procedure therefore serves to isolate the age effect plus the crossed age-cohort effect.

The same problem arises when we want to isolate the role of cohorts. For example, we can try to recalculate the predicted rates by replacing the cohort effect expressed as $\beta_1 j + \beta_2 j^2 + \theta_1 a_{jt} j + \theta_2 a_{jt} j^2$ by the cohort effect obtained if the cohort is set at an arbitrary level such as $j = 70$ ($\beta_1 70 + \beta_2 70^2 + \theta_1 a_{jt} 70 + \theta_2 a_{jt} 70^2$). This gives us the cohort effect plus the crossed age-cohort effect, in the same way as in the case of age. Once we have all this information, however, it is possible to separate the overall effect of age and cohorts into three components: a "pure" age effect that indicates how participation rates would evolve over time if all cohorts exhibited the same participation profile; a "pure" cohort effect that indicates how the participation rate for a population with a uniform age distribution would evolve over time as older cohorts get replaced by younger ones; a joint age-cohort that residual interactions between age and cohort effects.[21]

These different effects are illustrated in panels b-d of Chart 9. Let us take the example of panel b: here, we use the cohort that entered the labour force in 1970 as a reference cohort for the decomposition. The cohort effect thus indicates the difference between the observed rates and the rates that would have prevailed had all the cohorts followed the same age profile as cohort 17, other factors being kept constant. This cohort effect is therefore the "pure" effect mentioned earlier. The panel also shows the "pure" age effect (for a given cohort and other factors) as well as the combined age and age-cohort effect (the age effect for the observed cohorts in each year).

While it can be rather difficult to grasp all the details of these decompositions, the results speak for themselves: it is really the cohort effect that dominates the evolution of the participation and employment rates in 1976-94. The results are similar, regardless of which cohort is used as a reference for the decompositions (1970, 1980 or 1990). We find that cohort effects account for an increase of about 20 percentage points in participation and employment rates over the period. At the same time, the charts clearly show this phenomenon seems to be coming to an end. By comparison, age effects play a relatively small role in recent changes.

To sum up, our results indicate that the stagnation of female participation and employment rates is primarily a structural phenomenon related to the stabilization of the cohort effects, which were responsible for the remarkable increase in these rates in the 1970s and 1980s. The unfavourable macroeconomic situation amplified this phenomenon but was not the root cause. The relative performance of the participation and employment rates during the recessions of 1981-83 and 1989-94 clearly illustrates this phenomenon: in 1981-83, the downward pressure on the rates from the macroeconomic effect was offset by the cohort effects, pushing the rates up by one percentage point a year, whereas in 1989-94, because of the stabilization of cohort effects, macroeconomic effects comparable to those of 1981-83 resulted in lower participation and employment rates.

To clarify the role of cohort effects, we illustrate their magnitude at age 24, 34, 44, 54 and 64 in Chart 10. Take for example panel c, which shows the cohort effect at age 44 by year of entry into the labour force. The vertical line indicates the cohort aged 44 in 1994. The curve to the left of the line describes the evolution of cohort effects in 1976-94. The curve to the right shows the predicted evolution for the coming years.

Robustness Analysis

Education

The results indicate a general slowing trend for most of the ages under consideration, attributable to cohort effects. This is particularly true for the younger groups (ages 24 and 34), which explains why the levelling off and declining trend is more pronounced for the 25-44 age bracket than for the 45-64 bracket (Beaudry and Lemieux, 1998).

In our robustness analysis, we re-estimated the models separately for women who have pursued post-secondary studies and those who have only a high school diploma or less. The highlights of the results are:[22]

- The decline in the participation and employment rates in the 1990s is more pronounced among poorly educated women than for the entire female population. The growth in employment and participation in 1976-89 is also weaker within each education group than for the population as a whole. A significant portion of the rise in the rates for the population as a whole therefore seems to be attributable to the increase in average education levels.
- Similarly, the cohort effects exhibit a decline for the most recent cohorts in most cases. This suggests the average quality of cohorts is declining, since a high education level is a less selective characteristic than it was in the past.
- The younger cohorts have high and flat age profiles for women who pursued post-secondary studies. These profiles are similar to the ones for men with the same education levels.

Employment insurance

We have also re-estimated the regressions presented in Table 2 when the employment insurance subsidy rate is also used as a macroeconomic variable.[23] The results are not very conclusive, since the effect on the employment rate is negative when we also control for the unemployment rate among men aged 25-44.[24] The decline in the subsidy rate during the 1980s should therefore have increased the employment rate instead of lowering it. The effect on the participation rate is not significant. This being said, including the subsidy rate as a macroeconomic variable has little impact on the model's other coefficients. Our conclusions about the role of cohort effects versus macroeconomic effects during the 1990s therefore remain unchanged.

Alternative measures of the cycle

In another attempt to examine the robustness of our results, we have re-estimated our main models with an alternative cyclical variable, the after-tax real wage of women, instead of the unemployment rate of men aged 25-44.[25] This variable did not perform as well as a measure of the cycle as the unemployment rate of men in the same age group. For example, the estimated effect of the after-tax real wage of women was typically not statistically significant when the unemployment rate of men was included. The other parameters of the models were similar to those reported in Table 2.[26]

Predictions

Finally, we will now attempt to predict the future evolution of the participation and employment rates under two different macroeconomic scenarios: a 8.2 per cent unemployment rate for men aged 25-44 (the average for 1976-89) and a 6.6 per cent unemployment rate for the same group (the 1989 level). To do so, we must make some assumptions about the cohorts entering the labour force after 1994. We first assume their age profile will be similar to those of the last cohort

Chart 11 Participation and employment rate forecast for females 25-64

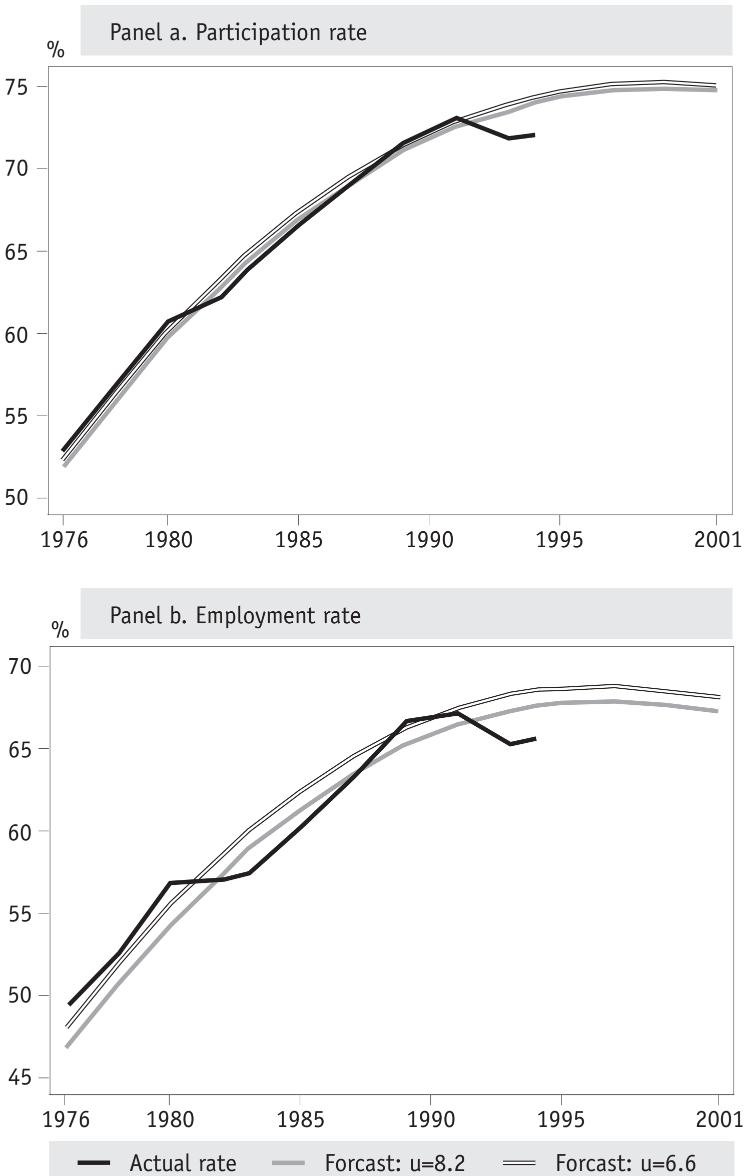

observed (which entered the labour force in 1992). This is a natural assumption since our empirical results indicate the cohort effects have been relatively stable for those cohorts entering the labour force in the 1980s and early 1990s (see panel a of Chart 6).[27] Since the age distribution is relatively uniform for cohorts entering the labour force throughout the 1990s, we also assume these new cohorts have the same size as the cohort that entered the labour force in 1992.[28]

The results of the simulations are presented in Charts 11. The conclusions are the same in both cases: large increases in the participation and employment rates are clearly a thing of the past; in the future, these rates can be expected to hold relatively stable. However, there is still room for a two-three percentage point increase in the rates if the macroeconomic situation continues to improve. It is illusory, though, to think that the rates could rise five-10 percentage points during the next period of expansion as they did in 1983-89. The cohort effects that prevailed then no longer exist.

Conclusion

This study's main finding is that the levelling off of female participation and employment rates is primarily a structural phenomenon related to the stabilization of the cohort effects, which accounted for the remarkable increase in these rates in the 1970s and 1980s. The unfavourable macroeconomic situation has amplified this phenomenon but is not the root cause. The relative performance of the participation and employment rates in the 1981-83 and 1989-94 recessions clearly illustrates this phenomenon. In 1981-83, the downward pressure on the rates from the macroeconomic effect was offset by the cohort effects, pushing the rates up by one percentage point a year, whereas in 1989-94, because of the stabilization of cohort effects, macroeconomic effects comparable to those of 1981-83 resulted in lower participation and employment rates.

This result is strongly dependent on the amount of flexibility used to capture cohort effects. It is essential the age profile as a whole, and particularly its slope, be allowed to vary from one cohort to another. This makes it possible to accurately trace both the rise and the flattening of the employment and participation profiles by age. These phenomena are consistent with a convergence in the behaviour of men and women in the labour market: men exhibit high and flat (at least until age 55) employment and participation profiles. The profiles of recent female cohorts are therefore closer to those of men than to those of older female cohorts.

Finally, the recent evolution of participation rates in the United States seems to corroborate our findings: Chart 12 shows that, despite more favourable macroeconomic conditions in the United States than in Canada, the U.S. female participation rate grew much slower in the 1990s than in the 1970s and 1980s. By contrast, the participation rate of U.S. men in the 1990s followed closely the pre-1990 trends. The fact there was a sharp break from previous trends in both countries suggests the poor performance of the Canadian economy in the 1990s was not the main factor behind the stagnation in female participation rates. This reinforces our conclusion on the importance of cohorts effects, rather than cyclical factors, as the main explanation for this phenomenon.

Chart 12 Participation rates in the United States Men and women age 25 to 64

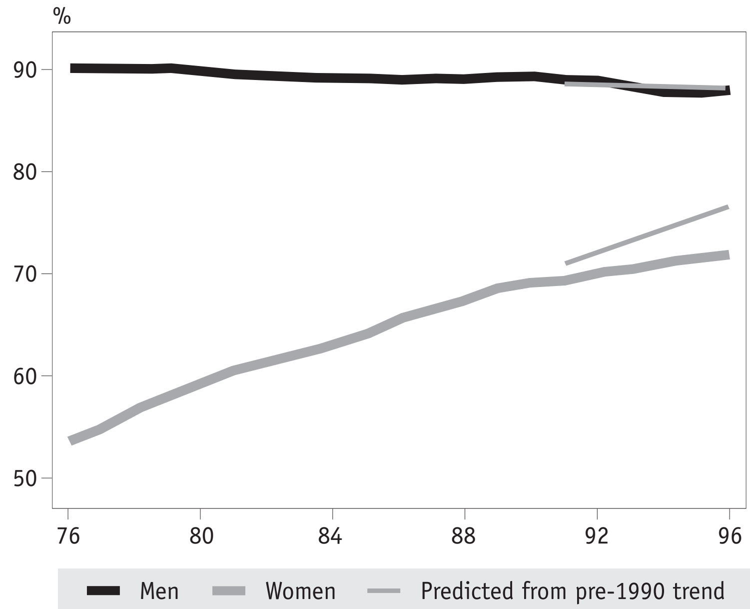

Paul Beaudry teaches in the Department of Economics at the University of British Columbia and is a Research Associate with CIRANO (beaudry@econ.ubc.ca). Thomas Lemieux teaches in the Department of Economics at the Université de Montreal and is a Research Associate with CIRANO.

Notes

1. The dotted line represents the values of the participation rate (from 1990 to 1996) predicted by a linear trend model fit to the 1976 to 1989 period.

2. SCF data are also available for 1972 to 1975. We are not using them, however, due to a problem of non-comparability of the questions on participation and employment, which were revamped as part of major changes to the Labour Force Survey in 1976.

3. The public use samples of the Labour Force survey have bigger sample sizes and are available for every year since 1976. Unfortunately, this data set in unsuitable for a detailed cohort analysis because respondents are pooled into 5-year age groups.

4. For odd-numbered years (1983 to 1993), the women included are between 25 and 65 years old.

5. Beaudry and Green (1996) point out that this restriction has little impact on individuals aged 25 and older, who are almost all heads of households or wives.

6. Before 1990, LFS questions on education did not specifically mention trade certificates. Since that time, those who have finished high school and also hold a trade certificate have been placed in the "post-secondary" category. This explains the break in the series between 1989 and 1991. Note also that changes to the LFS questionnaire make it difficult to use a more detailed classification than "high school or less" versus "post-secondary".

7. The "log-odds" specification is in fact only an application of the "logit" model to grouped data (grouped by cohort).

8. Consider $q = \ln(p/(l-p))$. This results in $p = \exp(q)/(1 + \exp(q))$ where $0 \le p \le 1$, since $\exp(q) > = 0$. Thus, no matter what the value of q predicted by a regression such as equation 1, the predicted value of p will therefore always be between 0 and 1.

9. In theory, the output gap should be more representative of the economic climate than unemployment among men aged 25 to 44. In fact, the validity of the output gap depends on the accuracy of a number of difficult-to-verify hypotheses on the structure of the macroeconomic model used to obtain this measurement.

10. This particular functional form was arrived at using specification tests. As a rule, we add terms of a higher and higher power until the terms added are no longer significant. The same procedure was used for equation (1).

11. We also assume that age composition is uniform (one fortieth of the population at each age).

12. The specification proposed in equation 2 fits the data much better than a specification more directly in the spirit of the illustrative example (an "older" and a "newer" cohort with completely separate age profiles). The problem with the illustrative example is that it implies a sharp discontinuity across cohorts which is inconsistent with the observed data. Equation 2 also implies some restrictions since only the linear term of the age profile is allowed to vary across cohorts. We chose that latter specification because allowing for more general interaction terms did not affect our main conclusions.

13. This version of the Chow test compares the estimated models for a prediction sample (1976 to 1989) and for the full sample (1976 to 1994). It can either be interpreted as a model stability test (are the estimated models the same for the two sample periods?) or as a prediction test (does the model estimated for the prediction sample predicts accurately for other years?).

14. Note that the slope of the age profile in the participation rate (or employment rate) and age age space may change even without interaction terms due to the log-odds specification.

15. The actual data shows a small increase in the participation rate from 1989 to 1991, followed by a small decrease from 1991 to 1993. However, these variations are not statistically significant. It is thus accurate to describe the period from 1989 on as one of stagnation in the participation rate. The stagnation in the participation rate can be seen more clearly in Chart 1a which is based on the larger and, therefore, more accurate samples of the Labour Force Survey.

16. An alternative set of "prediction tests" consists of verifying whether these average prediction errors are significantly different from zero. These tests yield the same inference as the Chow tests reported in Table 2, namely that only model 2 for the employment rate predicts accurately the post-1990 data.

17. To see this point, consider the simple linear participation model $p = \alpha + \beta j + \gamma a$, where the time trend has been set arbitrarily to zero. Since $t = j + a - 25$, if there is a true time trend δ, an equivalent specification can be rewritten as $p = \alpha + \beta j + \gamma a + \delta t - \delta t = \alpha + \beta j + \gamma a + \delta t - \delta(j + a - 25) = (\alpha + 25\delta) + (\beta - \delta)j + (\gamma - \delta)a + \delta t$. In other words, the cohort effect becomes $(\beta - \delta)$ instead of β. If β is equal to zero (stable cohort effects as for the post-1970 cohorts), this means that, controlling for age and temporal effects, participation rates are lower for younger (high j) than older (low j) cohorts, which is unappealing on a priori grounds.

18. The residual macro-economic effect is the average of the residuals for each year.

19. The predicted rate (p) is derived from the predicted log-odd (l) by the following equation: $p = \exp(l)/(1 + \exp(l))$.

20. All these results are presented in a longer version of the paper (Beaudry and Lemieux, 1998).

21. Technical details available upon request.

22. See Beaudry and Lemieux (1998) for detailed results.

23. The subsidy rate is the product of the replacement rate and the maximum number of weeks of eligibility for a worker who worked the minimum number of weeks required for eligibility, divided by the minimum number of weeks required for eligibility (qualification period). Since the weeks of eligibility and the length of the qualification period depend on the local unemployment rate, we use a national weighted average of the subsidy rates in each UI economic regions. The detailed regression results are reported in Beaudry and Lemieux (1998).

24. EI may have a different effect on the participation rate of different groups of women. For example, married women may react differently to changes in the EI subsidy rate than single women. The estimated effect in Table 3 may be thought as an average affect of the subsidy rate for the different subgroups of the population. Since the overall effects were not very encouraging, we have not explored how the results differed by subgroups.

25. This wage variable used is the logarithm of after-tax annual earnings of full-year full-time female workers.

26. The results from these regressions are available on request.

27. Since an entering cohort represents only a small fraction of the population, our results would not change very much if cohort effects were to increase again for the newer cohorts.

28. Tabulations from the SCF indicate that the 1992 cohort (women who turned 25 and 26 in 1992) represents 0.85 percent of the female adult population (15 and more). The four newer cohorts used in the simulations (age 17-18, 19-20, 21-22, and 23-34 in 1992) represent on average 0.83 percent of the female adult population, which is very similar to the 1992 cohort.

References

Beaudry, Paul, and David Green (1997) "Cohort Patterns in Canadian Earnings: Assessing the Role of Skill Premia in Inequality Trends," National Bureau of Economic Research Working Paper No. 6132, August.

Beaudry, Paul, and Thomas Lemieux (1998) "L'évolution du taux d'activité des femmes au Canada, 1976-1994: Une analyse de cohortes," CIRANO Working Paper No. 98s-06, March.

Card, David, and W. Craig Riddell (1993) "A Comparative Analysis of Unemployment in Canada and the United States," in D. Card et R. Freeman (eds.) *Small Differences that Matter: Labor Markets and Income Maintenance in Canada and the United States*, Chicago: University of Chicago Press for NBER, pp. 149-189.