Chapter 8

Sex, age, and transitions - the cohort effect

8.1Introduction

This chapter discusses some characteristics of the participants in the NESPD. Although summary statistics for the NES have been presented in a number of places, the new feature is that the basic unit of analysis is the cohort. Two definitions of cohort are used. The more usual type of cohort, of a group of individuals of the same age followed over their working life, is referred to as the *employee cohort*. In this chapter the *data cohort* is also used, where individuals are grouped according to their first appearance in the dataset. The advantage of the data cohort, constructed from the observation histories, is that patterns of observation can be tracked in detail to illuminate simple counts of numbers employed.

In this chapter the changing nature of the dataset and, by implication, the labour market is examined. As the labour market is followed over time, the individuals who make up the labour market (and population drawings such as the NES) will also change. Even on the simplest assumption that these changes are due simply to ageing (old workers leave, young workers join), the job characteristics for cohorts will still differ from those of the market as a whole. Suppose that wages rise with age. Then the growth in average market wages will understate the total wage growth experienced by individual workers as older, higher-earnings workers are replaced by younger, low-earnings workers.

The situation where the composition of the panel changes on a regular basis with selection and deselection being based on random criteria is called a rotating panel. If the distribution of characteristics of the constituent individuals remains constant as these constituents change then the rotating panel presents few statistical difficulties. An adequate description of the labour market can be achieved by studying one cohort over time or by making simple allowances for the time of entry into the labour market. "Snapshot" descriptions of the labour market will

accurately reflect trends in the labour market. Finally, models of employment, earnings, et cetera may be justifiably simplified by the assumption of time-invariant coefficients.

The assumption of a constant distribution of characteristics does not stand up to a casual analysis of the UK labour market over the period in question. For example, the UK has experienced a shift from manufacturing into services and from full-time to part-time work; union membership has fallen, as has employment in the public sector.

Moreover, the probability of appearing in the panel may not be random. If the experience of unemployment lowers the probability of being employed (Phelps (1972)) then fluctuations in demand can cause persistent differences in the probability of being employed between cohorts. Alternatively, the full-time and part-time participation rates for women have been rising and the full-time employment of men has fallen. Such fluctuations in labour supply may affect both wages and unemployment rates, and could discourage or encourage potential employees in the long term; for example, Dolton and Mavromaras (1994) argued that the choice of a teaching or non-teaching career for two graduate cohorts was significantly influenced by prevailing labour market conditions.

If individuals joining the labour market have different characteristics from their predecessors, two problems arise. Firstly, snapshot descriptions of the labour market may no longer represent dynamic effects accurately; secondly, econometric analyses of the labour market which do not allow for this changing structure are likely to be inefficient and may give misleading or inaccurate results. The aim of this chapter is describe the characteristics of the labour market to enable the validity of simplifying assumptions to be assessed.

The data presented here is used to show the composition of the NESPD, and to draw some inferences on the UK labour market in general. There are two important caveats to this. Firstly, the data on part-timers is suspect, given the NES's presumed ignorance of those earning below the National Insurance (NI) limit. Secondly, while appearance in the NES

implies employment during the survey week, and therefore participation in the labour market, the opposite is not true. Participation in the labour market does not necessitate employment; employment does not necessitate appearance in the Survey. The results of this chapter are mainly compared against Robinson(1994), a recent survey of UK labour market changes, to provide an alternative base for comparison[1].

Because of the enormous amount of information on cohorts available from the micro-data, the results presented here are necessarily selective. Before discussing these results, a brief description of the cohorts is useful.

8.1.1Employment cohorts

Employment cohorts were constructed for individuals born in every year from 1910 to 1975, so covering those who were 65 in 1975 and those who were 15 in 1990[2]. For each year from 1975 to 1990, numbers observed, both in total and in each of the dummy variable categories listed in Table A9.3, were stored for those working part-time or full-time in that year. In addition, gross weekly wages were stored for each observation. Thus, for example, it is possible to tell how many individuals born in 1928 were working full-time in London in 1976 and what their average weekly wage was. The data is split into males and females.

8.1.2Data cohorts: the observation histories

The data cohorts are constructed from OHs for the period 1977-1990. OHs were calculated separately for males and females, and were further broken down into full-time workers, part-

[1] Comparative results are based on alternative datasets. Robinson (1994) primarily uses the population census and the Labour Force Survey (LFS); Elias and Main (1982) uses the National Training Survey (NTS); Main (1988a, 1988b) the Women and Employment Survey (WES); Elias (1988) the LFS and WES. Coverage and survey methods for these sources of information differ from the NES and so provide effective comparators.

[2] Those born outside these dates were treated as being born in 1975 or 1910.

time workers, and those who had experience of both full- and part-time work[3]. Thus this last group, the "cross-over" workers, is likely to share the characteristics of both full-time and part-time workers. However, it is not possible to identify from these OHs the precise split between full-time and part-time working.

For males, the part-time and cross-over OHs are almost empty. Those males who have only full-time employment experience (in the NESPD) account for around 95% of the male sample, and so for the most part the other two male groups are dropped.

One useful side-effect of the OHs is the ability to identify and eliminate individuals with single observations. These tend to have different characteristics from those observed for longer periods, and given the potential for erroneous data entry (which is not corrected retrospectively by the DE) a working assumption is that a significant proportion of these are due to errors[4]. Single observations can account for anything from 5% to 40% of entrants in any one year, but only comprise around 5% of the total sample at any one time. Single observations are therefore generally excluded from the data cohorts.

For each OH, information is available on the age when joining the NESPD, the number of years spent in the private sector or covered by collective agreement or Wages Board, and wages for each year of observation. This information is disaggregated down to the level of observation patterns.

8.2Numbers in the dataset

Figure 8.1 shows numbers observed in full-time and part-time employment using the

[3] The individuals in the third category could not be broken down further without drastically shortening the survey period, because of the number of potential part-time/full-time work patterns.

[4] This assumption is probably unrealistic for the later years of the survey as we may expect more people to return to the NES after 1990. However, the bulk of entry to the survey occurs in the first few years.

employment cohorts, and figure 8.2 the proportion of total observed. Note that, although the NES is intended to be a 1% sample of those in employment, the numbers in figure 8.1 only show around 0.75% of employees as measured by the LFS, for example. This is partly because the NES does not include the self-employed, the armed services, occupational pensioners and so on. In theory, it should also exclude those earning below the NI limit[5]. Finally, the NES does have difficulties in tracing those who have changed jobs around the survey period (Adams and Owen (1988)). Thus it is not entirely surprising that the NES falls short of its target numbers; however, it still remains a significant survey of the labour market and far larger (in terms of participants) than alternative surveys[6].

Part-time employment is rising for both sexes; however, full-time employment is falling for males but rising for females. This is in line with other studies of the labour market (see Johnes and Taylor (1989); Robinson (1994); DE Gazette (DEG) and LFS). Over the survey period, male full-time employees have fallen from 65% of those employed to 55%, the growth in total employment coming from females. Note that, in contrast to LFS and Census data (Robinson (1994)), these figures do not show part-time male employment increasing its share of total employment; this may be due to the NES's difficulties tracking part-timers.

Table 8.1 depicts the numbers of new entrants to the dataset and those making their last appearance. The rise in the numbers leaving at the end of the period is due to the closure of the survey period. The number of males joining and leaving has always exceeded the numbers of full-time females joining and leaving; however, for females the number joining generally exceeds the number leaving whilst for males the reverse is the case. Thus the changing gender composition of the full-time labour force appears to be due to both a high rate of attrition among males in full-time employment and a net inflow of females into all types of jobs. It is

[5] In fact, a notable number do get through the NI contributions barrier, possibly due to inertia in the Inland Revenue records system.

[6] See Adams and Owen (1988) for details of the missing 25% of the target sample. Bell and Ritchie (1994) consider the implications for analysis of this missing data.

worthwhile noting that, since 1981, more females than males appear to be joining the labour force.

Table 8.1 Numbers joining and leaving the NESPD

		1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989
Males	joining	20581	1078	9337	6632	4816	4770	4958	4158	5068	5110	6552	4675
full-time	leaving	7598	8317	9374	7718	7058	6661	6611	6296	7055	6920	9809	12571
Females	joining	6808	4879	4572	3653	2900	3132	3309	3166	3657	4076	5213	3839
full-time	leaving	3553	3872	4281	3407	3222	3105	3144	3005	3227	3198	5240	6138
Females	joining	3730	2842	2485	1792	1448	1453	1645	1386	1852	2073	3053	2091
part-time	leaving	2378	2463	2654	2126	2061	1828	1893	1608	1793	1794	3217	2957
Females	joining	4592	3006	2348	1947	1565	1604	1343	1337	1383	1606	768	407
cross-over	leaving	210	457	650	818	956	1073	1131	1278	1402	1616	2305	3954

8.3Probabilities of observation

Figure 8.3 shows the probability for an individual observed in period t of not being observed in period t+1, calculated from the OHs. The initial leap in the probability of absence is due to an initialisation effect - moving the start date back to 1975 raises the chance of dropping out in 1977 and 1978 (Bell and Ritchie (1994)). As may be expected from empirical patterns of participation for men and women (Elias and Main (1982), Main (1988a), Elias (1988)), males have the lowest probability of becoming absent in the next year. Females employed part-time have highest chance, which is unsurprising given the NI earnings limit and the transient nature and high turnover rate for part-time jobs.

The downward trend in the disappearance rate may reflect a number of factors: improved administration of the NES, or the filtering out of those with poor employment records to give a dataset composed of "stayers". However, the sharp fall in the disappearance rate from 1979 to 1981 is unexpected, given the changes in the labour market. Over this period the unemployment inflow rate increased by 50% while outflows stayed roughly constant (Layard and Nickell (1987) pp132-333) and unemployment rose sharply for both males and females (Robinson (1994)). If the NES represents a truly random sample of the population, a rise in the disappearance rate rather than a fall is the more likely outcome.

In the absence of any information from the DE about improved data collection, the implication is that those in the dataset cut down on changes in jobs. The two biggest causes of missing observations in the NES are employees moving "out of scope" (eg self-employment, armed forces, occupational pension) or the last recorded employer having no record of the employee's current situation (Adams and Owen (1988)). Missing observations due to these causes will all fall if employees change jobs less frequently.

Figure 8.4 depicts the proportion in each year who have held their current post for over twelve months[7]. It is clear that there is a strong cyclical element, in that job changes fall markedly in the early 1980s when unemployment is rising and employment falling, and then rise from the mid-1980s as the economy picks up. Note, though, that although total employment was rising from 1983, unemployment rates did not start to fall until 1986/7 (Johnes and Taylor (1989); Robinson (1994)). Thus the turnover variable appears to be tracking employment rather than unemployment trends[8].

[7] Both those who join a new company and those who change jobs whilst remaining within the same company are deemed by the DE to have "changed jobs". Movements within companies are unlikely to affect the probability of observation.

[8] Turnover according to the LFS also appears to show this pattern. The DE Gazette (December 1989, table 1.6 pS11) records that, apart from a sharp fall in 1975, leaving rates over the period in question were fairly constant apart from a slight downward trend; however, engagement rates followed changes in employment.

These results can be interpreted in the context of a number of hypotheses. For example, a job-search theorist could argue that falling employment levels reduce the expected alternative wage by lowering the probability of finding acceptable job offers, which leads to a reduction in search activity and fewer changes in jobs (see Theodossiou (1992), for example). Alternatively, several models of the two-tier or segmented labour markets postulate an "inner" (insider, career, or primary) part of the labour force which has significant ability to protect its employment position against the "outer" (outsider, non-career, secondary) workers. The apparent increase in average tenure is consistent with this latter group (who, by definition, have poor employment records and prospects) bearing the brunt of employment shocks. Finally, a simple last-in-first-out seniority model will also lead to the result shown in figure 8.4.

Figure 8.5 extends this analysis by another year, giving the observation rates for the four possible combinations of observation patterns in the two years following an observation. Figure 8.5(a) shows that males are most likely to be observed for three consecutive periods, as a proportion of those observed in the first period. For all groups this proportion is rising. Figures 8.5(b) and (c) show that the female part-timers are most likely to have two consecutive observations and then to miss an observation, whereas they are least likely to miss an observation and then return the following year. The cross-over female group is most likely to return to the NES after an absence - consistent with a view that this group is at least partly due to women taking breaks from full-time work to raise families and returning to part-time work[9]. Finally, figure 8.5(d) shows that (ignoring single observations) part-timers have much the highest probability of not reappearing in the dataset in the two years following an observation.

[9] Although only a one-year gap is allowed for here, and the median time out taken for raising families varies from 3.7 to 9.7 years (Dex and Puttick (1988), p136) depending on the birth cohort. An alternative explanation would be that a large part of employment for this group is in temporary work and so is often missed by the Survey, which could be supported by Dex and Puttick's assertion that increasingly women are taking on short-term jobs between births.

Figure 8.5 complements figure 8.3; increasing observation frequency in the former reflects falling disappearance rates in the latter. The near-constant probability of not reappearing in the two years following an observation (figure 8.5(d)) suggests that the improvement in observation in figure 8.3 is due to a fall in the numbers of intermittent missing observations. Again, this may indicate fewer changes in employment, as moves between jobs should manifest themselves as one missing observation at most[10].

It is straightforward to reproduce figure 8.5 for each data cohort but this enormous amount of information is unlikely to be enlightening. Instead figure 8.6 presents, for each group, the numbers observed in each year following the first appearance. Thus, for example, male full-timers have an 80% chance of being observed one year after their first appearance which falls to 50% in the fourteenth year[11]. These are similar to survivor functions except that individuals absent one year are allowed to return in a subsequent year. All four groups follow similar patterns, and the "hazard rate" for both males and females can be shown to be very similar and almost constant for most data cohorts (Bell and Ritchie 1994).

Interestingly, the highest rate of continuing participation is achieved by the female cross-over group. It is also relatively flat, suggesting that the distribution of gaps in the employment record of this group is fairly even. If this group does reflect women moving into and out of part-time employment as family circumstances change (Main and Elias (1986)), then a high reappearance rate is to be expected. This raises the question of why the "survivor function" of female part-timers does not flatten out as well. One possibility is that a significant proportion of those who only work part-time are nearing the end of their working lives and leaving the labour force at a constant rate; age profiles presented in the next section provide some support for this view.

[10] Micklewright and Trinder (1981) comment on the problems caused for the NES by changing jobs near the end of the tax year. As data is not collected or corrected retrospectively, a missing observation is permanently recorded.

[11] These figures are averaged over each cohort in each group, as the disappearance rates are similar for all the data cohorts within a group. Ritchie (1995) presents the disaggregated result.

One result not apparent from figure 8.6 or from employment cohorts is that, for all but the longest periods, females full-timers are more likely than men to have few or many periods in the dataset. Moreover, it appears that, from the end of the 1970s, females joining the dataset are more likely to have a complete, or almost complete, set of observations in the dataset[12]. Equating appearance in the dataset with employment, this result is unexpected. Elias and Main (1982) and Main (1988a) characterise the employment patterns of men and women, using the NTS and WES. These studies indicate that men undertake paid work for more of their life and change jobs less frequently than women (relative to years of participation), which should be reflected in men appearing in the dataset for long periods if the NES accurately represented work patterns.

In this case the NES and alternative studies of the labour market appear to diverge. However, the NES results are conditioned on the fact that these women are only working full-time; those moving between full-time and part-time employment (the cross-over group) have fewer observations than full-time males (Ritchie (1995) table 10.2). Those women with some or all part-time working comprise two-thirds of the female observations, and so the overall effect is that the NES shows women with fewer observations. However, the suggestion that women only working full-time may have better employment records than their male counterparts is an interesting qualification to this result.

Robinson (1994) notes that female employment was hit harder than male employment in the early 1980s recession, but mainly in the manufacturing industries; in services, which weathered the recession relatively well, the impact on female employment was small. The net effect was that the loss in female employment was in low paid, possibly part-time employment. Similarly, several authors (for example Layard and Nickell (1987); Robinson (1994); Sloane

[12] Ritchie (1995). Bell and Ritchie (1993b). using a similar dataset grouped over part-time and full-time workers and including missing observations, argued that women were more likely to have long continuous periods of observation in the dataset and to return to the dataset after a long period of absence.

and Theodossiou (1994)) have noted the increase in demand for female-intensive professional services. The evidence from the NES supports Robinson's claim that "high-ranking" jobs, full-time by assumption, are increasing and improving the prospects for female employment.

8.4Age profiles

The employees in the NES are getting younger, on average. Figure 8.7 uses the employment cohorts to build up a picture of the average age of employees in the dataset, and shows that the average age for both full-time and part-time workers for both sexes is declining over time. This figure masks two opposite influences. Robinson (1994, population census) and Johnes and Taylor (1989, DEG) both note that the participation rate for older males is decreasing but for older females is rising. However, the participation rate for young females is rising even faster (Johnes and Taylor(1989)) and this group is relatively numerous, thus lowering the average age of female employees.

Figure 8.8 displays the age profiles of the dataset in each year for all females and for male full-timers. Figure 8.8(a) appears to show a significant "cohort" effect on males aged 25-30 in 1975. There are significantly higher numbers employed in these cohorts than in either younger or older groups of males. These are the workers who joined the labour force before the early 1980s recession, which would seem to support the Phelps (1972) argument that there is an element of hysteresis in employment.

In contrast, figure 8.8(b), showing the age profiles for females working full-time, indicates a strong "life-cycle" effect: that employment is strongly influenced by family circumstances, including breaks from full-time employment to raise children. Peak employment in full-time jobs occurs between sixteen and twenty, whereupon the numbers employed fall until around thirty; finally, there is some return to full-time employment over the remaining working life, but not to the levels observed for teenagers. There is a wide body of theoretical and empirical evidence supporting this outcome (see, for example, Elias and Main (1982); Joshi (1986);

Main (1988a); Main and Elias (1986)). Note that there is no real cohort effect here; unlike the males, no group of cohorts has consistently higher rates of full-time employment than any other. The counterpoint of figure 8.8(b) is the numbers employed in part-time work, given in 8.8(c). This indicates that full-time and part-time work are complementary for women, as the numbers in part-time employment mirror those for full-timers fairly closely. Again, the theoretical and empirical work on this area would predict this result, but the congruency between figures 8.8 (b) and (c) is pleasing.

The implication of these results is that women's participation in paid work is strongly influenced by the "life-cycle" effects mentioned above; moreover, for women part-time and full-time work are substitutes. In contrast, males are more prone to "cohort effects", where being employed at the right time has a significant effect on future employment prospects. Finally, there is no significant substitution between part-time and full-time work for males[13].

These results are very similar to the "M-shaped" participation patterns for females described by Main (1988a) and Elias(1988) using the WES: women typically work full-time, leave employment in their twenties, and then return to full-time employment by way of part-time employment, with some increase in part-time employment around retirement. What is striking is the complete absence of any change in the structure over time. While Main notes that apparent participation rates in employment have been shifting over time, the proportions of a cohort working full-time has barely changed over three decades (Main, pp43-46), reflected in figure 8.8b. Elias and Main both note some small increase in part-time employment, which is less clear in the NES data; however, figure 8.8c provides strong support for their contention that part-time employment is still dominated by family factors[14].

Figures 8.7 and 8.8 have detailed the "snapshot" distribution of ages for those currently

[13] The profile for part-time male employees is virtually zero except for the very young and very old.

[14] Elias does show that part-time employment in "low grade" occupations has been virtually constant.

working part-time or full-time. Using the data cohorts enables some analysis of individuals based on their whole job history. Figure 8.9 shows the average age at which each cohort joined the dataset, for each data group[15]. The initial fall in the starting ages is probably due to individuals who should have been included in the first cohorts having missing observations and so being included in later cohorts. By the early 1980s the average age of each cohort joining the dataset has settled down at around twenty-five to thirty, with full-timers joining at an earlier age. For female part-timers the starting age is much higher (in accordance with figure 8.8d) and falls steadily[16].

The stability of the starting ages is a promising finding, implying that inclusion in the dataset is primarily determined by the easily measurable characteristic of age (and simplifying the problem of selection bias arising from initial inclusion in the dataset). However, the *level* of joining ages is worrying. These ages are averages over all those joining the dataset in a particular year. On the assumption that most workers begin employment between the ages of 16 and 23 (school- and University-leaving), the evidence of figure 8.9 is that these are outweighed by large numbers of much older individuals also joining the dataset for the first time. While young workers both change jobs more frequently (and thus are less likely to be traced by the NES) and are paid less (and so may fall below the NI limit), these seem poor explanations.

Ritchie (1995) breaks down figure 8.9 into separate starting ages for each year for those remaining in the data cohort, to show how the age profiles of the cohorts change. As would be expected, the average starting age of each data cohort falls as older workers leave the labour force; this is most significant in the data cohorts beginning in the late 1970s. This effect is most significant for part-time employees. The high rate of attrition may be due to the poor

[15] The reason for the concentration on starting ages here is to remove the effect of ages rising over time.

[16] It is not clear whether this reflects the age profile of part-timers in the labour market, or whether this is due to the construction of the dataset (such as the initial-observation effect or the possibility that younger-part-timers are now earning above the NI limit.)

observation records for this group, but is more likely to result from the relatively high age of part-time female employees as the attrition rate does fall with the average age of the cohort.

Interestingly, the cross-over group shows almost no decline in average starting ages over time, indicating that observations of the members of this group are largely unrelated to age. Missing observations appear to be evenly spread over all ages. Again, this is consistent with the view that the cross-over group is composed of women who have taken or will take career breaks. Although there may be some loss of older workers due to retirement, there is a counterbalancing loss of younger women leaving the dataset for family reasons[17].

Figure 8.10 uses the data cohorts to construct an alternative age profile to that of figure 8.7. The contradistinction is that the former classifies females into part-time or full-time employees depending on their whole observation histories rather than a snapshot[18]. There are a number of effects to consider

- the age of those observed rises over time
- older cohorts retire, and younger ones join; the numbers may not be the same
- the average starting age of the data cohorts is constant or falling (figure 8.9)
- within cohorts, older members are more likely to leave (Ritchie (1995))

Figure 8.10 illustrates that the decline in average age for both part-time and full-time females of figure 8.7 has two different components. For those who only work part-time or full-time the average age is falling noticeably, whilst for those who hold both types of job the average age is rising. It has been noted that the age profile of this group is relatively flat, implying that disappearance rates are fairly constant for all ages. Thus these data cohorts age by one year

[17] This argument is consistent with the results in figures 8.8(c) and (d). However, until the cross-over group can be broken down into precise periods of part-time and full-time employment then the hypothesis of young women leaving full-time work and taking up part-time work after a break remains untested.

[18] Figure 8.10 only has figures from 1977 onwards and excludes single observations.

every year, and the effect of workers leaving or joining the dataset is muted as new cohorts replace a range of ages and not just older workers. However, as the average age of those starting is less than the age of those currently in the dataset, this damps down the rise in average ages.

The fall in the average age of full-time and part-time female employees is to be expected given the net inflows into the dataset (table 8.1) and the fact that these new employees are young on average. The age profile for males only working full-time shows little variation over the period, indicating that the various influences are finely balanced: although those currently in the dataset are growing older, this is balanced by the oldest workers leaving the dataset and being replaced by young workers.

8.5Wages and the cohort effect

We now consider levels and growth rates of wages in the NES, adjusted for a real-wage deflator[19]. Figure 8.11 gives the average (adjusted) wage observed for each group in each year, calculated from the employment cohorts. Figures calculated from the data cohorts show very similar patterns. It is apparent that overall the adjusted wages change very little over the period. Within cohorts, wage growth is almost always positive (Ritchie (1995)); this is to be expected, as those who remain in the dataset are likely to be the most successful in their experience of employment.

One interesting feature is that wages for part-time males appeared to rise in the mid-1980s. At a time of rising unemployment among males, the increase in supply of males might be expected to drive down part-time wages (Layard and Nickell (1987)). The rise in real wages and employment (measured by the LFS) suggest that demand for part-time male employees increased by more than supply during the mid 1980s. However, the number of men working

[19] DE Gazette, Table 5.3 January 1979/82/89/91. Figures are for April.

part-time is small and so this result may be an aberration.

In figure 8.12 the age earnings profiles for the NES are presented. These are averaged over all years at the same age, to give, for example, the average wage for a sixteen-year-old in all years. This is a valid simplification because, apart from that of male part-timers, the profiles are astoundingly stable over time; the only significant variation is in the sixty-plus category. This in itself is remarkable, given the strong cohort effect on the numbers of men employed mentioned above; for women, this result is less surprising given the strong evidence for a stable "life-cycle" pattern participation.

The numbers themselves indicate that males have a predictable concave age-earnings profile, with relative wages peaking in the late thirties. For females in full-time employment the profile is also vaguely concave, but peaks in the late twenties. Moreover, for males relative wages stay at a plateau for some years before tailing off; for females relative wages start falling almost immediately the peak has been passed. However, at sixteen both sexes appear to be earning the same wage. From a human capital point of view, the implication is that women acquire human capital at a lower rate throughout their working life; that they invest in less human capital in total; and that relative human capital deteriorates faster than for men. The fact that both sexes appear to start from the same point but diverge immediately supports arguments that attitudes and attributes prior to and on entering the labour market are an important determinant of both initial employment and subsequent work history (Polacheck(1981); Dolton and Kidd(1994); Vella(1994)). It does not support Becker and Lindsay's (1994) claim that female age-earnings profiles should be steeper as women should bear more of the cost of firm-specific investment because of their higher quit rate (although if men and women invest in different amounts of human capital there is an identification problem).

For part-timers, it is interesting to note that both sexes again start work at the same pay. The concave profile for male workers is a little surprising, implying that age and/or experience is an

important determinant of part-time earnings; this contrasts with a theoretical background that part-timers have little or no incentive to invest in human capital (for example, Theodossiou (1992)). The profile for part-time female employees is more in line with theoretical and empirical results (for example, Main (1988b)). There is some wage growth in the teens, possibly indicating some basic training. However, the complete absence of growth (positive or negative) between the ages of twenty and sixty suggests that the benefit from further experience or training runs out very quickly.

As Murphy and Welch (1990) show, the typical age-earnings profile is poorly approximated by the popular quadratic form. In the case of the NES this is clearly the case, and for females any continuously differentiable specification is likely to perform badly. In subsequent chapters a flexible dummy-variable specification has been used.

8.6Sector and agreement

This discussion of the cohort characteristics of the NESPD ends by looking briefly at two other areas of interest. Figure 8.13 shows the proportion of observations in the private sector for 1975 and 1990, from the employment cohorts. Three different trends are visible for the three different groups. For full-time male employees, employment in the private sector in 1990 is substantially higher than in 1975 for all ages. For females, all but the youngest part-timers spent more time in the public sector in 1990 than in 1975, whilst for full-timers there was little change for older age groups but the youngest became more likely to work in the private sector. The overall effect is that the proportion of employment in the private sector has increased between 1975 and 1990, as younger female employees are an increasingly large part of the NES; the lower public sector participation rates for this group reduces the overall proportion of female full-timers in the public sector (see figure 10.9).

Figure 8.14 shows the number of individuals whose wage bargains are affected by national collective agreements (but who are not necessarily union members). For males, the levels of

coverage have fallen fairly uniformly across all ages, suggesting that the decline in union membership throughout the 1980s fell across existing employees rather than new workers joining non-union recognised businesses[20]. For female full-timers, the largest decline in union membership appears to be in the young. For part-timers, there appears to be a small fall in the level of agreement coverage, with the young once again the least likely to be in positions covered by collective agreements. Again, the increasing prevalence of young female full-timers means that overall levels of coverage have fallen (see figure 10.7).

8.7Quasi-complete cohorts

One issue touched upon earlier is how closely the NESPD relates to the UK labour market. As other evidence suggests those who are unemployed or who withdraw from the labour force have different characteristics from those continually employed (eg low wage growth; relatively young or old; working in declining industries, and so on), then a simple test of whether absence from the dataset can be reasonably associated with non-employment may be to compare those who appear in the dataset all the time with those who have occasional absences. The more similar the characteristics of the two, the more likely it is that absence from the dataset is a statistical error rather than a reflection of non-employment.

To this end, similar statistics to those above were calculated for "quasi-complete cohorts" (QCCs): that is, those who had no missing observations between the first observation and leaving the dataset for good. Unfortunately, the results do not clarify the issue. The wage profile of the QCCs is almost identical to that of the full dataset. On the other hand, the QCCs have a higher overall within-cohort wage, and lower total "snapshot" wage. Moreover, the QCCs tend to be much older and have a disappearance rate of around 30%-40% - twice as high as for the full dataset.

[20] This is a qualified result, as the variable in question only checks for national collective agreements.

The overall suggestion is that those with missing observations do perform differently to QCCs; however, this is a very tentative result. The reason is probably that there is not a simple employment-non-employment question: absence from the dataset could have a number of causes, which may all act in different ways and have different effects on those who do not go missing. To identify some of these effects requires rather more information than is available from the NESPD.

8.8Summary

A few aspects of the NESPD have been discussed, and some difference between cohort and dataset characteristics has been noted. Although using these results to make inferences about the labour market as a whole is not necessarily justified, some sensible inferences may be drawn, particularly with respect to the participation of women. Some worries about the composition of the NES have emerged, most notably the high starting ages and the difficulty of determining how closely appearance in the dataset approximates to labour market experience. Many of these results are tempered by the difficulty of making assertions about the cross-over group, and must wait for a new set of OHs allowing for multiple destinations (full-time, part-time, unemployed, and so on).

With respect to transitions, there appear to be significant differences between both gender groups and cohorts. Bell and Ritchie (1993b) constructed hazard and survivor functions from similar observation histories which indicated very little difference between the sexes. This chapter suggest that these results need to be reconsidered in the light of the heterogeneity of the female groupings.

The age profile of the database was shown to vary significantly between the groups: the average age is constant for males, falls for female part-timers and full-timers, and rises for the cross-over group. Age-earnings profiles suggested the concave structure of human-capital theory. However, for women in particular there are clear indications that the returns to human

capital come to an abrupt halt and for part-timers, are negligible.

Some of the other information available was briefly discussed and it was noted that the widely reported fall in union membership over the 1980s appears to be spread over all male full-time workers but is chiefly found amongst the younger female employees.

Finally, a brief discussion of quasi-complete cohorts suggested that conditioning on a complete set of observations prior to finally leaving the dataset may lead to different results compared to the case where the whole dataset (including those who have intermediate missing observations) is used. This implies that selection bias is a significant issue for the NESPD.

Although cohort statistics over dataset figures have been emphasised, this does not imply that the former are necessarily any "better" than the latter. However, the view that some analysis of cohorts is necessary has been reinforced by results indicating that dataset statistics may not adequately reflect the characteristics of individuals. These results are supported in subsequent chapters which argue that the dynamic structure of the NESPD should not be ignored lightly.