CENTRE
for
ECONOMIC PERFORMANCE

DISCUSSION PAPER NO. 380

January 1998

MOVIN' ON UP: INTERPRETING THE EARNINGS-EXPERIENCE PROFILE

A. Manning

ABSTRACT

Human capital theory provides the generally accepted interpretation of the relationship between earnings and labour market experience, namely that general human capital tends to increase with experience. However, there are other plausible interpretations e.g. search models generally predict that more time in the labour market increases the chance of finding a better match and hence tends to be associated with higher earnings. In this paper we show how a simple search model can be used to predict the amount of earnings growth that can be assigned to search with the residual being assigned to the human capital model. We show how a substantial if not the larger part of the rise in earnings over the life-cycle in Britain can be explained by a simple search model and that virtually all the earnings gap between men and women can be explained in this way. Overall, the evidence suggests that we do need to rethink our interpretation of the returns to experience in earnings functions.

This paper was produced as part of the Centre's

Industrial Relations Programme

MOVIN' ON UP: INTERPRETING THE EARNINGS-EXPERIENCE PROFILE

Alan Manning

JANUARY 1998

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

ISBN 0 7530 1206 5

MOVIN' ON UP:
INTERPRETING THE EARNINGS-EXPERIENCE PROFILE

Alan Manning

The Centre for Economic Performance is financed by the Economic and Social Research Council.

MOVIN' ON UP:
INTERPRETING THE EARNINGS-EXPERIENCE PROFILE

Alan Manning

INTRODUCTION

Since at least the work of Mincer (1958) there has been a generally accepted interpretation of the correlation between earnings and labour market experience, namely that it reflects the fact that more experienced workers have larger amounts of general human capital (possibly, depending on the version of the theory, after current investments in human capital have been deducted **S** see Ben-Porath (1967)). So deeply engrained is this belief that apparent increases in the returns to experience observed in the United States (Katz and Murphy, 1992) and the United Kingdom (Schmitt, 1995) are interpreted as evidence of an increase in the returns to human capital and play an important role in the diagnosis of the cause of the rise in wage inequality as being a shift in relative labour demand against the less-skilled. But, there are few direct tests of the hypothesis for the simple reason that independent measures of productivity rarely exist.

This would not be a potential problem if the human capital model was the only plausible explanation for the earnings-experience profile but it is not. Burdett (1978) presents a search model in which the distinctive feature (at the time) was that employed as well as unemployed workers engaged in job search. Although it was not the main subject of his paper, he pointed out that his model predicted that "older workers... receive higher wages rates, on average, because they have obtained more job offers and the more job offers a worker receives, the greater the probability a 'high' wage rate job will be found" (p.219). Some evidence supportive of this view is to be found in the literature on the earnings losses suffered by displaced workers which typically finds that more experienced workers suffer greater wage losses even after controlling for job tenure (Ruhm, 1991; Jacobson, Lalonde *et al*, 1993). There is one part of the existing literature that is worried that estimates of earnings functions might be biassed by search or ‘job-shopping’. Topel (1986), Abraham and Farber (1987), Altonji and Shakotko (1987), Marshall and Zarkin (1987), Topel (1991) and Altonji and Williams (1992) present different ways of trying to derive ‘true’ earnings functions. But this literature is not very satisfactory as the answers obtained seem very sensitive to the precise method used and it has concentrated on possible bias in the estimated returns to tenure when the underlying logic also suggests a potential problem with the estimated return to experience.

In this paper we take a different approach from that literature. As we show below, the shape of the earnings-experience profile predicted by the search model depends on labour market transition rates: the rate at which workers move out of employment, into employment and the rate at which they change jobs. As these transition rates are observable one can compute the earnings-experience profile predicted by the search model and compare it with the actual, interpreting the residual as the part of the profile that can be ascribed to the human capital explanation (or any other theory that takes one’s fancy). What we show, using data from 20 years of the UK General Household Survey and Labour Force Survey, is that the search model can explain a large part of the earnings-experience profile for both men and women, that it can explain all the gap in the profiles between men and women and that the rise in the returns to experience seems to be the result of a rise in the returns ascribed to search as much as the residual.

The plan of the paper is as follows. In the next section we outline a simple search model of the labour market and derive the predicted relationship between earnings and experience. The second section describes the British earnings experience profile and the third through fifth sections then consider how well the model can explain UK data. Finally, we consider some extensions to the model.

1. THE RETURNS TO EXPERIENCE IN A SIMPLE SEARCH MODEL

Let us assume that the labour market has a distribution of wage offers and that firms can be characterised by their position in the wage offer distribution which we will denote by F. We are deliberately vague about the origin of this wage distribution e.g. one could justify it using a model of equilibrium wage dispersion along the lines of Burdett and Mortensen (1989) or a model of rent-sharing in the presence of employer heterogeneity along the lines of Pissarides (1994). But what is important is that there is a wage distribution facing a given worker. In the interests of simplicity we assume that jobs never change their position in the wage offer distribution. One can think of the wage offered by a job at position F in the wage distribution as being given by w(F,a) where a is the labour market experience of the worker[1]. In a pure search model where there is no 'true' return to labour market experience the function w(F,a) will be independent of a.

A couple of points are in order here. First, we will follow the bulk of the empirical literature in measuring a as potential labour market experience i.e. age minus age when left full-time education. Both the human capital and the search models would suggest that 'true' labour market experience would provide more explanatory power: we do not pursue this here as our data sets (in common with most others available to researchers) do not contain such information[2]. Second, we will refer to the dependence of w(F,a) on a as the 'true' returns to experience with the implication that they represent the returns to general human capital. But it should be remembered that there is no direct evidence for this and there are other stories one could tell based on incentives and monopsonistic discrimination that could also explain why firms make wage offers that depend on experience.

Individuals are assumed to enter the labour market at experience 0 and they exit the labour market at a rate $d_r(a)$. We assume that unemployed individuals of experience a receive job offers at a rate $?_u(a)$, and employed workers at a rate $?_e(a)$ and that all job offers are drawn at random from the wage offer distribution. We assume that all job offers are acceptable to the unemployed (assuming heterogeneity in reservation wages would make matters more realistic but much more complicated) but that employed workers only accept a new job offer if it is from a firm that is at a position in the wage distribution above the position of their current firm. This means that the rate at which workers in a firm at position F leave for other firms is given by $?_e$(a)(1-F). We also assume that workers leave employment for unemployment at a rate d_u(a). Because of the existence of on-the-job search the distribution of wages among workers will differ from the distribution of wage offers so let us denote by G(F,a) the fraction of workers of experience a who are employed in firms at position F or lower in the wage offer distribution.

To close the model, we need to make some assumption about labour market entrants. From the theoretical point of view perhaps the easiest assumption is to assume that these workers all initially enter unemployment so that u(0)=1 in which case we will have G(F,0)=F (G should be interpreted as a limit in this case). But, in practice not all labour market entrants are unemployed because of job search before they enter the labour market so that we treat u(0) as exogenously given. However, we maintain the assumption that G(F,0)=F so that the distribution of wages among labour market entrants is the wage offer distribution. The justification for this is that labour market entrants have had so little time in the labour market that the probability of them having received more than one wage offer is negligible. Given this set-up, let us derive G(F,a) for other experience levels.

It is useful to start with unemployment rates. Define U(a) to be the level of unemployment among workers of wage a and u(a) to be the unemployment rate among workers of age a. Obviously u(a)=[U(a)/L(a)] where L(a) is size of the labour force of age a. We must have:

$$\dot{u}(a) = \frac{\dot{U}(a)}{L(a)} - \frac{U(a)}{L(a)}\frac{\dot{L}(a)}{L(a)} = \frac{\dot{U}(a)}{L(a)} + d_r(a)u(a) \qquad \textbf{(1)}$$

As the rate at which workers exit unemployment is $(\lambda_u(a)+\delta_r(a))$ and the rate at which they enter unemployment is $\delta_u(a)$, we must have:

$$\dot{U}(a) = \delta_u(a)[L(a)-U(a)] - [\lambda_u(a)+\delta_r(a)]U(a) \qquad \textbf{(2)}$$

Substituting this into (1) leads to;

$$\dot{u}(a) = \delta_u(a)[1-u(a)] - \lambda_u(a).u(a) \qquad \textbf{(3)}$$

which simply says that the change in the rate of unemployment is the difference between inflows and outflows. One can see from (3) that the labour market exit rate, δ_r, plays no role. As this will continue to be the case in all the distributions derived below (as it is assumed to be the same for all workers, irrespective of employment status or wage), to economize on algebra we will assume it is zero in what follows: this is without loss of generality.

The following Proposition derives $G(F|a)$.

Proposition 1: The distribution function of wages conditional on experience $G(F,a)$ satisfies the following differential equation:

$$\frac{\partial G(F|a)}{\partial a} = -\lambda_e(a)[1-F]G(F|a) + \frac{\lambda_u(a)u(a)}{1-u(a)}.[F-G(F|a)] \qquad \textbf{(4)}$$

which has as a unique solution:

$$G(F|a) - F = -F.\frac{\int_0^a \Gamma(s)\beta(s)ds}{\Gamma(a)} \qquad \textbf{(5)}$$

where:

$$\alpha(s) \equiv \frac{\lambda_u(s)u(s)}{1-u(s)}$$

$$\beta(s) \equiv \lambda_e(s)(1-F) \qquad \textbf{(6)}$$

$$\Gamma(s) \equiv e^{\int_0^s [\alpha(s')+\beta(s')]ds'}$$

Proof: See Appendix A.

One can give a fairly simple intuitive explanation for the differential equation (4). The first term on the right-hand side represents the rate at which currently employed workers receive wage offers which would take them to a firm with a position above F in the wage offer distribution. The second term represents the employees who have come from unemployment. The term $\lambda_u u/(1-u)$ is the number of new entrants relative to the employed. [F-G] represents the difference in the distribution of wages between the new entrants and the existing workers.

Turning to the closed-form analytical solution (5) one can readily check that $G(F|a)$ is a strictly increasing function of F with $G(0|a)=0$ and $G(1|a)=1$ for all a (the last result follows because $\beta(s)=0$ when F=1). For a>0 and 0<F<1 one can check by inspection that $G(F|a)<F$ so that the distribution of workers across firms always strictly dominates (in the sense of first-order stochastic dominance) the distribution of wage offers. Given our assumption that the distribution of wages among labour market entrants is equal to the wage offer distribution this implies that the least experienced workers always do worse then more experienced workers: as we shall see this is consistent with the data.

We might be interested in how the function $G(F|a)$ varies with the labour market transition rates. We can prove the following

Proposition.

Proposition 2:
i)

$$\frac{MG(F^*a)}{Ma(v)} \;'\; 0 \qquad for \quad a\#v$$

$$\frac{MG(F^*a)}{Ma(v)} > 0 \qquad for \quad a>v \qquad \textbf{(7)}$$

ii)

$$\frac{MG(F^*a)}{M\text{ß}(v)} \;'\; 0 \qquad for \quad a\#v$$

$$\frac{MG(F^*a)}{M\text{ß}(v)} < 0 \qquad for \quad a>v \qquad \textbf{(8)}$$

Proof: See Appendix A.
What this proposition says is that anything that increases the arrival rate of job offers when employed will improve the wage distribution of workers and anything that increases the ratio of new entrants to existing workers will decrease it. As we shall see below this result will be of use in explaining why women have lower returns to experience than men as the rate at which women enter employment from non-employment is higher than for men.

So far we have shown that the wage distribution must improve on entry to the labour market. But, one might also wonder if one can prove a stronger result, namely that the wage distribution is monotonic in experience. A sufficient (but by no means necessary) condition for this is provided in the following Proposition.

Proposition 3: If a(a) is non-increasing ß(a) is non-decreasing in a, then MG(F*a)/Ma<0 for all 0<F<1 and for all a$0.

Proof: See Appendix A.

Put in intuitive terms, this means that we can be sure that the wage distribution is improving monotonically with experience if the arrival rate of job offers does not fall with experience and the entrant rate does not rise. This Proposition also provides clues for when the monotonicity might fail. If women of particular experience levels tend to leave employment to have children and then re-enter the labour market at some later date we might expect to see the proportion of entrants rising for some experience levels and hence we might expect to see the wage distribution worsening at that time for the reason that they will tend to re-enter the labour market at low wages.

One might wonder how well a pure search model can explain the broad features of the earnings-experience profile. Burdett (1978) showed, for a model in which ?$_u$=0 and the job offer arrival rate for the employed was constant, that the pure search model predicted earnings to be an increasing, concave function of experience. We have already provided a generalisation of the increasing part of this result in the previous proposition so one might wonder whether one can prove anything about concavity. If a and ß are constant then it is simple to show that the earnings function in the pure search model must be concave. But, in general the model does not predict concavity and, as we shall see, this is very desirable as the empirical earnings profiles are not everywhere concave.

So, let us now move on to the data. We will start by simply looking at some stylized facts about the earnings-experience profile in the UK.

2. THE EXPERIENCE EARNINGS PROFILE IN THE UK

The earnings data we will use in this study come from the UK General Household Survey for 1974–92. From the point of view of studying earnings functions, the GHS has one serious disadvantage, namely that it is not possible to construct a true hourly wage after 1978 because the questions on overtime hours were discontinued. Different researchers have dealt with this problem in different ways. Some have used only weekly earnings as their wage measure and, to avoid the problems caused by the importance of part-time work among women, have restricted attention to men. Others have used weekly earnings and restricted attention to workers who categorize themselves as full-time: this has the disadvantage of ignoring part-time women, a group who are growing in importance. In this paper we use as our earnings measure the usual weekly earnings divided by usual weekly hours excluding overtime i.e. we use an incomplete adjustment for hours. This may seem rather curious but we will argue that it is the best of a set of imperfect alternatives for the following reasons. First, the job search model suggests that we need a ranking of jobs: while it is reasonable to assume that all workers prefer, other things equal, jobs with higher hourly wages, it is not reasonable to think that all workers prefer jobs with longer hours (85% of part-time workers who are not sick or students in the 1995 Labour Force Survey say they do not want a full-time job). So the hourly wage is likely to give a better indication of the relative attractiveness of jobs than weekly wages. The other justification for the use of this earnings variable is that the bias induced by ignoring overtime hours seems to be small. For the years before 1978 when we can compute a ‘true’ hourly wage measure and compare it with ours, the correlation between the two is 0.975 after taking out time effects. The reason for the very close correlation is that overtime hours are a relatively small fraction of total hours (usual hours average about 35 hours, overtime hours under 3) and even then the correlation of overtime hours with hourly earnings is very weak. Hence, we would argue that the advantages in being able to include part-time women workers and doing some correction to weekly earnings for hours variation is better than doing no correction at all even though it would obviously be desirable to have a better hours measure.

Figure 1 presents the earning-experience profile for men and women in a single year; 1983. To show that this is a fairly typical year, Figure 2 presents (on a smaller scale) the profiles for each year. In this and all later work, log wages have been adjusted to be on the same scale by taking out time effects (time being measured by the month) and normalized so that the earnings of labour market entrants are zero. There are several features of these profiles worthy of discussion.

First earnings grow rapidly in the years after labour market entrance and then decline more slowly as is generally drawn. But the quadratic relationship between log earnings and experience that is commonly estimated is not a good fit to the profile. A similar point has been made by Murphy and Welch (1990) for the United States. They found that a quartic provided an adequate fit to the data. If one insists on fitting a polynomial to the data (and inspection of Figure 1 does not suggest that this is likely to be a particularly good specification) this does not seem to be a high enough order for British data; for each year, one typically finds that powers up to the eighth are statistically significantly different from zero. The reason for this difference is that there do seem to be significant differences in the profile between the UK and the US. In the US male earnings seem to peak at around 30 years of experience (a similar finding being reported by Mincer, 1974, for an earlier period). In the UK earnings seem to peak somewhat earlier (at around 15 years of experience for men and 10 for women) and earnings growth prior to that point is very fast.

The second feature of Figures 1 and 2 worth noting is the difference between men and women. On labour market entry the earnings growth of men and women are very similar but the profiles diverge as we increase years of experience, though less so in the later years than the earlier. In addition, the earnings profile for women seems to have a ‘bump’ period from 10 to 15 years of experience when earnings seem to fall very rapidly though this bump is less pronounced in later years. One implication of this is that although the profile is concave for men over the whole range of experience, it is convex over part of the range for women.

These figures do not tell us about the comparison of the level of wages of men and women on labour market entry (as these are normalized to zero for both groups). There was a significant positive difference in mean earnings between men and women on labour market entry in 1974S5 (consistent with the fact that the Equal Pay Act only came fully into force in 1975) but the difference is insignificantly different from zero after that date and actually slightly negative in the later years.

There have also been changes in the returns to experience over time. This is rather hard to see from Figure 2, so Figure 3 presents the earnings for different experience levels relative to labour market entrants over time. Because this is all relative to workers of one experience level we also present in Table 1, average earnings by grouped experience categories. From this table one can see the rise in the earnings of more experienced workers (both male and female) relative to those with 0S5 years of experiences in about 1984[3] and also the rise in female earnings relative to men and the fact that the peak in women’s earnings now seems to be noticeably later in their labour market career.

One can summarize these facts as:

i) For men the profile is concave with very rapid earnings growth in the first years after labour market entry and a gentle decline in the later years.

ii) For women the profile shows less of an increase than for men and has a convex portion with a ‘bump’ at around 10S15 years of experience.

iii) The gap between the male and female profiles has fallen over time.

iv) There has been a rise in the returns to experience.

Armed with these stylized facts, let us now attempt to evaluate how well they can be explained using the search approach.

3. EXPLAINING THE EARNINGS-EXPERIENCE PROFILE

The theoretical discussion above has all been in terms of the distribution of workers of particular experience levels across firms at different positions in the wage distribution. But if we want to discuss how well the theoretical model explains the earnings-experience profile we need to introduce the wages actually paid. As before we denote by w(F,a) the log wage paid to a worker of experience a by a firm at position F in the wage offer distribution. Given this the expected log wage will be given by:

$$E(w|a) = \int_0^1 w(F,a)\,dG(F|a) \qquad \textbf{(9)}$$

From (9) it should be apparent that there are two reasons why earnings might vary with experience. First, there are ‘true’ experience effects, namely that wages offered by firms vary with experience i.e. an effect through w(F,a) and, secondly, there are the effects associated with job search i.e. the effects through $G(F|a)$. There is no reason why the same approach could not be used to look at other moments or percentiles and we do consider the variance briefly below.

The aim of this paper is to try to disentangle the pure experience effect from the search effect. The way we will do this is to write (9) as:

$$E(w|a) = \int_0^1 [w(F,a) - w(F,0)]\,dG(F|a) + \int_0^1 w(F,0)\,dG(F|a) \qquad \textbf{(10)}$$

where we will interpret the first term as the ‘true’ return to experience and the second term (which can be interpreted as the average wage received by workers of experience a if the wage offer distribution remained the same) is the contribution of search. If w(F,a)=w(F)+?(a) then the first term will simply be ?(a): this is the case where the ‘true’ return to experience is the same at all points in the wage offer distribution so that there is a single measure of the 'true' return to experience. Where w(F,a) does not satisfy this condition the formula in (10) estimates the 'true' return to experience is a weighted average using the distribution of workers across the wage offer distribution as weights.

The strategy for the decomposition in (10) is the following. We will compute the second term using information on the wage distribution of labour market entrants[4] to estimate w(F,0) as it is our assumption that G(F*0)=F and estimates of G(F*a) based on labour market transition rates. The first term on the right-hand side of (10) (which we will interpret as the 'true' return to experience) will then be estimated as the residual part of the actual profile not explained by the search model. Note that we cannot estimate w(F,a) directly as what we observe is the distribution of wages across workers not across job offers[5].

At this stage it is worth considering the implications of assuming (as is commonly done) that the wage distribution depends not just on (F,a) but also on job tenure. The important point to note is that this does not alter the interpretation of the second term in (10) as the wage one would expect if there were no true returns to experience or tenure. The reason for this is that labour market entrants must, by definition, have tenure zero so that the wage distribution used to estimate w(F,0) can be interpreted as the wage offer distribution for workers with zero experience **and** tenure. However, the interpretation of the residual in the earnings profile must now be interpreted as the joint effect of experience and tenure and there is no way to disentangle the two components (though see Hartog and Teulings, 1996; Manning, 1997, for such attempts).

Our estimates of the transition rates are taken from the annual UK Labour Force Survey conducted in the spring every two years from 1975 to 1983 and annually thereafter. This contains information on current labour market status and also a retrospective question on labour market status a year ago and whether there has been a change in employer for those who report being in employment in both years. Since the start of the panel aspect of the Labour Force Survey in 1992 one can check the accuracy of the retrospective questions: overall it seems to be of the order of 95%. So, they seem to be fairly accurate.

We used these questions to estimate the labour market transition rates for each year of experience in each year. As the information on labour market status is only available at discrete points in time and the model we have presented has been in continuous time there is an issue of how we infer the underlying transition rates from the data that we have. One possibility is to use the predictions of a continuous time model for the discrete observations to estimate the continuous time transition rates making some assumption about how the transition rates vary as experience varies through the year. For estimating $(?_u,d_u)$ there is no problem with doing this but matters are more difficult for the arrival rate of job offers when employed, the reason being that $?_e$ is the arrival rate of all job offers, not just the arrival rate of better job offers so that one needs to adjust the rate at which workers change jobs for the distribution of workers across firms which will itself change through time. This adjustment is computationally extremely cumbersome in continuous time. So, we adopt a simpler approach and use a discretized version of the model which is described in Appendix B where we also describe how we estimate the transition rates. For $(?_u,d_u)$ we have worked out the difference between this simple method of computing transition rates and the rates derived from a model based on continuos time and the discrepancies are small. The reason for this is that the discrete time approximation will be good if a year is a period of time over which it is rare to have more than one labour market transition and, while there are some people who have a large number of moves (and these people are obviously much more important in statistics based on flows) most people move only rarely. Evidence from the British Household Panel Study for the early 1990s suggest under 2.5% of individuals have more than two labour market states in the course of a year (and this overstates the mobility for our purposes as it distinguishes between unemployment and inactivity). Hence the discrete time approximation is likely to be a good one.

Figure 4 presents information on the transition rate from employment to non-employment for men and women at different experience profiles over the years. The main features of the data are that male employment to non-employment rates are quite constant over the life-cycle though were noticeably higher in the early 1980s. For women the transition rates in the earlier years show a ‘bump’ associated with having children but, by the early 1990s this has all but disappeared and the transition rates are very similar to those for men. Figure 5 presents similar information for the transition rate from non-employment to employment: this tends to decline with labour market experience and be higher for women than men though again the convergence between the sexes over time is very striking. Finally Figure 6 presents information on the job-to-job mobility rate: these seem to always have been very similar for men and women and decline with experience. Note that this job mobility rate is not the same thing as the job offer arrival rate of the model because it is not corrected for the wage offer distribution.

We used these transition rates and the discrete time version of (4) (which is given in Appendix B) to compute G(F*a). In doing this we are making an implicit steady-state assumption. In computing the G(F*a) for a worker with 30-years of experience this year what is relevant is the transition rate for a worker with 29 years of experience last year rather than this year. For some years we obviously do have this information but the biannual nature of the LFS in the early years together with the fact that we only have 20 years of data mean that there would be much information that is missing. The steady-state assumption may not be too bad as one can see from Figures 4S6 that the pattern of transition rates is broadly the same over the sample period but one should be aware of the potential problems caused by it.

To illustrate the predictions of the model Figure 7 presents estimates of G(0.5*a) i.e. the proportion of workers of experience a paid below the median wage for labour market entrants. As can be seen the model predicts a decline for men which is very sharp in the first few years after labour market entry and then flattens off and a ‘bump’ for women caused by their tendency to leave paid employment for domestic reasons.

We now turn to the predictive power of the model. Figure 8 presents the actual earnings-experience profile compared with the predicted from the pure search model for women and men (note that these graphs only show the actual profile in 1976, 1978, 1980 and 1982 as the LFS was not conducted in these years). In order to increase the sample size (and hence the precision of our estimates) the actual earnings distributions is computed using two pooled years of the General Household Survey and the year given in Figure 8 refers to the second of the two years[6]. Everything is measured relative to labour market entrants (as we have nothing which explains the level of their earnings) so the predicted and actual are constrained to coincide for labour market entrants. The overall impression is how much of the rise in earnings is explained by the pure search model. Particularly striking is the fact that the pure search model traces out the ‘bump’ in the profile for women which it predicts is the result of women leaving paid employment to have and bring up children. One might wonder about the precision of these estimates. Appendix B describes how we computed the standard error for the predicted wages. Typically the standard error for the estimate of the actual wage is in the range of 0.04 to 0.06 and the standard error for the predicted wage starts off around 0.06 and then rises to a maximum of 0.1 (some actual estimates are reported in Table 2).

Although the overall fit is quite good, there are still substantial residuals in certain parts of the profile which should be interpreted as the ‘true’ returns to experience. These rise very sharply in the early part of the worker’s career but then decline gradually actually ending up negative implying that the oldest workers have a level of human capital below that of labour market entrants.

To provide a better idea of the contribution of the search component and the residual to earnings growth, Tables 2a and 2b present estimates of the actual gain, the predicted gain and the residual at different years of experience over the sample period together with the computed standard errors. Even at 5 years of experience the search model explains at least 40% of the earnings gains. We can also use the information in this table to examine the hypothesis of an increase in the 'true' returns to experience. In the conventional analysis the actual returns are used which, as we saw earlier in Table 1 and can be seen again here, did seem to rise in the early 1980s. But with our decomposition we want to look at the residuals. These do seem to rise in the early 1980s (which is when the search model performs worst) but, by the early 1990s the residual seems to be back where it was in the late 1970s. It is in the component predicted by the search model that the rise in the returns to experience can be found.

One can also ask how well the search model can explain differences in the returns to experience between men and women. Figure 9 shows the difference in earnings at different experience levels and the predictions of the pure search model (the decomposition is constructed so the actual and predicted coincide at zero experience). What is striking is that virtually all the differences can be explained by the search model i.e. by differences in transition rates. This also suggests that the reason for the narrowing of the earnings gap between men and women is largely, if not wholly, the result of the narrowing in transition rates that we saw earlier. Note that this is not saying that the model can explain earnings differences between men and women on labour market entry but, as noted earlier, these disappeared in the mid-1970s so that the difference in the returns to experience is now virtually all of the gender pay gap.

4. FURTHER RESULTS: DIFFERENT EDUCATION GROUPS

So far, we have lumped all education groups together. In this section we look at different levels of education. Because of the need to have consistency over time we divide into three levels of education: those who left full-time education at age 16 or before, those who left after 16 but before 21 and those who left at 21 or later (the Labour Force Survey has no information on qualifications until 1979). We will refer to these education groups as low, medium and high respectively. For all of the sample period there are relatively small number of observations in the GHS in the medium and high groups respectively. So, we group 6 years together, making three periods which we will refer to as early (1974–79 inclusive), middle (1980–85) and late (1986–92). With separate results for men and women this leads to the six figures which are in Figures 10a–10f.

The main differences in the profiles across education groups is that the profile seems to be flattest for the medium education group and steepest for the lowest education group. In terms of the prediction of the search model one can see that the model does worst in the early 1980s for all education groups (as we saw before) but that the search model does better at explaining the profile the higher the education level. For the highest education group, the search model seems to do an amazingly good job in fitting the actual profile. This difference in fit across education groups is largely because of the differences in the shape of the profiles rather than because of differences in the predictions of the search model.

What possible explanation for this is there? One interpretation might be that the higher the age at which an individual leaves full-time education, the less of their human capital is provided at work. For many of the workers who left education early, what skills and qualifications they have achieved will often have been achieved while in work (e.g. apprenticeships). However it is much rarer, for example, for university graduates to attain further qualifications once they are in employment. If this is correct, our results are consistent with the fact that human capital does increase rapidly in the first years after labour market entry for those workers who left school at the earliest opportunity.

5. THE ACTUAL AND PREDICTED VARIANCE PROFILE

So far we have concentrated on how well the search model can explain the variation in the first moment of the wage distribution with experience as this is the aspect of the earnings-experience profile that is most commonly explained. But, as already mentioned, one could apply the same methodology to trying to explain other aspects of the wage distribution and doing so might be expected to give further insight into the adequacy of the search model. So, in this section, we look at the actual and predicted variance of earnings. The actual variance can obviously be written as:

$$Var(w|a) = \int w(F,a)^2 dG(F|a) - \left[\int w(F,a) dG(F|a)\right]^2 \qquad \textbf{(11)}$$

The variance predicted by the pure search model, on the other hand, will be given by:

$$\int w(F,0)^2 dG(F|a) - \left[\int w(F,0) dG(F|a)\right]^2 \qquad \textbf{(12)}$$

The difference between the two (the residual) does not have as simple an expression as is the case for the mean so is not written here. Figure 11 presents the actual and predicted variance-experience profiles for our year. There seem to be no particularly striking differences between men and women. The actual profile also does not seem to vary greatly with experience though there is some evidence of a fall in the first years after labour market entry. This contrasts with the findings of Dickens (1996) who, using a different data set, found that the variance of earnings is first increasing and then constant over the life-cycle. There are a number of possible explanations for the differences. First, Dickens restricts attention to those aged 22–59 so misses the first 6 years of the labour market for those who left education at the earliest opportunity. Secondly, Dickens uses age alone as the experience variable as there is no information on education in his data set. This does appear to make some difference: if we draw the profile in our data using age as the experience variable there is some slight increase over the life-cycle after an initial dip though it is not as large as in Dickens' study. The intuition for this is the following: university graduates entering the labour market at the age of 21 do so at wages quite similar to those who have been in the labour market since the age of 18 so the variance at age 21 is quite low. But the graduates rapidly overtake the earnings of the less-educated group and the variance rises. But, if one conditions on experience this does not occur as people of all education groups have rapid earnings increases on entry into the labour market (though some more than others **S** see previous section).

Turning to the predicted variance, there overall predicted variance is pretty close to the actual though there are a number of years in which the predicted variance is substantially above the actual . The predicted variance tends to have a shallow u-shape for men but a slightly more complicated shape for women. Overall, the pure search model does not do too badly in explaining the variance profile. This is consistent with a model in which we can write w(F,a)=w(F)+?(a) in which case the search model will accurately pick up movements in the variance profile over the life-cycle even if there are substantial 'true' returns to experience.

6. EXTENSIONS AND MODIFICATIONS

One should be aware of certain assumptions in our simple search model that may make it appear to perform better or worse than really does so let us consider likely sources of bias. First, our assumption that jobs never change their position in the wage offer distribution means that job mobility is probably given an exaggerated role in wage growth (though it is unclear whether the model would perform better or worse with some within-job wage mobility as jobs will go up as well as down the wage offer distribution). Secondly, the assumption that workers only care about wages mean that the ‘only way is up’: if workers also care about non-wage attributes then some job-to-job mobility will mean workers move down the wage offer distribution **S** this effect will tend to reduce the prediction of earnings growth in the search model and, given that the main weakness of the search model is its inability to predict very rapid earnings growth in the earliest years, make it fit the data worse. Going in a similar direction is the fact that any ‘noise’ in the wage distribution of labour market entrants e.g. because of unmeasured worker heterogeneity or measurement error will tend to increase the search model’s prediction for wage growth as variance in the wage offer distribution is exploited by the process of job mobility.

As this is an issue of some importance, let us consider how we might try to get some idea of the potential problems caused by it. The discussion that follows should be interpreted broadly as it involves a number of approximations. Suppose that we can write the log wage of labour market entrant I as:

$$w_{i0} \quad ' \quad ?_{i} \; \% \; ?_{i0} \; \% \; e_{i0} \qquad \textbf{(13)}$$

where $?_i$ is a fixed effect, $?_{i0}$ is measurement error and e_{i0} measures the position in the wage offer distribution of the job that the worker has. In this set-up the distribution of e can be exploited by the individual in the process of job search but the individual heterogeneity and the measurement error cannot be. Suppose that each of the individual components of (13) is independently normally distributed with mean μ_w for the first component and mean zero for the others (this is without loss of generality) and standard deviations ($s_?$,$s_?$,s_e) respectively. These assumptions mean that the overall log wage distribution for labour market entrants will be normal as well, an assumption that is not too at variance with the facts.

Suppose that we assumed (mistakenly) that all the cross-section variance in wages can be exploited by the process of job search. An alternative way of representing the wage distribution for labour market entrants in (13) is:

$$w(F,0) = \mu_w + \sigma_w \Phi^{-1}(F) \qquad \textbf{(14)}$$

where μ_w is the mean of the wage distribution for entrants, σ_w is the standard deviation, $\Phi^{-1}(F)$ is the inverse of the distribution function for the standard normal and F the position of the individual in the overall wage distribution.

The wage growth predicted by the search model is then given by:

$$\int w(F,0)dG(F,a) - \int w(F,0)dF = \sigma_w \int \Phi^{-1}(F)dG(F,a) \qquad \textbf{(15)}$$

Now consider the correct estimate of wage growth due to the result of the search model. One can think of an individual as being characterised by a position in the ability distribution, $F_?$, a position in the measurement error distribution, $F_?$, and a position in the wage offer distribution F_e. The first two will always be uniformly distributed across the population while the distribution of the latter will change with search. Analogously to (14), we can then write:

$$w(F_?,F_?,F_e,0) = \mu_w + \sigma_? \Phi^{-1}(F_?) + \sigma_? \Phi^{-1}(F_?) + \sigma_e \Phi^{-1}(F_e) \qquad \textbf{(16)}$$

The first three components will not vary with experience so that the ‘true’ wage growth predicted by search will be:

$$\int w(F,0)dG(F,a) - \int w(F,0)dF = \sigma_e \int \Phi^{-1}(F_e)dG(F_e,a) \qquad \textbf{(17)}$$

Comparison of (15) with (17) shows that (conditional on the normality assumption) there is a very simple relationship between the true returns to experience predicted by the search model and the returns predicted by the model which assumes that all wage variation can be exploited through search. One is a simple multiple of the other, the ratio being given by the ratio $\sigma_\varepsilon / \sigma_w$ i.e. the ratio of the standard deviation of the wage offer distribution to the standard deviation of total wages.[7]

Now let us consider how we can get an estimate of this ratio. If wages for labour market entrants are given by (13) then for workers in employment in the next year who have had no intervening period of unemployment we will have:

$$w_{i1} = ?_i + ?_{i1} + \varepsilon_{i1} \qquad \textbf{(18)}$$

where the distribution of ε_{i1} will, from our assumed model of search, have a very particular relationship to the distribution of ε_{i0}, namely it is censored at ε_{i0} with:

$$\varepsilon_{i1} = \varepsilon_{i0} \quad \textit{with probability} \left[1 - ?_e(0) + ?_e \Phi\left(\frac{\varepsilon_{i0}}{\sigma_\varepsilon}\right)\right]$$

$$= \varepsilon, \ \varepsilon \geq \varepsilon_{i0} \quad \textit{with probability} \ \phi\left(\frac{\varepsilon_{i0}}{\sigma_\varepsilon}\right) \qquad \textbf{(19)}$$

Appendix C proves certain facts about the distribution which can be used to derive the following:

$$Var(w_0) = \sigma_?^2 + \sigma_?^2 + \sigma_\varepsilon^2 \equiv \sigma_w^2$$

$$Var(w_1) = \sigma_?^2 + \sigma_?^2 + \sigma_\varepsilon^2\left(1 - \frac{?_e(0)^2}{\pi}\right) \qquad \textbf{(20)}$$

$$\Rightarrow \quad Var(w_1) - Var(w_0) = \frac{\sigma_\varepsilon^2 ?_e(0)^2}{\pi}$$

so that knowledge of the variances of log wages for labour market entrants at entry and one year later (for those in continuous employment) can, together with knowledge of the job-to-job mobility rate from transition rates, be used to get an estimate of $s_{e.}$ This can then be compared with the overall variance in entrants' wages. Data on the variances was obtained from the British Panel Study (as it is important to have a sample of individuals in continuous employment) which gives the variance at labour market entry in the early 1990s as 0.49 and one year later as 0.42. A reasonable value for $?_e(0)$ is about 0.8 when (20) then implies that s_e=0.58 as compared to a value of s_w=0.7. While these computations have a back-of-the-envelope quality to them, this would suggest that there our computed returns to experience as a consequence of search are possibly not a wild over-estimate of the truth. So far, we have emphasized a bias in our approach which will tend to overstate the contribution of search. But we have also made simplifying assumptions which will tend to worsen the contribution of the search model. The assumption that job arrival rates for the employed are the same at all points in the wage offer distribution implicitly assumes that search intensity is the same. Yet, we would expect search intensity to be concentrated in workers with the worse jobs as the potential for gain is greatest there. Embodying this in our search model would tend to increase the predicted rate of earnings growth. It is also quite likely that the job destruction rate varies with the position in the wage offer distribution as lower-paid workers do seem to be at greater risk of job loss. Again, this would tend to increase the predicted rate of earnings growth in the early years in the job search model as workers who reach the higher rungs on the ladder are less likely to fall off.

This discussion has hinted that there are useful ways in which the search model could be modified but it is unclear whether these changes would make the model fit better or worse.

7. CONCLUSIONS

In this paper we have argued that simple search or ‘job-shopping’ model can do surprisingly well in interpreting the earnings-experience profile. It does extremely well in explaining differences in the profile between men and women. We should not be surprised by these results as the literature on displaced workers (see, inter alia, Jacobson *et al* 1993, and Ruhm, 1991) suggest that losing your job leads to lower earnings even once one has controlled for lost job tenure and the search model simply formalizes this unsurprising fact. These results suggest that labour economists should be much more cautious than they commonly are in interpreting returns to experience purely as returns to human capital: perhaps we should think of part of it as returns to search capital.

However, there are also dangers in exaggerating the differences between the two views. In the model we have presented above search capital is destroyed completely whenever a worker loses a job while, in contrast, human capital theory would emphasize that general human capital would not be lost in this situation. But both views would emphasize the importance in differences in transition rates between men and women in accounting for the gender gap: the difference is more in the ‘spin’ that would be put on the mechanism behind the story. The human capital story would emphasize that women who have taken time out of the labour market have less human capital, hence are less productive and receive lower wages. The search model would suggest they are no less productive themselves but that they are tend to be in the less well-paid jobs (which could be less productive employers).

There are useful extensions that could be done to the approach used in this paper. Earnings functions usually condition not just on experience but also on job tenure (as well as other factors) which is often interpreted as the returns to specific human capital so that the wage paid should perhaps be written as w(F,a,t) where t is job tenure. The function w(F,a) we have used here should be interpreted as being this function after we have marginalised with respect to job tenure. But it would be very useful to work with job tenure included and one can use the approach derived in this paper to compute the predictions of a search model for the correlation between wages experience and tenure. We have not done this here for the reasons that it complicates matters very considerably and that the data we have used contains only banded information on job tenure. But, it would be interesting to know how well the search model does in explaining the correlations between wages and tenure and Manning (1997) considers this extension.

ENDNOTES

1. One might wonder about the possible dependence of wages on tenure: we return to this issue below.

2. However, the two theories do differ in their predictions about what function of experience might be most useful in predicting wages e.g. the human capital model would look for some measure of the stock of human capital while the pure search model would suggest that the current length of time in employment might be more useful.

3. It should be noted that Gosling, Machin and Meghir (1994) suggest that the apparent rise in the returns to experience is in fact a cohort (i.e. birth date) effect. It is difficult to distinguish the two hypotheses as once one has a term in cohort squared this is equivalent to including experience multiplied by a time trend.

4. If the pure search model was correct then one could also use the distribution of wages among those workers entering employment after a period of non-employment to estimate w(F,0). We do not do this here as there is no particular reason to believe that the pure search model is correct and we do not have the relevant data for all years in the GHS.

5. Though one should note that given an estimate of G(F*a) one could use this function and the distribution of wages across workers to provide an estimate of w(F,a).

6. This corresponds, approximately, to the fact that the Labour Force Survey, taken in the spring of the year refers to transitions since the previous spring.

7. I am grateful to Coen Teulings for drawing my attention to the possibility of doing this.

TABLE 1: The Changing Returns to Experience

	Men					Women				
	6S10	11S15	16S25	26S35	36S45	6S10	11S15	16S25	26S35	36S45
1974	0.44 (0.02)	0.53 (0.02)	0.56 (0.02)	0.56 (0.02)	0.45 (0.02)	0.29 (0.03)	0.19 (0.03)	0.17 (0.02)	0.13 (0.02)	0.08 (0.02)
1975	0.41 (0.02)	0.49 (0.02)	0.54 (0.02)	0.50 (0.02)	0.44 (0.02)	0.20 (0.02)	0.13 (0.02)	0.12 (0.02)	0.11 (0.02)	0.05 (0.02)
1976	0.43 (0.02)	0.49 (0.02)	0.53 (0.02)	0.52 (0.02)	0.43 (0.02)	0.23 (0.02)	0.14 (0.02)	0.13 (0.02)	0.12 (0.02)	0.08 (0.02)
1977	0.44 (0.02)	0.53 (0.02)	0.57 (0.02)	0.55 (0.02)	0.43 (0.02)	0.29 (0.02)	0.20 (0.02)	0.16 (0.02)	0.18 (0.02)	0.12 (0.02)
1978	0.42 (0.02)	0.49 (0.02)	0.56 (0.02)	0.52 (0.02)	0.40 (0.02)	0.27 (0.02)	0.26 (0.02)	0.15 (0.02)	0.17 (0.02)	0.09 (0.02)
1979	0.43 (0.02)	0.49 (0.02)	0.52 (0.02)	0.52 (0.02)	0.41 (0.02)	0.27 (0.02)	0.18 (0.02)	0.14 (0.02)	0.13 (0.02)	0.09 (0.02)
1980	0.46 (0.02)	0.50 (0.02)	0.57 (0.02)	0.56 (0.02)	0.44 (0.02)	0.29 (0.02)	0.20 (0.02)	0.11 (0.02)	0.15 (0.02)	0.06 (0.02)
1981	0.49 (0.02)	0.59 (0.02)	0.61 (0.02)	0.57 (0.02)	0.51 (0.02)	0.29 (0.02)	0.23 (0.03)	0.11 (0.02)	0.10 (0.02)	0.07 (0.02)
1982	0.45 (0.02)	0.62 (0.02)	0.61 (0.03)	0.62 (0.02)	0.46 (0.02)	0.34 (0.03)	0.23 (0.03)	0.16 (0.02)	0.14 (0.02)	0.06 (0.02)
1983	0.42 (0.02)	0.56 (0.03)	0.56 (0.02)	0.54 (0.02)	0.47 (0.02)	0.34 (0.03)	0.30 (0.03)	0.17 (0.02)	0.14 (0.02)	0.09 (0.02)
1984	0.51 (0.03)	0.66 (0.03)	0.67 (0.02)	0.65 (0.02)	0.57 (0.03)	0.39 (0.03)	0.34 (0.03)	0.24 (0.02)	0.21 (0.02)	0.17 (0.03)
1985	0.54 (0.02)	0.67 (0.03)	0.72 (0.03)	0.71 (0.02)	0.62 (0.03)	0.32 (0.03)	0.36 (0.03)	0.23 (0.02)	0.15 (0.02)	0.13 (0.03)
1986	0.56 (0.03)	0.67 (0.03)	0.73 (0.02)	0.71 (0.02)	0.59 (0.03)	0.35 (0.03)	0.38 (0.03)	0.26 (0.02)	0.22 (0.02)	0.13 (0.03)
1987	0.50 (0.03)	0.65 (0.03)	0.71 (0.02)	0.70 (0.02)	0.61 (0.03)	0.33 (0.03)	0.31 (0.03)	0.26 (0.02)	0.20 (0.02)	0.12 (0.03)
1988	0.49 (0.03)	0.63 (0.03)	0.75 (0.02)	0.68 (0.03)	0.57 (0.03)	0.32 (0.03)	0.26 (0.03)	0.34 (0.02)	0.15 (0.02)	0.13 (0.03)
1989	0.53 (0.03)	0.65 (0.03)	0.74 (0.02)	0.69 (0.02)	0.61 (0.03)	0.30 (0.03)	0.32 (0.03)	0.25 (0.02)	0.18 (0.02)	0.11 (0.03)
1990	0.53 (0.03)	0.63 (0.03)	0.68 (0.03)	0.69 (0.03)	0.60 (0.03)	0.31 (0.03)	0.33 (0.03)	0.28 (0.03)	0.21 (0.03)	0.13 (0.03)
1991	0.49 (0.03)	0.62 (0.03)	0.71 (0.03)	0.68 (0.03)	0.56 (0.03)	0.35 (0.03)	0.36 (0.03)	0.27 (0.03)	0.20 (0.03)	0.15 (0.03)
1992	0.51 (0.03)	0.63 (0.03)	0.69 (0.03)	0.67 (0.03)	0.55 (0.03)	0.38 (0.03)	0.31 (0.03)	0.34 (0.03)	0.23 (0.03)	0.13 (0.03)

TABLE 2a: Actual, Predicted and Residual Returns to Experience

Women

	5 year			15 year			25 year			35 year			45 year		
	a	p	r	a	p	r	a	p	r	a	p	r	a	p	r
1975	0.61 (.04)	0.39 (.06)	0.21 (.07)	0.49 (.04)	0.38 (.06)	0.10 (.07)	0.45 (.05)	0.40 (.06)	0.05 (.07)	0.45 (.03)	0.49 (.06)	-0.04 (.07)	0.43 (.04)	0.54 (.06)	-0.11 (.08)
1977	0.60 (.03)	0.30 .(05)	0.30 (.06)	0.52 (.05)	0.33 (.05)	0.18 (.07)	0.48 (.04)	0.37 (.05)	0.10 (.06)	0.48 (.04)	0.44 (.05)	0.04 (.06)	0.39 (.03)	0.47 (.05)	-0.07 (.06)
1979	0.57 (.03)	0.27 (.03)	0.30 (.04)	0.48 (.03)	0.28 (.03)	0.20 (.05)	0.43 (.04)	0.30 (.03)	0.13 (.05)	0.47 (.03)	0.37 (.04)	0.10 (.05)	0.35 (.03)	0.37 (.04)	-0.02 (.05)
1981	0.69 (.03)	0.20 (.04)	0.49 (.05)	0.52 (.05)	0.25 (.04)	0.27 (.07)	0.58 (.04)	0.33 (.05)	0.25 (.06)	0.52 (.04)	0.38 (.05)	0.14 (.06)	0.42 (.06)	0.40 (.05)	0.02 (.08)
1983	0.59 (.03)	0.22 (.04)	0.37 (.05)	0.52 (.07)	0.22 (.04)	0.30 (.08)	0.49 (.05)	0.28 (.04)	0.21 (.06)	0.43 (.07)	0.31 (.04)	0.13 (.08)			
1984	0.70 (.03)	0.26 (.04)	0.43 (.06)	0.68 (.05)	0.25 (.04)	0.43 (.07)	0.58 (.05)	0.27 (.04)	0.31 (.06)	0.55 (.06)	0.31 (.04)	0.24 (.07)	0.54 (.05)	0.29 (.04)	0.26 (.07)
1985	0.61 (.03)	0.31 (.05)	0.30 (.06)	0.62 (.05)	0.28 (.05)	0.34 (.07)	0.55 (.05)	0.35 (.04)	0.20 (.06)	0.58 (.06)	0.40 (.04)	0.19 (.07)	0.40 (.05)	0.36 (.04)	0.03 (.07)
1986	0.66 (.04)	0.34 (.05)	0.32 (.06)	0.72 (.06)	0.33 (.05)	0.39 (.08)	0.61 (.05)	0.37 (.05)	0.24 (.07)	0.64 (.06)	0.41 (.05)	0.23 (.08)	0.38 (.05)	0.38 (.05)	-0.00 (.07)
1987	0.66 (.04)	0.36 (.05)	0.30 (.06)	0.69 (.05)	0.37 (.05)	0.32 (.07)	0.54 (.05)	0.38 (.05)	0.16 (.07)	0.63 (.05)	0.46 (.06)	0.17 (.07)	0.48 (.09)	0.43 (.06)	0.06 (.10)
1988	0.67 (.04)	0.39 (.05)	0.28 (.06)	0.86 (.08)	0.36 (.05)	0.50 (.09)	0.63 (.05)	0.39 (.05)	0.24 (.07)	0.63 (.05)	0.45 (.05)	0.18 (.07)	0.43 (.03)	0.38 (.05)	0.05 (.06)
1989	0.65 (.04)	0.47 (.05)	0.19 (.07)	0.65 (.05)	0.49 (.06)	0.16 (.08)	0.60 (.05)	0.55 (.06)	0.05 (.08)	0.54 (.05)	0.64 (.06)	-0.11 (.08)	0.41 (.05)	0.61 (.06)	-0.19 (.08)
1990	0.76 (.04)	0.46 (.05)	0.30 (.06)	0.66 (.05)	0.47 (.05)	0.18 (.07)	0.58 (.04)	0.53 (.06)	0.05 (.07)	0.63 (.05)	0.61 (.06)	0.02 (.08)	0.43 (.08)	0.57 (.06)	-0.14 (.10)
1991	0.57 (.05)	0.41 (.05)	0.16 (.07)	0.68 (.05)	0.49 (.05)	0.20 (.07)	0.65 (.06)	0.54 (.06)	0.11 (.08)	0.51 (.06)	0.61 (.08)	-0.10 (.08)	0.40 (.08)	0.56 (.06)	-0.16 (.10)
1992	0.64 (.04)	0.37 (.05)	0.27 (.07)	0.65 (.06)	0.38 (.05)	0.28 (.08)	0.64 (.06)	0.44 (.05)	0.20 (.08)	0.49 (.04)	0.48 (.06)	0.00 (.07)	0.35 (.06)	0.41 (.06)	-0.06 (.08)

TABLE 2a

Men

	5 year			15 year			25 year			35 year			45 year		
	a	p	r	a	p	r	a	p	r	a	p	r	a	p	r
1975	0.60	0.41	0.19	0.83	0.96	-0.13	0.93	1.18	-0.25	0.80	1.14	-0.34	0.72	1.14	-0.42
	(.03)	(.05)	(.06)	(.03)	(.08)	(.08)	(.03)	(.10)	(.10)	(.03)	(.13)	(.13)	(.03)	(1.95)	(1.95)
1977	0.86	0.35	0.51	1.04	0.85	0.19	1.08	1.05	0.03	1.07	1.11	-0.04	0.86	1.11	-0.24
	(.03)	(.06)	(.06)	(.02)	(.08)	(.08)	(.03)	(.09)	(.09)	(.03)	(.09)	(.10)	(.04)	(.09)	(.10)
1979	0.80	0.34	0.47	0.98	0.75	0.23	0.94	0.90	0.04	0.90	0.93	-0.02	0.75	0.94	-0.19
	(.02)	(.04)	(.05)	(.03)	(.06)	(.06)	(.03)	(.06)	(.07)	(.03)	(.07)	(.07)	(.03)	(.07)	(.07)
1981	0.91	0.24	0.68	1.17	0.60	0.56	1.15	0.79	0.36	1.15	0.87	0.28	1.04	0.91	0.14
	(.03)	(.05)	(.06)	(.03)	(.07)	(.07)	(.04)	(.08)	(.09)	(.04)	(.09)	(.10)	(.03)	(.07)	(.07)
1983	0.81	0.23	0.58	1.13	0.47	0.66	1.05	0.57	0.48	0.98	0.57	0.41			
	(.03)	(.05)	(.06)	(.05)	(.06)	(.07)	(.05)	(.06)	(.08)	(.04)	(.06)	(.08)			
1984	0.89	0.27	0.61	1.21	0.53	0.68	1.15	0.62	0.52	1.10	0.64	0.46	0.97	0.57	0.39
	(.03)	(.05)	(.06)	(.04)	(.06)	(.07)	(.05)	(.06)	(.08)	(.05)	(.06)	(.08)	(.05)	(.06)	(.08)
1985	0.96	0.32	0.65	1.24	0.60	0.64	1.23	0.70	0.53	1.21	0.74	0.47	1.10	0.72	0.39
	(.04)	(.05)	(.07)	(.04)	(.06)	(.07)	(.04)	(.07)	(.08)	(.05)	(.07)	(.09)	(.06)	(.07)	(.09)
1986	1.01	0.30	0.70	1.24	0.60	0.65	1.27	0.71	0.56	1.29	0.76	0.54	1.02	0.71	0.31
	(.04)	(.05)	(.07)	(.03)	(.06)	(.07)	(.04)	(.07)	(.08)	(.05)	(.07)	(.08)	(.05)	(.07)	(.08)
1987	0.90	0.31	0.58	1.23	0.61	0.62	1.20	0.69	0.51	1.18	0.73	0.45	1.05	0.70	0.34
	(.04)	(.05)	(.06)	(.04)	(.05)	(.07)	(.04)	(.06)	(.07)	(.05)	(.06)	(.08)	(.06)	(.06)	(.09)
1988	0.91	0.37	0.54	1.32	0.61	0.72	1.30	0.64	0.66	1.16	0.67	0.49	1.06	0.62	0.44
	(.04)	(.05)	(.07)	(.05)	(.07)	(.08)	(.05)	(.07)	(.08)	(.06)	(.07)	(.09)	(.05)	(.07)	(.09)
1989	0.93	0.40	0.52	1.24	0.73	0.51	1.26	0.83	0.42	1.18	0.88	0.29	1.06	0.83	0.23
	(.04)	(.05)	(.06)	(.04)	(.06)	(.07)	(.05)	(.07)	(.09)	(.06)	(.07)	(.09)	(.05)	(.07)	(.09)
1990	0.85	0.39	0.47	1.07	0.69	0.38	1.10	0.79	0.31	1.08	0.80	0.28	0.92	0.79	0.13
	(.05)	(.06)	(.08)	(.04)	(.08)	(.09)	(.05)	(.09)	(.11)	(.06)	(.10)	(.11)	(.08)	(.09)	(.12)
1991	0.86	0.39	0.47	1.13	0.78	0.35	1.10	0.89	0.21	1.13	0.93	0.20	0.86	0.90	-0.04
	(.04)	(.07)	(.08)	(.04)	(.09)	(.10)	(.06)	(.10)	(.12)	(.06)	(.10)	(.12)	(.06)	(.10)	(.12)
1992	0.83	0.31	0.52	1.13	0.58	0.55	1.15	0.66	0.49	1.00	0.71	0.29	0.79	0.68	0.11
	(.04)	(.06)	(.07)	(.04)	(.07)	(.08)	(.05)	(.08)	(.09)	(.06)	(.08)	(.10)	(.07)	(.08)	(.11)

Notes:

1. A is the actual average wage, p is the predicted and r the residual. Standard errors computed by the method described in Appendix B are reported in parentheses.

APPENDIX A

Proof of Proposition 1

Denote by J(F*a) the proportion of the population of experience a that is in a firm of position F or less. Obviously [1-u(a)]G(F*a)=J(F*a) so that knowledge of J(.) and (4) allows us to derive G(.).

The derivative of J(F*a) with respect to a will be equal to the inflow to this group minus the outflow. As the instantaneous separation rate is [d_u(a)+$?_e$(a)(1-F)] the outflow will be given by [d_u(a)+$?_e$(a)(1-F)]J(F*a). The inflow comes from those workers who are unemployed who get a wage offer that is less than F. The rate at which this happens is $?_u$(a)Fu(a).

Putting this information together we have that:

$$\frac{\partial J(F|a)}{\partial a} = -[\delta_u+\lambda_e(1-F)].J(F|a) + \lambda_u.F.u(a) \qquad \textbf{(21)}$$

Now, as (1-u)G=J we have that:

$$\frac{\partial G(F|a)}{\partial a} = \frac{u'(a)}{1-u(a)}G(F|a) + \frac{1}{1-u(a)}\frac{\partial J(F|a)}{\partial a} \qquad \textbf{(22)}$$

which, using (3) and (21) can be re-arranged to yield (4).

To derive the closed-form solution note that given the definitions of (a,ß,G) one can write the differential equation (4) as:

$$\frac{\partial[G(F|a)-F]}{\partial a} = -F\beta(a) - [\alpha(a)+\beta(a)][G(F|a)-F] \qquad \textbf{(23)}$$

which can be written as:

$$\frac{\partial\left([G(F|a)-F].e^{\int_0^a[\alpha(s)+\beta(s)]ds}\right)}{\partial a} = -F\beta(a)e^{\int_0^a[\alpha(s)+\beta(s)]ds} \qquad \textbf{(24)}$$

which, given the definition of Γ can be written as:

$$\frac{\partial\left([G(F|a)-F].\Gamma(a)\right)}{\partial a} = -F\beta(a)\Gamma(a) \qquad \textbf{(25)}$$

Which, together with the initial condition G(F|0)=F leads to (5).

Proof of Proposition 2
As a preliminary, note from (6) that:

$$\frac{\partial\Gamma(s)}{\partial\alpha(v)} = \frac{\partial\Gamma(s)}{\partial\beta(v)} = 0 \quad if \quad v>s$$

$$\frac{\partial\Gamma(s)}{\partial\alpha(v)} = \frac{\partial\Gamma(s)}{\partial\beta(v)} = \Gamma(s) \quad if \quad v\leq s \qquad \textbf{(26)}$$

By inspection of the formula in (5) we can see that α(v) and β(v) have no effect on G(F|a) if v≥a. So, let us restrict attention to v<a. Then by differentiating (5) we have (making extensive use of (26)).:

$$\frac{\partial G(F|a)}{\partial\alpha(v)} = \frac{-F}{\Gamma(a)^2}.\left(\Gamma(a)\int_v^a\Gamma(s)\beta(s)ds - \Gamma(a)\int_0^a\Gamma(s)\beta(s)ds\right) > 0 \qquad \textbf{(27)}$$

and:

$$\frac{\partial G(F|a)}{\partial\beta(v)} = \frac{-F}{\Gamma(a)^2}.\left(\Gamma(a)\int_v^a\Gamma(s)\beta(s)ds + \Gamma(a)\Gamma(v) - \Gamma(a)\int_0^a\Gamma(s)\beta(s)ds\right)$$

$$= \frac{-F}{\Gamma(a)}.\left(\Gamma(v) - \int_0^v\Gamma(s)\beta(s)ds\right) \qquad \textbf{(28)}$$

Let us consider the term in the large brackets in the final line. When v=0 this is strictly positive as Γ(0)=1. Taking the derivative of it with respect to v we have that:

$$\frac{\partial\left(\Gamma(v)-\int_0^v \Gamma(s)\beta(s)ds\right)}{\partial v} = [\alpha(v)+\beta(v)]\Gamma(v) - \beta(v)\Gamma(v) > 0 \qquad \textbf{(29)}$$

so that the term is increasing in v. Hence it must always be strictly positive making the expression in the last line of (28) as a whole negative.

Proof of Proposition 3

If we differentiate (5) with respect to a we obtain:

$$\begin{aligned}\frac{\partial G(F|a)}{\partial a} &= \frac{-F}{\Gamma(a)^2}.\left(\Gamma(a)^2\beta(a) - [\alpha(a)+\beta(a)]\Gamma(a)\int_0^a \Gamma(s)\beta(s)ds\right) \\ &= \frac{F}{\Gamma(a)}.\left([\alpha(a)+\beta(a)]\int_0^a \Gamma(s)\beta(s)ds - \Gamma(a)\beta(a)\right)\end{aligned} \qquad \textbf{(30)}$$

Let us define the term in the large brackets in the last line of (30) as Z(a). It is the sign of Z(a) which determines the sign of the derivative of G with respect to a. Obviously Z(0)<0 (which is simply our earlier result that initially the wage distribution must improve). So if we can show that $Z'(a) \leq 0$ then we have proved the result. Differentiating Z(a) we have:

$$Z'(a) = [\alpha' + \beta'] \int_0^a G(s)\beta(s)ds + [\alpha+\beta]G\beta - \beta' G - \beta G'$$

$$= [\alpha' + \beta'] \int_0^a G(s)\beta(s)ds - \beta' G \qquad \textbf{(31)}$$

$$= \frac{\alpha' + \beta'}{\alpha+\beta} Z(a) + \frac{\alpha'\beta - \beta'\alpha}{\alpha+\beta} G(a)$$

where the second line follows from the fact that $G' = (\alpha+\beta)G$ and the final line from (30). As $Z(0)<0$ if $Z(a)>0$ for some a it must cross the axis from below. But if $Z(a)=0$, the final line shows that, under the conditions stated in the Proposition, $Z'(a) \leq 0$ so this is not possible.

APPENDIX B
The Discrete Time Version of the Model

There are two types of discretization that we are going to do to implement the theoretical model. First, we are going to model time in discrete periods (a year) and, secondly (for computational reasons) we can only consider a finite number of values of F. Let us start with considering the implications of discrete time.

In the discretized version of the model we assume that workers get at most one job offer per year or get their job destroyed once so that ($?_u$,d_u,$?_e$) now refer to probabilities and hence must be less than one.

For each cell (differentiated by year, sex, experience and, possibly, education) we have an estimate of the number of people in each of the cells (employed, employed), (non-employed, employed), (employed, non-employed) and (non-employed, non-employed) where the first state refers to the situation a year ago and the second refers to the situation now. Obviously these numbers enable us to compute the non-employment rate for this experience group both now (which we will denote by u(a)) and a year ago (which we will denote by u(a-1)). Let us denote by L_{ne} the number of individuals non-employed a year ago but employed now with similar notation for the other four states and use L to denote the total sample size. Then our estimates of ($?_u$,d_u) are given by:

$$?_u \; ' \; \frac{L_{ne}}{L_{ne}\%L_{nn}}, \qquad d_u \; ' \; \frac{L_{en}}{L_{ee}\%L_{en}} \qquad \textbf{(32)}$$

Now consider how we can work out $?_e$, the arrival rate of job offers when employed. For those workers who are in employment in both years we have information on whether they have changed employers or not. Denote by ?(a) the proportion of these workers who have changed jobs. So the probability of worker in employment a year ago changing jobs but being in employment in both years is given by [1-d_u(a)]?(a). If the distribution of workers with experience a is given by G(F*a-1) then it must be the case that:

$$[1-\delta_u(a)]\lambda(a) = \lambda_e(a)\int_0^1 (1-F)\,dG(F|a-1) \qquad \textbf{(33)}$$

as only a fraction (1-F) of job offers result in a change of employers. (33) shows that knowledge of $\lambda(a)$ and $G(F|a-1)$ is enough to enable computation of $\lambda_e(a)$. But how can we compute $G(F|a-1)$? As in the proof used in Proposition 1 let us denote by $J(F|a)$ the percentage of the labour force employed at a firm less than or equal to F. In the discrete version of (21) we have:

$$J(F|a) = [1-\delta_u(a)-\lambda_e(a)(1-F)][1-u(a-1)]G(F|a-1) + F\lambda_u(a)u(a-1) \qquad \textbf{(34)}$$

from which we can derive $G(F|a)$ by $[1-u(a)]G(F|a)=J(F|a)$. Substituting (32) into this we have that:

$$G(F|a) = [1-\delta_u(a)]\left[1-\frac{\lambda(a)(1-F)}{\int_0^1 (1-F')\,dG(F'|a-1)}\right]\frac{[1-u(a-1)]}{[1-u(a)]}G(F|a-1) + F\frac{\lambda_u(a)u(a-1)}{[1-u(a)]} \qquad \textbf{(35)}$$

Now using (32) and the definition of the unemployment rates we have that:

$$\frac{[1-\delta_u(a)][1-u(a-1)]}{[1-u(a)]} = \frac{L_{ee}}{L_{ee}+L_{en}}.\frac{L_{en}+L_{ee}}{L}.\frac{L}{L_{ne}+L_{ee}} = \frac{L_{ee}}{L_{ne}+L_{ee}} \qquad \textbf{(36)}$$

$$\frac{\lambda_u(a).u(a-1)}{[1-u(a)]} = \frac{L_{ne}}{L_{ne}+L_{nn}}.\frac{L_{nn}+L_{ne}}{L}.\frac{L}{L_{ne}+L_{ee}} = \frac{L_{ne}}{L_{ne}+L_{ee}} \qquad \textbf{(37)}$$

so that (35) can be written as:

$$G(F|a) = s(a)\cdot\left[1-\frac{\lambda(a)(1-F)}{\int_0^1 (1-F')\,dG(F'|a-1)}\right]G(F|a-1) + [1-s(a)]F \qquad \textbf{(38)}$$

where $s(a)=L_{ee}/(L_{ne}+L_{ee})$ which has the natural interpretation as the proportion of current workers who have not entered from non-employment in the past year.

Equation (38) has a recursive structure and we can start it off from the assumption that $G(F|0)=F$ and then run it forward.

An additional difficulty from the computational point of view is the need to evaluate $G(F|a)$ at a finite number of points of F. We used one hundred values of F evenly spaced from 0.005 to 0.995 (so effectively we use percentiles). The integral in (36) is then approximated by:

$$\int_0^1 (1-F)\,dG(F|a-1) = \sum_{i=1}^{100} F_i[G(F_i|a-1)-G(F_{i-1}|a-1)] \qquad \textbf{(39)}$$

where F_i denotes the value of F at the ith point in the partition and we adopt the notation that $F_0=0$. Experimentation showed that not much was gained from moving to a finer partition.

We also use the same percentiles when computing the wage distribution. We do this in the following way. All workers in the relevant cell are ordered by their wage and then assigned a position in the wage distribution from 0 to 1. We then estimate the wage at position F_i by linear interpolation using the two wages adjacent to this point. The outcome of this procedure is an estimate of $w(F_i)$ in each cell. We use this distribution for computing all statistics in the paper so that the mean wage we use is $(\Sigma w(F_i))/100$ and so on. There is obviously an approximation error in doing this but it seems to be small. For example, the correlation across all cells in all years of the true mean wage and the mean computed from the discretized version is 0.99

and for the variance it is 0.97 (note that the variance computed from the discretized version must always be smaller than the true variation because the discretized model effectively assumes there is no wage variation within percentiles).

Computing covariances

In computing the covariance matrices below it is important to be aware of the independence assumptions we are making. We will assume that the information on wages is independent of that on transition rates (they do come from different data sets) and that data on transition rates and wages for different experience levels are independent of each other.

Given the finite number of values of F considered we can represent G(F*a) as a vector **G**(a). Let us denote the asymptotic covariance matrix of this vector as $\Sigma_G(a)$. As **G**(0)=F we must have $\Sigma_G(0)=0$. As we can see from (36) **G**(a) is a function of **G**(a-1) and the scalars s(a) and ?(a). s(a) and ?(a) are Bernouilli random variables with the variance of s(a) given by s(a)(1-s(a)) divided by the sample size and similarly for ?(a). Let us denote the variance of s(a) and ?(a) by $\sigma^2_s(a)$ and $\sigma^2_?(a)$ respectively. s and ? are independent because s is the fraction of current workers who were employed last year and ? is the fraction of workers employed in both years who have changed jobs. Then, from (39), the asymptotic covariance matrix is going to follow the following recursive formula:

$$\Sigma_G(a) = \sigma^2_s(a)\left(\frac{\partial \mathbf{G}(a)}{\partial s(a)}\right)\left(\frac{\partial \mathbf{G}(a)}{\partial s(a)}\right)' + \sigma^2_?(a)\left(\frac{\partial \mathbf{G}(a)}{\partial ?(a)}\right)\left(\frac{\partial \mathbf{G}(a)}{\partial ?(a)}\right)'$$
$$+ \left(\frac{\partial \mathbf{G}(a)}{\partial \mathbf{G}(a-1)}\right)\Sigma_G(a-1)\left(\frac{\partial \mathbf{G}(a)}{\partial \mathbf{G}(a-1)}\right)' \qquad \textbf{(40)}$$

which can be computed using the initial condition $\Sigma_G(0)=0$.

When we compute the expected wage from the search model using the decomposition in (10) we obviously need the covariance matrix for w(F,0). We compute this in the following way. When we discretize, the function w(F,0) is represented by the vector **w** (which is of the same dimension as **G**). As each element in this vector represents percentiles of the wage distribution there is a well-known formula for the covariance matrix for **w** which we will denote by Σ_w and whose individual elements are given by:

$$Cov(w_i,w_j) = \frac{F_i(1-F_j)}{Nf(w_i)f(w_j)} \quad i \leq j \qquad \textbf{(41)}$$

where f is the value of the density at w and N is the sample size. In computing this the only problem is to estimate the density. We have proceeded by assuming that the distribution of entry wages is normal (an assumption that is accepted using standard normality tests for most though not all of the years) and computed the density at different percentiles accordingly.

When we apply the decomposition in (10) we calculate:

$$\sum w_i(G_i-G_{i-1}) = \boldsymbol{w}'A\boldsymbol{G} \qquad \textbf{(42)}$$

where A is a square matrix of the same dimension as the rows of **w** and **G**, with zeroes everywhere except ones on the main diagonal and -1 on the diagonal immediately to the left of the main one. Given this we can compute:

$$Var(\sum w_i(G_i-G_{i-1})) = (\boldsymbol{w}'A)\Sigma_w(\boldsymbol{w}'A)' + (A\boldsymbol{G})'\Sigma_G(A\boldsymbol{G}) \qquad \textbf{(43)}$$

and it is this that is reported in the text.

APPENDIX C
Justifying Equation (20)

It is simpler to derive (20) using a standardised normal so let us define $?=e/s_e$. For economy of notation let us denote $?_e(0)$ by $?$. Given the assumption on the distribution of e_1 in (19) we have that:

$$E(?_1 * ?_0) = [1 - ? + ?F(?_0)]?_0 + ?\int_{?_0} ?f(?)d? \qquad \textbf{(44)}$$

$$= [1 - ? + ?F(?_0)]?_0 + ?f(?_0)$$

using well-known results on the censored normal distribution. Taking the distribution of (44) with respect to $?_0$, we obtain:

$$E(?_1) = ?\int F(?)?f(?)d? + ?\int f(?)^2 d?$$

$$= -?\int F(?)f'(?)d? + ?\int f(?)^2 d? \qquad \textbf{(45)}$$

$$= 2?\int f(?)^2 d? = \frac{?}{\sqrt{p}}$$

where the second line follows from the first as $-?f(?)=f'(?)$ for the normal distribution, the third line follows from the second by integrating by parts.

We must also have:

$$E(?_1^2 | ?_0) \; = \; [1-? + ?F(?_0)]?_0^2 + ? \int_{?_0} ?^2 f(?)\,d? \qquad \textbf{(46)}$$

$$= \; [1-? + ?F(?_0)]?_0 + ?[1-F(?_0)]$$

Taking the expectation of (46) with respect to $?_0$, we obtain:

$$E(?_1^2) \; = \; (1-?)? \int F(?)?^2 f(?)\,d? + ? \int [1-F(?)] f(?)\,d?$$

$$= \; (1-?)-? \int F(?)?f^{)}(?)\,d? - \frac{?}{2} \int [1-F(?)]^{2)}\,d? \qquad \textbf{(47)}$$

$$= \; (1-?) + ? \int [F(?) + ?f(?)] f(?)\,d? + \frac{?}{2}$$

$$= \; (1-?) + \frac{?}{2} \int F(?)^{2)}\,d? + \frac{?}{2} \; = \; 1$$

where the third line follows from the second by integrating by parts. Combining (45) and (47) and the relationship between ? and e yields (20).

Figure 1
The Earnings-Experience Profile in 1983

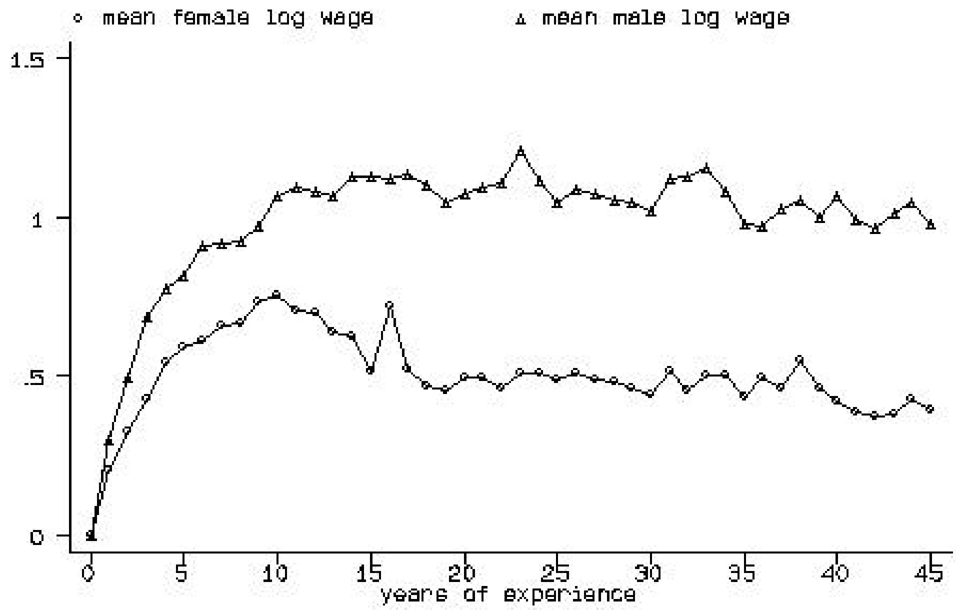

Figure 2
The Earnings-Experience Profile in All Years

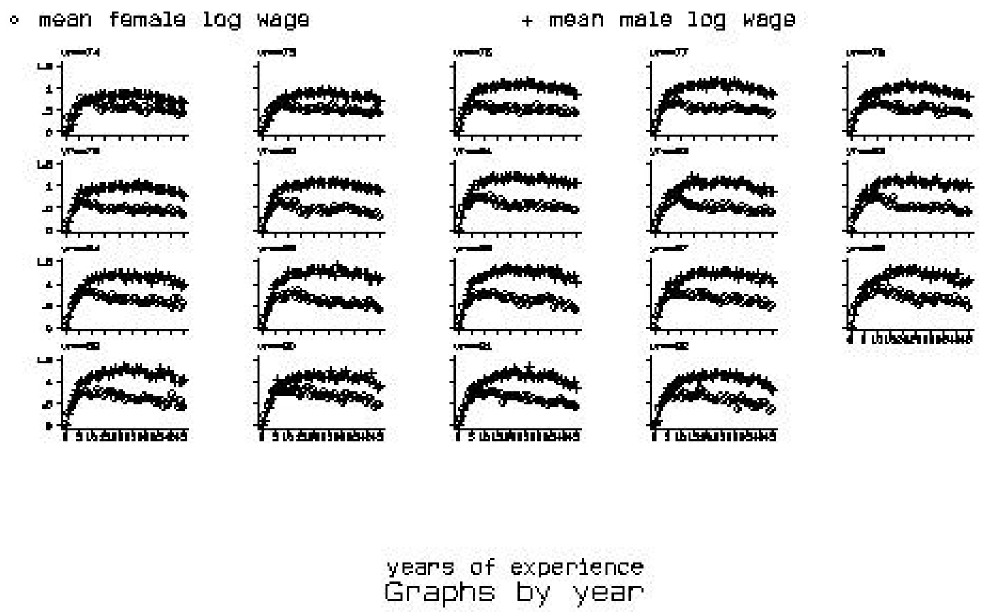

Figure 3
Changes in the Returns to Experience Over Time

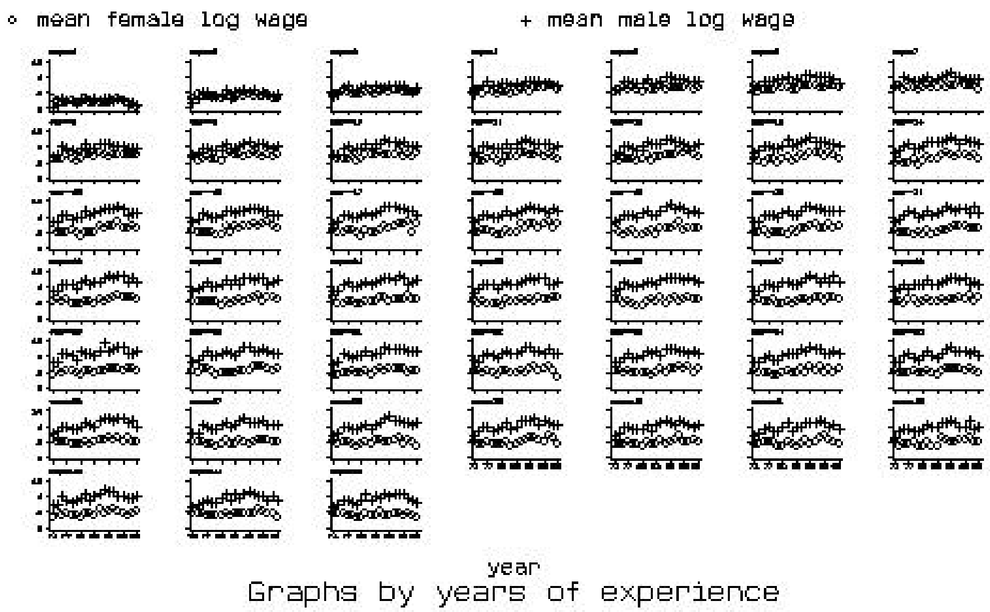

Figure 4
Transition Rates From Employment to Non-Employment

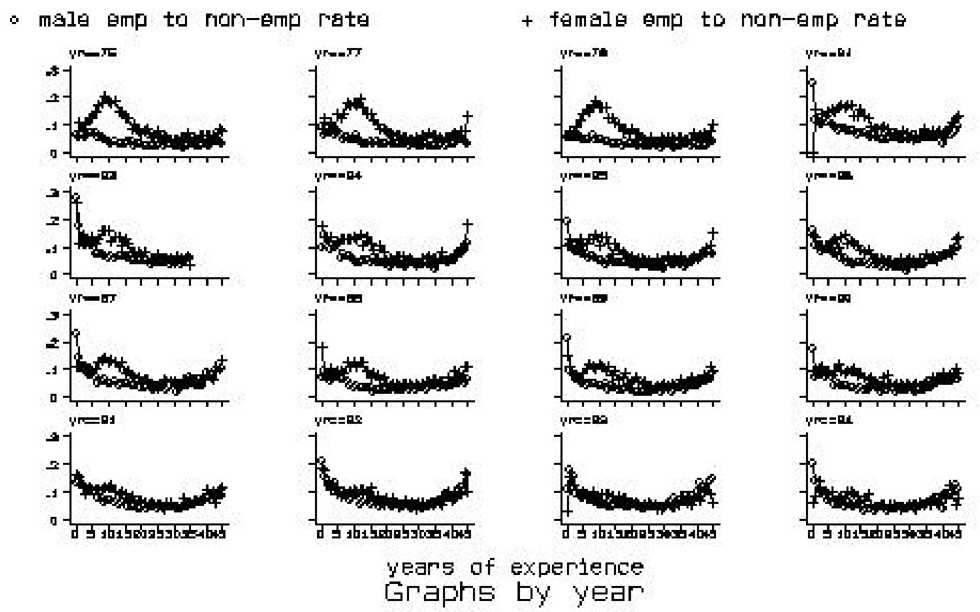

Figure 5
Transition Rates from Non-Employment to Employment

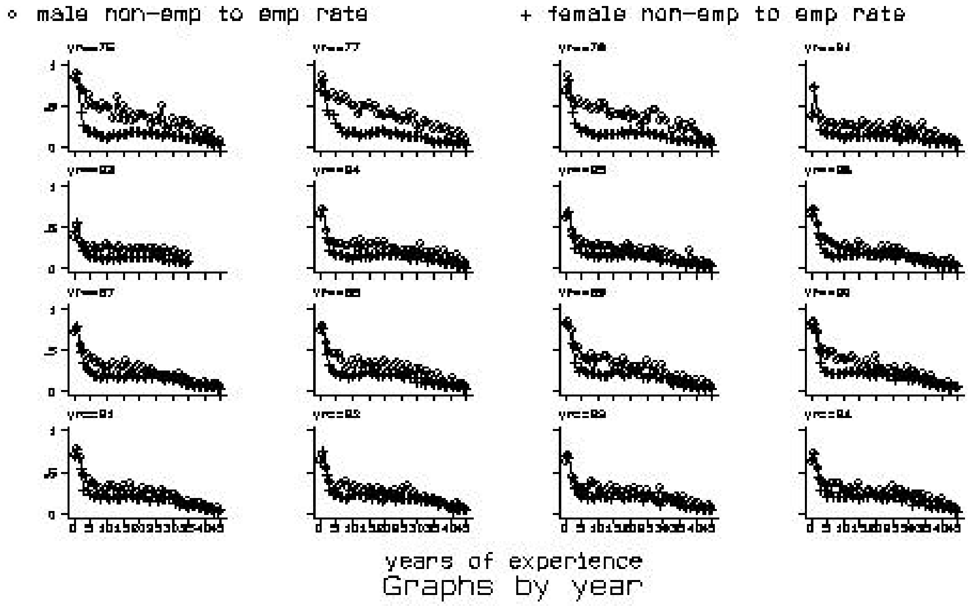

Figure 6
Job-to-Job Mobility Rates

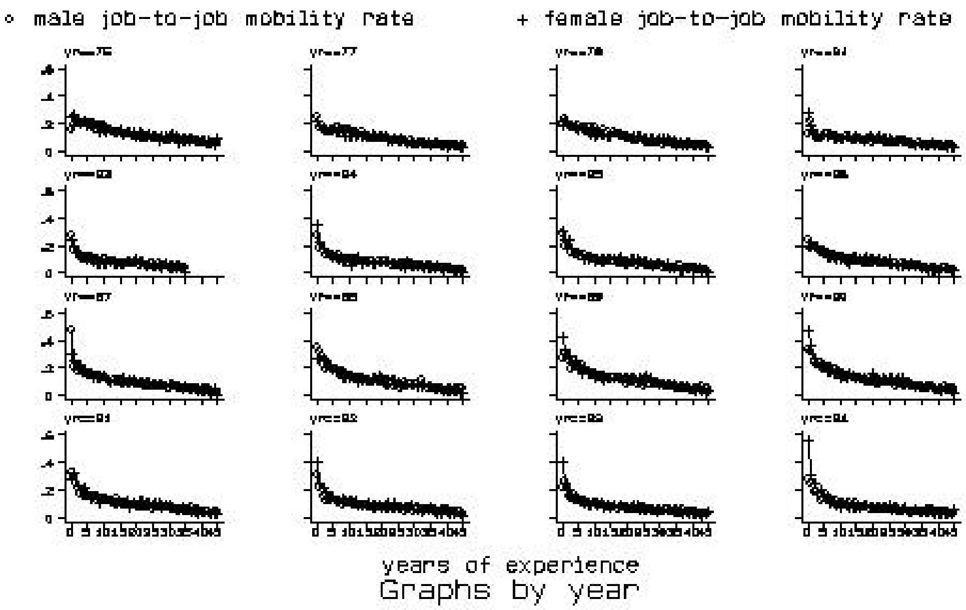

Figure 7
The Predictions of the Search Model for the Proportion Paid Below the Median

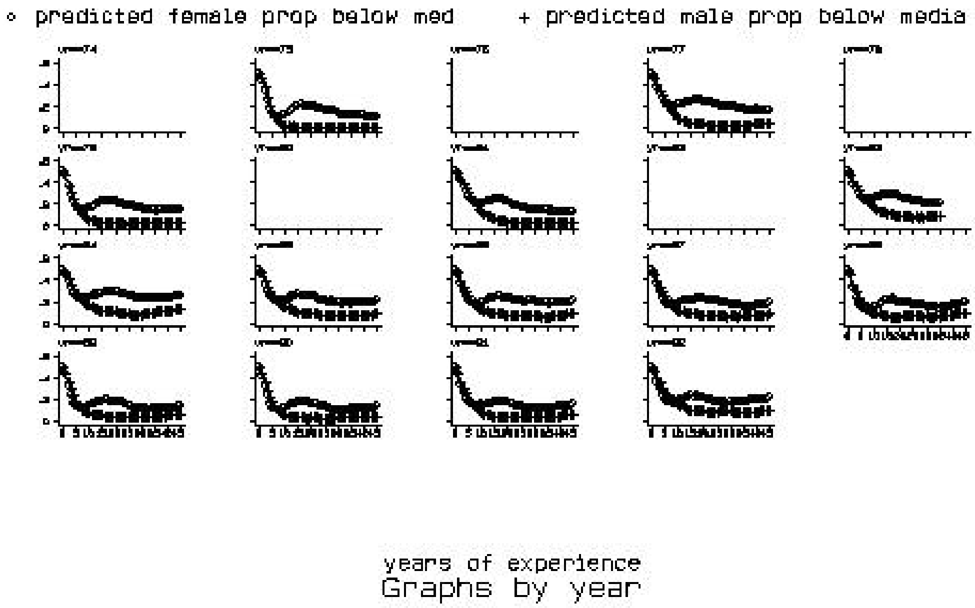

Figure 8
Actual and Predicted Earnings-Experience Profiles

Women

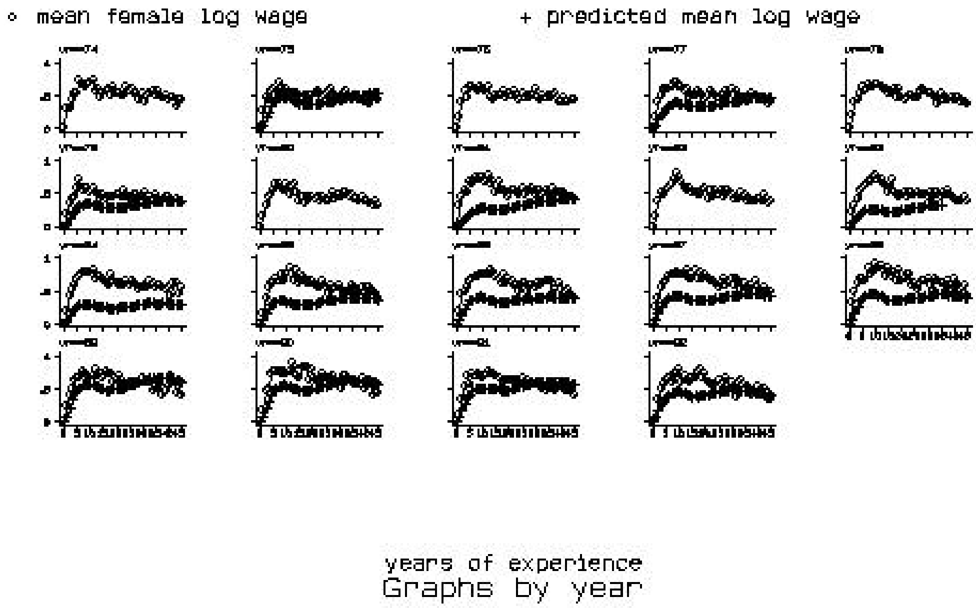

Men

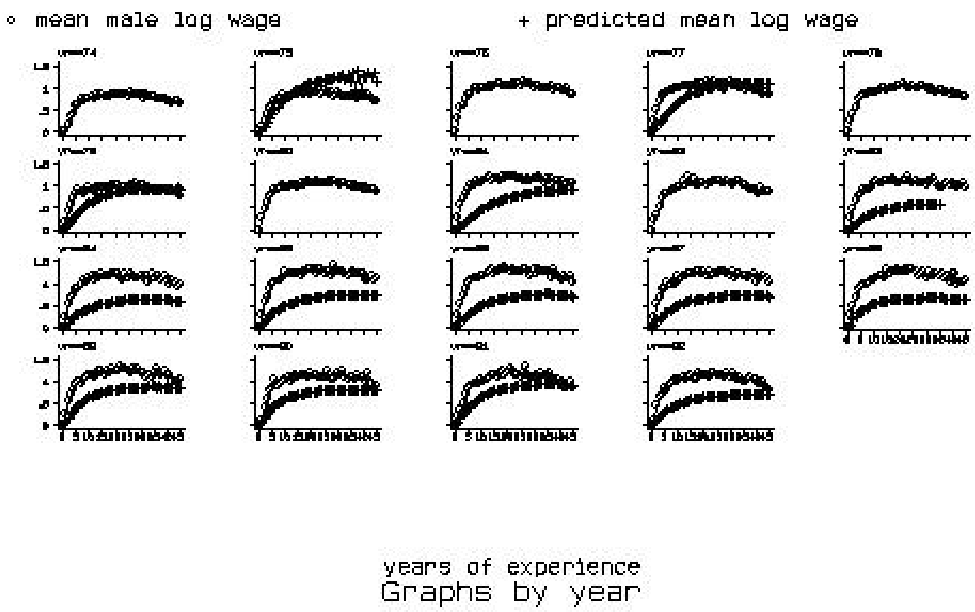

Figure 9
Actual and Predicted Male-Female Wage Differential

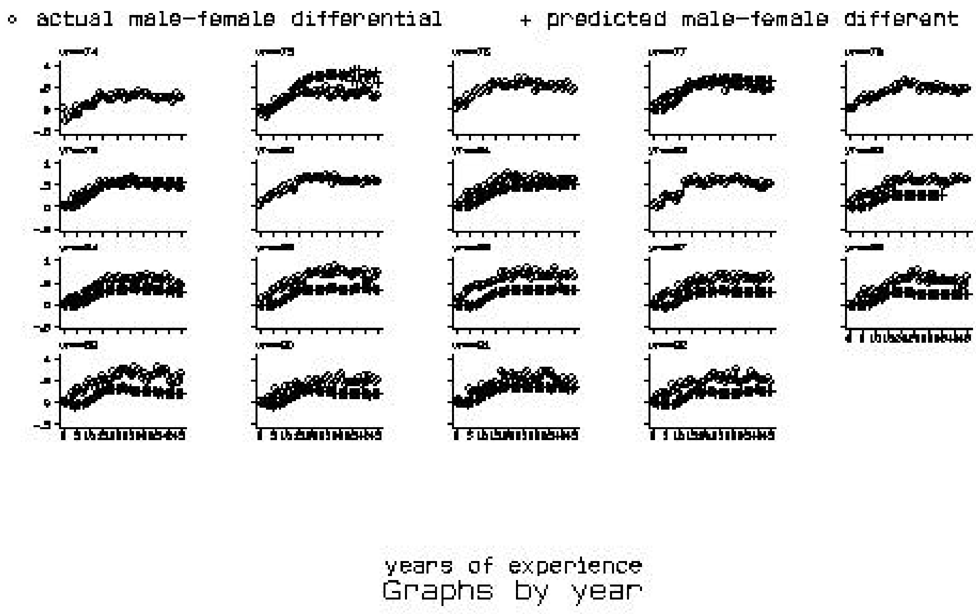

Figure 10a
Actual and Predicted Earnings-Experience Profiles by Education: 1974S79

Women

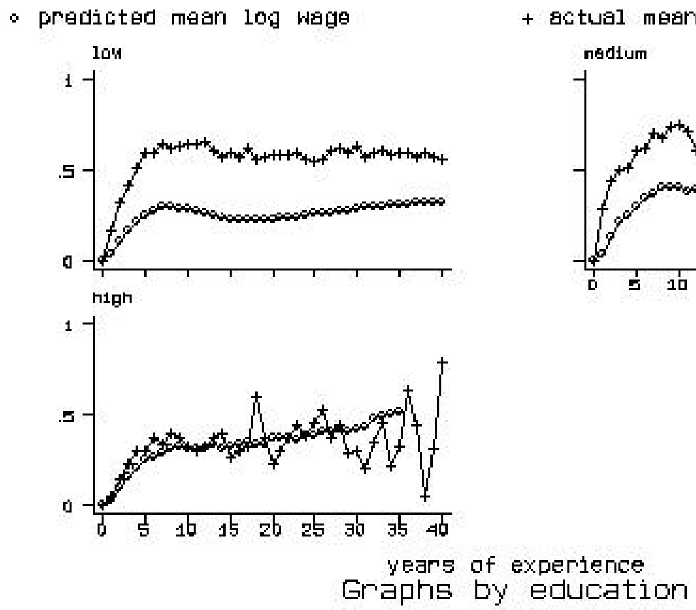

Men

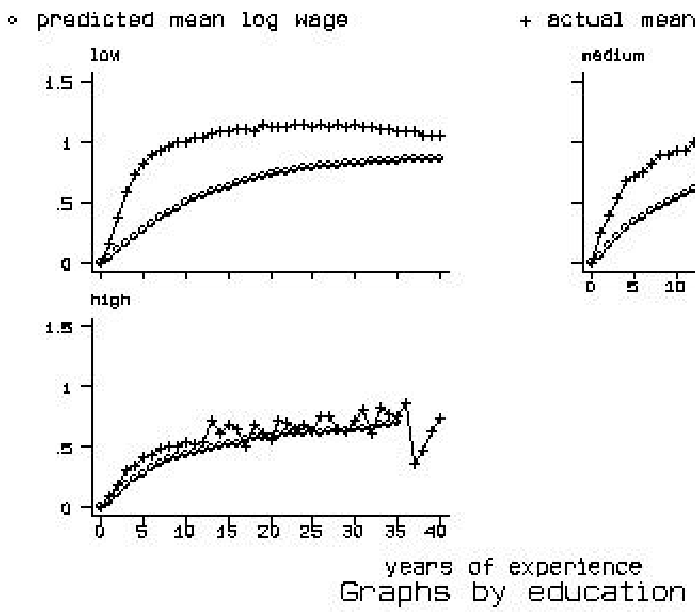

Figure 10b
Actual and Predicted Earnings-Experience Profiles by Education: 1980S85

Women

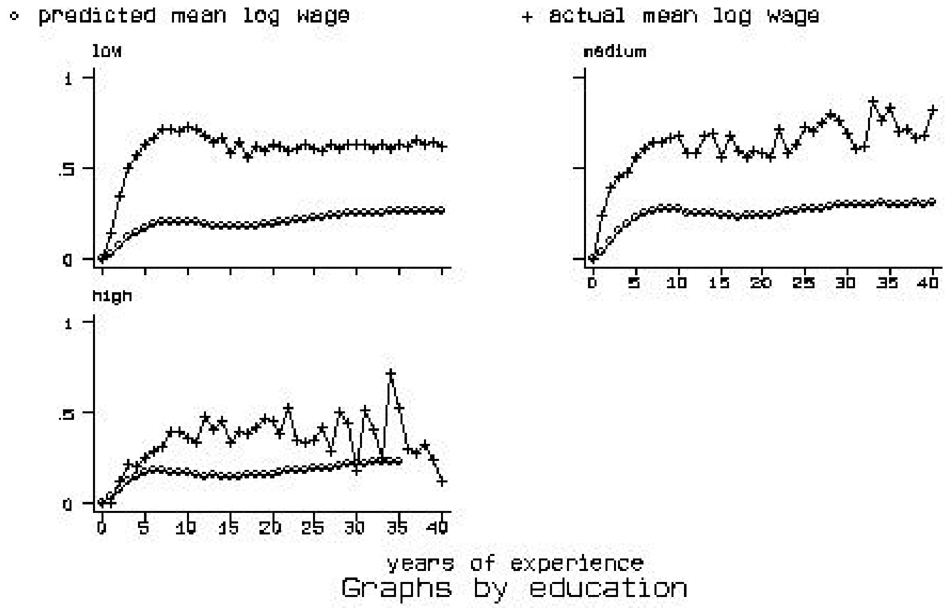

Men

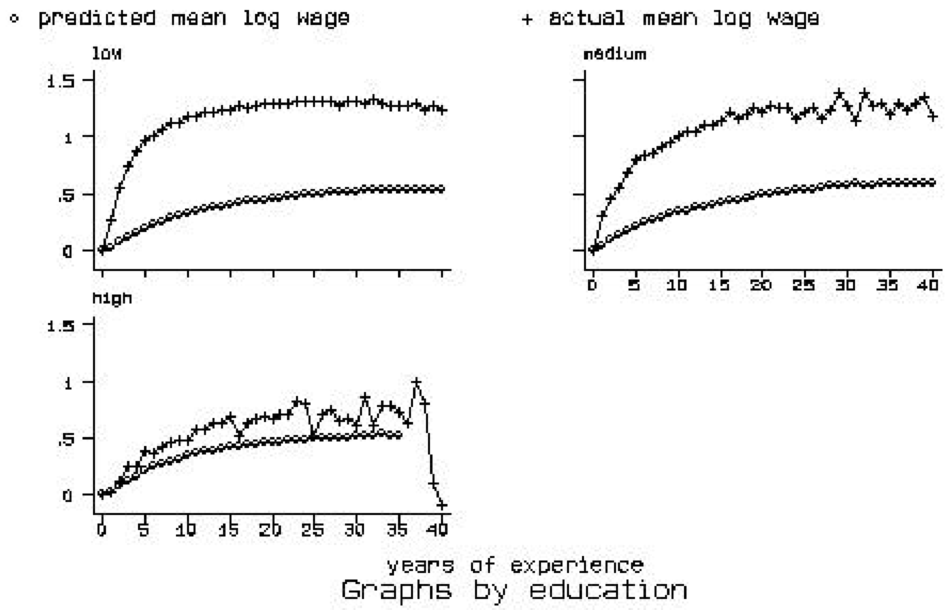

Figure 10c
Actual and Predicted Earnings-Experience Profiles by Education: 1986-92

Women

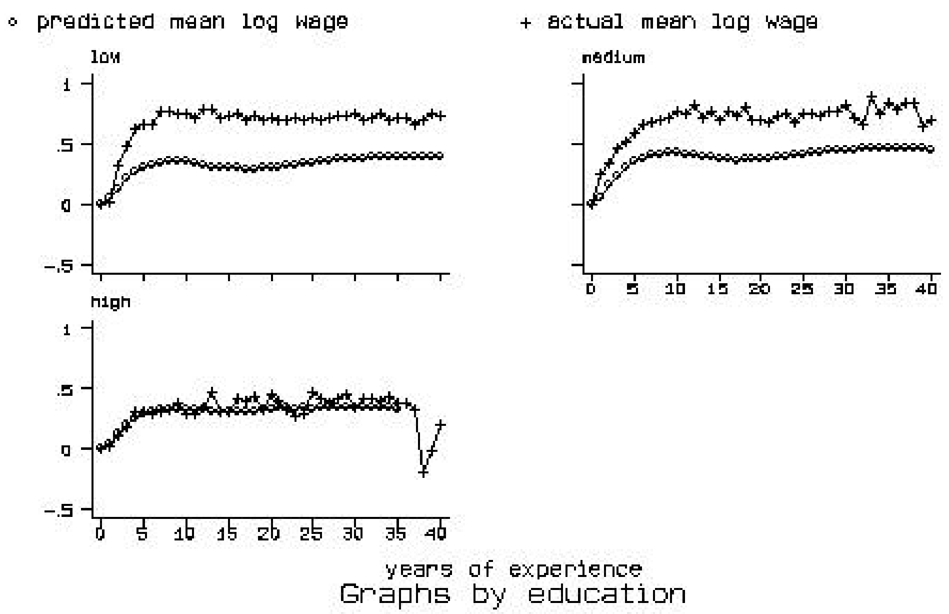

Men

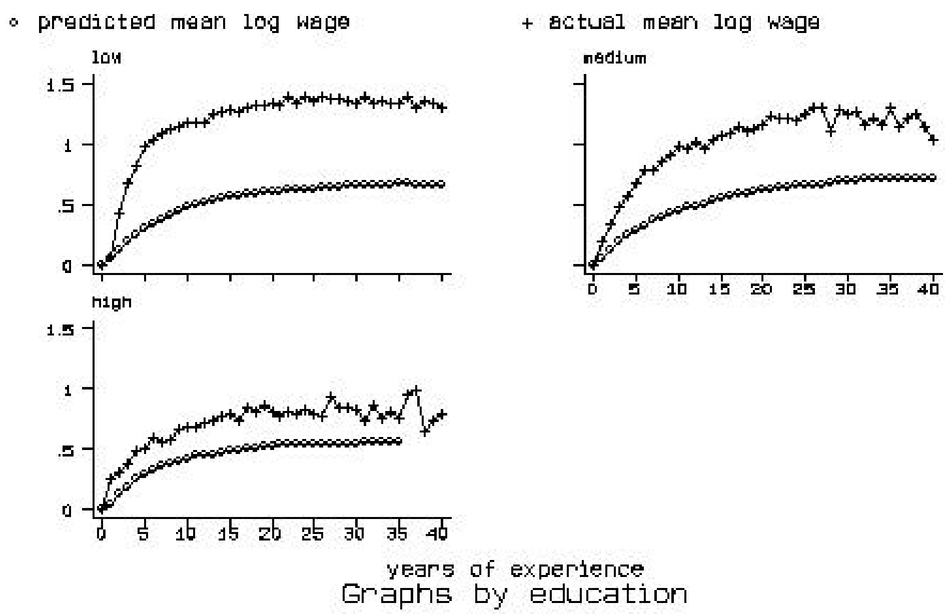

Figure 11a
Actual and Predicted Variance of Log Wages

Women

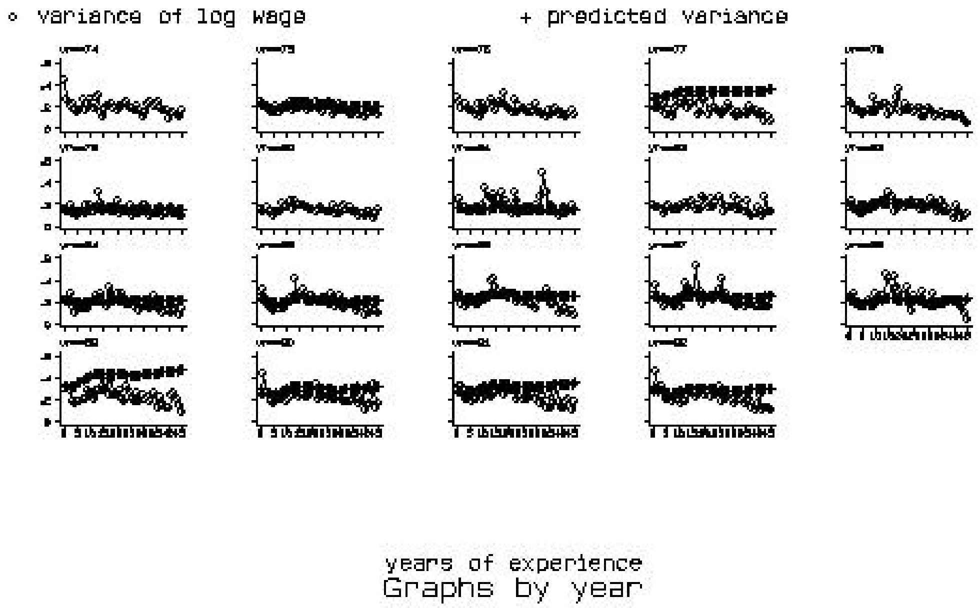

Men

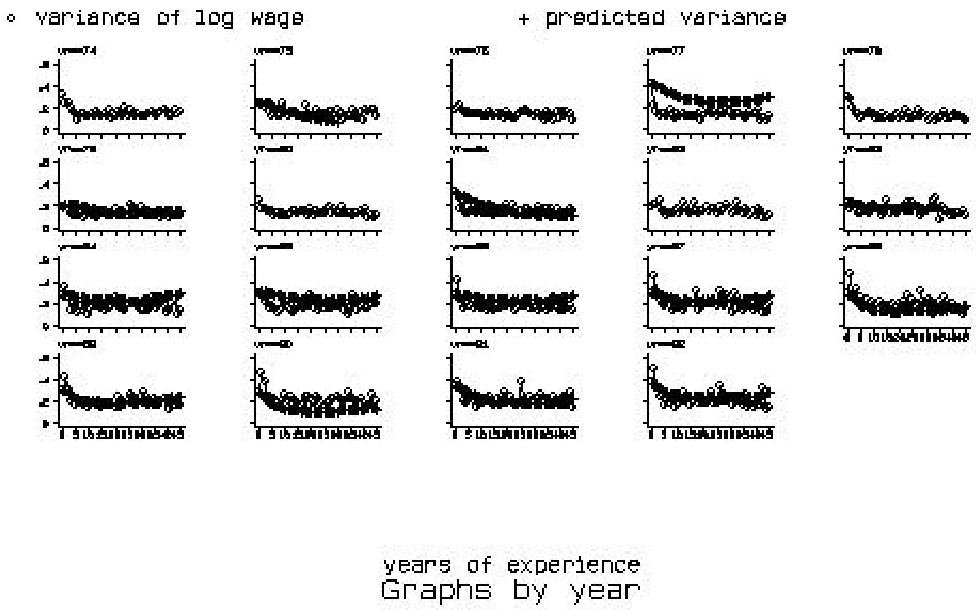

REFERENCES

Abraham, K. G. and H. S. Farber (1987), 'Job Duration, Seniority, and Earnings', American Economic Review, 77(3), p.278.

Altonji, J. G. and R. A. Shakotko (1987), 'Do Wages Rise with Job Seniority?', Review of Economic Studies, 54(3), p.437.

Altonji, Joseph G. and N. Williams (1992), 'The Effects of Labor Market Experience, Job Seniority, and Job Mobility on Wage Growth', National Bureau of Economic Research Working Paper 4133, August 1992.

Ben-Porath, Y. (1967), 'The Production of Human Capital and the Life-Cycle of Earnings', Journal of Political Economy, 75, p.352.

Burdett, K. (1978), 'A Theory of Employee Job Search and Quit Rates', American Economic Review, 68(1), p.212.

Burdett, K. and D. Mortensen (1989), 'Equilibrium Wage Differentials and Employer Size', unpublished, University of Essex.

Dickens, R. (1996), 'The Evolution of Individual Male Earnings in Great Britain, 1975-94', unpublished, CEP, LSE.

Gosling, A., S. Machin and C. Meghir (1994), 'The Changing Distribution of Male Wages in the UK, 1966-92', unpublished, Institute for Fiscal Studies.

Hartog, J. and C. Teulings (1996), 'Corporatism or Competition: Labour Contracts. Institutions and Wage Structures in International Competition', unpublished manuscript, University of Amsterdam.

Jacobson, L. S., R. J. LaLonde, *et al* (1993), 'Earnings Losses of Displaced Workers', American Economic Review, 83(4), September 1993, pp.685S709.

Katz, L. F. and K. M. Murphy (1992), 'Changes in Relative Wages, 1963-1987: Supply and Demand Factors', Quarterly Journal of Economics, 107(1), February 1992, pp.35S78

Marshall, R. C. and G. A. Zarkin (1987), 'The Effect of Job Tenure on Wage Offers', Journal of Labor Economics, 5(3), July 1987, pp.301S24.

Manning, A. (1997), 'Got a Mighty Good Thing, Ain't Gonna Give It Up: Wages, Experience and Tenure in a Search Model', unpublished, LSE.

Mincer, J. (1958), 'Investments in Human Capital and Personal Income Distribution', Journal of Political Economy, 66.

Mincer, J. (1974), *Schooling, experience and earnings*, National Bureau of Economic Research, New York.

Murphy, K. M. and F. Welch (1990), 'Empirical Age-Earnings Profiles', Journal of Labor Economics, 8(2), April 1990, pp.202S29.

Pissarides, C. (1994), 'Search Unemployment with On-the-Job Search', Review of Economic Studies, 61, pp.457S75.

Ruhm, C. J. (1991), 'Are Workers Permanently Scarred by Job Displacements?', American Economic Review, 81(1), March 1991, pp.319S24.

Schmitt, J. (1995), 'The Changing Structure of Male Earnings in

Britain, 1974S88', *Differences and Changes in Wage Structures*, University of Chicago Press: Chicago.

Topel, R. (1986), 'Job Mobility, Search, and Earnings Growth: A Reinterpretation of Human', in Ehrenberg, Ronald G. (*ed*), *Research in Labor Economics*, Volume 8 (Part A).

Topel, R. H. (1991), 'Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority', Journal of Political Economy, 99(1), February 1991, pp.145S76.