

A NONPARAMETRIC ANALYSIS OF CANADIAN EMPLOYMENT PATTERNS

LUKE IGNACZAK AND MARCEL VOIA¹

*Department of Economics, Carleton University, 1125 Colonel By Drive, Ottawa ,
Ontario K1S 5B6, Canada.*

E-mails: `lignacza@connect.carleton.ca`, `mvoia@connect.carleton.ca`

February 2009

Abstract. Popular perception holds that employment stability has decreased in recent decades. However, no conclusive evidence exists on secular declines in the length of jobs held. Furthermore, most studies conclude that the proportion of long term jobs has remained remarkably stable over the last few decades. To shed light on this discrepancy we use distribution analysis to systematically track changes in Canadian employment durations over an extended period. This is done in order to reconcile popular perception with recent studies and nest the existing literature in a broader historical context. Using finite mixture decomposition on successive cohorts of workers starting from the 1950s we identify worker types within cohort-based distributions. Then, using tests of stochastic dominance, we show that the distribution of employment has indeed changed. The finite mixture decomposition reveals that earlier cohorts were more likely to have longer tenure than later cohorts and that there are shifts in proportions between longer and shorter work episodes. Our results also indicate that after the 1960s employment durations declined sharply for men, while for women the results were mixed.

Classification codes: J01, C14, C12, C16, C41.

Key words and phrases: Nonparametrics, Stochastic Dominance, Kolmogorov-Smirnov type statistic, Bootstrap, Heterogeneous Distribution, Censored Distributions, Finite Mixtures, Employment Duration.

¹The data used for this study were provided by Statistics Canada, but the results obtained and any errors are solely those of the authors.

It is commonly perceived that employment stability has declined, that is, workers are more likely to have a larger number of shorter term jobs over the course of their career than had been the case in the past. However, international evidence on the issue of declining job tenure is mixed. Furthermore, there have been only a limited number of comprehensive studies on employment tenure in the Canadian context and existing studies showed no long term secular declines in stability. The aim of this paper is to fill a gap in the long term analysis of employment durations in Canada using evidence gleaned from the General Social Survey (GSS).

Using evidence from distributions of completed and censored employment spells we will show that employment durations have indeed changed over time. Specifically, we find evidence consistent with these durations declining over the latter half of the twentieth century. This will be, as far as we know, the first paper to do so using such a long term perspective. It is a particularly relevant question as recent studies in the U.S. have shown that employment stability has indeed decreased over the last forty years, contrary to much of the earlier research. No such determination has yet been made for Canada.

The importance of better understanding historical patterns of employment stability is twofold. Firstly, for workers, employment stability is important in the maintenance of an individual's human capital and is generally associated with higher-quality work and higher levels of satisfaction. This has a number of policy implications when combined with information on which factors affect stability. Secondly, by looking at the long term it is possible to reconcile the puzzling discrepancy between popular opinion and the academic literature. We explain the discrepancy by pointing out that many of the major changes in stability occurred before the 1970s. This is important because most studies only begin looking at stability at that time. This is done by using data about the 1950s, 1960s and 1970s which does not appear in other studies of employment stability.

In section one we will discuss the literature on tenure and outline a basic worker heterogeneity model which will be our basic point of reference. Section two will detail the empirical methodologies employed in the identification and the tests employed in this study. The third section focuses on the retrospective data used and describes some of the problems that are encountered in using it. The fourth section presents the results of our analysis. Finally, section five concludes and discusses further avenues for research on this topic.

1. THEORY AND EVIDENCE

The majority of studies on employment stability are based on data from the United States. This is because of the large number of datasets available in the United States. One of the seminal papers in this field was that of Hall (1982) who, using the Current Population Survey (CPS) of 1968 and 1978, constructed synthetic cohorts, and found three characteristics which continue to define job durations in the U.S.. These three characteristics are:

- (1) Most jobs do not last a long time, but those that last for five years are likely to last twenty years.
- (2) A substantial portion of workers will, at some point in their careers be in a “lifetime job,” this being defined as one which lasts 20 years or more.
- (3) Men are more likely than women, and whites are more likely than blacks to be in jobs of a longer duration.

Many authors have written on the topic of job stability. The results are sometimes conflicting, but more often than not the results show a good deal of stability in employment durations, particularly in the earlier literature. With some minor caveats, Hall (1982), Diebold et. al. (1994) and Farber (1995) using CPS data all found that job durations have not changed greatly between the 1970s and early 1990s. Furthermore, Jaeger and Stevens (1999), using the Panel Study of Income Dynamics (PSID) and CPS, found little evidence of a trend in job stability.

Conversely, using CPS data Swinnerton and Wial (1995) showed that there had been a decline in stability over the 1980s. Rose (1995), corroborated Swinnerton and Wial’s results using the PSID, and argued that stability decreased between the 1970s and 1980s for men. On a related matter, it has been surmised that worker displacement may have seen significant increases between the 1970s and 1980s (Fallick, 1998). Furthermore, evidence from the PSID shows that job losses stemming from plant closings rose between the 1960s and 1980s (Hamermesh, 1989). More recently, Farber (2007) using a birth-cohort approach found an increase in short-tenure positions and declines in long-term employment relationships using the 1973-2006 CPS. In particular, Farber’s results indicate that mean employment tenure for men has deteriorated by about 50% after controlling for education between the 1914 and the 1975 birth cohorts. For women the declines were smaller. Many of these results point to declines occurring well before the 1980s.

In the Canadian context, researchers have primarily focussed on the period following 1976, when the Labour Force Survey (LFS) became available. In this literature it has been largely shown that employment is generally stable and declines in job stability are largely pro-cyclical. Green and Riddell (1997) found that, between 1979 and 1991, the distribution of job durations “hollowed out”, such that both, short and long tenure jobs increased at the expense of mid-length jobs. Picot, Heisz, and Nakamura (2001), using the LFS, the Longitudinal Worker File (LWF), Survey of Consumer Finances (SCF) and Survey of Labour and Income Dynamics (SLID) showed that job duration declined in the 1980s but rose again over the 1990s, more than making up for the declines the previous decade. Heisz (2005), using the LFS and CPS, corroborates these findings and points out that stability has declined to a greater extent in the United States than in Canada over the 1990s. Conversely, Brochu (2008) found that a secular increase in stability has occurred for workers with low initial job tenure.

In general, because of data availability, the literature generally focuses on analyzing either in-progress job lengths or retention rates. Using a cross-cut of the data to analyze in progress job lengths allows for a clear picture of tenure at a particular point in time. However, this method leads to the problem of length bias where the number of short-term jobs are under reported. Similarly, because lower mobility workers remain in their positions longer, any snapshot in time will show a larger share of long-term jobs in the labour market than is actually the case. Retention rates solve this problem by analyzing the conditional probability of jobs continuing for a certain length given their current duration. A complication for some studies using this method is that estimates rely on synthetic cohorts in repeated cross-sections data and may be prone to errors generated by uncontrolled changes in the composition of the cohorts.

Both of the above methods are prone to the problems of telescoping recall errors as years in the current job are generally recorded at the time of the survey. Optimally, it would be desirable to follow workers from the start to termination of employment but data for this purpose is usually lacking. Our method will use retrospective data to generate distributions of completed employment spells.

The theoretical basis for the ensuing analysis will be a model of worker heterogeneity. This model is a modified version of the pure mover-stayer model of worker heterogeneity introduced by Blumen et. al. (1955) as described by Farber

(1999). In this section we will briefly describe this model and move to discuss its implications and the current state of evidence on employment durations.

In the model there are multiple types of workers with distinguishing turnover probabilities denoted by λ_i . For example, in a two type model we could have high mobility workers (with mobility λ_H) and low mobility workers (with mobility λ_L) where $\lambda_H > \lambda_L$. If turnover probabilities remain fixed, then θ_i will represent the proportion of workers with mobility λ_i . The total probability of turnover in the labour market can then be described by a weighted mean of the turnover probabilities of all worker types. This can be written as

$$P = \sum_{i=1}^n \theta_i \lambda_i, \sum_{i=1}^n \theta_i = 1 \quad (1.1)$$

Thus, the total degree of turnover in the labour market is entirely dependent on the proportion of individuals within the different groups. In reality, proportions can be variable over time. If we analyze any particular group starting work at the same point in time the empirical distribution of these workers employment durations should allow us to identify the underlying worker types. We will not assume that individual types are fixed, only that the worker is associated with one type over their employment length. Indeed, in their subsequent employment the same worker can be a different type. The types will be identified by the data rather than vice-versa.

In effect, our goal is to reverse engineer a mover-stayer model to identify worker types by looking at the distribution of completed job durations. Hence, we will be able to determine whether job durations have changed over time due to the composition of worker types within the distributions. To find the worker types within employment duration distributions we will use a finite mixture decomposition. To connect these tests to the worker heterogeneity model, we must make the assumption that different worker types can be identified by the different durations of employment shown by the data. It should be understood that the types identified below will be the result of past labour market equilibria. Once the decompositions are done we will move on to stochastic dominance testing which will show a broad decline in the length of employment spells over the latter half of the twentieth century.

2. METHODOLOGY

We will use two primary methods in order to determine whether employment durations have declined. The first method will be a finite mixture decomposition of

completed employment durations by decadal cohort in order to identify worker types by their spell lengths. The second will be a series of tests of stochastic dominance, again, by decadal cohort to test for an overall decline in spell length using both completed and ongoing job spells. To decide what types of parametric specification to use for the mixture model, a non-parametric graphical representation is used. The fit of the estimated mixture distributions is checked using tests of equality of distributions. Our arguments about the non-parametric estimation follow Silverman (1986), the subsequent heterogeneity analysis uses the method of McLachlan and Peel (2000), while the stochastic dominance testing are an extension of Linton, Maasoumi and Whang (2005) to account for censored distributions as in Ignaczak and Voia (2009). The proposed extension has the advantage of allowing testing only on the range over which restricted dominance can be inferred. Further, the finite mixture decomposition suggests the presence of nuisance parameters in the distributions, which require a parametric bootstrap to estimate the critical values of the test statistics as in Huynh and Voia (2008).

2.1. Heterogeneity Testing.

2.1.1. *Nonparametric Density Estimation.* Before assuming what type of decomposition is suitable we first examine the distribution using a kernel density estimator. Varying bandwidths are required when long-tailed or multi-modal density functions are estimated using kernel methods. Therefore, we use adaptive kernel density estimation to avoid the potential problems of using kernels with a fixed bandwidth, such as undersmoothing in areas with only sparse observations and oversmoothing in others. Adaptive kernels were introduced and discussed by the following authors: Silverman (1986), Bowman and Azzalini (1997), or Pagan and Ullah (1999). Salgado- Ugarte et al. (1993), Salgado-Ugarte et al. (1995) and Salgado-Ugarte and Perez-Hernandez (2003). The method employed is described in the technical annex.

2.1.2. *Testing for internal heterogeneity.* To identify the potential number of types of workers in each cohort we use a finite mixture decomposition of the cohort-specific outcome duration distribution. The results of the finite mixture decompositions are used to understand the changes in worker heterogeneity over time. The finite mixtures induce nuisance parameters in the distributions of interest. Therefore, to deal with this potential problem, the results of the finite mixture decompositions

are used to improve the size and power of the stochastic dominance test statistics by employing a parametric bootstrap in computing critical values.

For a random variable Y (our outcome variable of interest), the finite mixture models decompose a probability density $f(y)$ into the sum of K class probability density functions. In this case, if $f_k(y)$ is the k^{th} class probability density function and p_k denotes the proportion of the k^{th} class, we can define the finite mixture model with K components as $f(y) = \sum_{k=1}^K p_k f_k(y)$. The proportion p_k can be interpreted as the prior probability of observing a sample from class k , with the property that is greater or equal to zero and that they sum to one ($p_k \geq 0$ and $\sum_{k=1}^K p_k = 1$).

The objective is to estimate the parameters of the class probability densities and the proportions p_k of each class, while fixing an upper bound on the number of possible classes. The class of probability densities are assumed to have parametric components that can be estimated. The adaptive kernel density plots are examined to determine what parametric specifications we use to estimate the probability densities. The plots suggest that the outcome duration is approximated with mixtures of log-normal distributions. The parameters of such mixtures are estimated by maximum likelihood and guided by model selection methods. The following likelihood functions are used:

$$f_{Duration}(y, \theta) = \sum_{k=1}^K p_k \frac{1}{y \sigma_k \sqrt{2\pi}} \exp\left(\frac{-(y - \mu_k)^2}{2\sigma_k^2}\right), \quad (2.1)$$

The parameters of interest are: $\theta = \{K, p_k, \mu_k, \sigma_k\}$ with $k = 1, \dots, K$ and $\sum_{k=1}^K p_k = 1$. All the parameters of interest with the exception of the number of types are estimated by the likelihood. The number of types are estimated using model selection based on the AIC criteria. The following AIC criteria is minimized:

$$AIC_k = -2 \ln l(\theta|y) + 2d_k, \quad (2.2)$$

where d_k is equal to the dimension of the model and acts as a correction term without which one will choose the model that maximizes the unconditional log-likelihood.

2.2. Stochastic Dominance Testing. To formally determine whether employment durations have declined we use tests of stochastic dominance and equality of distributions. However, because a complete distribution of employment durations

for any particular group of workers can range in lengths of up to many decades it is often impossible to compare complete distributions of recent groups of workers. Because of this constraint the stochastic dominance tests will be limited to some feasible length, longer for earlier cohorts and shorter for more recent groups.

These tests will allow us to answer the question of whether evidence of a secular decline in employment durations exists. The tests are explained formally below.

Consider that we observe the duration of employment of different groups (cohorts) at different time periods. The outcome variable of interest is $Y^{(C_t)}$, where $Y^{(C_t)}$ is the length of employment duration for cohort C_t , where t is the starting period of a given cohort. To insure correct comparison of the cohort specific outcome distributions we censored our cohort data at different comparison levels (13 years, 23 years and 33 years). This censoring of the distributions induces some difficulties that are addressed using a simulation study by Ignaczak and Voia (2009). Define the associated censored cumulative distribution functions as $F^{(C_t)}$. We shall be interested in various properties of the conditional distribution functions

$$F^{(C_t)}(y \leq y_{max}) = \mathbf{P} [Y^{(C_t)} \leq y].$$

Let

$$G_1^{(C_t)}(y \leq y_{max}) = F^{(C_t)}(y),$$

and define the higher orders of G_1 by

$$G_s^{(C_t)}(y \leq y_{max}) = \int_0^y G_{s-1}^{(C_t)}(x) dx.$$

We consider three possibilities for $F^{(C_t)}$:

- (1) The distributions of the tested comparison cohorts are equal. In this case we write the null hypothesis as

$$H_0^{(1)} : F^{(C_{t_i})} \equiv F^{(C_{t_j})}, \text{ where } i \neq j.$$

- (2) One of the distributions first order stochastically dominates another. We shall consider the case when $F^{(C_{t_i})}(y) \leq F^{(C_{t_j})}(y)$ for all $y \leq y_{max}$. We formulate the corresponding null hypothesis as

$$H_0^{(2)} : F^{(C_{t_i})} \leq F^{(C_{t_j})}.$$

- (3) The two distributions intersect, but we have that one distribution second order stochastically dominates the other. In this case we write the null

hypothesis as

$$H_0^{(3)} : \int_0^y (y - x) dF^{(C_{t_i})}(x) \leq \int_0^y (y - x) dF^{(C_{t_j})}(x),$$

for all $y \leq y_{max}$.

These tests as well as the parametric bootstrap used to obtain the critical values for the stochastic dominance tests are fully described in the technical annex.

3. DATA

The survey data used in this analysis is a combination of the 15th and 20th cycles of the Canadian General Social Survey (GSS). These surveys were conducted throughout 2001 and 2006 respectively. They have a combined 47,918 records and yield a total of 59,892 employment spells for analysis. Because the GSS is primarily concerned with social issues it lacks some desirable economic information. For example, only current occupational and industrial detail is provided. Furthermore, only current income is recorded. However, despite these shortcomings the 15th and 20th cycles of the GSS offer a wide range of retrospective data on a number of interesting economic phenomena.

The variable of interest is the duration of an employment spell. It should be noted that the data obtained focusses on work interruptions of 3 months or more. The data do not permit us to verify whether the interruption ended with a return to the same employer or a new employer. Hence, the operational concept of employment will be one of an employment spell rather than a job with the same employer. In effect, we focus on employment stability rather than job stability. An additional issue with the data is that many respondents reported interruptions of less than 3 months as a break. This may be because a change of employer was interpreted as an interruption by workers, despite the lack of a 3 month break. These data were left in the study as we are primarily concerned with self-reported employment spells rather than interruption length.

In the GSS, respondents are asked to retrospectively identify up to 5 work episodes over their life course starting with the first. They are asked, using the 15th cycle questionnaire, “How old were you when you first started working for a period of six months or longer?” (WH01_Q20A) followed by “How old were you when your FIRST absence from work started?” (WP01_Q50A) with an absence defined as having “been away from work for more than 3 months because of a lack of work,

sickness, maternity/paternity leave, retirement, or any other reason?" (WP01_Q45). This information allows us to extract details on the duration of employment for individuals. Exit from employment is identified by the respondent providing an exit date. The year of retirement is used as an end date if one is not otherwise provided. In other cases, when respondents refuse to answer, don't know or are not asked the data is dropped. Furthermore, any employment spells of immigrants before immigration are dropped. Data for ongoing spells, ones which have started but not ended, are retained for testing if the spell length is less than 60 years in duration.

Using the information above employment spell durations were constructed by obtaining the year in which a particular spell ended and subtracting it from the year it began. We consider employment spells for all employees, whether full-time or part-time, the self employed as well as unpaid family workers. One of the difficulties, encountered with retrospective surveys is that respondents will often not be able to remember the exact timing of an event, particularly with the passage of time. Hence, monthly data has a tendency to occur in unrealistic patterns for particular months. Ureta (1992) found that the type of question asked results in different answers, questions regarding the length of work suffer from less heaping than those regarding when work started. To preclude these so-called heaping effects, employment durations were aggregated to an annual level. Employment begun in a particular year was recorded as beginning in that calendar year while employment ending in a particular calendar year was recorded as ending in the subsequent calendar year. Hence, a job which began in June of 1987 and ended in August of 1987 would be counted as beginning in 1987 and ending in 1988, producing a duration of 1 year. Farber (1995), in his study of job durations in the U.S. has similar problems with monthly data but opts to use a correction. Because we will be focussing on a much longer timeframe the loss of monthly information will not significantly affect our conclusions.

Forward telescoping and memory effects are two serious recall problems that can emerge with retrospective data. Descriptions of these effects can be found in Torelli and Trivellato (1993). Forward telescoping occurs when respondents remember past events as occurring nearer the time of the survey than they actually had. While, memory effects occur when respondents forget that an event had actually taken place. The first effect would make ongoing employment spells seem shorter than

they actually were. As for completed spells, it would depend on the time pattern of the telescoping error. Interestingly, constructing retention rates from LFS data would be subject to the same type of telescoping errors because the survey asks for time workers were in their current position. The second error is more serious and difficult to adjust for; if workers forget an employment spell entirely (likely an earlier and shorter spell) it would bias down the number of short duration spells. In a subsequent paper we will address this issue and generate retention rates to compare GSS results to the LFS. Despite these caveats, it can be readily identified from Figure 2 that for the major change in durations occurring between spells begun in the 1960s and 1970s to be rejected we would require some form of unrealistic nonlinearity in memory effects.

To test for telescoping errors the dataset is broken down into the 15th cycle and 20th cycle of the GSS. As both are representative of the Canadian population at their respective points in time we can test whether the five year gap between them has caused people to recall the year their job began differently. Specifically, for forward telescoping, we test, for each year of birth, whether the year the first job begins is significantly further along the time continuum for those surveyed in 2006 than for those surveyed in 2001. This is done by gender, by class of worker and, for the employed, by full-time and part-time work status. No clear pattern emerges from this series of tests, signifying that forward telescoping is not a serious concern, at least over that 5 year interval.

4. RESULTS

In this section we will present the results of the estimation by ten year employment cohort. To determine whether employment spell durations, have declined in length we select a ten year grouping of spells which start in a particular decade. This is far from an obvious choice, as a birth cohort approach would call for a grouping of workers who were born at a particular time. But to reflect the focus on work rather than workers the cohort grouping uses the employment spell as the unit. This allows for the same individual to appear multiple times in the same or different cohorts if his employment spells began within or between cohorts. This method is a natural remedy to the problem of length bias mentioned above. It can also adequately capture the period effects that can answer our question of interest. Furthermore, it ensures our testing regions will not contain any ongoing employment spells.

When comparing empirical CDFs the use of employment cohorts, birth cohorts or even generations yield very similar results. To ensure that the results are robust to the cohort selection length alternate lengths at various intervals were constructed. The results yielded broadly similar conclusions to the ten year employment cohort concept chosen. The cohort length was selected because it allows for an adequate number of observations to occur in the earliest usable cohorts. This is important because there is currently no information available on these groups in the existing literature. Furthermore, the concept of decades lends itself well to discussion of the topic at hand.

When performing the finite mixture analysis and dominance testing, complete distributions could not be analyzed. This is because there were a number of respondents with ongoing employment spells who began working in those decades. Hence, the upcoming analysis will focus on the portion of the distribution where all employment spells are observed. That is, for the 1980-89 cohort observed spells only up to 13 years can be tested. This is because anyone who began employment in 1989 and not completed after 13 years would still be employed after the 2001 survey was conducted. Hence, for the 1970-79 cohort we have 23 testable years and so on and so fourth until the 1950-59 cohort where we would have 43 testable years. In both, the finite mixture analysis and the dominance testing we, however account for all observed data. Therefore, for the mixture analysis we report the shares of the observations that were not accounted for in the analysis (the incomplete spells) and, for the dominance testing we account for censored observations (we used all the observed data, but censor the incomplete spells using cut off points that match the 13, 23 and 33 year testing regions).

4.1. Distribution Analysis. Kernel density plots reveal that the employment duration distributions by cohort are multimodal. One prominent hump always occurs near the beginning of the distribution with a second significantly farther out. This characteristic pattern shows the so-called movers and stayers respectively. The movers are those who had short durations while the stayers had lifetime work, which lasted 20 years or more. As can be seen in Figure 1, the distributions shown often had complex shapes.

[Figure 1]

4.1.1. *Identifying Worker Types.* Worker types were identified using finite mixture decomposition for various truncated portions of the distribution. Tables 1, 2 and 3 provide a decomposition of the conditional distributions by cohort for the 13, 23 and 33 year testing regions by gender. The 23 and 33 year testing region will allow for an identification of actual movers and stayers while the 13 year testing region decomposes the former into earlier or later movers.

Operationally, worker types are identified by their contribution to the mixture of log-normal distributions which forms the total employment duration distribution. Each type of worker has their own share of the total distribution, showing how common they are and their own mean duration describing their average tenure. The standard deviation is presented to show the breadth to the variation within the group but has no immediate interpretation on worker types.

In Table 1 we consider an analysis of cohorts with durations of at most 33 years. The durations presented are for the 1950-59 cohort and the 1960-69 cohort. The proportion of long duration spells of twenty years or more increased between the 1950s and 60s for both men and women. It should also be noted that the share of shorter-term spells declines for both genders between the 1950s and 1960s, however the mean duration of these spells also fall. For example, note the share of the type I men decreases from 45% to 43% and their mean durations fall from 4.6 to 3.9 years of continuous employment between the 1950-59 and 1960-69 cohorts. This is indicative of “early movers” becoming slightly rarer but leaving employment earlier.

[Table 1]

The degree of heterogeneity, as represented by the number of types of workers stays constant for men. For women the degree of heterogeneity declines in the 1960s cohort, but the share of employment durations greater than 20 years increases. On the whole, women have a much larger share of leavers than do men. In both cohorts at least 70% of women had spells with a mean duration below 10 years. For men, only about 45% of the 1950s and 1960s cohorts showed a mean duration below 10 years. This difference is largely a reflection of women’s role in family formation.

In Table 2 we present the shares of ongoing spells for the two cohorts described above. The results show that for both groups (males and females) the share of workers with ongoing tenures greater than 33 years decreased for those beginning work in the 1960s compared to those who began in the 1950s.

[Table 2]

In table 3 we present the same results but for a 23 year testing region for all but the 1980-89 cohort. The proportions of worker types and mean tenures stay somewhat constant for men but another worker type emerges for the 1970-79 cohort. For women, the mean tenures and proportions change in no discernable pattern and the number of types remains constant at 2. One interesting development in this group is the large share of women working for a short duration in the 1960-69 cohort.

[Table 3]

In Table 4 we present the shares of ongoing spells for these three cohorts. The results show that for males there is a large drop continuing employment of greater than 23 years, especially for the 1970-79 cohort. For females, the share of ongoing spells are almost the same for the 1950-1959 and 1960-1969 cohorts, but are begin decreasing for the 1970-1979 cohort.

[Table 4]

In Table 5 we present the results obtained from all cohorts up until 1980-89. Here we condition the durations to be at most 13 years in order to remove all ongoing spells. The analysis reveals that an increase in the proportion of men who were employed for a very short period emerged between the 1950s and 1970s cohorts, and then declined for the 1980s cohort. This decline in short term employment for men who began work in the 1980s was coupled with a declining mean tenure for both shorter and longer-term workers with durations less than 14 years. That is, the length of tenures for early and late movers in the 1980s were shorter than in previous decades. The number of types for men did not change over the decades for this portion of the duration distribution. For women the mean tenures remained relatively stable across cohorts. For the 1960s cohort it should be noted that only 1 type emerged. These results are consistent with the 23 and 33 year results.

[Table 5]

In Table 6 we present the share of ongoing spells for the four cohorts described above. The results show that for males there is a continuous tendency to have fewer longer duration employment spells over the observed cohorts. For females, the long duration spells show almost no change until the 1970s cohort, after which they decline substantially. This change may be related to the expansion of maternity benefits which occurred as part of the unemployment insurance changes instituted in the early 1970s.

[Table 6]

Both men and women exhibit no clear increase or decrease of heterogeneity over time. Men do display an increase in the 1970s cohort, but only for the 23 year decomposition. While the proportion of early leavers generally holds stable for men, their mean tenure tends to decline secularly. For women, the share of early leavers declines and the mean time before departure falls in the 13 year decompositions. Additionally, for women the share of later leavers also tends to increase over time, reflecting the increasing age at which women had children.

Thus, there are a few indications which point to increasing numbers of shorter duration worker types as well as increases in longer-type workers within the 13 year distribution. This fits with Green and Riddell's assertion of a hollowing of the duration distribution, but over a longer timeframe. However, the best evidence for decreasing tenure does not come from the decomposition but rather from the increasing share of total spells within the sample which are of a short duration.

It has been shown that changes in proportions among the completed employment spells have occurred. However, the changes among the types have not shown any unambiguous movement towards shorter spell lengths. From Tables 2, 4 and 6 it can be seen that many employment spells were still ongoing at the time of the survey. Among men, these proportions were quite high. Including these ongoing spells in the analysis will allow us to better judge whether the the overall distribution of spell durations has changed over the latter half of the twentieth century. To capture this effect the complete distribution must be examined without truncation.

4.1.2. Testing Employment Duration. To test the hypothesis of decreasing employment tenure we focus employment duration distributions by cohort. Stochastic dominance testing was performed at intervals of 13, 23 and 33 for successive employment cohorts constructed from a mixture of the worker types defined above. A cascading format was used to maximize the possible testing regions across cohorts. To make the test valid for the entire distribution all continuing jobs were allowed to remain in the density functions. Hence, for cohort 4 (1980-1989) only 13 years were tested and all ongoing employment spells were censored at 14 years. This was done to weight the testable portion of each distribution distribution appropriately downwards on the right. The same logic applies to all other cohorts with the exception that the testing period is extended by ten years each time we move back one cohort.

[Figure 2]

To understand the concept being employed it is useful to look at the empirical CDFs of all employment durations for men, as shown in Figure 2. The earliest cohorts have a somewhat typical shape. As the cohorts advance the last ten years of each distribution begin to slope sharply upwards. This is the result of a large number of incomplete ongoing durations. These portions of each distribution will not be tested but are necessary for an accurate comparison of completed spells. For men we can clearly see an increasing share of shorter term employment spells over time. But the most interesting result is the width of the gap between the 1960s and 1970s. Even if we impose some reasonable corrections for recall errors, this outsized gap will clearly remain.

[Table 7]

To test the changes more formally we use tests of stochastic dominance. The results shown in Table 7 underscore that there has been a shift over time to shorter average employment durations. The largest change appears to have occurred between the 1960s and 1970s. This led to a significantly larger share of shorter-term employment spells after the 1960s. Formal tests of stochastic dominance back this claim, clearly showing the distributions of cohorts 1 and 2 dominating cohorts 3 and 4 over all testable intervals. In general, each cohort dominates its successor for men.

[Figure 3]

For women the empirical CDFs in Figure 3 appear to have far less dramatic differences. Most noteworthy is the fact that a large share of employment spells, no matter the decade in which they began, are shorter in length than those of men.

[Table 8]

More formally, the results of stochastic dominance tests for the employment durations of women, seen in table 8, yield significantly different results than those of men. Clearly, the results show that cohort 4 is dominated by cohorts 1 at order 1 and by cohorts 2 and 3 at order 2. Cohort 3 is dominated by cohorts 1 at order 1 and by cohort 2 weakly at order 2. Cohort 2 is dominated weakly by cohort 1 at order 2. Hence, there are some indications that women who began work in the 1980s were more likely to have shorter spells than women in previous decades. This could be the result of more generous maternity leave provisions in the UI system or

because of the 1990s recession cutting short their tenures. Hence, for women there were some, but not necessarily secular, declines in employment tenures. This long-term divergence between the experience of women and men was also documented by Farber (2007) for the U.S.

5. CONCLUSIONS

The results of the distribution analysis showed that the heterogeneity of worker types has not increased or decreased substantially. Furthermore, decreases in the mean tenures of shorter-duration worker types have occurred when considering durations up to 13 years. The analysis of heterogeneity did not however lead to conclusive evidence on declining tenures.

Tests of stochastic dominance on successive employment cohorts starting from the early 1950s confirm that the distributions of completed employment spells have changed. Distributions have shifted towards an increasing concentration of shorter-term employment spells among men while for women the results were more mixed. The largest shift in distributions occurred between the 1960s and 1970s cohorts, perhaps as a result of increasingly generous unemployment insurance coverage or the often mentioned subsequent breakdown of employer-employee relationships in the 1980s and 1990s.

Although the use of retrospective data introduces some uncertainty into the results, the paper generally matches existing literature well. Moreover, the results conform to the popular consensus of a decline in employment stability. The use of this data thus allows us to provide evidence towards the resolution of a discrepancy in the literature that could not otherwise have been uncovered with more conventional data.

Further work in this area can easily be undertaken and can greatly enhance our knowledge of employment stability. A focus on individuals rather than jobs or employment cohorts could be undertaken to help understand some of the contributing factors to what makes a worker a mover or a stayer. Furthermore, an analysis of employment scaring can be performed on individuals who begin working during recessionary periods. The dataset is also amenable to work on better determining the effects of telescoping and recall errors in retrospective data. All these areas are interesting in their own right and can make significant contributions to the field. Furthermore, a more theoretical job search model of worker types can be built to

see if unemployment insurance effects can be used to model the shift in distributions between the 1960s and 1970s.

6. REFERENCES

Abramson, I. S. (1982), "On Bandwidth Variation in Kernel Estimates A Square Root Law", *Annals of Statistics* 10(4), pp. 1217-1223.

Blumen, I., M. Kogen and P. McCarthy (1955), "The Industrial Mobility of Workers as a Probability Process," Vol. 6, *Cornell studies in industrial and labour relations*, Cornell University, Ithaca, NY.

Bowman, A. W. and Azzalini, A. (1997), "Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations," Vol. 18 of Oxford Statistical Science Series, Oxford University Press, Oxford, UK.

Brochu, P. (2008) "Rising Job Stability in the 1990s: the Impact of Compositional Change," University of Ottawa Working Paper.

Diebold, F. X., D. Neumark and D. Polsky (1994), "Job stability in the United States," NBER Working paper no. 4859.

Fallick, B. C. (1998) "A Review of the Recent Empirical Literature on Displaced Workers", *Industrial and Labor Relations Review*, Vol. 50, pp. 5-16.

Farber, H. S. (1995) "Are Lifetime Jobs Disappearing? Job Duration in the United States: 1973-1993," NBER Working Paper No. W5014.

– (1999) "Mobility and Stability: The Dynamics of Job Change in Labour Markets," in Orley C. Ashenfelter and David Card eds., *Handbook of Labour Economics* Volume 3B, Elsevier Science B.V.

– (2007) "Is the Company Man an Anachronism? Trends in Long Term Employment in the U.S., 1973-2006," Princeton University Industrial Relations Section Working Paper #518.

Green, D. A., and W. C. Riddell (1997) "Job Durations in Canada: is Long Term Employment Declining?" in *Transition and Structural Change in the North American Labour Market*, ed. Michael G. Abbott, Charles M. Beach, and Richard P. Chaykowski (Kingston, ON: IRC press).

Hall, R. E. (1982) "The Importance of Lifetime Jobs in the U.S. Economy," *American Economic Review* 72, pp. 716-24.

- Hamermesh, D. (1989) "What do we Know About Worker Displacement in the U.S.?" *Industrial Relations* 28, pp. 51-59.
- Heisz, A. (2005) "The Evolution of Job Stability in Canada: Trends and Comparisons with U.S. Results," *Canadian Journal of Economics*, vol. 38(1), pp. 105-127.
- Huynh, K. P. and M. C. Voia (2008) "Finite Mixtures and Stochastic Dominance: Estimation and Inference" Carleton working paper.
- Ignaczak, L. and M. C. Voia (2009) "Censored Distributions and Stochastic Dominance: Estimation and Inference" Carleton working paper.
- Jaeger, D. A. and A. H. Stevens (1999), "Is Job Stability in the United States Falling? Reconciling Trends in the Current Population Survey and Panel Study of Income Dynamics," *Journal of Labor Economics* 17, pp. S1-S28.
- Linton, O., Maasoumi, E. and Whang, Y. (2003), "Consistent Testing for Stochastic Dominance under General Sampling Schemes," Manuscript.
- McLachlan, G., and D. Peel (2000), *Finite Mixture Models*, Wiley, New York.
- Pagan, A. and Ullah, A. (1999), *Nonparametric Econometric, Themes in Modern Econometrics*, Cambridge University Press, New York, USA.
- Picot, G. , A. Heisz, and A. Nakamura (2001) "Worker Mobility, Hiring and the Youth Labour Market During the 1990s," Analytical Studies Branch Research Paper Series No. 155, Statistics Canada.
- Rose, S. J. (1995), "Declining Job Security and the professionalization of opportunity," National Commission for Employment Policy Research report no. 95-04.
- Silverman, B. W. (1986), *Density Estimation for Statistics and Data Analysis, Monographs on Statistics and Applied Probability*, Chapman and Hall, London, UK.
- Swinnerton, K. A., and H. Wial (1995) "Is Job Stability declining in the U.S. economy?" *Industrial and Labor Relations Review* 48, pp. 293-304.
- Ureta, M. (1992) "The Importance of Lifetime Jobs in the U.S. Economy, Revisited," *American Economic Review* 82, pp. 322-34.

7. TECHNICAL ANNEX

7.1. Heterogeneity Testing. The adaptive kernel density estimator which utilizes a variable bandwidth. Varying the bandwidth along the support of the sample data reduces the variance of the estimates in areas with few observations, and reduces the bias of the estimates in areas with many observations.

The estimation procedure follows two steps: the first step computes an initial (fixed bandwidth) density estimate to get an idea of the density at each of the data points, and in the second step, this pilot estimate is used to adapt the size of the bandwidth over the data points when computing a new kernel density estimate.

Kernel density estimates are not unbiased, they are asymptotically biased, with a bias varying with the bandwidth and the shape of the true density function. For a given bandwidth, the bias does not tend to 0 as the sample size increases, therefore we should be careful about any inference that we make using this approach.

To estimate the densities using an adaptive kernel we use Abramson (1982). Let Y_1, Y_2, \dots, Y_n be *iid* random variables with continuous distribution function $F(y) = \Pr(Y_i \leq y)$.

The estimator constructs the local bandwidth h_i as a product of an estimated local bandwidth factor λ_i and a fixed bandwidth h at each sample point ($h_i = \lambda_i h$). The local bandwidth factor stretches or shrinks the bandwidths to adapt to the density of the data, while the fixed bandwidth controls for the overall degree of smoothing.

$$\widehat{f}_{h_i}(y) = \frac{1}{\sum_{i=1}^n w_i} \sum_{i=1}^n \frac{w_i}{h_i} K\left(\frac{y - Y_i}{h_i}\right)$$

where Y_i are data points associated to the weights w_i and $K(u)$ is the kernel (window) function.

The kernel function is a weight function that puts different weights on different points. Typically, it puts more weight on points near y and the weights decline as Y_i gets farther away from y . “Near” and “far” from y is determined by the bandwidth parameter h_i . The local bandwidth factors are proportional to the square root of the underlying density functions at the sample points:

$$\lambda_i = \lambda(Y_i) = \left(\frac{G}{\widetilde{f}(Y_i)}\right)^{0.5}$$

where G is the geometric mean over all i of the pilot density $\widetilde{f}(Y)$. The pilot density estimate is the kernel density estimate with fixed bandwidth h .

One can construct bands around the estimated density functions using the fact that the variance of the adaptive Kernel density estimator can be expressed as

$$V(\widehat{f_{h_i}}(y)) = \left(\sum_{i=1}^n \frac{w_i^2}{n^2} \right) \frac{f(y)}{h_i} \int (K(s))^2 ds.$$

7.2. Stochastic Dominance Testing. The next subsections are extensions to stochastic dominance tests as developed by Linton, Maasoumi and Whang (2005) to allow for censored distributions as in Ignaczak and Voia (2009).

7.2.1. *Testing $H_0^{(1)}$ vs $H_1^{(1)}$.* Considerations in this subsection are based on an extension to the classical Komogorov-Smirnov test by allowing for censored distributions. Namely, with the help of the parameter

$$\kappa = \sup_{y \leq y_{max}} \left| F^{(C_{t_i})}(y) - F^{(C_{t_j})}(y) \right|,$$

we rewrite the null and the alternative hypotheses under consideration as follows:

$$H_0^{(1)} : \kappa = 0 \quad \text{vs} \quad H_1^{(1)} : \kappa > 0. \quad (7.1)$$

An estimator of κ can be defined by

$$\widehat{\kappa} = \sup_{y \leq y_{max}} \left| \widehat{F^{(C_{t_i})}}(y) - \widehat{F^{(C_{t_j})}}(y) \right|,$$

where $\widehat{F^{(C_{t_i})}}(y) = \frac{1}{n} \sum_{i=1}^n 1 \{Y^{(C_{t_i})} \leq y\}$ and $\widehat{F^{(C_{t_j})}}(y) = \frac{1}{m} \sum_{j=1}^m 1 \{Y^{(C_{t_j})} \leq y\}$ are the corresponding empirical distribution functions. The estimator $\widehat{\kappa}$ is consistent. Based on its asymptotic distribution we obtain that

$$\widehat{K} = \sqrt{\frac{nm}{n+m}} \widehat{\kappa}$$

is an appropriate statistic for testing the null hypothesis of equality of distributions (EoD) $H_0^{(1)}$ against the alternative $H_1^{(1)}$ of first order stochastic dominance (FOSD). Here n and m and sample sizes for the two distributions. The corresponding rejection (i.e., critical) region is $R : \widehat{K} > k_\alpha$ and the acceptance region is $A : \widehat{K} \leq k_\alpha$, where k_α is the critical value. Under the presence of the nuisance parameters, the k_α -critical value is not distribution free, and is estimated using a parametric bootstrap. The steps of implementing the bootstrap are outlined in subsection 2.3.

7.2.2. *Testing $H_0^{(2)}$ vs $H_1^{(2)}$.* Considerations in this subsection follow those in Linton, Maasoumi and Whang (2005) and allows for censored distributions. Namely, with the help of the parameter

$$\delta = \sup_{y \leq y_{max}} \left(F^{(C_{t_i})}(y) - F^{(C_{t_j})}(y) \right),$$

we rewrite the hypotheses $H_0^{(2)}$ and $H_1^{(2)}$ as follows:

$$H_0^{(2)} : \delta = 0 \quad \text{vs} \quad H_1^{(2)} : \delta > 0. \quad (7.2)$$

The empirical estimator of δ is given by

$$\widehat{\delta} = \sup_{y \leq y_{max}} \left(\widehat{F^{(C_{t_i})}}(y) - \widehat{F^{(C_{t_j})}}(y) \right).$$

The estimator $\widehat{\delta}$ is consistent. Therefore,

$$\widehat{D} = \sqrt{\frac{nm}{n+m}} \widehat{\delta}$$

is an appropriate statistic for testing the null hypothesis $H_0^{(1)}$ against the alternative $H_1^{(1)}$. The corresponding rejection (i.e., critical) region is $R : \widehat{D} > d_\alpha$ and the acceptance region is $A : \widehat{D} \leq d_\alpha$, where d_α is the α -critical value of the maximum of a Gaussian stochastic process Γ which depends on both distributions $F^{(C_{t_i})}$ and $F^{(C_{t_j})}$. Since the distributions are not, in general, identical, the critical value d_α is not distribution free and has to be therefore estimated. For this we can use a parametric bootstrap method as in subsection 2.3.

7.2.3. *Testing $H_0^{(3)}$ vs $H_1^{(3)}$.* As before, we use Linton, Maasoumi and Whang (2005) and allow for censored distributions (see Ignaczak and Voia (2009)). Hence, if $F^{(C_{t_i})}$ second order stochastically dominates (SOSD) $F^{(C_{t_j})}$, then the parameter

$$\tau = \sup_{y \leq y_{max}} \left(G_2^{(C_{t_i})}(y) - G_2^{(C_{t_j})}(y) \right)$$

is strictly positive. Therefore, we shall test the null hypothesis using

$$H_0^{(3)} : \tau = 0 \quad \text{vs} \quad H_1^{(3)} : \tau > 0. \quad (7.3)$$

that one of the distributions SOSD another one.

Define an estimator of τ by

$$\widehat{\tau} = \sup_{y \leq y_{max}} \left(\widehat{G_2^{(C_{t_i})}}(y) - \widehat{G_2^{(C_{t_j})}}(y) \right),$$

The estimator $\hat{\tau}$ is consistent and we have that

$$\hat{T} = \sqrt{\frac{nm}{n+m}} \hat{\tau}.$$

The corresponding rejection (i.e., critical) region is $R : \hat{T} > \theta_\alpha$ and the acceptance region is $A : \hat{T} \leq \theta_\alpha$, where θ_α is the α -critical value of a distribution that depends on a transformation of $F^{(C_{t_i})}$ and $F^{(C_{t_j})}(x)$. Hence, θ_α is not distribution free and has to be estimated. For this we can use a parametric bootstrap method as in subsection 2.3.

7.3. Parametric Bootstrap. The parametric bootstrap proposed by Huynh and Voia (2008) is used to simulate the critical values for the EoD test in the following fashion:

- (1) Sample n -values from $Y_1^{(C_{t_i})}, \dots, Y_n^{(C_{t_i})}$ from the estimated distributions obtained using Canadian data:

$$\int_0^y \hat{f}_{Duration}(s) ds = \int_0^y \sum_{k=1}^K \hat{p}_k \frac{1}{s \hat{\sigma}_k \sqrt{2\pi}} \exp\left(\frac{-(\ln s - \hat{\mu}_k)^2}{2\hat{\sigma}_k^2}\right) ds,$$

- (2) Then, sample using the cohort i estimated distributions m values from $Y_1^{(C_{t_i})}, \dots, Y_m^{(C_{t_i})}$.
- (3) The distributions are adjusted to be stochastically equal under the null hypothesis (compare the pooled simulated mixtures obtained for the two distributions of interest with the simulated mixture distribution obtained by pooling the data of the two distributions).
- (4) With the use of the resulted empirical distribution functions, $\hat{F}^{(C_{t_j})^*}(y)$ and $\hat{F}^{(C_{t_i})^*}(y)$, define

$$\hat{K}^* = \sup_{y \leq y_{max}} \sqrt{\frac{nm}{n+m}} \left| \hat{F}^{(C_{t_i})^*}(y) - \hat{F}^{(C_{t_j})^*}(y) \right|.$$

- (5) Repeat steps 1-3 B times and define the critical value k_α^* as the smallest value of y subject to at least $100(1 - \alpha)\%$ of the obtained B values of \hat{D}^* are at or below y .
- (6) The rejection region is $\hat{K}^* > k_\alpha^*$.

To estimate the critical values for the FOSD test the same steps are followed as in the EoD case, but construct the estimator

$$\hat{D}^* = \sup_{y \leq y_{max}} \sqrt{\frac{nm}{n+m}} (\hat{F}^{(C_{t_i})^*}(y) - \hat{F}^{(C_{t_j})^*}(y)).$$

The critical value d_α^* is defined as the smallest value of y subject to at least $100(1 - \alpha)\%$ of the obtained B values of \widehat{D}^* are at or below y . The rejection region is $\widehat{D}^* > d_\alpha^*$.

To estimate the critical values for the SOSD test, the same steps are followed as in EoD and FOSD cases, but the following operator is constructed: $\widehat{G}_2^{(C_{t_i})^*}(y)$ and $\widehat{G}_2^{(C_{t_j})^*}(y)$ to define

$$\widehat{T}^* = \sup_{y \leq y_{max}} \sqrt{\frac{nm}{n+m}} \left(\widehat{G}_2^{(C_{t_i})^*}(y) - \widehat{G}_2^{(C_{t_j})^*}(y) \right).$$

The critical value θ_α^* is defined as the smallest value of y s.t. at least $100(1 - \alpha)\%$ of the obtained B values of \widehat{T}^* are at or below y . In this case, the rejection region is $\widehat{T}^* > \theta_\alpha^*$.

8. TABLES AND FIGURES

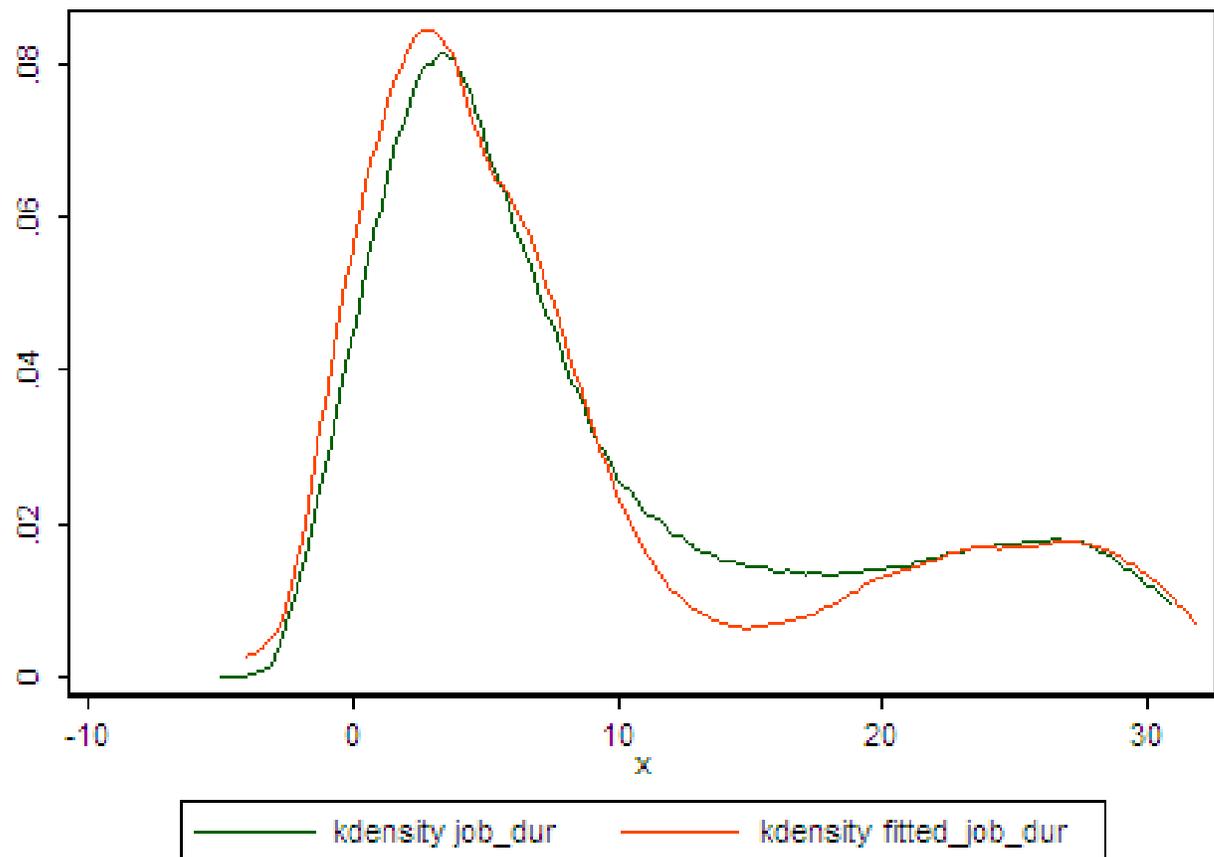


FIGURE 8.1. Adaptive Kernel Density Plot (Total- 1960-69 Cohort)

TABLE 8.1. Tabulated decompositions (by cohort) for complete employment spells of 33 years or less

Cohort	Type	male					female				
		obs	out of:	Share	mean	sigma	obs.	out of:	Share	mean	sigma
1950-59	I	238	858	0.45	4.57	5.09	764	1048	0.73	4.96	4.38
	II			0.3	19.18	5.87			0.18	13.58	5.66
	III			0.25	29.66	2.03			0.09	27.85	3.01
1960-69	I	532	1245	0.43	3.89	3.95	1508	1876	0.82	6.67	7.58
	II			0.24	16.37	5.92			0.18	26.02	4.12
	III			0.34	27.92	3.35			-	-	-

TABLE 8.2. Tabulated shares (by cohort) for ongoing employment spells of 34 years or more

Cohort	male			female		
	obs	out of:	Share	obs.	out of:	Share
1950-59	620	858	0.72	284	1048	0.27
1960-69	713	1245	0.57	368	1876	0.19

TABLE 8.3. Tabulated decompositions (by cohort) for complete employment spells of 23 years or less

Cohort	Type	male					female				
		obs	out of:	Share	mean	sigma	obs.	out of:	Share	mean	sigma
1950-59	I	157	858	0.68	4.54	4.96	685	1048	0.82	4.86	4.17
	II			0.32	16.79	3.73			0.18	13.2	4.48
1960-69	I	344	1245	0.7	4.23	4.45	1269	1876	0.93	5.87	6.1
	II			0.3	15.99	4.12			0.07	19.54	1.71
1970-79	I	1060	2199	0.68	3.9	3.75	2435	3277	0.61	3.92	3.51
	II			0.22	13.33	3.36			0.39	12.8	5.33
	III			0.1	19.87	1.59					

TABLE 8.4. Tabulated shares (by cohort) for ongoing employment spells of 24 years or more

Cohort	male			female		
	obs	out of:	Share	obs.	out of:	Share
1950-59	701	858	0.82	363	1048	0.34
1960-69	901	1245	0.72	607	1876	0.32
1970-79	1139	2199	0.52	842	3277	0.26

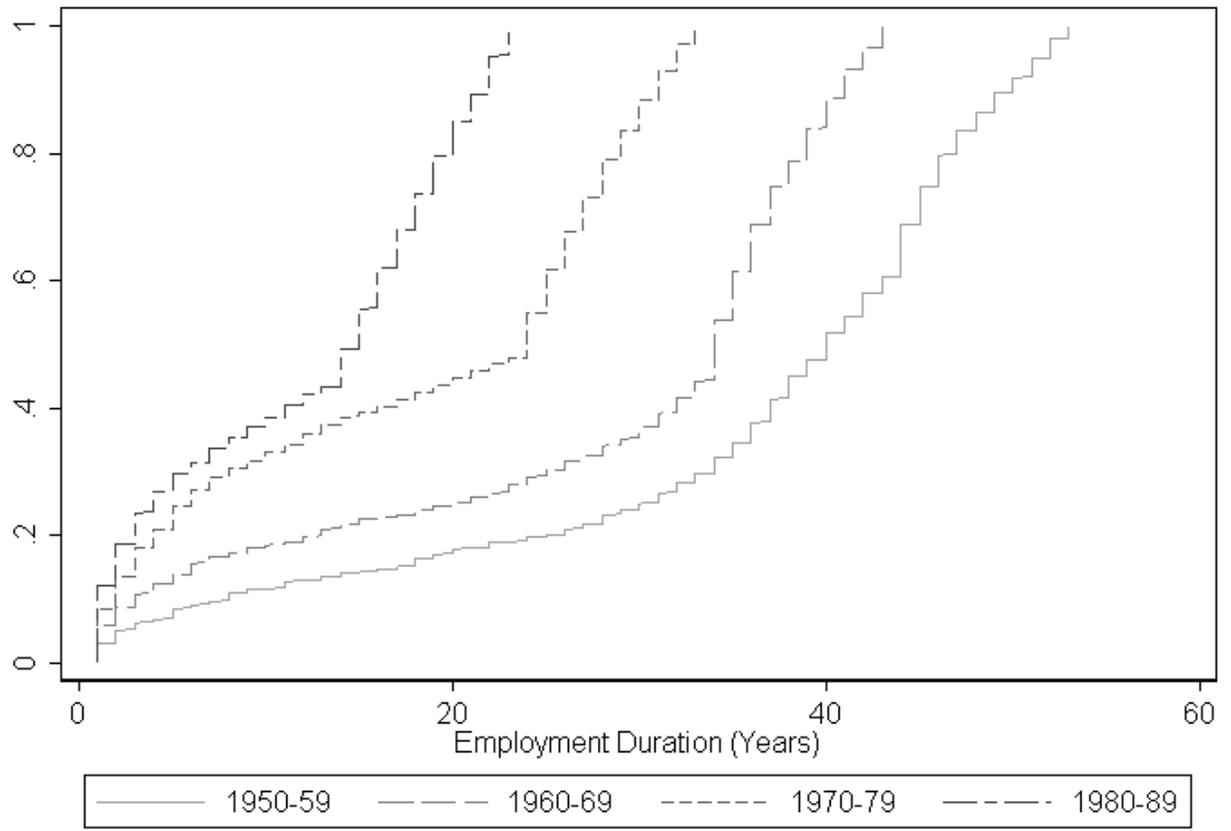
TABLE 8.5. Tabulated decompositions (by cohort) for complete employment spells of 13 years or less

Cohort	Type	male					female				
		obs	out of:	Share	mean	sigma	obs.	out of:	Share	mean	sigma
1950-59	I	116	858	0.74	2.27	2.06	594	594	0.78	3.93	2.92
	II			0.26	9.97	1.24			0.22	9.22	2.45
1960-69	I	266	1245	0.91	2.91	2.31	1063	1063	1	4.86	4.45
	II			0.09	12.59	1.04			-	-	-
1970-79	I	803	2199	0.87	2.78	2.19	1942	1942	0.58	2.87	2.1
	II			0.13	11.37	1.12			0.42	8.5	2.9
1980-89	I	1292	2872	0.69	1.96	1.89	3092	3092	0.53	2.55	1.73
	II			0.31	8.74	1.33			0.47	8.08	2.88

TABLE 8.6. Tabulated shares (by cohort) for ongoing employment spells of 14 years or more

Cohort	male			female		
	obs	out of:	Share	obs.	out of:	Share
1950-59	742	858	0.86	435	1048	0.42
1960-69	979	1245	0.79	791	1876	0.42
1970-79	1396	2199	0.63	1262	3277	0.39
1980-89	1580	2872	0.55	1552	4644	0.33

FIGURE 8.2. Empirical CDF: Men by Cohort



Data source: General Social Survey Cycles 15 and 20

TABLE 8.7. Stochastic Dominance Test Results for Employment Durations by Cohort: Men

	Test Period	Cohort 1	Cohort 2	Cohort 3
Cohort 1 1950-1959		-		
Cohort 2 1960-1969	33 Years	Order 1 (p<0.01)		
Cohort 3 1970-1979	23 Years	Order 1 (p<0.01)	Order 2 (p<0.01)	
Cohort 4 1980-1989	13 Years	Order 1 (p<0.01)	Order 1 (p<0.01)	Order 1 (p < 0.01)

Note: p-values are in parentheses. The table is negatively symmetric, therefore we show only one side of the table. The table should read along the columns.

Considering the first column, the table reads that Cohort 1 FOSD dominates Cohorts 2, 3 and 4 for different testable years periods. It is assumed that a FOSD domination for a longer testable period implies a FOSD dominance for a shorter testable period.

FIGURE 8.3. Empirical CDF: Women by Cohort



Data source: General Social Survey Cycles 15 and 20

TABLE 8.8. Stochastic Dominance Test Results for Employment Durations by Cohort: Women

	Test Period	Cohort 1	Cohort 2	Cohort 3
Cohort 1 1950-1959	-			
Cohort 2 1960-1969	33 Years	Order 2 ($0.1 < p < 0.05$)		
Cohort 3 1970-1979	23 Years	Order 1 ($p < 0.01$)	Order 2 ($0.1 < p < 0.05$)	
Cohort 4 1980-1989	13 Years	Order 1 ($p < 0.01$)	Order 2 ($p < 0.01$)	Order 2 ($p < 0.01$)

Note: p-values are in parentheses. The table is negatively symmetric, therefore we show only one side of the table. The table should read along the columns.

Considering the first column, the table reads that Cohort 1 SOSD dominates Cohort 2, and Cohort 1 FOSD dominates Cohorts 3 and 4 for different testable years periods. It is assumed that a FOSD domination for a longer testable period implies a FOSD dominance for a shorter testable period.