Canadian International Labour Network

Labour Market Outcomes:

A Cross-National Study

CILN is a collaberative research venture between the Social Sciences and Humanities Research Council (SSHRC) and McMaster University. Additional funding is provided by the University of British Columbia, the University of Toronto, Queen's University, York University and Human Resources Development Canada (HRDC).

McMaster University

DEPARTMENT OF ECONOMICS

COHORT, YEAR AND AGE EFFECTS IN CANADIAN WAGE DATA

J.B. Burbidge, L. Magee, A.L. Robb

Department of Economics
McMaster University
Hamilton, Ontario
L8S 4M4

email: BURBIDGE@MCMASTER.CA

17 JULY 1997

COHORT, YEAR AND AGE EFFECTS IN CANADIAN WAGE DATA

ABSTRACT

We use Canadian SCFs 1971-1993 to study the wages of full-time, full-year male and female workers. Median real wages of 24-year-old males without a university degree fell by 25% between 1978 and 1993. For 24-year-old females the decline was more modest and reversed in 1987, but real wages in 1993 were still significantly lower than they were in 1978. We investigate whether these changes are permanent "cohort" effects or more temporary "year" effects. Graphs of median wages against year and age indicate some periods where year effects are more prominent than cohort effects and other periods where the reverse is true. We then compare the results from two models, one assigning the trends to year effects, the other assigning them to cohort effects, and use these models to produce real wage projections.

COHORT, YEAR AND AGE EFFECTS IN CANADIAN WAGE DATA[1]

1. Introduction

Any country's ability to make intergenerational transfers either to its young or to its elderly depends not only on the relative sizes of successive cohorts but also on the real wages of active workers. Given the age distribution of the population, higher real wages mean that any particular menu of programs for the young and the old can be financed with lower tax rates, and vice versa. Had real wages continued to grow at the same rate after 1977 as in the two decades before 1977 perhaps Canadian policy makers now would have been debating the merits of payroll tax decreases rather than increases. And clearly the country's ability to pay pensions to the baby-boom generations will depend on the real wages of the generations that have entered the labour market more recently. Even if the real wages of prime-age workers are sufficient to fund unchanged pension formulas for those who are retired, intergenerational equity may point to decreases in tax rates and benefits when the lifetime real wages experienced by older cohorts are significantly greater than the real wages of their successors.

In earlier work with Bar-Or, we employed data from the Canadian Survey of Consumer Finances to study the return to a university education over the period 1971 to 1991 (see Bar-Or et al. 1995). We concluded that, unlike the U.S., the university-high-school wage premium did not rise sharply during the 1980s. Only for males and females with less than six years of experience was there a strong upward trend over the 1980s and even with very large data sets the premium exhibits great instability late in the data period. In a subsequent paper, we used the same data to examine changes in wage inequality amongst full-time, full-year workers. Applying non-parametric statistical methods we found statistically significant and large increases in inequality for those with low levels of education and experience, combined with more modest changes in inequality for those with medium levels of education and experience, and with actual declines in inequality for older more experienced workers with a university degree (see Burbidge, Magee and Robb, 1997).

In this paper we show that real wages of full-time, full-year male workers in Canada have been falling since the late 1970s. These declines have been more pronounced at younger ages and lower levels of education. Real wages have oscillated for young full-time, full-year female workers, falling from the late 1970s until the late 1980s and then rising. One interpretation is that these changes are "year" effects; as the "bad" years fade into history, today's younger workers will experience abnormally large rises in wages with age. Thus their average real

---

[1]An early draft of this paper was presented at the Savings Workshop, Tilburg University, July 1995, and at the meetings of the Canadian Association on Gerontology, Quebec City, October 1996. We wish to thank participants in these workshops for helpful comments. We are also grateful to the Social Sciences and Humanities Research Council of Canada and the Seniors' Independence Research Program of Health Canada, which funds the program on the Independence and Economic Security of the Older Population at McMaster University, for generous financial support.

lifetime wages may be a bit lower than those of older cohorts, but not dramatically lower. An alternative view is that these changes are "cohort" effects which means that the entire lifetime path of real wages for a typical younger worker may lie below that for a typical worker from an earlier cohort. Beaudry and Green (1997), in a paper that closely parallels this paper, use Canadian SCF data to argue that the cohort effect interpretation is more plausible than the year effect interpretation.

Opinions on the relative importance of year and cohort effects differ sharply in the literature on earnings. For example, Juhn, Murphy and Pierce (1993) use U.S. data from the March Current Population Survey to argue that an increase in the demand for skilled workers is the primary factor causing an increase in the return to a university degree and increasing inequality in male wage distributions. Table 3 in their paper studies wage inequality in 6-year cohorts (years at which the cohort enters the labour market run from 1929-34,...,1983-88) for 1964, 1970, 1976, 1982 and 1988. They difference inequality measures within a cohort and find that the changes in inequality across successive cohorts are much smaller than the changes in inequality over time (pp. 425-426), concluding that year effects are more important than cohort effects. On the other hand, Gosling, Machin and Meghir (1995, GMM) use different statistical methods from Juhn, Murphy and Pierce to conclude there are very strong cohort effects and only weak year effects in earnings data drawn from 27 years of U.K. Family Expenditure Surveys, 1966-1992. GMM assign a primary role to age and cohort variables, by constructing the year variables to be orthogonal to the age and cohort variables. (In the language of vector autoregressions, cohort variables are placed "higher in the ordering" than year variables.) In addition, GMM assume year-age interactions are not present in the data. As a result, they do not find much evidence of year effects. Underlying the GMM approach is the view that year effects are merely temporary movements around some long-term trend; trend effects are incorporated into cohort variables.

It is well known that the identity linking age, year and cohort variables makes the identification of three separate effects extraordinarily difficult. Is it differences in the methods used to identify these effects that gives rise to the disparate results or are the U.S. and U.K. experiences very dissimilar? In this paper we analyse the same Canadian data set using various methods, so that any differences between our sets of results are due to differences in method, not the data. We begin by graphing the smoothed median wages data against year and age, and we go on to argue that these graphs may shed a different light on the identification of year and cohort effects from the more common parametric methods. Specifically, we think we can "see" periods where year effects are trending more than cohort effects and other periods where the reverse is true. It may be too restrictive to force either the cohort or year effects to be "trendless" over the entire 1971-93 period. However, allowing both effects to have trends would require a different method of identification. We do not have such a method to suggest, so instead we conduct a kind of sensitivity analysis where we compare the results from two models, one assigning the trends to year effects, the other assigning them to cohort effects. We then use these models to produce real wage projections for 2003, ten years beyond our last observations, to see what differences emerge between the "year" and "cohort" specifications.

The paper is structured as follows. In section 2, we describe the data we use to study the shifting patterns of wages by age, year and cohort. In section 3 we present and discuss various graphs of wages. Section 4 presents the two statistical models and compares and contrasts projections from these models. The final section summarizes the paper's main results and sketches the next step in this research program.

2. Data

The data used in this paper are taken from the Canadian Survey of Consumer Finances (SCF) and are described in detail in our earlier work. Here we provide only an outline of our extracts.

We employ seventeen years of Canadian SCF data that are publicly available as microdata sets for the calendar years 1971-1993. Biennial household surveys from 1971 to 1979 and subsequent releases of public use sample tapes based on census families were replaced in 1981 by annual surveys and releases of microdata files on individuals. This means that surveys prior to 1981 do not report wages and other characteristics of working children who live with their parents. To help minimize the differences between the household and individual surveys we restrict our attention to those aged 24 to 60 (see the discussion below). In addition, the SCF for wages during the 1983 calendar year (conducted in April and May of 1984) focused on assets and debts, so that wage data comparable to those of other years are unavailable for 1983. Our time series is thus biennial from 1971 to 1981 and annual from 1981 with the exception of 1983. Another major change in the nature of the data during the period was in the educational classifications. Below, we discuss two such changes, one of which occurred in 1975, the other in 1989.

Table 1 shows the cumulative effects of our exclusion restrictions in two steps. First we eliminate observations we are not interested in, namely those outside our age cuts (24 to 60) and those who were not working at all, so had no wage observation. Second, we eliminate observations for which the calculated wage variable may be inaccurate. In this case, we eliminate the self employed and those not working full time, full year. The self-employed[2] are eliminated because their income comprises wages and profits. We selected only full-time full-year workers in order to get an accurate measure of wages. A key reason for the full-time restriction is that hours of work information are not available for the previous year, to which the wages data apply. The full-year restriction is primarily to select a homogeneous group of workers, so that we are not mixing part-time and full-time workers. Also, we have some concern about the accuracy of reported weeks worked, and, in any event, one would probably want to study full-time and part-time workers separately. For full-time workers, we selected only those

[2]To eliminate the self-employed, we dropped the occupations of farmers, fishermen, loggers, etc. and those who reported their major source of income was not wages and salaries.

working 52 weeks per year, though, given the almost total absence of workers reporting 50 or 51 weeks, it would have made no difference had we selected on the basis of 50 or more weeks.

Considering first the males, the first two columns of Table 1 show the cumulative percentage of the sample (using the SCF universal weights) that is dropped due to the age and not working restrictions while the next two columns show the percentage of the remaining sample (after the first two restrictions are imposed) that is dropped due to the elimination of the self-employed and those who are not working full time, full year. The next four columns show the corresponding information for females. For males, the age-restriction percentage jumps from 26% in 1979 to 37% in 1981 and for females the corresponding numbers are 31% and 39%. These increases reflect the switch from the census-family data of the 1970s to the individual data for 1981 and subsequent years. As we observed above, working children living at home with their parents are excluded from the 1971-1979 data sets. Below, in section 4, we estimate models with separate year dummies for each year. That the year dummies don't exhibit abrupt changes between 1979 and 1981 is some evidence that our selection criteria are effective in smoothing the change in data sets between 1979 and 1981. Apparently there were few children aged 24 or older living with their parents in this part of the data period. Excluding those who are not working eliminates about another 6% of males and about another 40% of females in the early 1970s, and another 12% of males and another 20% of females by the end of the period. Females labour force participation rates have risen while male participation rates have fallen. Of those remaining, dropping the self-employed usually eliminates another 10% of both males and females. There appears to be some tendency in the data for males for self-employment to rise during and shortly after recessions. There is also only a weak upward trend in the "drop self emp." column which may reflect the offsetting trends out of farming and fishing and into self-employment in contracting and services. Finally, eliminating those employed less than full-time, full year, drops another 20% of males and another 31% of females. Note that the effect of this restriction jumps upward between 1973 and 1975. After 1973 "full-time" means the person typically worked 30 or more hours per week. For the 1971 and 1973 surveys it was left to the household to determine what "full time" meant. It would appear that some households used the term to describe a typical work week of less than 30 hours. Finally, Table 1 shows that, for both males and females, the cumulative effects of our exclusions leave us with about 65% of working males of the "right" age and about 50 to 55% of working females (again of the appropriate ages).

Annual wages and salaries from employment for the selected workers is divided by weeks worked (52) to form the weekly wage. The wages are adjusted by the Canadian CPI to the 1993 base year. Table 2 for males and Table 3 for females report the mean and three quantiles of weekly wages in 1993 dollars, for all education levels grouped together and then for university graduates and all others (non-university), as well as mean age.[3] Both tables use the universal weights supplied by Statistics Canada. Even though average age is remarkably constant over the

---

[3]The reason for this grouping of education categories is explored below.

data period, Table 2 shows that real weekly wages exhibit a hump-shape.[4] The median weekly wage for male non-university graduates grew at 2.3% per annum between 1971 and 1977 but then fell at an average rate of .5% per annum between 1977 and 1993. For male university graduates wages grew at .8% per annum between 1971 and 1979 and fell at .3% per annum thereafter. For females wages trend upwards over this period but the average growth rate was higher during the 1970s than subsequently.

One of the themes explored in this paper is that changes in wages over this period vary importantly with education. Accordingly, we would have liked to have various education categories that are consistent over time. For reasons discussed in earlier papers (see especially Bar-Or et al., 1995), however, one cannot overcome all discontinuities in the data. The difficulties of disaggregating the non-university category further (for example, into high school and other) led us to leave the non-university (NONU) group as one broad aggregate and university (UNIV) as the other. Nevertheless, some definitional discontinuities remain. Inspection of columns 4 and 5 in these tables reveals that there are at least two major breaks: (a) 1973 to 1975, which saw a large increase in the number of individuals reporting university degree (UNIV); (b) 1988 to 1989 which displays decreases in the UNIV category following several years of increases. Starting in 1975 the education question did not permit one to separate those who had taken some courses at university from university graduates and, as a consequence, the fraction indicating "university graduate" for 1975 jumps in both Tables 2 and 3. The education question was changed in the other direction in 1989 and thus the UNIV fraction drops in this year. In the subsequent analysis we keep track of the 1975 and 1989 breaks in educational classification. Our reading of events is that major switches in trends do not occur in either 1975 or 1989, and the changes we observe in wages are thus not the result of changes in definitions.

3. Eye ID

While the problem of identifying cohort, year and age effects is well known to researchers working in this area, some recapitulation of the issue and approaches typically adopted will clarify subsequent discussion. Suppose log wages (W) of an individual are influenced by events associated with the cohort birth year (C), the particular year in which wages are observed (Y), the age in the observation year (A), and by other variables (X). We can write W = f(C,Y,A,X). The identification problem arises because we do not have very good models for the influences associated with each of A, Y and C. A natural approach in such a case is to model these as fixed effects (that is, by using dummy variables) or perhaps trend effects treating age, year and birth year themselves as variables. The identification problem then arises because A, Y and C are

[4]It seemed odd to us that the average age was so constant over the period given that the period covers the entry of the baby boom into the labour market. In checking, though, we found a similar constancy of average age among the census population for the same age cuts.

linked by the identity A = Y - C. If we use the age, year and cohort variables themselves as regressors, we can use this identity to eliminate any one of C, A, or Y and we cannot estimate independent effects for all three. The problem is more complex if one admits the possibility of interactions, say, between cohort and year effects. However, the basic identification issue remains.

One way to resolve the problem is to model the effect of one or more of the independent influences (cohort, year or age) as functions of other variables rather than simply including them as fixed effects or trends. For example, cohort effects might be modeled as a function of cohort size, or year effects might be assumed to be a function of the real interest rate, the growth rate of real output; or the unemployment rate. Either route would "solve" the identification problem but one would need to be convinced the model was a reasonable one. We have experimented with modeling the year effects along these lines but have not found the results very satisfying.

A second approach to solving the identification problem is to experiment with different functional forms for the three effects. Thus, some researchers have employed a step function for cohort effects by creating n-year age cohorts (typically 5 or 10-year), while retaining, for example, year fixed effects and a polynomial in age. Without a clear theoretical foundation, however, this approach appears rather arbitrary at best. It may also be quite misleading since different n-year groupings may yield quite different results. A variation on this theme is to use a step function (or grouping) for cohorts but use some prior knowledge to group the years. For example, those entering the labour market during war years, or during depression years, might be assumed to constitute different cohorts. Again, it is hard not to suspect that such groupings are likely to be arbitrary.

We think there is an alternative way to separate out the three effects that is worth exploring. Examining three-dimensional graphs, along with more conventional two-dimensional ones, may help to disentangle the relative importance of cohort, year and age effects in wage data. The approach we suggest here is as follows. We place Years along the X-axis, Age along the Y-axis and the variable of primary interest (e.g., weekly wages) along the Z-axis. In such three-dimensional graphs, "pure" (or additive) year effects that dominate the other effects should stand out as valleys or ridges lined up with some particular year or group of years (wages should be affected the same way at all ages). Likewise, pure age effects that dominate should be valleys or ridges lined up with some particular age or group of ages. Finally, pure cohort effects that dominate should appear as valleys or ridges running from southwest to northeast with a slope of unity in the XY-plane. Obviously if a year effect interacts with an age effect in just the "right" way it could look like a cohort effect. This is another way of thinking of the identification problem discussed above. Clearly any data set admits more than one interpretation.

To fix ideas about pure cohort effects we graph estimates of the Canadian population,[5]

---

[5]We obtained these numbers from thirty-seven CANSIM matrices, various numbers, beginning with C892355 and ending with C892487.

aged 24 to 60, for the years 1971 to 1993 in Figure 1. Publication Quality Graphics (PQG) in GAUSS allows one to set the dimensions of the rectangular block containing the surface (we used 1 by 1 by 1) and the vantage point from which one views the surface. The graph in Panel A is viewed from (-.5,-.5,2) and the graph in Panel B is a contour plot showing some of the elevations in panel A. As noted above, for data sets with strong cohort effects, such as the population data shown in Figure 1, the ridges and valleys run southwest to northeast, are parallel and tend to have slopes of unity in the XY-plane.

The idea behind eye-id should be clear from this Figure. It would be hard to imagine that a combination of age and year effects could combine to generate these pictures. It is the regularity of the ridge lines with unity slopes that leads one to the conclusion that these must be cohort effects. Alternate views of the population data are provided in Panels C, D and E of Figure 1. Panel C shows population by age for various years, panel D shows population by year for various ages, and panel E population by age for various cohorts. Panels C and D will be more familiar to the reader as they are different ways of looking at the cross-sections. Panel E is a 45 degree cut through the data shown in Panel A. The 1930 birth cohort, for example would have been 41 in 1971 and the graph in Panel E represents the 45 degree cut through the surface in Panel A from that starting point (1971, age 41). The dashed line labelled 1930 is slightly downward sloped showing a declining 1930 birth cohort with age (presumably due to death by age). The 1945 cohort we observe only until age 48 and the line is pretty well horizontal. The 1960 cohort, still only in their early thirties at the end of the data period, shows marked increases which must be due to immigration.[6]

Panel C, if translated 90 degrees counter clockwise, would be similar to an age pyramid. It shows a growing population. Only in 1990 do we begin to see the clear age bulge associated with the baby boom of the 1950's. The baby boom is seen as well in Panel D for the 24 year olds whose size peaks in 1987 (birth year 1963).

Turning now to the wages data, we create for each sex, for each year (1971 to 1993), for each age (24 to 60), and for each education sub-group (NONU and UNIV), estimates of the weighted median (using the SCF universal weights) and an estimate of its variance. Three-dimensional graphs of raw quantile wages (even at the medians) are quite noisy. To make patterns easier to discern we smoothed the median estimates using local regressions; the details are in the appendix. Most of the surfaces have a smooth appearance, except for those areas with many small-sized cells. In order to avoid biasing the procedure in favour of showing short-term year or age effects, the amount of smoothing is restricted to be the same in both directions. In effect, the smoothed median for any year/age pair places greatest weight on the raw median for that year/age pair but it places some weight also on median wages up to four years, or four ages,

---

[6]Beaudry and Green find their main results are insensitive to immigration status but as noted above their sample is quite different from ours and we plan to look into this in further work.

away from it.[7]

The five panels (A through E) of Figure 2 comprise the graphs for all males while the five of Figure 3 comprise the corresponding graphs for those without a university degree (NONU) and the five of Figure 4 are for those with a university degree. The corresponding graphs for females are contained in Figures 5, 6 and 7. In each case, panel A graphs the smoothed median, panel B plots the associated contour lines, panel C shows median weekly wages by age for various years, panel D gives wages by year for various ages, and panel E shows wages by age for various cohorts.

Inspection of panel A (and panel D; see below) of Figure 2 reveals that, for all males, real wages generally rose during the early 1970s at all ages but have been declining since about 1978 for younger males. The contours in panel B for the 1980s and 1990s are not far off the 45 degree slope associated with cohort effects for those of middle ages. The almost vertical lines in the 70s from about age 30 to 55 suggest pronounced year effects in that period. While we recognize that one can never be sure how to decompose wage changes into age, year and cohort effects, it appears that year effects may have dominated cohort effects during the 1970s and that this was reversed somewhat during the 1980s. While the three-dimensional graph in panel A and the contour plot in panel B provide useful overviews of the data they are somewhat unconventional and may obscure important details of what has happened. Accordingly, we supplement these with panels C through E. C uses a slice of the three-dimensional surface shown in panels A and B to draw cross-sectional age-wages profiles for 1971, 1976 and 1990. In Figure 2C, the 1976 profile lies above that for 1971 at all ages, which points to the generally strong real wages growth rates for this period. While the 1990 profile is higher than the other two after age 38 it lies below both of them at low ages.

Figure 2A illustrates the increase in real wages during the 1970s and the subsequent decline, particularly at younger ages, since the late 1970s, but the magnitudes of these changes are difficult to discern. Figure 2D studies wages by year at four ages: 24, 33, 42 and 50. For full-time full-year male workers aged 24, real weekly wages declined by over 20% between 1978 and 1993, and the pace of the decline accelerated in the 1990s. Even for somewhat older workers, there was some reduction in real wages over this period - 12% for those aged 33 and 3% for those aged 42. Only the line for those aged 50 trends upward between 1978 and 1993.

---

[7]Another way to smooth the quantiles would be to combine the individual observations from adjacent cells and form quantiles, for example using a two-dimensional extension of the kernel-smoothing algorithm in Magee, Burbidge and Robb (1991). An advantage of this method is that it avoids possible distortions resulting from inappropriate weighting that might occur using the above method, especially for thinly populated cells. Disadvantages of such a method include greater computational requirements and the possibility that straightforward kernel-smoothing could lead to larger “edge effects”, or bias at and near the boundary of the (A,Y) surface.

One can also use the surface depicted in Figure 2A to follow particular cohorts by slicing along a 45 degree line in the Years-Age plane. We graph wages by age for the 1930, 1945 and 1960 cohorts in Figure 2E. Clearly, with observations over only a twenty-three-year period, we cannot compare completed wages histories of different cohorts. We can compute, however, the average ratio of real wages for overlapping age ranges of particular cohorts to give some sense of how different lifetime real wages might be. For example, we can observe median real wages of both the 1930 and 1945 cohorts between the ages of 41 (= 1971-1930) and 48 (= 1993-1945) and we can compute the ratio of the sum of wages for the two cohorts for ages between 41 and 48. On this basis, real wages of the 1945 cohort were 10% higher than those of the 1930 cohort and, over the age range 26 to 33, they were 8% higher than the real wages of the 1960 cohort. For males without a university degree, those born just after World War II have done better, to this point at least, than those born in the 1960s.

The five panels of Figure 3 display the information for those without a university degree. Not surprisingly, since they make up the largest fraction of the overall males group, the results are fairly similar to those for all males. Again, there is the suggestion of year effects in the early 1970s but little since that time, at least for older workers.

The five panels in Figure 4 (males, with a university degree) provide a striking contrast to Figure 3 and underline the importance of controlling for education when studying wage data. We can see real wage growth in the 1970s at some ages but declines at other ages; indeed the overall trend at many ages is downwards for the entire period. This is perhaps clearer in panel D than panel A. At age 24, real wages peaked in 1978 and then fell by about 15% to 1993; real wages at age 42 decline continuously throughout the data period by about 14%. Since real wages fell even more quickly at age 24 for those without a university degree, the return to a university degree has increased somewhat at younger ages (see Bar-Or et al., 1995). As is well known, age-wage profiles are much steeper for those with a university degree; contrast the C panels of Figures 3 and 4 (note the different scales on the vertical axes). The erratic behaviour of the 1971 line in Panel 3C reflects the small number of older university-educated males in that year's sample. Year effects appear to be less important and cohort effects more important in Panel B of Figure 4 than in the corresponding panel in Figure 3. Panel E of Figure 3 shows that younger cohorts have lower real wages.

For full-time, full-year females, real wages at many ages trend upwards over the data period. This can be seen in various ways from all panels of Figure 5. Only at the youngest ages do we see no upward trend over the data period. For example, at age 24, real wages look a bit like a sine curve; they rose from 1971 to 1978, then fell to 1986 and then rose to 1993, ending up in 1993 at approximately the level they started in 1971 (see Panel D). In contrast, real wages of 50 year olds rose by about 35% over the data period. These trends imply large estimates of real-wage divergence between certain cohorts. For example, the numbers underlying panel E can be used to calculate that, for ages 41 to 48, real wages of the 1945 cohort were 21% larger than for those of the 1930 cohort. As in the case of males, there appear to have been large year effects in the 1970s (note the almost vertical lines from about age 28 on in that period).

The charts for full-time, full-year women without a university degree (Figure 6) are not very different from those for all females, as was the case for males. For women with a university degree the patterns are different. It is less easy to discern what is going on because many cells are thinly populated, particularly in the 1970s. In fact, we are able to estimate the model only for ages 24 to 50. Generally speaking, real wage growth for this group has been more modest than for the non-university group. Panel D of Figure 7 highlights the contrast across ages; at age 24 real wages fell by about 10% between 1971 and 1993, whereas, for 50 year olds, wages rose by 60% over the same period. Predictably, these trends imply differences in what we observe across cohorts. For example, the real wages of the 1945 cohort are only 11% higher than those of the 1930 cohort and the 1930 cohort's wages are 4% higher than those of the 1960 cohort.

To this point we have allowed the data to speak with little encumbrance. We now explore some models with more structure.

4. Predicting Cross-Sections using Cohort and Year Models

In section 3, we referred to the difficulty in identifying age, year and cohort effects. This identification is not necessary for descriptive exercises such as the one in that section. However, it becomes necessary if one wishes to make out-of-sample predictions, as we do in this section. Consider an additive model that describes the mean or median log wages of a person aged *a* in year *t* as

$$W_{at} = A_a + Y_t + C_{t-a},$$

where W_{at} is the mean or median log wage of persons aged *a* in year *t*, A_a is the age effect at age *a*, assumed to be constant over time, Y_t is the year effect at year *t*, assumed to be constant across age, and C_{t-a} is the cohort effect for those born in year *t-a*, assumed to be constant as the cohort ages over time.

Predicting W_{at} at a time period $t = t'$ beyond t_{max}, the end of the sample period, requires estimating A_a, $Y_{t'}$ and $C_{t'-a}$. An important consideration is how to identify these three effects. Typically this involves the ability to estimate regression coefficients that describe the three effects as functions of polynomials in age, year, or birth year, or dummy variables and other variables. Some restrictions are necessary to achieve this.

Several authors identify the separate effects by forcing the year effects to display no trends. Deaton and Paxson (1994) and Baker and Benjamin (1995) restrict the year effects to sum to zero and be orthogonal to a time trend variable, in effect assigning any trend effects to C_{t-a} rather than Y_t. We will use "trend effects" or "trends" to refer to any steady increase or decrease in W_{at} over time, during part or all of the time period. Gosling, Machin and Meghir (1995) identify the separate effects by defining the year effects to be orthogonal to the age and cohort effects. Again, this assigns any trends to the cohort effect, leaving only detrended year effects.

All of these authors recognize the basic identification issue and the unavoidable arbitrariness of their identification method.

This identification issue has important consequences for out-of-sample prediction. For example, suppose there is a negative trend effect late in the sample, which could be assigned as a cohort effect, a year effect, or some combination. Further, suppose that the predictions are based on the assumption that future cohort and year effects will be similar to the last few estimated cohort and year effects. If the trend is treated as a cohort effect, some of the $W_{at'}$ predictions will decrease over time as the low-wage cohorts grow older. This does not happen, though, when the trends are assigned to year effects, unless the downward trend in Y_t is assumed to continue out-of-sample. For example, Baker and Benjamin's (1995) pessimistic conclusions about reduced savings and reduced work in recent Canadian cohorts are in part due to their having assigned all trends to cohort effects.

There is a second issue in predicting out-of-sample $W_{at'}$'s. Even if the three separate effects are identified, meaning that regression coefficients associated with the three effects are identified, it still is necessary to predict year, cohort and age effects that have not yet been observed. It seems natural to assume the age effects will be the same as in the sample, but the cohort and year effects are more difficult to predict. If Y_t and $C_{t\text{-}a}$ are estimated by polynomials in year and birth year respectively, for example, then out-of-sample predictions can be obtained by extrapolation. However, researchers are unlikely to place much faith in an extrapolation of these trends. If Y_t and $C_{t\text{-}a}$ are modeled using year and cohort-specific dummies, it is even more difficult to extrapolate convincingly.

For our predictions, we need to deal with these two issues. We examine the identification issue by estimating two models. The dependent variable is weekly wages. Both models have as regressors a quartic polynomial in age. The "cohort model" also includes a set of cohort dummies, while the "year model" includes a set of year dummies. The cohort model could have included orthogonalized year dummies, but when this was considered, it did not change the predictions much. The estimated direct effect of the orthogonalized variables on the predictions is small, because by construction these variables have no trend. Because of the orthogonalization, they are not very correlated with the other variables, so those other coefficients are not affected much either. To facilitate exposition, then, we consider only the simple cohort and year models described above.

There is necessarily some arbitrariness in handling the second issue, that is, predicting the future year and cohort effects. In the year model, we predict future year effects by setting them equal to the most recent estimated year effect in our sample: 1993. In the cohort model, we choose not to do the analogous thing. The most recent cohort effect, for the 1969 birth cohort, is not estimated with much precision since it is entirely determined by the 24-year-olds in the 1993 sample. Instead, we predict future cohort effects by setting them equal to a weighted average of the four most recent estimated cohort effects in our sample, with weights .4, .3, .2 and .1 assigned to the 1969, 1968, 1967, and 1966 cohorts respectively.

We wish to estimate conditional medians using a large number of observations. One approach (Gosling, Machin and Meghir (1995)) is to estimate the median for each age/year cell, then use these along with some cell weights to estimate median regression functions by a minimum chi-square or GLS procedure. MaCurdy and Mroz (1995) also use this method as part of a more sophisticated approach. This technique requires an adequate number of observations in each cell for the results to be reliable, which would be a concern here. Instead, we conduct median regressions using an iterated weighted least squares algorithm suggested by Fair (1974) and Schlossmacher (1973), adjusted for the sampling weights. The algorithm iterates the following two calculations until convergence:

$$\hat{\beta}_m = (X^T W^{-1} X)^{-1} X^T W^{-1} y, \text{ where } W = \mathrm{diag}(\max(.00001, |e_{m,i}|)/w_i),$$

$$e_{m+1} = y - X\hat{\beta}_m.$$

$e_{m,i}$ is the i^{th} element of the residual vector e_m, and w_i is the sampling weight. The "max" operation ensures that there is no numerical problem when inverting W. Upon convergence, where $e_m = e_{m+1}$, apart from this small "max" adjustment, β minimizes $\sum[e_{m,i}^2/(|e_{m,i}|/w_i)] = \sum w_i|e_{m,i}|$, which is a weighted LAD estimator of the population conditional median.

Figure 8 shows four predicted 2003 cross-sections resulting from applying two models, the year model and the cohort model, to two subsamples, males with and without a university degree (UNIV and NONU, respectively). The cohort model predicts a substantially lower profile for UNIV than does the year model. For NONU, the two models' predictions for 2003 cross around age 50. The cohort effects themselves are discussed in more detail below. Note that the cohort model predicts a much smaller return to a university education for older workers in 2003 than the year model does.

Figure 9 shows the same plots for females. For both the UNIV and NONU groups, the year and cohort models give similar 2003 predictions.

Although education differences are often of interest, one could argue that the changing nature of education over such a long period of cohorts can make the results misleading, and that it would be better to look at all education groups together. Figures 10 and 11 show the 2003 predictions for all males and females. To get a sense of the sensitivity of total predicted income (and tax base) to the model, we used predicted population counts by age and gender for 2003[8] to form a weighted sum of median wages. For males, the year model predicted 3.28% higher median wages than the cohort model. For females, the year model predicted 5.15% lower median wages than the cohort model. The models used to produce these estimates are very crude and do not employ modern macroeconomic forecasting techniques. We do not wish to argue that these estimates should be taken seriously but it is clear that year and cohort models are capable of

---

[8]We are very grateful to Frank Denton, Chris Feaver and Byron Spencer for providing us with these projections from their MEDS model.

generating different predictions. Policy makers reshaping programs such as the Canada Pension Plan may wish to consider scenarios in which the real wages tax base is quite different from what it is today.

Figures 12 and 13 show the year effects estimated from the year model for males and females. Male wages clearly peaked in the 1970s, with the downward trend since then being larger for UNIV than NONU in absolute size. For females, UNIV shows a similar pattern as males, with a less pronounced downward trend since the late 1970s. NONU females, however, shows an increasing trend throughout the sample.

Figures 14 through 19 show cohort effects, expressed as indexes set at 100 for birth year 1945. The plot labelled “regression” is based on the coefficients from the cohort model. The “direct” and “chain” plots are based on the following procedure. Both use the smoothed weekly wages shown in the A panels of Figures 2 through 7. The 1927 value of the “direct” index, for example, is 100 times the sum of estimated median wages of people aged 44-48 who were born in 1927 (i.e. their wages during the years 1971-1975) divided by the sum of estimated median wages of people aged 44-48 who were born in 1945 (i.e. their wages during the years 1989-1993). The index is calculated only for those cohorts where the age overlap is at least five years; thus the index can be calculated for cohorts born between 1927 and 1963. The “chain” index is based on the overlap (again, at least 5 years) between one cohort and the next. For example, the 1915 cohort overlaps the 1916 cohort at ages 56 to 60 and one can calculate the ratio of the sum of median wages, 1915 over 1916. Likewise, for ages 55 to 60, one can obtain the ratio for 1916 to 1917; and so on. Multiplying the ratios for 1915 to 1944 together, and then multiplying the result by 100, one obtains an index of the real wages of the 1915 cohort on a scale where the 1945 cohort equals 100. This “chain” index can be calculated for the 1915-1964 cohorts, for all groups except females with a university degree where the range is 1925-1964 (here thin data sets forced us to estimate smoothed median wages for ages 24 to 50, not 60).

Both all males and NONU males show inverted-U shaped cohort effects, peaking for the 1940s birth cohort. UNIV males have a downward cohort effect during most of the sample. The sparse sample for early birth year cohorts results in some discrepancies across indexes and a very noisy plot for the regression index. There is a possible rebound for the most recent cohorts showing up in the regression index, but it could be noise. Female cohort effects show a regular upward trend, although the all females plots appear to stop increasing after about 1950. The NONU female plot coming from the cohort model regression continues to increase well beyond the 1940s cohorts, whereas the other indexes level off. The trends are not as clear for UNIV females - the sample is very sparse for early cohorts.

For all full-time, full-year males the 1960s cohorts have experienced real wages that are 10-20% lower than those of the 1945 cohort. Depending on the measure chosen, full-time full-year females born in the 1960s have about 10% higher real weekly wages than the 1945 cohort. These numbers may not appear to be significant to some observers but they are “large” by the standards used in the public finance literature. There, much attention has been focused on the dynamic efficiency gains of switching from income to consumption taxation, where the steady-state gains are estimated to be about 5% of lifetime wages (see Auerbach, Kotlikoff and Skinner (1983)). Many commentators have talked about the intergenerational equity of changing taxes and transfers but the differences across generations cohorts in real wages have not received as much attention as perhaps they should have.

5. Conclusions

The main objective of this paper has been to bring some facts about the changing structure of real wages in Canada to the attention of policy makers. Among the more striking results, median real wages of full-time, full-year males aged 24 without a university degree fell by 25% between 1978 and 1993. For the corresponding group of female workers, the decline was more modest and reversed in 1987, but real wages in 1993 were still significantly lower than they were in 1978. If these changes are modeled as cohort effects as opposed to year effects, the prospects for younger males are very bleak. For younger females, the opposite is true. As suggested in Beaudry and Green, future research should focus on attaining a better understanding of what is causing these changes.

At this point it is difficult to state precisely why different researchers reach different conclusions about the relative importance of year and cohort effects. The data sets employed are different and the empirical approaches are different. Here we simply speculate on two factors that may lead to a reconciliation of these disparate results.

First, the evidence in favour of cohort effects depends quite heavily on the recent experience of younger workers over the late 1980s and early 1990s. That JMP's data set ends in 1988 while GMM's data and ours extends through to 1992 may explain some of the differences in results. Dropping the last four years of our data would lead one to downplay the role of cohort effects in Figures 2-7. Secondly, we suspect empirical technique may be important. By construction, the JMP argument against the importance of cohort effects tends to place less weight on cohorts that have entered the labour market recently. And certainly, our interpretation of GMM's way of estimating cohort effects is that it is biased in favour of finding cohort effects to be more important than year effects.

Using a descriptive graphical approach with Canadian data, we observe, particularly for males, what appear to be strong trending year effects in the 1970s and trending cohort effects in the 1980s. There may be some periods in other countries when year effects dominate and other periods when cohort effects are important. It may be a mistake to force the data into one mould or the other.

Appendix

Let x_i be log wages and w_i be the sampling weight for observation I belonging to a particular age/year cell. The x_i's are in ascending order. We obtain the αth quantile by interpolation. Let $k(\alpha) = \max_j(1,\max\{j|W_j < \alpha W_n\})$, where W_j, $j = 1,..,n$ are the cumulative sums $W_j = \sum_{i=1}^{j} w_i$ and n is the total number of observations in the cell. And let $b_\alpha = (\alpha W_n - W_J)/(W_{J+1} - W_J)$, where $J = \max\{j|W_j < \alpha Wn\}$. Then the αth quantile $m^{(\alpha)}$ is estimated as

$$m^{(\alpha)} = (1-b_\alpha)x_{k(\alpha)} + b_\alpha x_{k(\alpha)+1}.$$

Ideally, these local regressions would use the inverses of the variances of the cell medians when constructing the weight matrix for the GLS regression given below. But some of the cell sample sizes are quite small, and we fear that the noise resulting from sampling error in the variance estimates could be damaging. MaCurdy and Mroz (1995) handle this by weighting by the inverse of the cell sample size, n_{ay}. This is appropriate, for example, if one is working with sample means and each underlying observation has the same variance, σ^2. Then each cell mean would have a variance σ^2/n_{ay}, justifying weighting by $1/n_{ay}$. In other cases, when the spread of the distribution differs across cells, this weighting is not precisely valid. Still, if some cell sample sizes are small, and the spread of the distribution is not thought to vary much across cells, it seems preferable to weight by $1/n_{ay}$ instead of trying to estimate the variance of each cell mean or, as in our case, median, and this is the approach we adopt.

There is a further complication, however, since we use sampling weights. Instead of $1/n_{ay}$, we weight by:

$$v_{ay} = \sum_{i=1}^{n_{ay}} w_i^2 / \left(\sum_{i=1}^{n_{ay}} w_i\right)^2 .$$

This is motivated by analogy with the variance of a sample mean of homoskedastic data being given by σ^2/n_{ay}. The variance of a weighted mean in this case is given by $\sigma^2 v_{ay}$.

These quantile estimates and weights are obtained for each age/year cell. For a particular quantile, let them be m_{ay} and v_{ay}. The following local regression procedure gives a smoothed quantile estimate for the age/year cell (A,Y). It uses a weighted GLS regression with data in the neighbourhood of (A,Y) to predict the quantile at (A,Y). Let $k_{ay} = \max(0,1-[(a-A)^2+(y-Y)^2]/H)$, where H is a smoothing parameter. Using some rule for stacking the cells, let X be a matrix with typical row (1,a,y), y be a vector with typical element m_{ay}, and V and K be diagonal matrices $V = \mathrm{diag}(v_{ay})$, $K = \mathrm{diag}(k_{ay})$. The smoothed quantile estimate is the prediction from a GLS regression:

$$[1,A,Y](X^TKV^{-1}X)^{-1} X^TKV^{-1}y.$$

This procedure is repeated for any desired (A,Y) and quantile. This smoothing allows for the prediction of quantiles for years when data are nonexistent (1972, 1974, 1976, 1978, 1980 and 1984) as well. We set H equal to 20.

This procedure can be viewed as a variation on GMM's, which was sketched in the introduction. Our first stage is similar to theirs - both obtain cell-specific quantile and variance estimates. GMM's second stage fits a surface defined by cubic polynomials in age and cohort, and orthogonalized year dummies. Our second stage fits a surface defined by local linear functions of age and year (or equivalently, age and cohort, or year and cohort). The local nature of our fit allows age, cohort, and year effects to show up on the surface in a less restricted way than if they had been required to be identified by prior parametric restrictions. GMM's more parametric approach allows for more convenient hypothesis testing for the significance of the various effects than does our approach. The validity of GMM's tests, however, relies on the validity of the parameterization. The orthogonalization of the year effect in GMM, which is necessary for identification, results in any long-term trend in the year effect being assigned instead to a cohort effect, which might contribute to the apparent significance of their cohort effect.

References

Auerbach, A.J., L.J. Kotlikoff and J. Skinner 1983 "The Efficiency Gains From Dynamic Tax Reform". International Economic Review, 24, pp. 81-100.

Baker, M. and D. Benjamin 1995 "The Economics of Aging in Canada: A Survey of Some Issues and Preliminary Evidence". Department of Economics, University of Toronto, unpublished.

Bar-Or, Y., J. Burbidge, L. Magee, A. L. Robb 1995 "The Wage Premium to a University Education in Canada: 1971-1991". Journal of Labor Economics, 13, pp. 762-794.

Beaudry, P. and D. Green 1997 "Cohort Patterns in Canadian Earnings: Assessing the Role of Skill Premia in Inequality Trends". Presented at the Canadian Economics Association Meetings in St. John's, Newfoundland, June.

Burbidge, J.B., L. Magee and A.L. Robb 1997 "Canadian Wage Inequality over the Last Two Decades". Empirical Economics, 22, pp. 181-203.

Deaton, A. and C. Paxson 1994 "Saving, Growth, and Aging in Taiwan". Studies in the Economics of Aging, D. Wise (ed.), University of Chicago Press, Chicago, pp. 331-357.

Fair, R.C. 1974 "On the Robust Estimation of Econometric Models". Annals of Economic and Social Measurement, 3/4, pp. 667-677.

Gosling, A., S. Machin and C. Meghir 1995 "The Changing Distribution of Male Wages in the UK". Institute for Fiscal Studies and Department of Economics, University College London.

Juhn, C., K.M. Murphy and B. Pierce 1993 "Wage Inequality and the Rise in Returns to Skill". Journal of Political Economy, 101, pp. 410-442.

MaCurdy, T. and T. Mroz 1995 "Measuring Macroeconomic Shifts in Wages from Cohort Specifications". Presented at 1995 Canadian Econometric Study Group, Montréal.

Magee, L., J.B. Burbidge and A.L. Robb 1991 "Computing Kernel-Smoothed Conditional Quantiles from Many Observations". Journal of the American Statistical Association, 86, pp. 673-677.

Schlossmacher, E.J. 1973 "An Iterative Technique for Absolute Deviations Curve Fitting". Journal of the American Statistical Association, 68, pp. 857-859.

Table 1: This table shows the cumulative percentage effects of our exclusions in two stages; age and dropping not working first and then dropping self-employed and those not working full-time, full-year next; for both the male and female data sets, by year.

Cumulative effects of selection criteria (%)

	MALES					FEMALES			
Year	Age rest'ns	Drop not work.	Drop self emp.	Full-time Full-year		Age rest'ns	Drop not work.	Drop self emp.	Full-time Full-year
1971	26	32	11	28		30	75	5	41
1973	27	32	12	26		31	72	6	42
1975	27	31	11	30		31	64	8	51
1977	26	30	12	26		31	61	9	49
1979	26	30	13	28		31	59	9	49
1981	37	46	7	25		39	62	9	47
1982	37	46	10	30		38	62	11	46
1984	36	46	10	30		38	61	11	46
1985	35	46	9	26		37	59	10	43
1986	35	45	9	26		37	58	10	45
1987	34	45	9	26		36	56	9	45
1988	34	44	8	25		36	55	10	43
1989	34	44	8	25		36	54	9	41
1990	34	45	8	28		36	54	10	42
1991	34	46	10	29		36	55	10	42
1992	34	46	11	30		35	55	10	42
1993	34	46	10	30		35	55	11	42

Table 2: This table shows summary statistics for extracts drawn from the Canadian SCFs: full-time, full-year males; 1971-1993.

			Per-cent		REAL WEEKLY EARNINGS (1993 dollars) All education levels				Not University				University			
Year	Num. obs.	Avg. Age	NU	UN	Mean	.25Q	.50Q	.75Q	Mean	.25Q	.50Q	.75Q	Mean	.25Q	.50Q	.75Q
1971	9474	40	89	11	742	527	676	865	698	513	660	824	1106	719	973	1336
1973	10378	40	89	11	780	557	716	918	737	543	697	871	1117	763	1032	1342
1975	9820	39	84	16	825	587	755	969	771	569	727	910	1115	767	1010	1339
1977	14126	40	85	15	831	601	782	997	791	587	758	953	1063	763	1022	1312
1979	13922	39	85	15	829	595	780	987	783	578	748	935	1086	770	1036	1294
1981	13897	39	83	17	815	590	764	993	764	565	731	921	1067	780	1023	1288
1982	12854	40	81	19	815	570	765	986	757	542	725	918	1069	772	1019	1272
1984	12434	39	81	19	810	576	773	987	753	548	730	924	1051	795	1018	1249
1985	13024	39	81	19	807	556	774	1005	752	526	731	938	1044	758	1016	1274
1986	11394	39	80	20	807	565	757	997	746	532	727	917	1059	771	1003	1254
1987	14869	39	79	21	817	558	769	996	756	533	721	937	1052	721	985	1246
1988	13380	39	80	20	823	567	769	1002	769	542	730	929	1043	693	981	1247
1989	15159	39	82	18	813	555	764	990	753	535	715	924	1079	749	996	1317
1990	15792	40	81	19	829	557	757	1018	775	534	726	944	1060	713	976	1259
1991	14285	40	80	20	808	547	755	994	740	517	695	914	1078	739	994	1290
1992	12921	40	78	22	822	554	764	1003	759	528	711	932	1052	724	998	1254
1993	12876	40	79	21	801	538	750	1000	738	510	693	923	1037	731	980	1261

Table 3: This table shows summary statistics for extracts drawn from the Canadian SCFs: full-time, full-year females; 1971-1993.

			Per-		REAL WEEKLY EARNINGS (1993 dollars)											
	Num.	Avg.	cent		All education levels				Not University				University			
Year	obs.	Age	NU	UN	Mean	.25Q	.50Q	.75Q	Mean	.25Q	.50Q	.75Q	Mean	.25Q	.50Q	.75Q
1971	2975	39	92	8	449	312	409	546	422	301	395	511	764	582	701	903
1973	3479	39	92	8	468	334	437	574	444	325	420	547	729	563	699	902
1975	3888	38	88	12	497	357	463	595	468	348	449	554	720	547	702	877
1977	5765	38	87	13	511	357	479	626	477	348	461	587	737	489	744	948
1979	6217	38	88	12	518	366	485	638	482	354	463	592	771	560	754	995
1981	7216	38	86	14	520	357	484	644	482	339	465	598	763	565	761	948
1982	7339	38	84	16	516	358	479	646	473	335	452	584	751	554	749	941
1984	7466	38	83	17	527	357	494	666	481	339	467	600	750	543	746	935
1985	8074	38	83	17	521	343	481	653	474	321	457	588	753	527	738	941
1986	6787	38	82	18	532	351	502	665	484	332	476	602	752	502	749	931
1987	9075	38	81	19	537	360	499	673	483	328	471	601	769	567	749	956
1988	8464	38	81	19	535	353	492	670	479	332	462	593	777	537	750	970
1989	9887	38	84	16	531	352	486	660	479	330	459	594	794	555	770	990
1990	10711	38	84	16	550	376	514	691	506	357	483	630	783	567	742	986
1991	9917	39	82	18	555	369	517	695	503	348	481	620	797	564	789	994
1992	9114	39	80	20	581	392	540	728	524	367	499	652	813	587	800	1013
1993	9068	39	80	20	569	385	520	708	511	363	481	635	799	560	770	1000

Figure 1: Canadian Population by Age(24-60) and Year (1971 - 1993)

Panel A
View from (-0.5,-0.5,2.0)

Panel B
Contour Plot

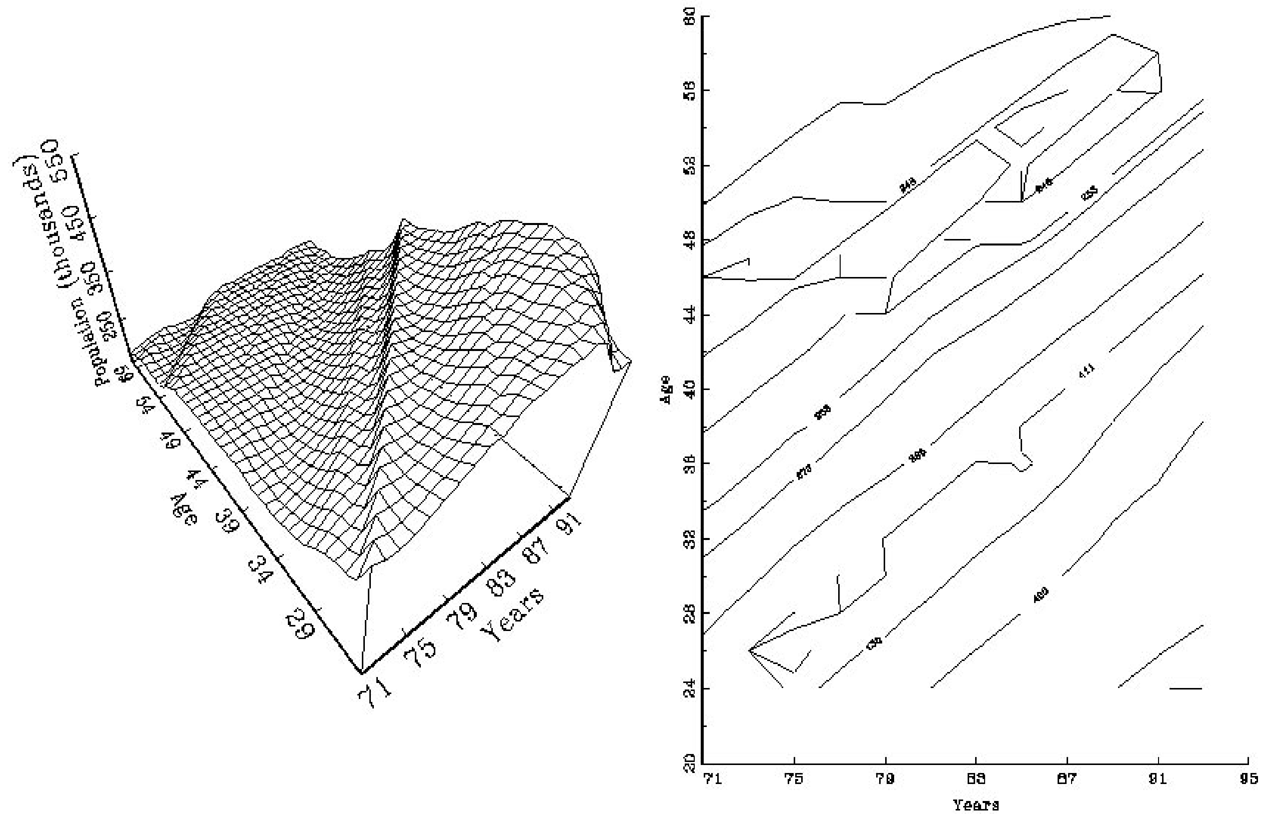

Figure 1: Canadian Population by Age (24-60) and Year (1971 - 1993)

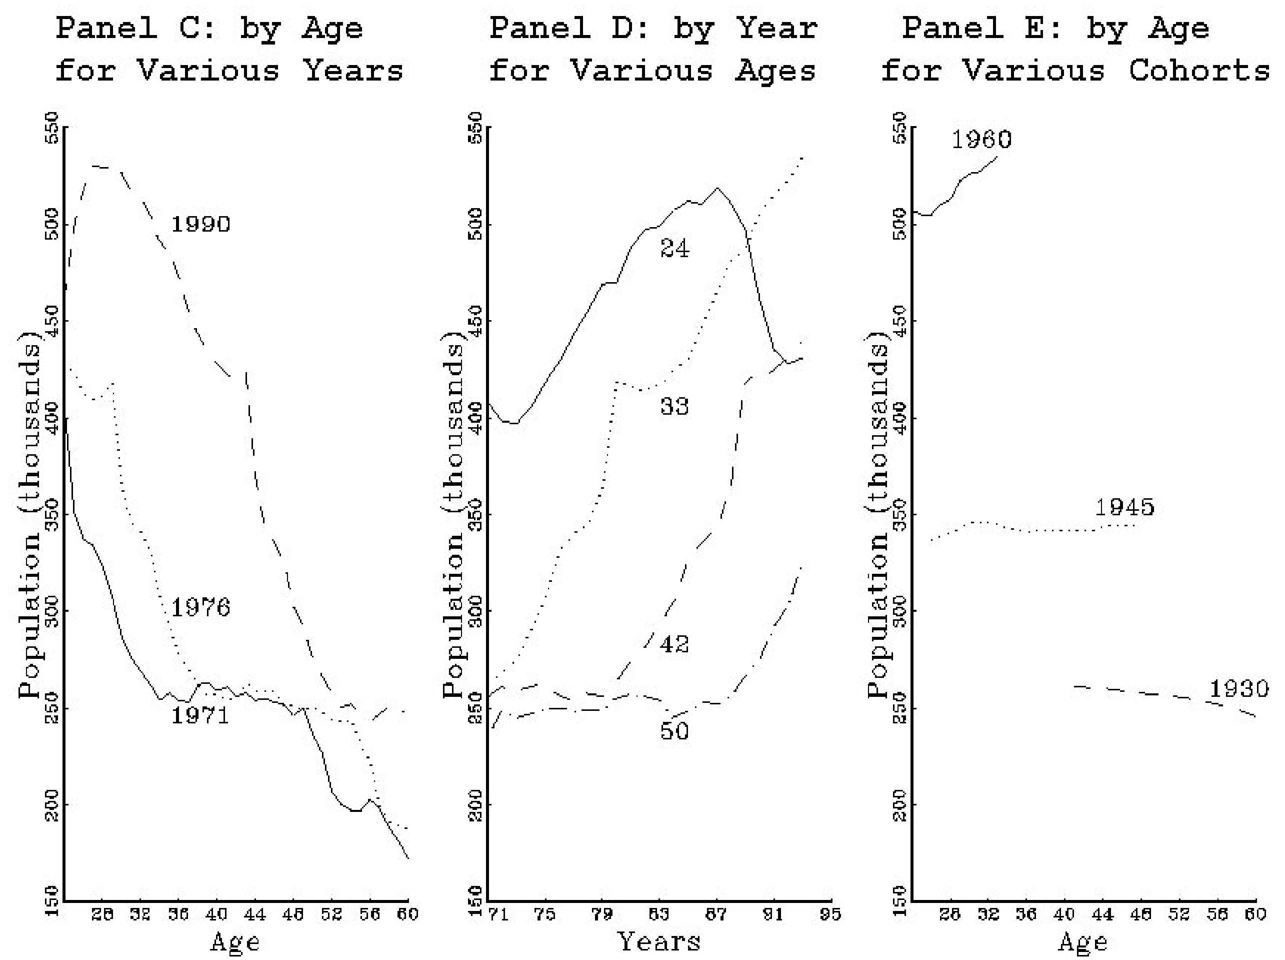

Figure 2: Median Weekly Wages:
All Males Aged 24-60

Panel A View from (-0.5,-0.5,2.0)

Panel B Contour Plot

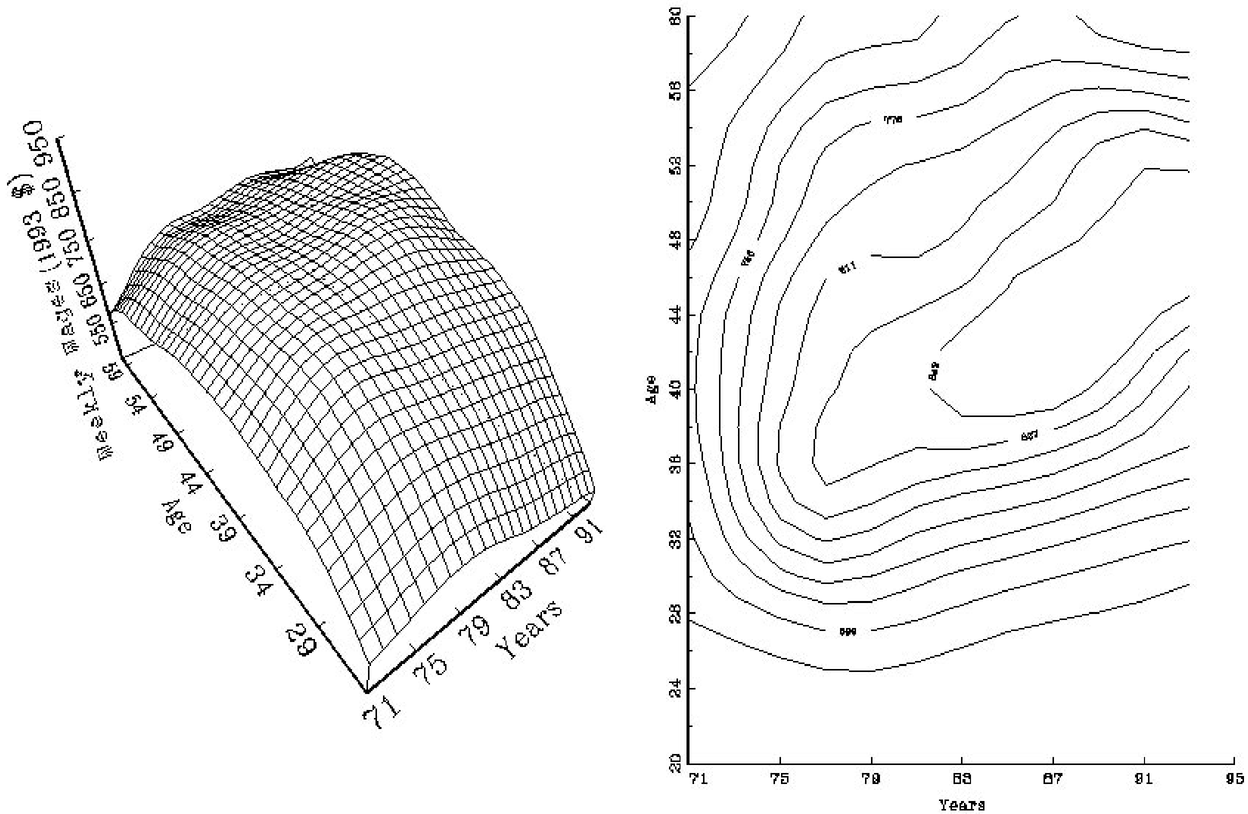

Figure 2: Median Weekly Wages:
All Males Aged 24-60

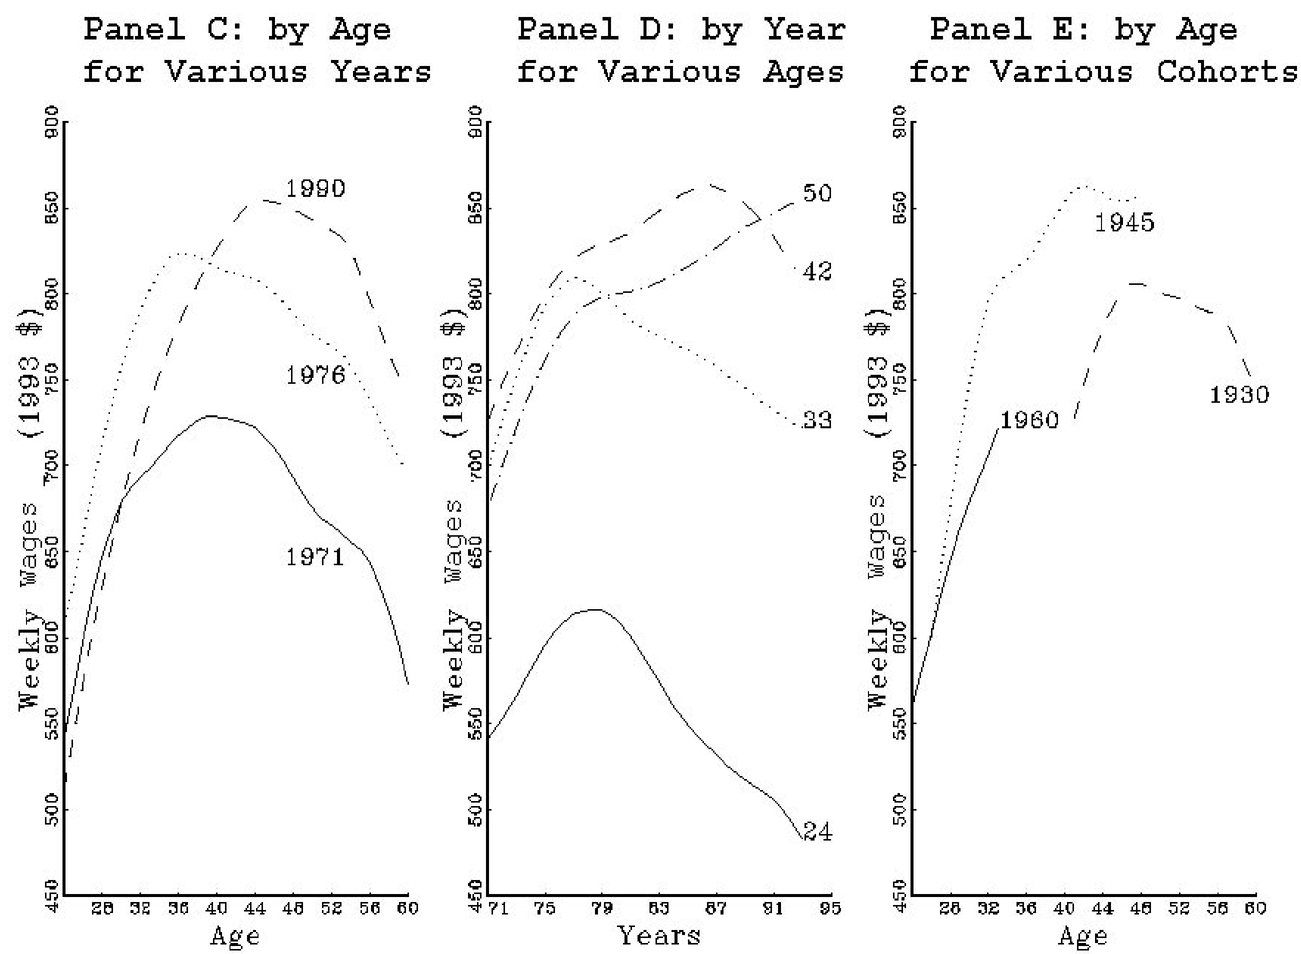

Figure 3: Median Weekly Wages:
Males Without a University Degree

Panel A
View from (-0.5,-0.5,2.0)

Panel B
Contour Plot

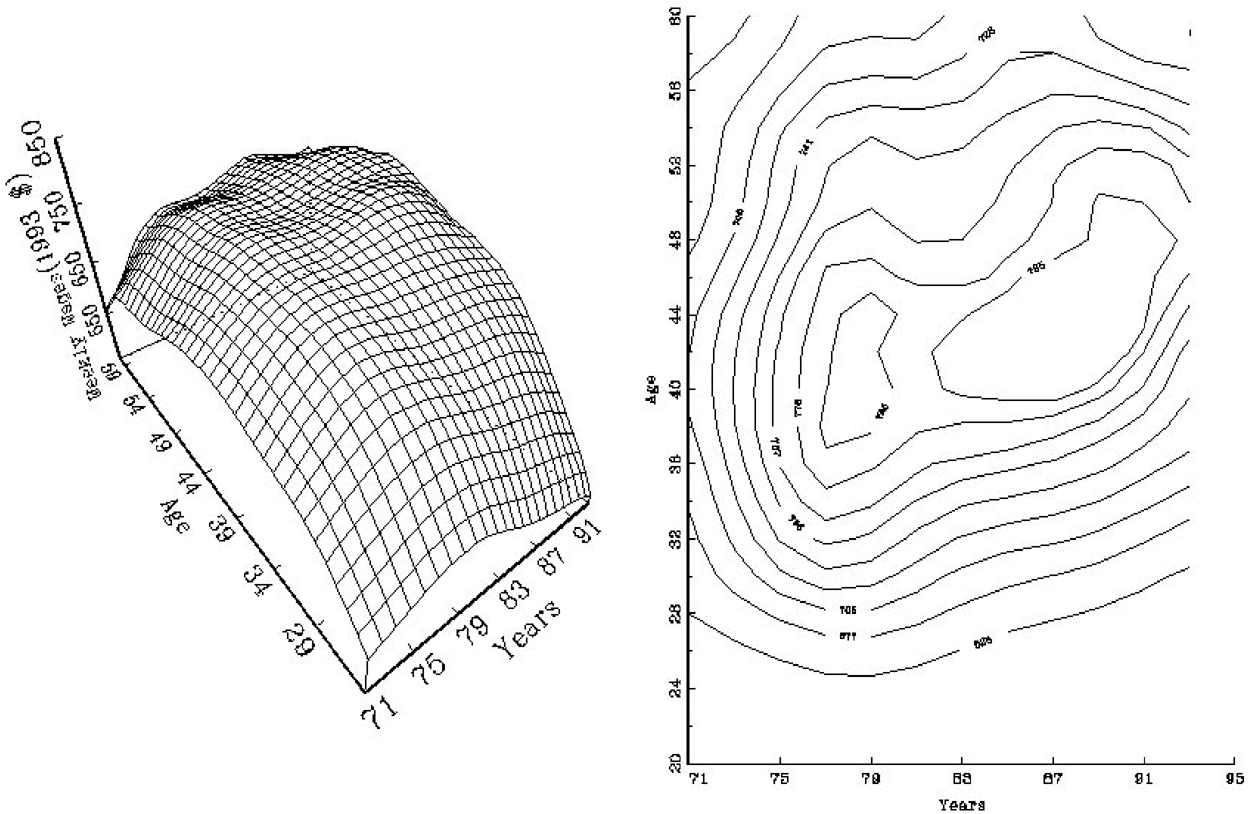

Figure 3: Median Weekly Wages: Males Without a University Degree

Panel C: by Age for Various Years | Panel D: by Year for Various Ages | Panel E: by Age for Various Cohorts

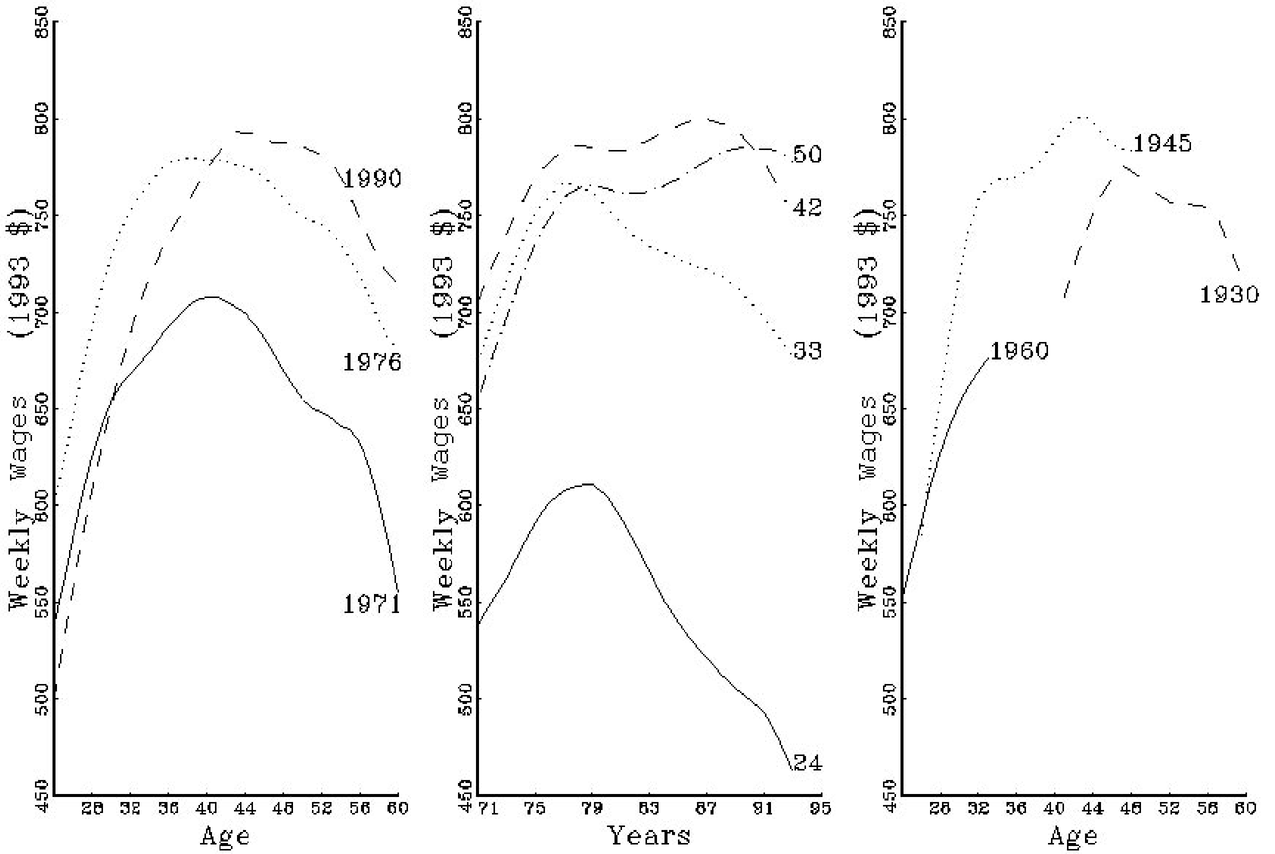

Figure 4: Median Weekly Wages: Males With a University Degree

Panel A
View from (-0.5,-0.5,2.0)

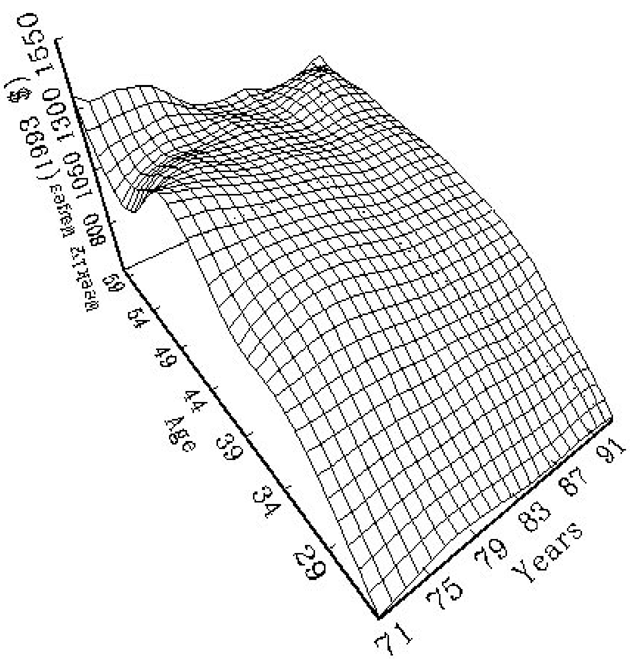

Panel B
Contour Plot

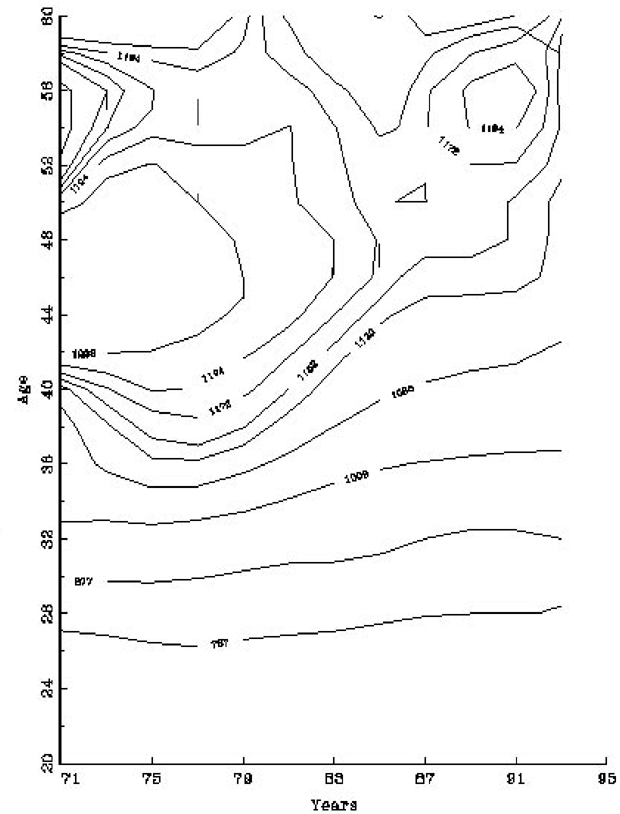

Figure 4: Median Weekly Wages: Males With a University Degree

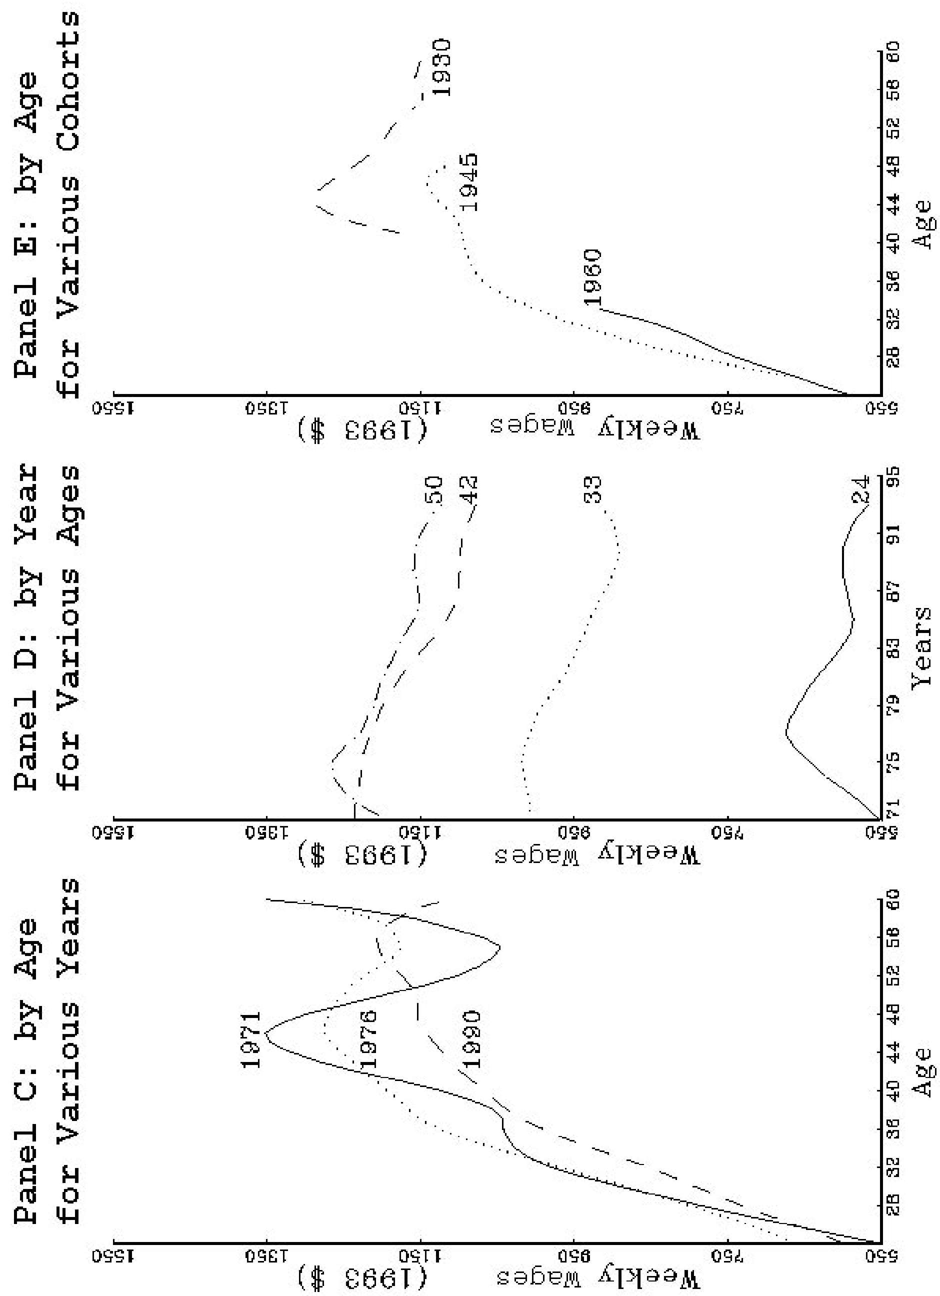

Figure 5: Median Weekly Wages: All Females Aged 24-60

Panel A
View from (-0.5,-0.5,2.0)

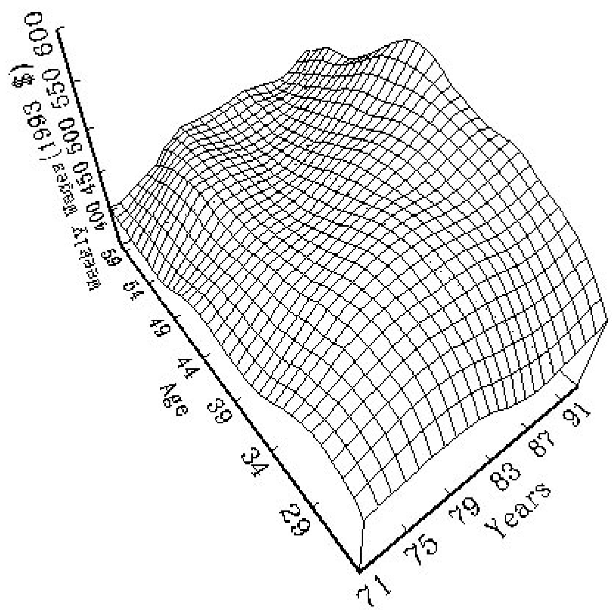

Panel B
Contour Plot

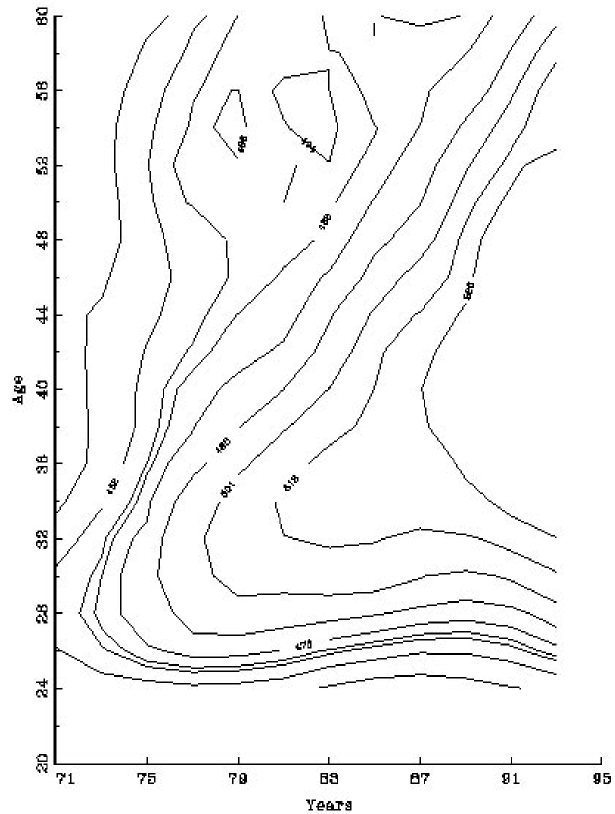

Figure 5: Median Weekly Wages:
All Females Aged 24-60

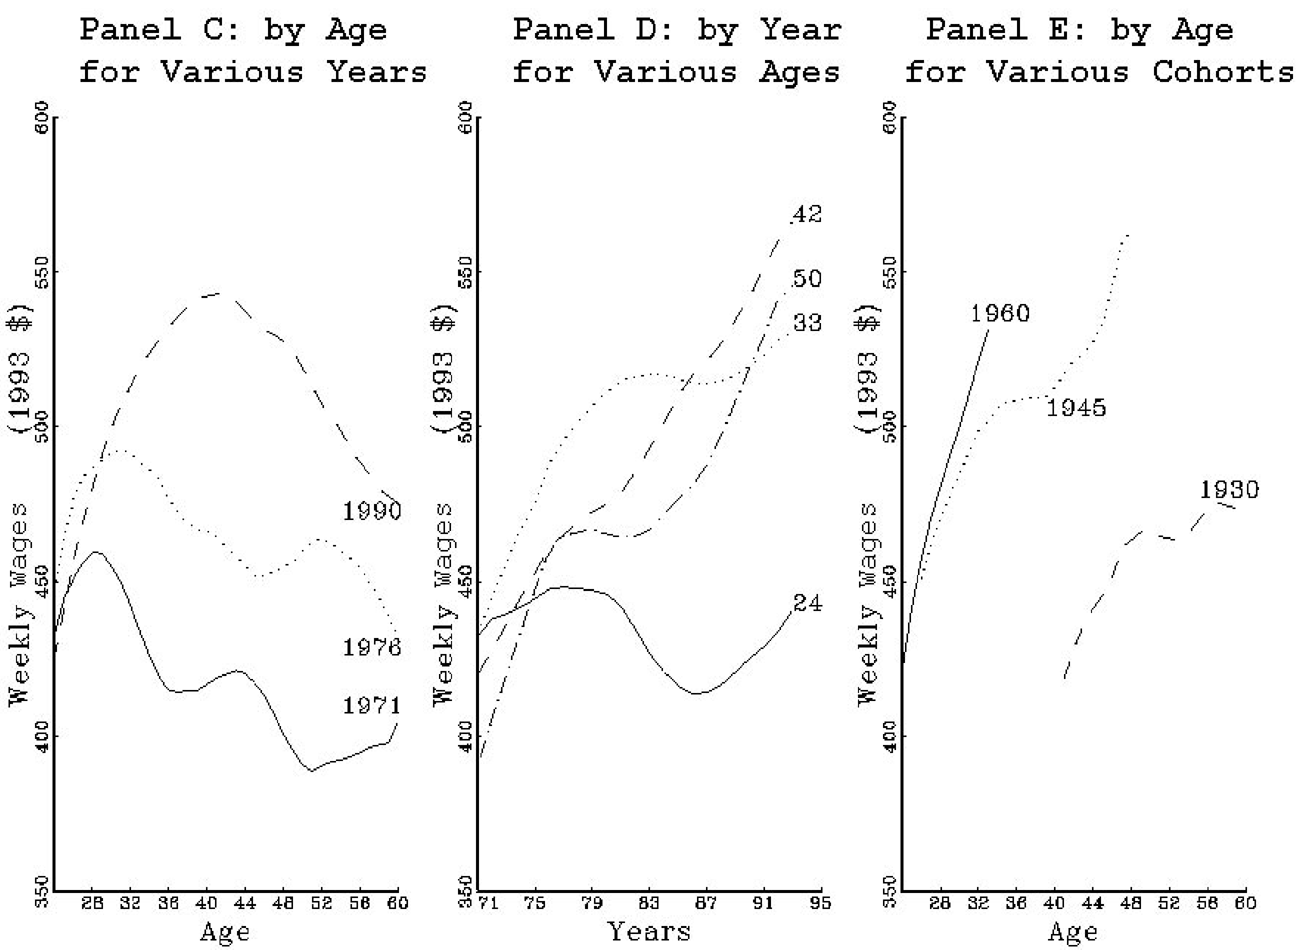

Figure 6: Median Weekly Wages: Females Without University: Aged 24-60

Panel A
View from (-0.5,-0.5,2.0)

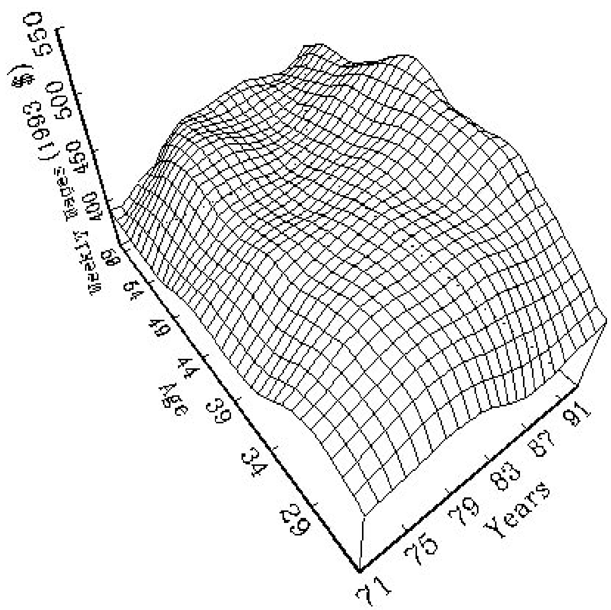

Panel B
Contour Plot

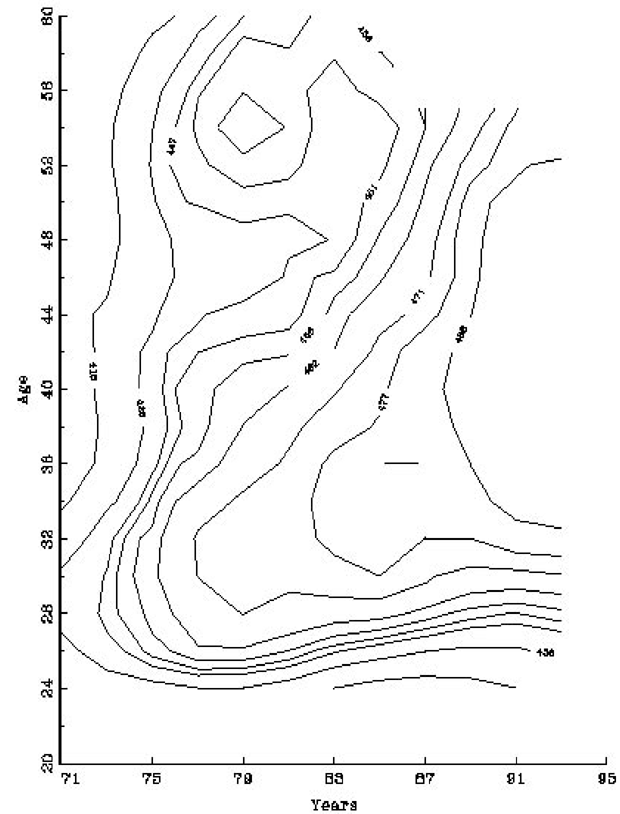

Figure 6: Median Weekly Wages: Females Without University: Aged 24-60

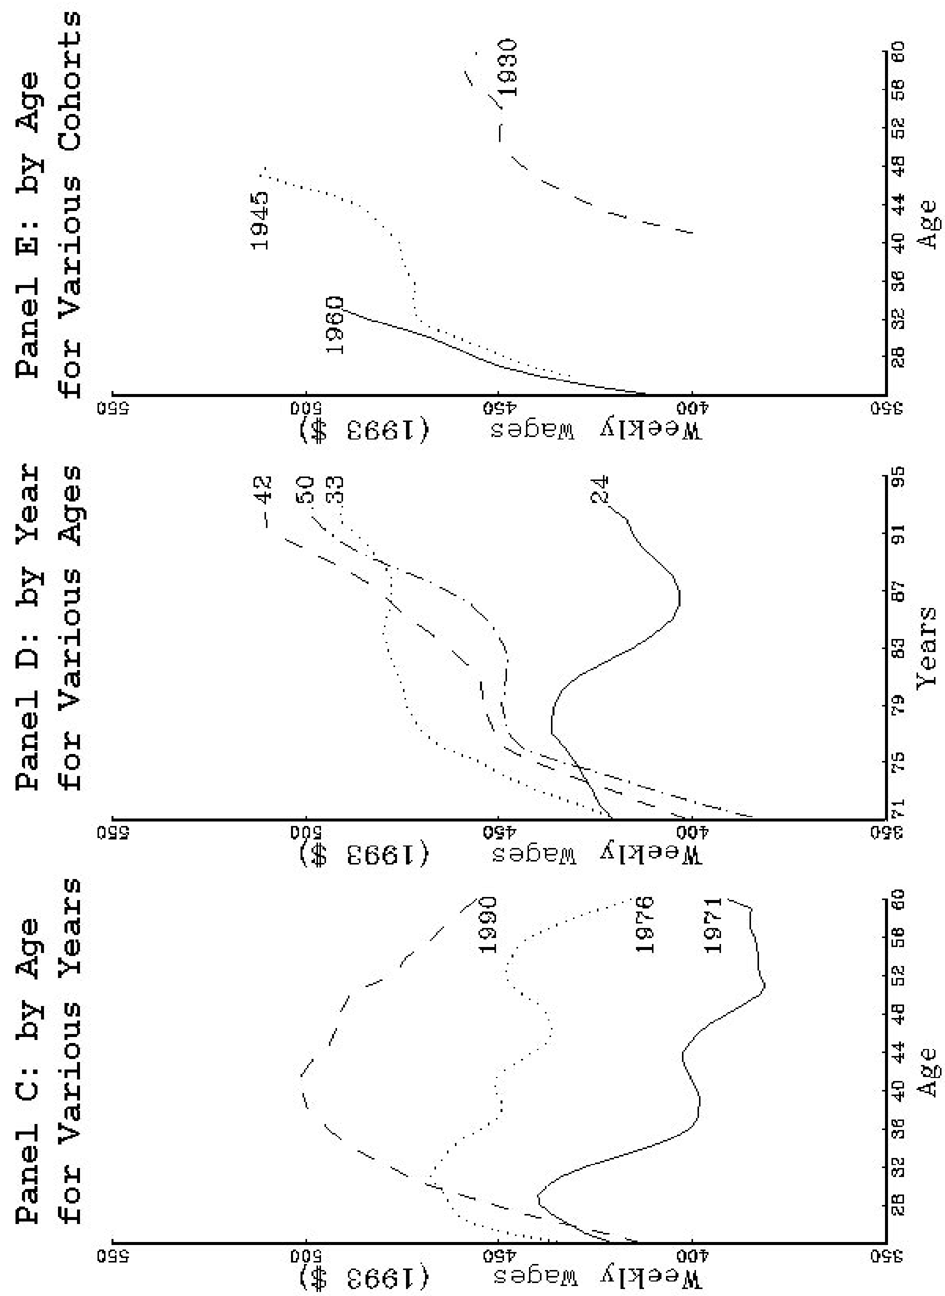

Figure 7: Median Weekly Wages:
Females With University: Aged 24-50

Panel A
View from (-0.5,-0.5,2.0)

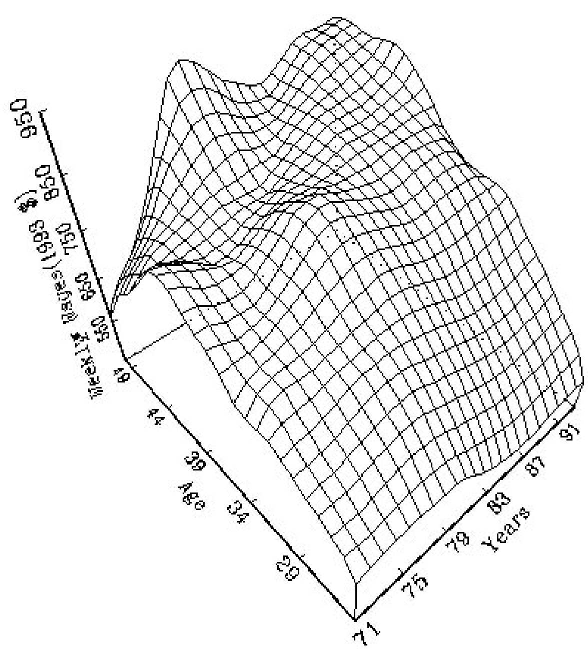

Panel B
Contour Plot

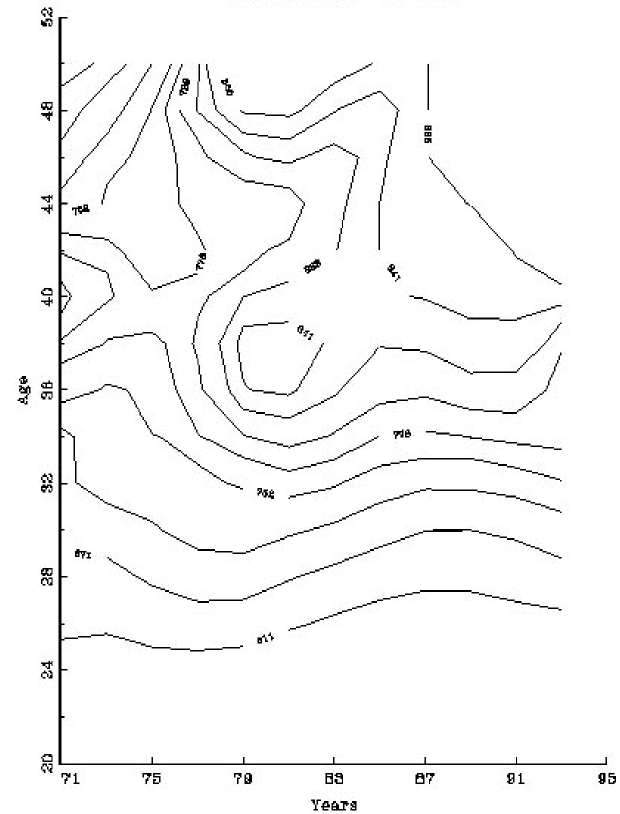

Figure 7: Median Weekly Wages:
Females With University: Aged 24-50

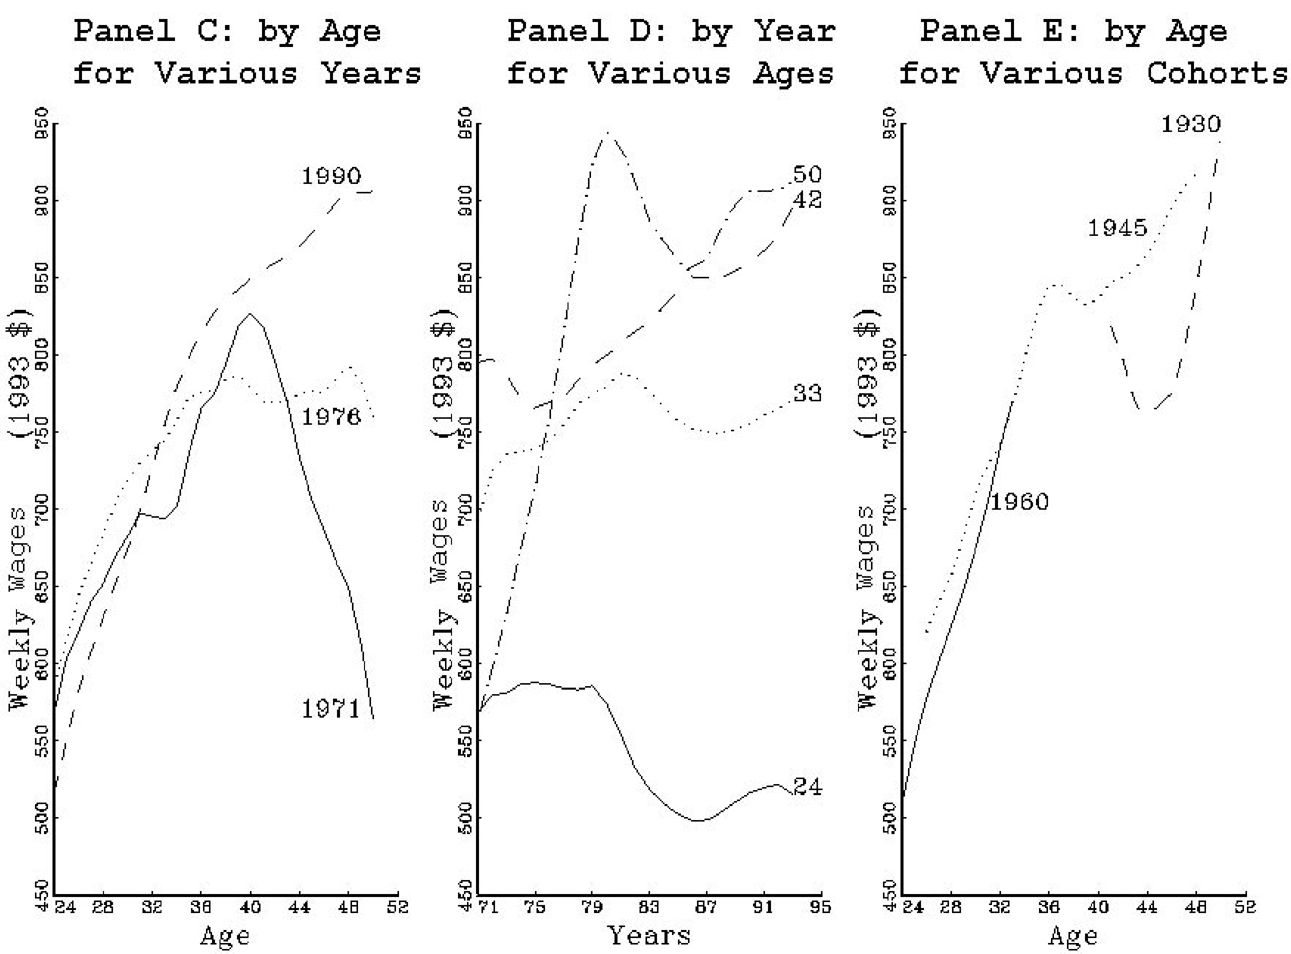

Figure 8

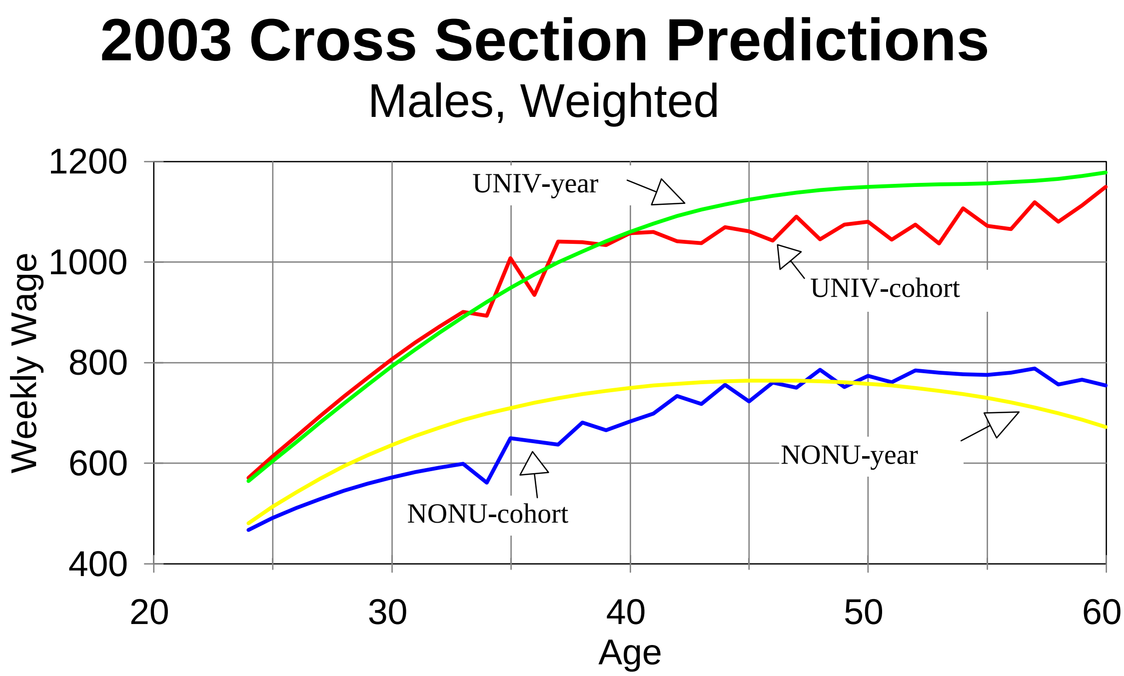
2003 Cross Section Predictions
Males, Weighted
UNIV-year
UNIV-cohort
NONU-year
NONU-cohort
Weekly Wage
1200
1000
800
600
400
20
30
40
50
60
Age

Figure 9

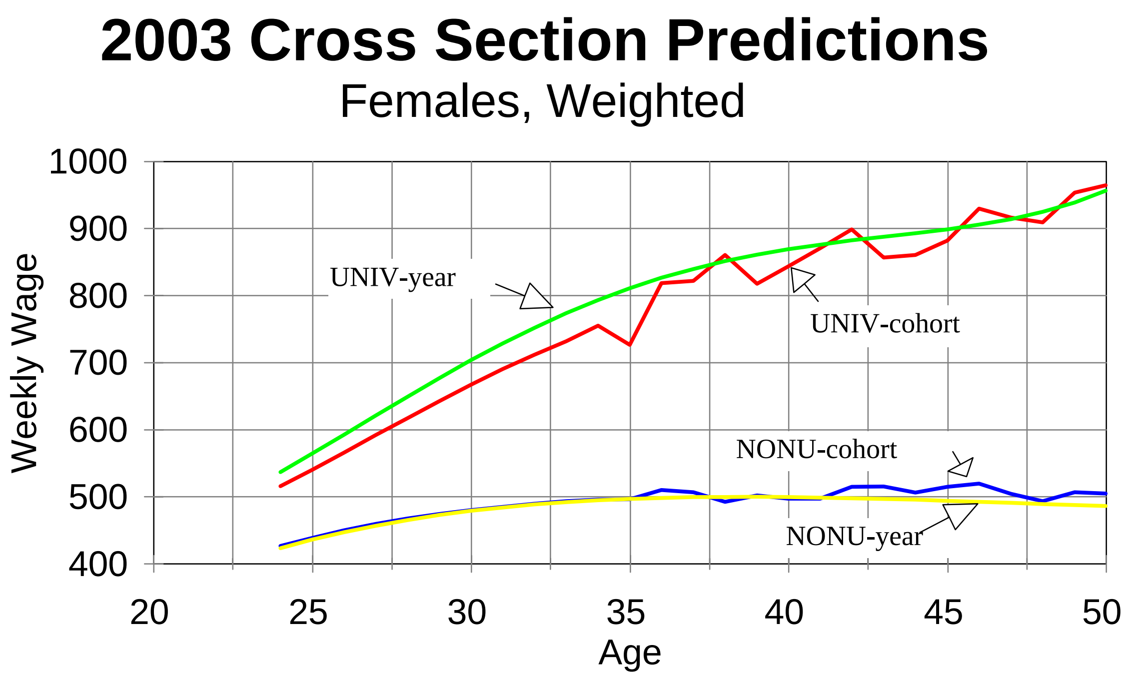
2003 Cross Section Predictions
Females, Weighted
Weekly Wage
1000
900
800
700
600
500
400
20
25
30
35
40
45
50
Age
UNIV-year
UNIV-cohort
NONU-cohort
NONU-year

Figure 10

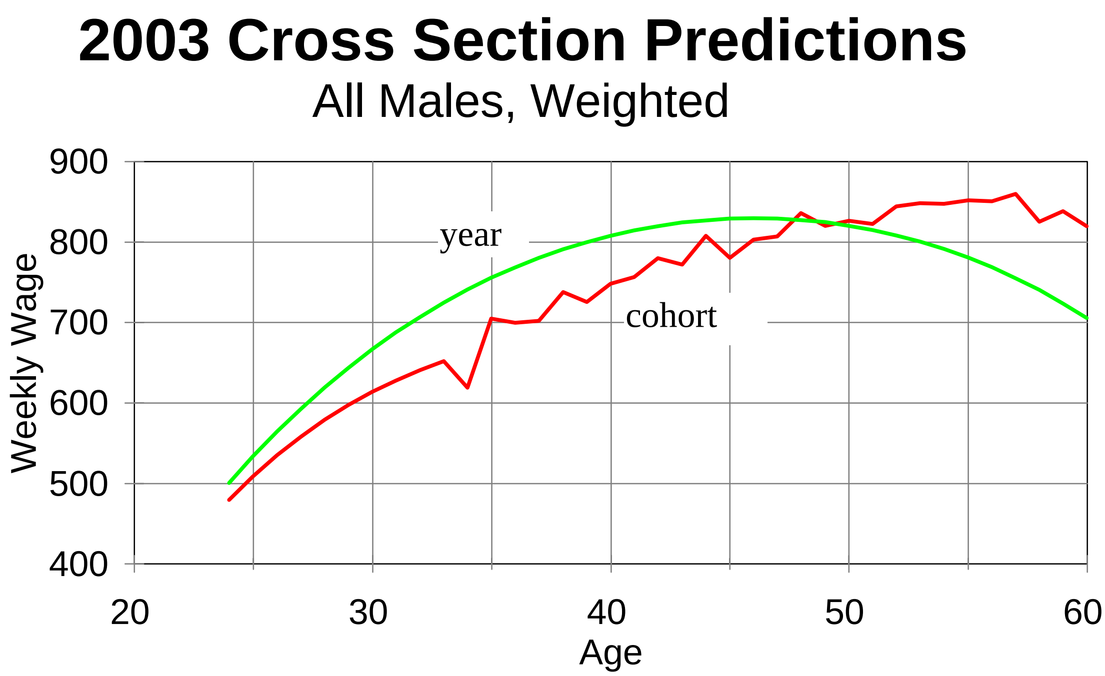
2003 Cross Section Predictions
All Males, Weighted
900
800
700
600
500
400
Weekly Wage
year
cohort
20
30
40
50
60
Age

Figure 11

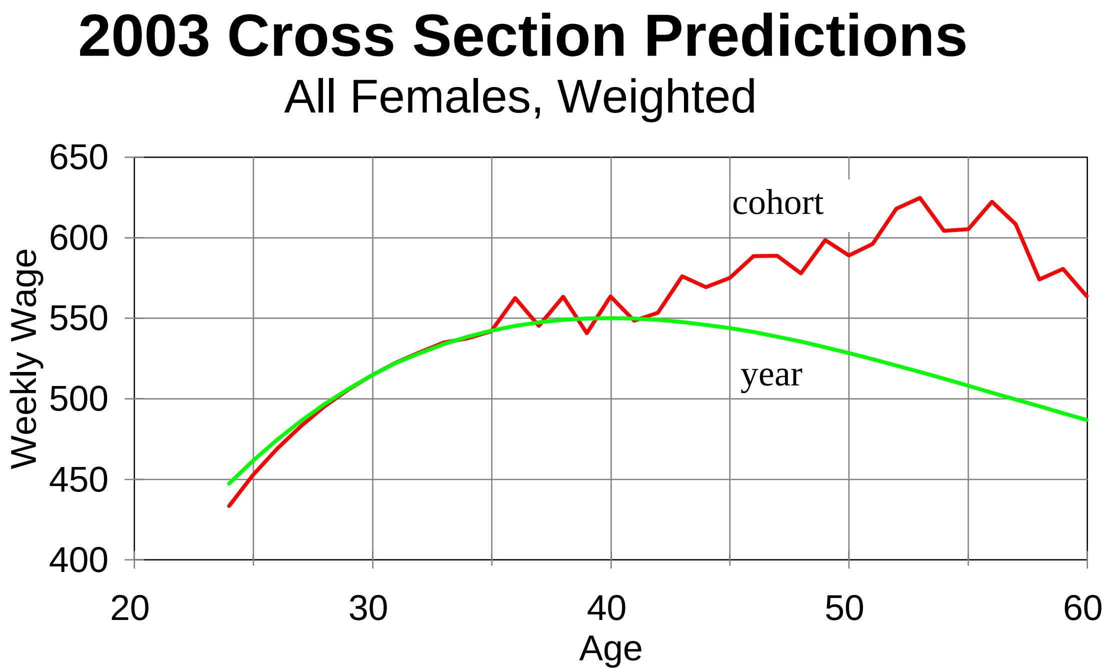
2003 Cross Section Predictions
All Females, Weighted
650
600
550
500
450
400
Weekly Wage
cohort
year
20
30
40
50
60
Age

Figure 12

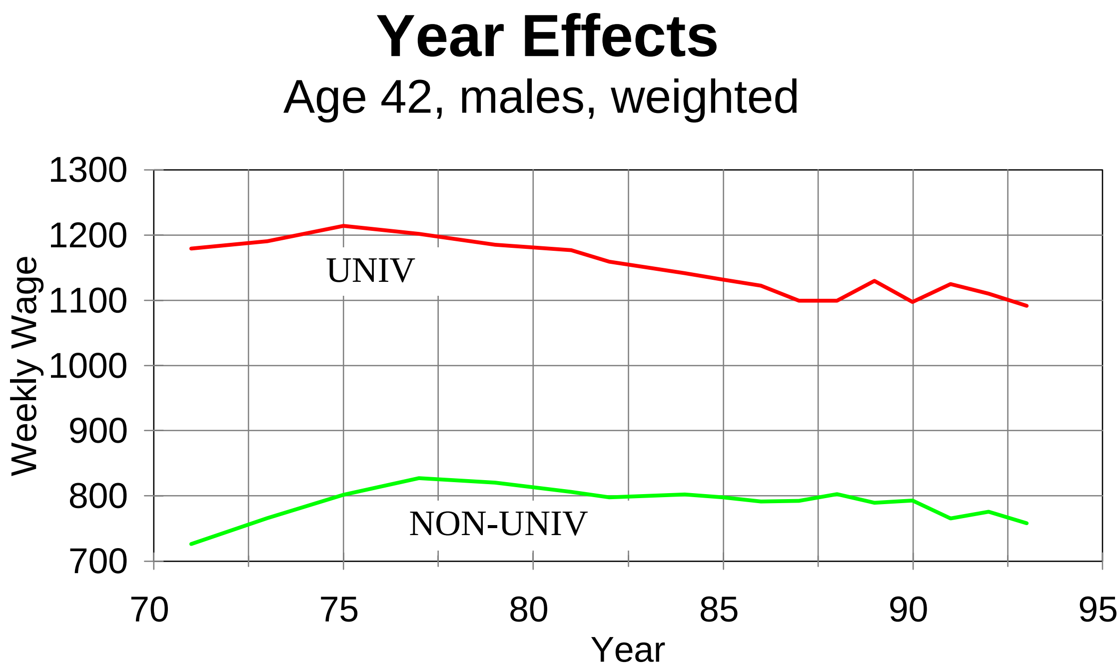
Year Effects
Age 42, males, weighted
1300
1200
1100
1000
900
800
700
Weekly Wage
UNIV
NON-UNIV
70
75
80
85
90
95
Year

Figure 13

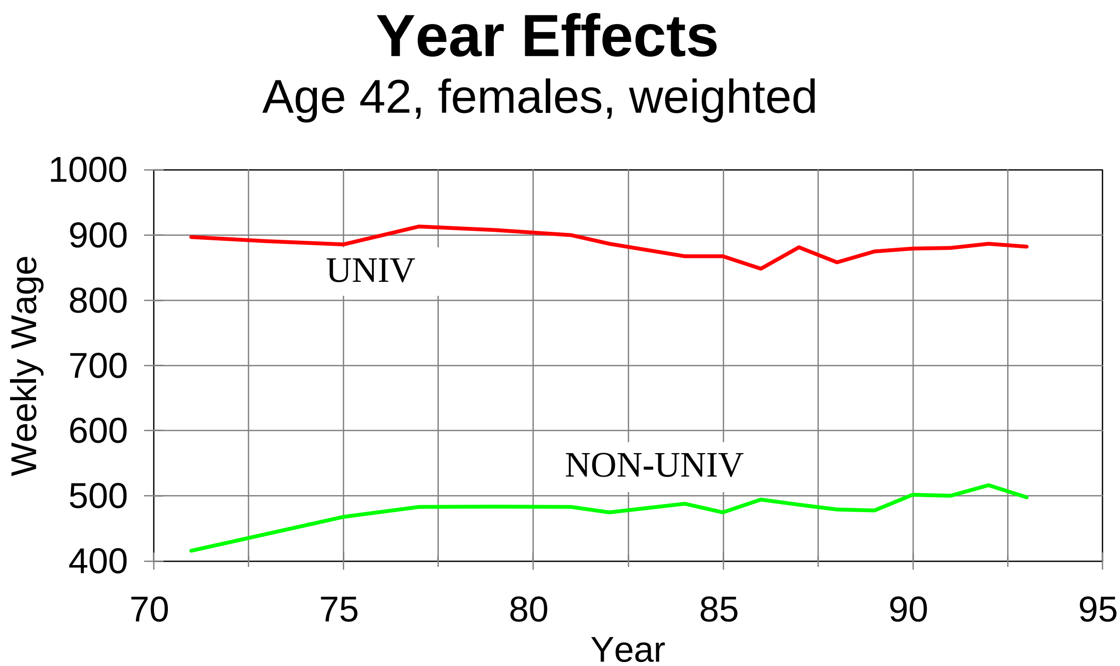
Year Effects
Age 42, females, weighted
Weekly Wage
1000
900
800
700
600
500
400
70
75
80
85
90
95
Year
UNIV
NON-UNIV

Figure 14

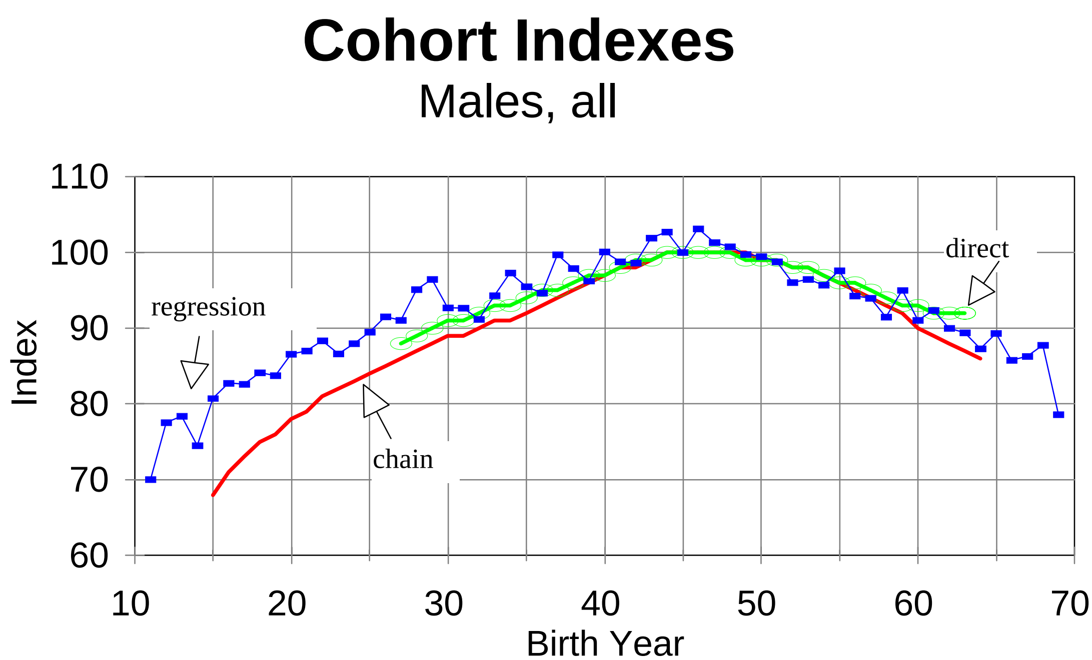
Cohort Indexes
Males, all
110
100
90
80
70
60
Index
10
20
30
40
50
60
70
Birth Year
regression
chain
direct

Figure 15

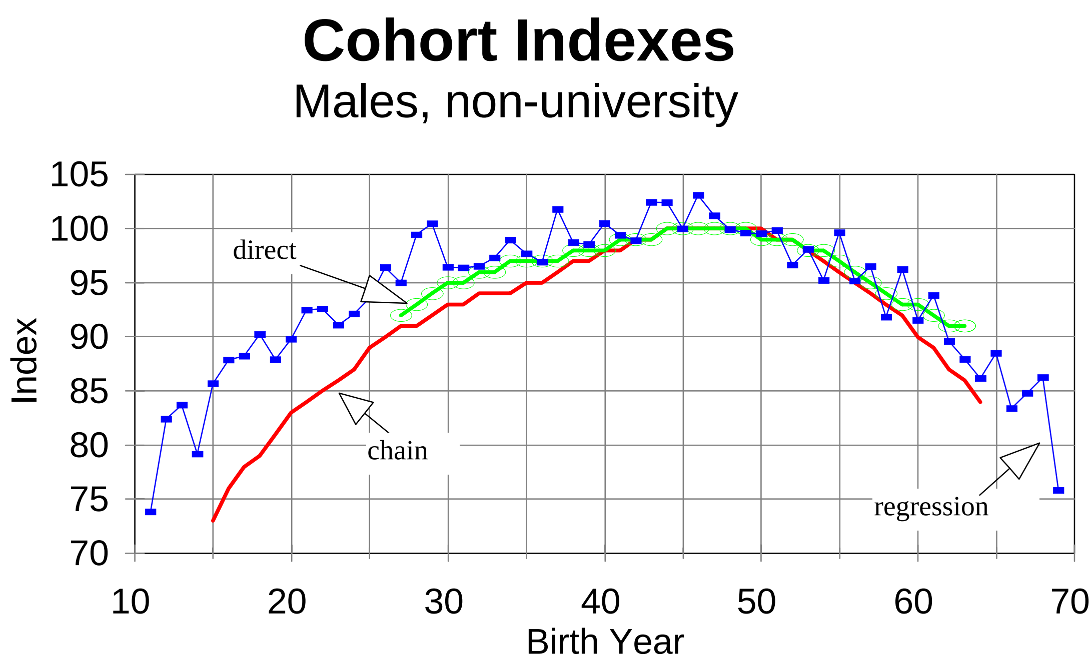
Cohort Indexes
Males, non-university
Index
105
100
95
90
85
80
75
70
10
20
30
40
50
60
70
Birth Year
direct
chain
regression

Figure 16

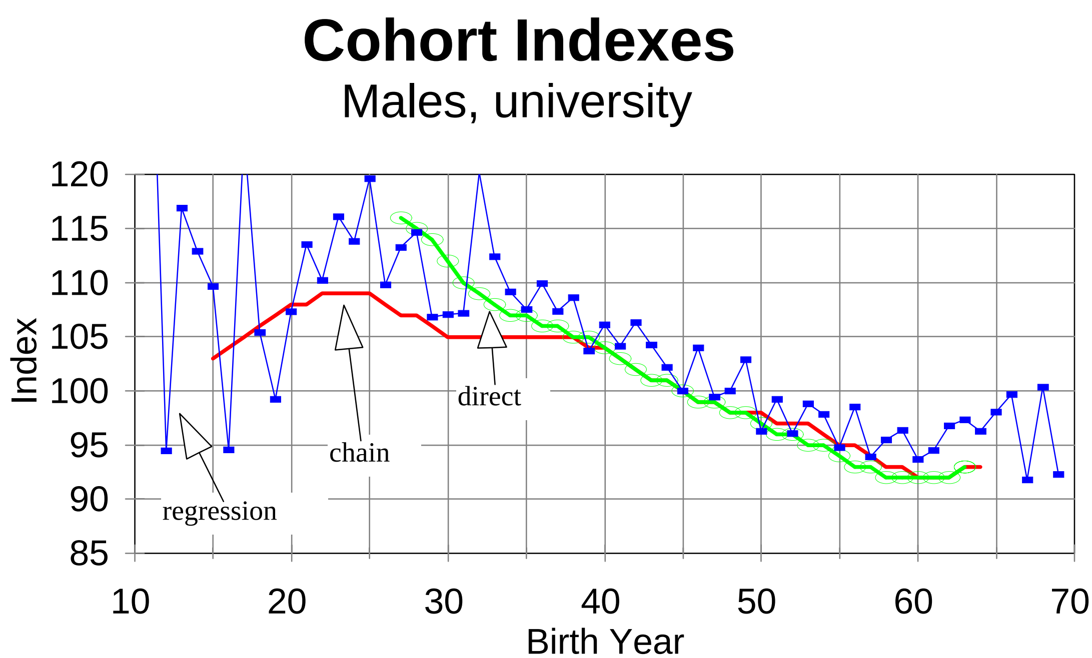
Cohort Indexes
Males, university
Index
Birth Year
120
115
110
105
100
95
90
85
10
20
30
40
50
60
70
regression
chain
direct

Figure 17

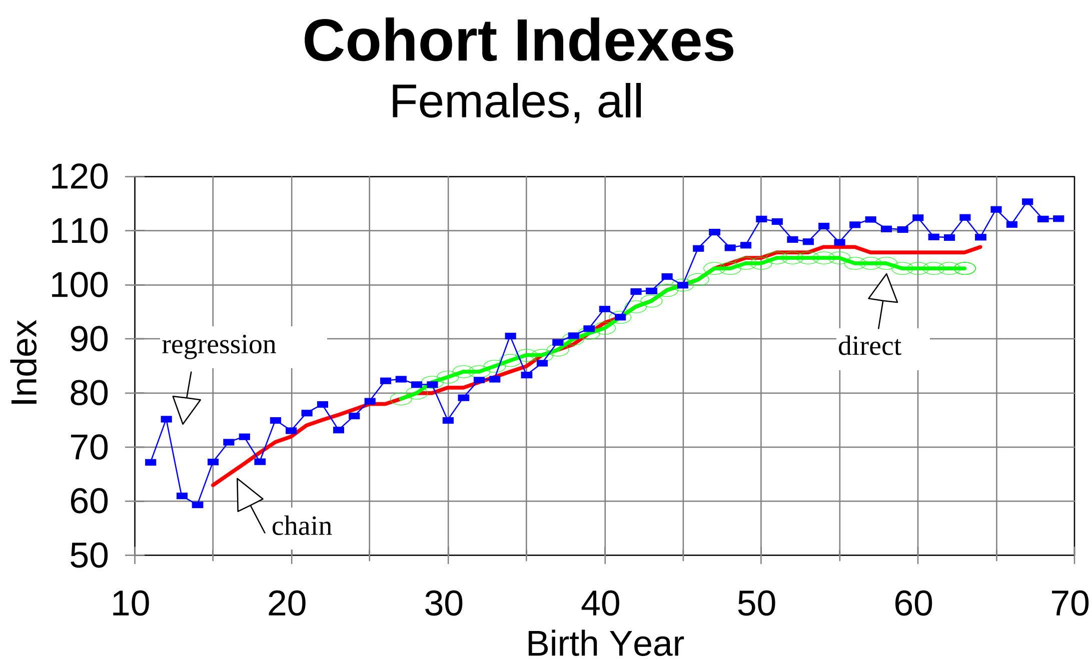

Cohort Indexes
Females, all
Index
Birth Year
120
110
100
90
80
70
60
50
10
20
30
40
50
60
70
regression
chain
direct

Figure 18

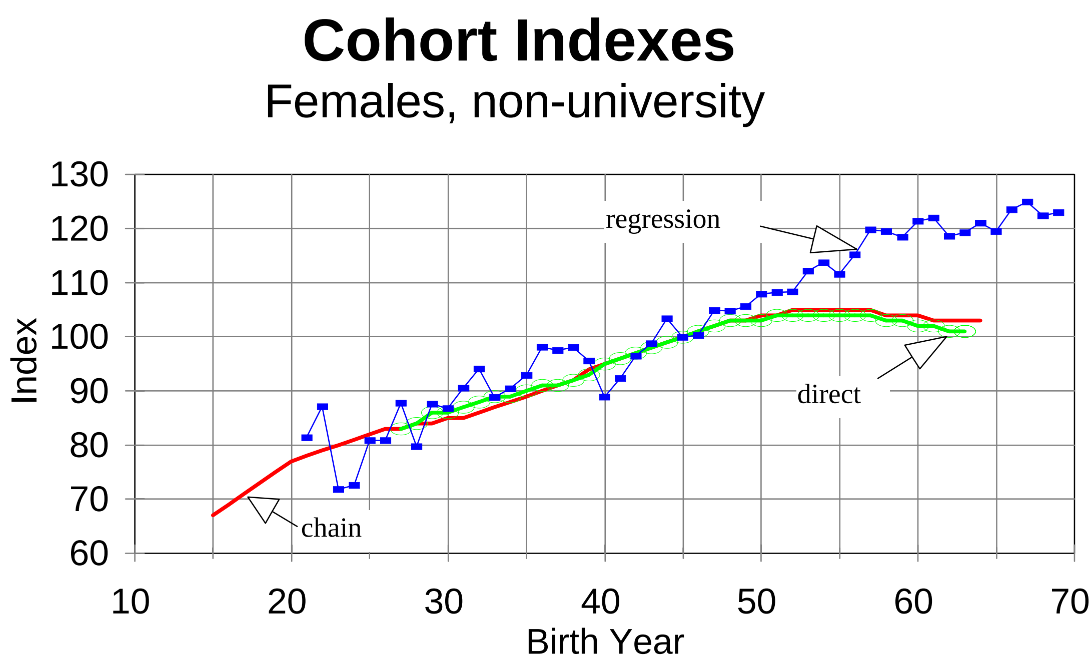
Cohort Indexes
Females, non-university
130
120
110
100
90
80
70
60
Index
regression
direct
chain
10
20
30
40
50
60
70
Birth Year

Figure 19

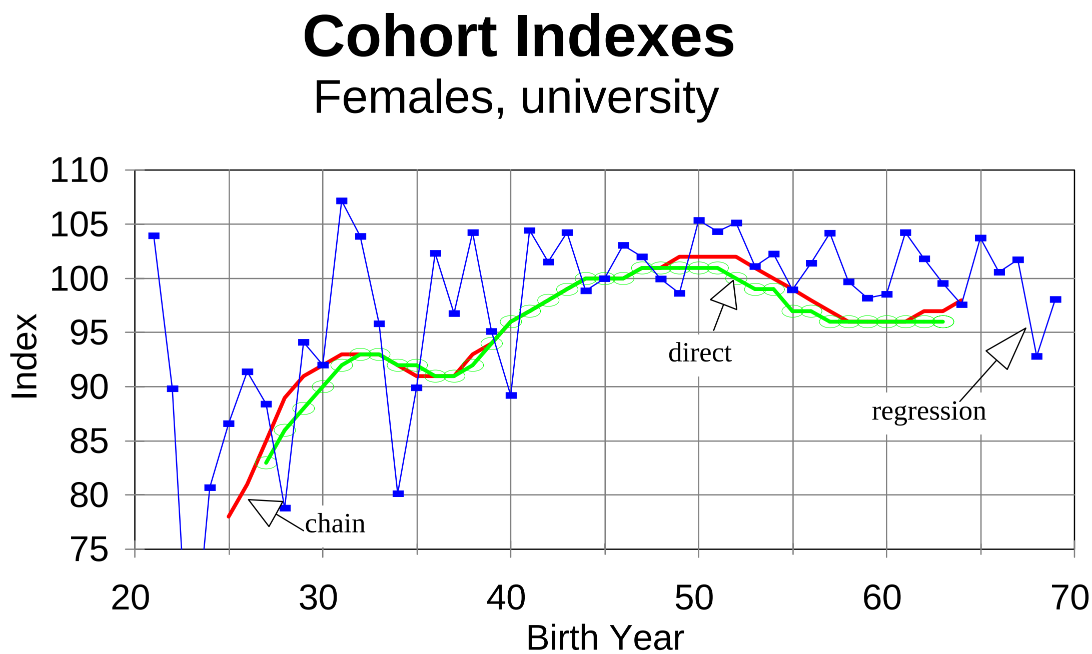
Cohort Indexes
Females, university
Index
110
105
100
95
90
85
80
75
20
30
40
50
60
70
Birth Year
direct
regression
chain

CILN Working Papers (downloadable)

wp43 **John Flemming John Micklewright**
Income Distribution, Economic Systems and Transition

wp42 **John Micklewright Kitty Stewart**
Is Child Welfare Converging in the European Union?

wp41 **W Bentley Macleod**
A Note on the Optimality of Bonus Pay

wp40 **Xin Meng Robert G Gregory**
Impact of Interupted Education on Earnings:
The Educational Cost of the Chinese Cultural Revolution

wp39 **Miles Corak**
Death and Divorce: The Long Term Consequences of Parental Loss on Adolescents

wp38 **Lori Curtis Martin Dooley**
Child Health and Family Socioeconomic Status in the Canadian National Longitudinal Survey of Children and Youth

wp37 **Heather Antecol**
An Examination of Cross-Country Differences in the Gender Gap in Labor Force Participation Rates

wp36 **W. Craig Riddell**
Canadian Labour Market Performance in International Perspective: Presidential Address to the Canadian Economics Association

wp35 **Shelley Phipps**
Economics and Well-Being of Canadian Children

wp34 **W. Craig Riddell**
Measuring Unemployment and Structural Unemployment

wp33 **Susan Johnson**
Automatic Certification or Mandatory Representation Votes? How the choice of union recognition procedure affects union certification success.

wp32 **James Andreoni Eleanor Brown Isaac C. Rischall**
Charitable Giving by Married Couples: Who Decides and Why Does it Matter?

wp31 **Herb Schuetze Peter Kuhn**

Self-Employment Dynamics and Self-Employment Trends: A Study of Canadian Men and Women, 1982-1995

wp30 **Isaac C. Rischall**
The Effect of High School Effort on Future Earnings

wp29 **Isaac C. Rischall**
The Roles of Education, Skill and Parental Income in Determining Wages

wp28 **Isaac C. Rischall**
The Value of a Fresh Start: Earnings Persistence and the Migration of Single Mothers

wp27 **Martin Browning Tom Crossley**
Shocks, Stocks and Socks: Consumption Smoothing and the Replacement of Durables During an Unemployment Spell

wp26 **Heather Antecol Kelly Bedard**
Against All Odds: The Surprising Labor Market Success of Young Mexican Women

wp25 **Heather Antecol**
Why is there Inter-Ethnic Variation in the Gender Wage Gap? The Role of "Cultural" Factors

wp24 **Martin Browning Tom Crossley**
Unemployment Insurance Benefit Levels and Consumption Changes

wp23 **Heather Antecol Peter Kuhn**
Employment Equity Programs and the Job Search Outcomes of Men and Women: Actual and Perceived Effects

wp22 **Thomas F. Crossley**
Firms and Wages: Evidence from Displaced Workers

wp21 **Jennifer Stewart Martin Dooley**
The Duration of Spells on Welfare and Off-welfare among Lone Mothers in Ontario

wp20 **Peter Kuhn Arthur Sweetman**
Vulnerable Seniors: Unions, Tenure and Wages Following Permanent Job Loss

wp19 **Kelly Bedard**
Human Capital Versus Signaling Models: University Access and High School Drop-outs

wp18 **Peter Kuhn Arthur Sweetman**
Assimilation and Economic Success in an Aboriginal Population: Evidence from Canada

wp17 **Martin D. Dooley**
The Evolution of Welfare Participation Among Canadian Lone Mothers From 1973 - 1991

wp16 **Lori Curtis Martin D. Dooley Ellen L. Lipman David H. Feeny**
The Role of Permanent Income and Family Structure in the Determination of Child Health in the Ontario Child Health Study

wp15 **LaDonna A. Pavetti**
A New Social Contract: Moving to a Work-Based Assistance System

wp14 **Gary Burtless**

The Job Prospects of U.S. Welfare Recipients: Lousier Pay but Bigger Earnings Supplements

wp13 **J.B. Burbidge L. Magee A.L. Robb**
Cohort, Year and Age Effects in Canadian Wage Data

wp12 **Martyn Andrews Alan Harrison**
Testing for Efficient Contracts in Unionized Labour Markets

wp11 **Herb J. Schuetze**
Taxes, Economic Conditions And Recent Trends in Male Self-Employment: A Canada-U.S. Comparison

wp10 **Peter Kuhn**
Canada and the "OECD Hypothesis": Does Labour Market Inflexibility Explain Canada's High Level of Unemployment?

wp9 **Stephen R. G. Jones W. Craig Riddell**
The Measurement Of Unemployment: An Empirical Approach

wp8 **Pierre Lefebvre Philip Merrigan Martin Dooley**
Lone Female Headship and Welfare Policy in Canada

wp7 **Heather Antecol Peter Kuhn**
Gender as an Impediment to Labor Market Success: Why do Young Women Report Greater Harm?

wp6 **John Myles Paul Pierson**
Friedman's Revenge: The Reform of "Liberal" Welfare States In Canada and the United States

wp5 **Jeff Borland**
Earnings Inequality in Australia: Changes and Causes

wp4 **Jeff Borland**
Union Effects and Earnings Dispersion in Australia, 1986-1994

wp3 **Robert Gregory Boyd Hunter**
The Macro Economy and the Growth of Income and Employment Inequality in Australian Cities

wp2 **Peter Kuhn**
Labour Market Polarization: Canada in International Perspective

wp1 **Peter Kuhn A. Leslie Robb**
Shifting Skill Demand and the Canada-US Unemployment Gap: Evidence from Prime-Age Men

Last updated March 27, 2000