



IRISS at CEPS/ INSTEAD

An Integrated Research
Infrastructure in the
Socio-Economic Sciences

EQUILIBRIUM JOB SEARCH AND GENDER WAGE DIFFERENTIALS IN THE UK

by

Audra J. Bowlus and Louise Grogan

IRISS WORKING PAPER SERIES

No. 2001-06



**Directorate General
XII**

Please pay a visit to <http://www.ceps.lu/iriss/iriss.htm>



IRISS-C/I

*An Integrated Research Infrastructure in
the Socio-Economic Sciences at
CEPS/Instead, Luxembourg*

A project supported by the European Commission, the
Ministry of Culture, Higher Education and Research
(Luxembourg) and the National Science Foundation
(USA)

RESEARCH GRANTS

for individual or collaborative research projects
(grants awarded for periods of 4-6 weeks)

Who may apply?

IRISS-C/I supports collaborative and/or internationally comparative research in economics and other social sciences. We encourage applications from all interested individuals who want to carry out research in the fields of expertise of CEPS/Instead. A separate funding agreement is in place for US researchers. Their access to IRISS-C/I is funded by the National Science Foundation.

What is offered by IRISS-C/I?

Free access to the IRISS-C/I research infrastructure (office, computer, library...);
Access to the CEPS/Instead archive of micro-data;
Technical and scientific assistance;
Free accommodation and a contribution towards travel and subsistence costs.

Research areas

Survey and panel data methodology; income and poverty dynamics; persistent and new poverty; gender, ethnic and social inequality; unemployment; segmentation of labour markets; new forms of work/atypical employment; education and training; social protection and redistributive policies; impact of ageing populations; intergenerational relations; effects of family-related policies; regional development and structural change.

Additional information and application forms

IRISS-C/I
B.P. 48 L-4501 Differdange (Luxembourg)
IRISS-C/I homepage: <http://www.ceps.lu/iriss/iriss.htm>

Equilibrium Job Search and Gender Wage Differentials in the UK

Audra J. Bowlus
University of Western Ontario

Louise Grogan
Canadian International Labour Network and McMaster University

June 2001

Abstract

The role of gender differences in labour market behaviour in determining the UK male-female wage differential is examined using the British Household Panel Study and the general equilibrium job search framework of Bowlus (1997). We find that search behaviour explains 30-35% of the gender wage differential. This is similar to US findings. Despite more generous maternity policies, females in the UK are more likely to exit to non-participation. Finally, we find the level of search friction is lower in the UK than in the US due to low job destruction rates in the UK.

Keywords: labour force participation, search models, gender wage differentials

JEL codes: J64, J62, J16

Correspondence to: Louise Grogan, Department of Economics, McMaster University, Hamilton, ON, Canada L8S 4M4. Tel: (905) 525-9140 x 23221, Fax: (905) 521-8232.

Bowlus is a CIBC Junior Research Fellow in Human Capital and Productivity. The authors are grateful to Brendan Halpin at the ISER, University of Essex, the staff of CEPS/INSTEAD and IRISS-CI facility (which sponsored the initial stages of the data preparation), Stephen Jones and participants at the September 2000 Canadian International Labour Network (CILN) Conference in Burlington ON, and the Social Sciences and Humanities Research Council of Canada.

1 Introduction

Gender wage differentials are pervasive across countries, ages, and skill groups. The UK is no exception. Data from the British Household Panel Survey (BHPS) reveal a female-male hourly earnings ratio of approximately 75% for the early 1990's. This ratio is similar in magnitude to that found for the US and other developed countries.

Gender wage differentials are often related to productivity differences between men and women. However, in reduced form regressions observed productivity differences rarely account for all of the observed differential (see, for example, England, 1982). The remainder is often attributed to discrimination against women in the labour market. Another possible source of gender wage differentials is differences in labour market behaviour between males and females. Search behaviour has long been noted as a potential source of wage differentials (see, for example, van den Berg, 1990a, 1990b). However, it is difficult to quantify in a reduced form econometric setting. The advent of equilibrium search models has allowed researchers to examine the role of search behaviour differences in determining wage differentials. For example, Bowlus et al. (2001) find that a large portion of the black-white wage differential in the US can be traced back to differences in search behaviour. With respect to male-female wage differentials, Bowlus (1997) finds that 20-30% of the US male-female wage differential can be explained by search behaviour differences across males and females.

Given the magnitude of the gender wage differential in the UK is similar to that in the US, one might expect the forces behind the differential to be similar in magnitude as well. However, with different institutional and labour market structures there is no reason to expect a similar level of search friction across the UK and the US nor similar differences between the labour market behaviour of males and females. In fact Ridder and van den Berg (1999) find that search friction levels do vary across countries with European countries exhibiting more search friction than the US. Our study is focussed, therefore, on determining the levels of search friction for males and females in the UK and the extent to which they contribute to gender wage differentials. To this end it closely follows the framework laid out in Bowlus (1997) enabling comparisons of the UK results to those for the US.

Some of the differences in search friction across countries that Ridder and van den Berg (1999) identify may well be related to differences in labour market institutions and policies affecting worker and firm behaviour. For example, Ridder and van den Berg look at the role of minimum wage policies. Given the focus of our study is on gender wage differentials, an important source of differences in behaviour across countries may be differences in maternity leave policies. In general the UK has a more generous legislated policy than the US, which until recently had no national policy on family leave.¹ However, it is difficult to predict how such policy differences may affect behaviour and hence gender wage differentials. The US system is often

¹Under the Family Leave Law, women in the US can now take up to 3 months of unpaid maternity leave without fear of losing their jobs.

interpreted as more flexible and less constraining to firms and may therefore reduce hiring frictions. In contrast, longer, legislated maternity leaves may allow women to remain employed around the time of childbirth and thus result in fewer exits to non-participation. In support of this notion Rønsen and Sundström (1996) find that the extension of maternity leave benefits in Norway reduced the level of exits to non-participation around childbirth, and increased the rate at which mothers returned to work after childbirth. Waldfogel et al. (1999) compare the effect of family leave policies in the US, UK, and Japan on the labour force attachment of new mothers. They find that family leave coverage increases the likelihood that women return to work after childbirth in all three countries. In the model we present these tendencies would result in higher wages for females.

Gender differences in labour market behaviour, especially with respect to non-participation and child rearing, have long been studied in labour economics.² The literature documenting such differences for the UK has emerged more recently, due to the fact that longitudinal, individual level data were not available before the 1990's. This changed in 1991 with the start of the British Household Panel Survey (BHPS), which continues today. Booth et al. (1999) use waves 1 to 5 of the BHPS to examine differences in the participation rates of males and females. They find the year-on-year persistence in paid work propensities is higher for males than for females indicating females have less labour force attachment than males in the UK. Booth et al. also find the year-on-year persistence of non-work is higher for females than for males. A key determinant of these propensities for females is household structure, in particular the presence of children.

The Booth et al. (1999) study establishes for the UK the existence of differences across men and women in their labour market behaviour, particularly amongst women with young children. In this paper we examine the role these differences and others play in determining gender wage differentials. To do so we adopt the model and estimation methodology in Bowlus (1997). Bowlus presents a three-state general equilibrium search model of the labour market and shows that within the framework a higher tendency to exit to non-participation by females will in itself result in a gender wage differential. This is because the higher exit rate for women leads to a lower reservation wage for women and lower wage offers from firms to women. The higher exit rate also prevents women from climbing the wage distribution via on-the-job search as fast as men.

To estimate the model we use data on labour force status, unemployment spells, non-participation spells, job spells and wages from the BHPS. While we try to stay as close to the Bowlus (1997) study in our sample construction as possible for comparison purposes, differences between the BHPS and the National Longitudinal Survey of Youth (NLSY), the data set used by Bowlus, dictate some deviations. These deviations, however, allow us to examine some aspects of the relationship between labour market behaviour and wage differentials that were not explored by Bowlus. For ex-

²Recent examples include Wright and Ermisch (1991) and Elias and Gregory (1994) for the UK and Wellington (1993) and Blau and Kahn (1997) for the US.

ample the BHPS contains a representative sample of the population. In contrast, the NLSY is restricted to a sample of young individuals just entering the labour force after finishing their education. Therefore, in our analysis we present estimates for a wider and more representative age range for the UK. Because Bowlus focussed on individuals just entering the labour force, the wage analysis focussed on accepted wages at first jobs. In contrast the BHPS gives a stock sample of individuals who are already in the labour force and thus our focus is on the cross-sectional distribution of earnings. In terms of the model the latter contains the effect that women will move up the wage distribution at a slower rate if they are more likely to exit to non-participation than men. Further, the estimates presented here using the BHPS stem from a more recent time period than those in Bowlus. The BHPS starts in 1991, whereas the NLSY began in 1979.

Examining the BHPS data reveals that, as expected, there are behavioural differences between males and females in the UK and these differences are similar to those in the US. Most importantly, in both countries females are more likely to enter non-participation at the time of a family concern, i.e. the birth of a child, than males. Females also have shorter job durations than males and males are more likely to make a job-to-job transition. There are, however, some differences across the countries. In the UK females with a university degree education stay in non-participation for shorter periods than those with a high school equivalent (O-level) education, while in the US the opposite is the case. Females in the UK also appear to be more likely to go into non-participation than those in the US. Finally, the UK exhibits a much higher job-to-job transition rate than the US for both males and females.

The estimation results indicate that the above behavioural differences between males and females in the UK do have an effect on the male-female wage differential. As found in Bowlus (1997) the majority of the wage differential in the UK is due to productivity differences for both education groups. However, 30% of the differential for higher educated workers is due to search behaviour differences, while the figure is 35% for O-level educated workers. These figures are similar to those in Bowlus: 25-40% for high school graduates and 20-27% for university graduates. Interestingly the level of search friction we find for the UK is lower than that in the US. This is primarily due to the low estimated job destruction rates in the UK.

The paper is organised as follows. Section 2 gives a brief overview of the model and estimation methodology taken from Bowlus (1997). Section 3 discusses the sample construction within the context of the model and presents evidence from the BHPS on important male-female differences in labour market patterns. The estimation results for the UK are presented in Section 4 and compared to those for the US. Conclusions are given in Section 5.

2 Model and Estimation Procedure

To study the effects of male-female behavioural differences on gender wage differentials we use the model and estimation procedure outlined in detail in Bowlus (1997).

Here we provide a brief overview of that framework. The search model used by Bowlus is a derivative of the Mortensen (1990) general equilibrium search model. It contains three labour market states: employment, unemployment and non-participation. Workers search for jobs while employed and unemployed, but not while they are out of the labour force. Thus to regain employment after a spell of non-participation one must first re-enter unemployment.³ The following transitions are allowed within the model: employment to unemployment, unemployment to employment, job to job, employment to non-participation, unemployment to non-participation and non-participation to unemployment.

The parameters governing these transitions include: λ_0 , the job offer arrival rate while unemployed; λ_1 , the job offer arrival rate while employed; δ , the job destruction rate; η_1 , the arrival rate of a family concern (i.e. birth); and η_2 , the exit rate out of non-participation. Following Bowlus (1997) events governed by η_1 consist of family concerns that raise the value of non-market time in the non-participation state such that all unemployed and employed workers choose to exit to non-participation; η_2 then governs the rate at which this value is lowered such that workers choose to return to unemployment and resume searching for a job. The timing of these family concerns is exogenous to the worker with the decision to exit effectively suppressed as well. In our study of gender wage differentials in the UK labour market such exits consist solely of caring for family members, primarily children. It is likely that these decisions are not exogenous, but rather that they depend on the current labour market state and wage rate. However, to facilitate a cross-country comparison we maintain the model in Bowlus and leave this important extension for further research.

In equilibrium workers adopt a state-dependent reservation wage strategy such that, while unemployed, they accept any wage offer above their reservation wage, r , and, while employed, they accept any outside wage offer higher than their current wage, w . The unemployed reservation wage is solved for by equating the value of unemployment and the value of employment evaluated at r and is given by (Mortensen and Neumann, 1988):

$$r = b + (\kappa_0 - \kappa_1) \int_r^\infty \left[\frac{1 - F(w)}{1 + \kappa_1(1 - F(w))} \right] dw \quad (1)$$

where b is the workers' value of non-market time while unemployed, $F(w)$ is the wage offer distribution and $\kappa_i = \lambda_i/(\delta + \eta_1)$, $i = \{0, 1\}$. The equation for the reservation wage reveals that r increases (decreases) when the arrival rate of offers while unemployed (employed) increases thus making unemployment (employment) more attractive. Note that, if the arrival rates are the same, then $r = b$. The parameters κ_0 and κ_1 are measures of search friction in the labour market as they measure the number of job offers expected during an employment spell. Lower levels of κ_0 and κ_1 indicate a greater presence of search friction with the level of search

³This is consistent with the spirit of job search models, which assume that observed transitions to employment are related to job search activities.

friction decreasing in the rate of job offer arrivals (the λ_i 's) and increasing in the level of job destruction (δ) and the exit rate to non-participation (η_1).

Firms maximise profits in this model by posting a wage. In equilibrium all firms earn the same profit level, but because of on-the-job search firms do not offer the same wage. Some firms offer lower wages and consequently have high per worker profits but small labour stocks, while other firms offer higher wages and make up for low per worker profits with large labour stocks. Via on-the-job search these high wage firms attract workers from the lower paying firms. In equilibrium the wage offer distribution is non-degenerate. The lowest wage offered is r , as all offers below are rejected, and the highest wage, w_H , is less than the highest productivity level in the market. Thus all firms have some monopsony power. Note that as on-the-job search becomes more effective, i.e. λ_1 increases, firms lose monopsony power and the wage distribution collapses to the competitive price, and without on-the-job search, i.e. $\lambda_1 = 0$, all firms offer only r . If the market contains only one firm type with productivity level P , profit maximisation implies the following solution for $F(w)$ (Mortensen, 1990):

$$F(w) = \frac{1 + \kappa_1}{\kappa_1} - \frac{1 + \kappa_1}{\kappa_1} \left[\frac{P - w}{P - r} \right]^{1/2}, r \leq w \leq w_H. \quad (2)$$

Because of on-the-job search the cross-section earnings distribution, $G(w)$, is not the same as the wage offer distribution. Over time workers move up the wage offer distribution such that the earnings distribution lies to the right of the offer distribution. The earnings distribution is given by the following formula:

$$G(w) = \frac{F(w)}{1 + \kappa_1(1 - F(w))}. \quad (3)$$

Gender wage differentials can be generated in this model easily by allowing firms to post gender-specific wage offers.⁴ For example, if the arrival rate of family concerns is higher for women than for men, women will earn, on average, less than men. This occurs for three reasons. First, the reservation wage for women will be less than that for men which will shift the wage offer distribution for women to the left of that for men. Second, firms will have more monopsony power over females due to the greater search friction levels for women and will respond by offering lower wages to women. Third, women will climb up the wage distribution at a slower rate than men and hence an even larger earnings differential will emerge. Of course, there can be other differences between males and females that contribute to the observed wage differential. Only through estimating the model can these different forces be sorted out.

Before turning to the estimation methodology we note that the homogeneous productivity version of the model presented above does a poor job of fitting observed wage data. Therefore we follow Bowlus et al. (1995, 2001) and estimate the model

⁴It is possible to generate gender wage differentials if wage offers are not gender-specific. However, the model is easier to solve and estimate under the gender-specific assumption.

assuming discrete productivity heterogeneity. A full description of the estimation methodology for the three-state model is given in Bowlus (1997). As in Bowlus the likelihood function is composed of pieces related to wages, non-employment and job durations, and transitions. Because the BHPS starts with a stock sample as compared to the NLSY flow sample, the data we collect differ slightly from that collected by Bowlus, and, therefore, the likelihood function must be modified.

First, because we observe the stock of employed and non-employed workers, we collect the labour force state of each respondent at the start of the survey in 1991. The likelihood of being in each of the three states - non-participation (family care), unemployment, and employment - is, respectively,

$$\Pr(N) = \frac{\eta_1}{\eta_1 + \eta_2}, \quad (4)$$

$$\Pr(U) = \frac{\eta_2(\delta + \eta_1)}{(\eta_1 + \eta_2)(\delta + \eta_1 + \lambda_0)}, \quad (5)$$

and

$$\Pr(E) = 1 - \Pr(N) - \Pr(U). \quad (6)$$

Second, because we sample from the stock we have a stock sample of durations – non-employment durations and job spells. This means that we have over sampled long spells and must account for this in the log likelihood function. As Bowlus (1998) points out stock sampled durations under the assumption of Poisson arrival rates are sampled from a gamma distribution whereas spells sampled from the flow are exponentially distributed. However, we use only the residual portion of each spell. That is, the duration of the spell after the survey date. With spells that have an underlying exponential distribution, residual durations are also distributed as exponential.⁵ Thus the job spells in our sample (residual and flow) are distributed exponential with parameter $\lambda_1(1 - F(w)) + \delta + \eta_1$. The non-employment spells are also distributed as exponential with the parameter depending on the state (non-participation or unemployment) at the start of the spell. Spells (residual or flow) that are in unemployment when first observed are exponential with parameter $\lambda_0\eta_2/(\eta_1 + \eta_2)$ while those in non-participation are exponential with parameter $\lambda_0\eta_2/(\lambda_0 + \eta_1 + \eta_2)$. Note that η_2 is identified only if the mean duration of spells starting in non-participation is larger than the mean duration of spells starting in unemployment.

Third, again because of stock sampling, the wages from the 1991 survey are sampled from the cross-section earnings distribution instead of the wage offer distribution. Thus these wages are distributed according to $g(w)$, the probability density function (pdf) of $G(w)$, while those accepted from unemployment are distributed according to $f(w)$, the pdf of $F(w)$. Finally, we enter into the likelihood function the transitions workers make following the completion of their job spells. These transitions can take

⁵This result is particular to the exponential distribution.

three forms: employment to non-participation, employment to unemployment and job-to-job. The transition probabilities are, respectively,

$$\Pr(E \rightarrow N) = \frac{\eta_1}{\delta + \eta_1 + \lambda_1(1 - F(w))}, \quad (7)$$

$$\Pr(E \rightarrow U) = \frac{\delta}{\delta + \eta_1 + \lambda_1(1 - F(w))}, \quad (8)$$

and

$$\Pr(E \rightarrow E) = \frac{\lambda_1(1 - F(w))}{\delta + \eta_1 + \lambda_1(1 - F(w))}. \quad (9)$$

The final likelihood function is then the product of all the above components after appropriately dealing with the censoring of durations. We follow Bowlus et al. (1995, 2001) and use simulated annealing to handle the discontinuities in the log likelihood function due to the form of the wage distribution. We refer the reader to these studies for details. We also follow Bowlus (1997) by estimating only a two-state model for males. As we show in the next section so few males are in the family care state or exit to the family care state that this simplification has little effect on our results.

3 Data

To estimate the above model and examine the relationships between gender wage differentials and labour market behaviour in the UK we use the BHPS. The BHPS is a representative survey at the national level of all private households in the UK. The first wave was conducted in 1991 and we have data through 1998. The BHPS includes information on current labour market status, remuneration from work, transitions made between interview periods, and the reasons for such transitions.

3.1 Sample

To create our sample we start with the stock of individuals present in the first wave of the BHPS. To be included in our sample individuals must be either working, unemployed or non-participants because of family care needs at the 1991 interview. Thus students, those on government training programs, those who do not report a labour market status, the self-employed, and those who are retired at the 1991 interview are excluded.⁶ As well, individuals who are observed to transit directly to retirement, training schemes, self-employment or higher education from their 1991 state are dropped from the sample. Finally, we select individuals who have had at least one work spell since completing their education.

⁶These are standard exclusion restrictions in the search literature, and allow us to focus on behaviour related to exit from and entry to non-participation for child-related family reasons.

We restrict the age range of the sample to 20-40 years of age in 1991 giving us an average age of 30. This draws in a considerably larger and more diverse group of workers than that used in Bowlus (1997) for the US where the mean age for high school graduates was 18.7 and 23.9 for college graduates. We are not able to restrict the age range further to conduct a more direct comparison, because of small sample sizes in the BHPS. However, we have the advantage of being able to observe behaviour differences amongst a wider age range of men and women and an age range where exits to and returns from non-participation may be more likely. For example, only 10% (19%) of high school (college) graduates are married at the start of their jobs in the Bowlus sample, while about one third of our sample is continuously married throughout the sample period.

We divide our sample into two education groups: those educated at the O-level and those with a university degree. We consider individuals to have achieved a university degree if they have completed a bachelor's degree or higher, a teaching degree, or a nurse's qualification. This classification is roughly analogous to the college graduates of the Bowlus (1997) study. Individuals with O-level qualifications are those that passed at least one of a set of subject-specific exams at age 16, and then stopped their education. Although they are slightly younger, this group is comparable to the group of high school graduates in the Bowlus study.

We attempt to follow each individual in the sample through one full job cycle. That is, we include information from the start of a job until the start of the next job. We define a job spell as a continuous period of work for a single employer. Thus we use an employer-based definition instead of a task or position-based definition. This definition is comparable with that used in the NLSY and requires us to combine job spells in the BHPS that occur at the same employer.

For those who are employed at the 1991 survey date, we follow them from the survey date until their job spell is censored or completed.⁷ A job spell can be censored in the BHPS data for three reasons: the end of the sample period, attrition from the sample, or the inability to link job spells across survey dates.⁸ If a job spell is observed to complete during the sample period, we record the type of transition that takes place after the job spell. The transition can take one of three forms: a job-to-job transition, a transition to unemployment due to job loss, or a transition

⁷It is not possible to follow them from the start of the job itself, because we only have information regarding the starting date of the current position, not the starting date at the employer, in 1991.

⁸Despite the practice of "feed forward" in the BHPS, where interviewers are given information about the labour market status of respondents in previous interviews to clarify apparent contradictions in responses across waves, there are still some substantial differences in year-on-year records of the timing of labour market transitions. This results in what is called a seam problem and can make linking job spells across survey years and timing the end of a job spell difficult. In our construction of the spells we treat such apparent recall errors according to a rule assuming that the labour market status changes reported closest to the date of the transition are correct. If the reported start dates of a given spell vary by more than three months across subsequent interviews, the spell is censored at the date of the earlier interview. If the inconsistency in reports is less than three months, the spell is not censored and the date of the labour market transition is set on or about the start date of the earlier interview.

to non-participation due to a family concern. If the individual makes a transition to unemployment or non-participation, we also record the length of the non-employment spell until the start of their next job or to the end of the sample period.

For those who are not employed at the 1991 survey date we record whether or not they are unemployed or non-participants and then follow them until their non-employment spells are censored or they make a transition to employment. If they do find employment, we repeat the above procedure and follow them through a full job spell cycle. As in the Bowlus (1997) study we only record the state at the start of the non-employment spell and do not record transitions between unemployment and non-participation.

Our treatment of temporary absences from a job is different from that in Bowlus (1997). Bowlus subsumed temporary lay-off spells of less than 3 months into the job spell. This is preferred if one is trying to identify more permanent separations and actual search activity. However, it is not possible to do this with the BHPS, because we can not determine whether or not the individual returned to the same employer after an unemployment spell. Thus we treat all lay-offs, no matter how short, as unemployment spells. Bowlus also treats all employment spells with less than 20 hours per week as non-employment. We are unable to conduct a similar treatment on our sample, because hours of work are not recorded for all job spells in the BHPS. Thus all job spells are treated as valid job spells. Our treatment of wages earned in part-time jobs is discussed below.

With regard to maternity leave absences from work Bowlus (1997) again subsumed all maternity leave spells that were shorter than 3 months into the job spell. This choice was arbitrary for Bowlus given the US did not have an established maternity leave policy in the late 1970's and early 1980's. However, during our sample period the UK did have an established maternity leave policy. From 1991-1994 pregnant women had the right to up to 40 weeks of maternity leave if they had completed two years of continuous employment by the beginning of the eleventh week before the expected week of childbirth. With the adoption of the European Union standards in 1994 the official maternity leave policy was augmented to include at least 14 weeks of leave for all women independent of job tenure, and a compulsory 2 week period after the birth.

Since the change in the policy extended coverage but for a shorter duration, one would expect more women to take maternity leave after the policy change but that average durations would shorten. In fact we find no effect of the policy on the likelihood of maternity leave and longer, not shorter, maternity leave spells after 1994. It appears likely that firm policies with respect to maternity leaves were already in place and dominated the government legislation.⁹ Since more than 70% of maternity leave spells observed amongst our sample members occurred before the adoption of EU maternity leave legislation, we subsume all maternity leave spells shorter than

⁹Dex et al. (1996) find that employers have increasingly provided additional career-break schemes, top-ups to maternity benefits, workplace nurseries and flexible hours to working mothers in the UK.

40 weeks into the job spell. Those longer than 40 weeks we treat as spells of non-participation. Because the additional coverage after 1994 was for only 14 weeks, our limit of 40 weeks will treat these women on maternity leave appropriately.

To complete the job spell cycle we must determine how to code the transitions after a job spell ends. For the most part this has been done in the literature by using the observed spell following the job spell. That is, if the observed spell is another job spell, then a job-to-job transition is recorded; if it is an unemployment spell, then a job loss is recorded; and if it is a non-participation spell, then a family concern transition is recorded. It is also possible to record the transitions using information provided by the respondent on the reason why they left the job. Thus, if the reason is to take another job, a job-to-job transition is recorded; if it is plant closure, lay-off or fired, a job loss is recorded; and finally if it is maternity leave or family care, a family concern transition is recorded. These two coding schemes do not necessarily give the same results. For example, some individuals report that they left their previous job due to lay-offs, but are observed to be employed at another job in the next spell. Thus they do not experience an intervening period of unemployment. As we must use the reason for leaving to code the non-employment spells into unemployment and non-participation, we have decided to also use this method to code the job-to-job transitions. This choice does not significantly impact observed job-to-job transition rates for any of the groups.

Finally, we collect wage information for each job spell. We use the BHPS composite of net earnings reported in the previous payment period, the time period that the previous payment period included, and the hours of work in the previous period to construct a full-time equivalent weekly wage, based on 37 hours of work per week. We deflate wages to September 1991 using the monthly Consumer Price Index. Individuals for which this information is missing are not dropped from the sample, but only their non-employment durations contribute to the estimation procedure. We convert wages to full-time equivalent levels so that wages reveal an hourly price and do not reflect labour supply decisions that are not modeled. Alternatively, we could have used weekly wages without adjusting for hours of work, with the implicit assumption that firms offer workers weekly wages that workers must accept or reject. Estimation using unadjusted wages would likely lower the lowest observed wages of women dramatically and would make the apparent male-female wage differential for each education group much larger. Due to extreme outliers in the data we trim the wage samples 5% at the top and bottom.¹⁰

Unfortunately many job spells are missing information on hours and thus do not have a wage associated with the spell. This is true of all job spells that occur entirely between two interviews. Surprisingly, this affects a large number of spells. For a small number of these spells we do know the wage at the interview immediately following the spell and we use the interview date wage to impute the missing wage. In general we have wage data for most working individuals at the start of the survey in 1991,

¹⁰This is a common practice in the estimation of search models, because it aids in the estimation of the productivity parameters.

but are able to collect only limited wage information for many later job spells.

3.2 Descriptive Statistics

Before estimating the model with the above data, we first examine the data and provide a brief overview of their salient features. We are particularly interested in the differences in labour market behaviour between males and females. Table 1 provides various sample statistics of interest.

As expected we find a higher employment rate amongst men than women, and a higher rate amongst more educated workers. The fraction of respondents who are unemployed at the start of the survey in 1991 is much higher for men than for women. Amongst men, unemployment is higher for the lower educated group. Surprisingly, the female unemployment rate is quite low and does not vary across the education groups. The remaining fraction of respondents is in non-participation (defined as family care) at the start of the survey. This fraction is effectively zero for both groups of men and is twice as high for lower educated women than for higher educated women.

The lower level of labour force attachment for lower educated females is also found in the work history files of the BHPS.¹¹ Amongst O-level females in our sample 59% has had at least one spell of non-participation since completing full-time education. This compares to 37% of university females. As well, the lower educated group is more likely to have children under the age of 16 at the beginning of the panel.¹² Interestingly the university females in our sample are more likely to have had a previous maternity leave spell than the O-level females, 24% versus 21%. Given that more O-level females have children, this result likely stems from the higher labour force participation rate of university females.

The fourth row in Table 1 shows the mean residual duration of job spells that are ongoing at the time of the 1991 interview. That is, we have calculated the average length of time spent in these jobs after the interview date. On average, employed men and women work 2.5 to 3 years after the survey date before making a transition. For both education groups males have, on average, longer job spells than females. The difference between the sexes is greater for the university group than for the O-level group. The next row shows the job spell mean following an initial spell of non-employment. These spells are shorter than those underway in the stock because they are more likely to be censored. As well, they are less numerous. These spells reinforce the relationships of longer durations for males than females in the low education group and similar mean durations amongst males and females in the high education group. The job spell data in the BHPS have high censoring levels, especially for the O-level sample.

¹¹This information is contained in a data set created by Brendan Halpin at the University of Essex using information from the BHPS working life history files. We thank Brendan for passing his data sets on to us.

¹²The BHPS collects information on the age of all household members, but it is not known whether or not these are biological children of one or both parents.

Row 7 indicates a very high rate of job-to-job transitions in the data. The fraction of completed job spells that end in a quit to another job ranges from 59% to 85%. This compares to figures ranging from 34% to 52% for the US in Bowlus (1997). Given that Bowlus was working with a younger sample, these differences are quite astounding and indicate a significant difference in labour market behaviour across the two countries. Across gender and education a similar pattern emerges in the UK as in the US. Men have a higher tendency to exit to another job than women. Higher educated workers are also more likely to make a job-to-job transition.

With regard to exits to non-employment, we see that almost no men exit to family care, while a substantial fraction of women do. The fraction of women exiting to family care is higher for university graduates. The greater tendency of university females to exit to family care is not consistent with the stock levels of non-participants found in 1991. This may indicate a greater tendency of higher educated females to exit for family care reasons as they get older; or it may indicate the lower stock levels stem from a higher exit rate for university females. In comparing women in the UK with those in the US, we find that of those going into non-employment a far greater fraction of women in the UK enter into family care. In the US the percentage was less than 20%; unemployment is a much more likely destination (Bowlus, 1997). Again this may be related to the older age group of the UK sample, but the difference is so large that it likely indicates a significant structural difference between the two markets.

Related to the transition to non-employment is the length of time spent in each state. The mean duration of spells starting in unemployment is shorter for females than males in both education groups.¹³ On average women spend just less than one year in unemployment while men spend slightly more. With respect to non-participation we find that O-level females have longer durations than university females, consistent with their higher stock percentages. Here again the censoring rates are quite high.

Finally we turn our attention to wages. At the start of the sample period in 1991 we find a substantial education premium as well as a gender wage gap. University females earn on average 84% of the salary of university males, while O-level females earn 78% of the salary of O-level males. Thus females appear to fare better in our sample than in the national statistics for the UK. This is likely due to our use of full-time equivalent wages. The second row of mean wages in Table 1 shows the mean wages of those individuals who find employment after being non-employed in 1991. We would expect the mean wage following a non-employment spell, either unemployment or non-participation, to be lower than the mean of the cross-section wage distributions. This is true for all of the groups except higher educated females. As shown by the large standard errors these means are not estimated with much

¹³Here we have included both the residual non-employment spells from the start of the survey and those that occur after a transition to non-employment during the course of the sample period. This is done because of small sample sizes and because the model treats these spells as from the same distribution.

precision due to small sample sizes.

4 Estimation Results

We now use the above data sample to estimate the model following the procedure outlined in Section 2. We estimate the model separately for males and females for both education groups. The parameter estimates for the four subgroups are shown in Table 2.¹⁴

For the O-level group we find the following relationships across males and females. First, females have a much higher job offer arrival rate while unemployed. Females are exiting unemployment almost twice as fast as males. Such a high exit rate for females helps to keep their unemployment rate and mean unemployment duration low. Second, females and males experience more similar offer arrival rates while employed, although again females have a higher job offer rate. Females also have a much lower job destruction rate than males. This would put females at an advantage in terms of search friction levels if not for their high rate of entrance into non-participation. The exit rate to non-participation for females is of a similar magnitude to the job destruction rate for men and thus overall females exit firms to non-employment at a faster rate than men. These differences result in females facing less search friction while unemployed (κ_0) than males but more while employed (κ_1). Since κ_1 is the parameter combination that enters the wage offer and earnings distributions, this difference between males and females helps to explain the presence of the wage differential.

Search friction, however, is not the whole story. We see that O-level females also have, on average, a lower average productivity level. The firms hiring females have a mean productivity level that is lower than those hiring males, and the difference is greater across workers in the cross-section. Thus the model is not able to explain the full gender wage differential through differences in search behaviour; productivity differences play a role as well. Finally we find that, as expected, the reservation wage is lower for females. Even though the lower reservation wage for females was expected, it is not necessarily consistent with the other parameter estimates for females. Since females have a much higher level of κ_0 than males and κ_0 raises the reservation wage, the model would predict (given similar productivity distributions) a higher reservation wage for females than males. Offers are arriving so quickly while unemployed that they should be pickier. In this case, the value of κ_0 is so high (and the productivity differences not so large) that the only way the model can explain the reservation wage strategy of the females is to give them a much smaller value for b than that for males.¹⁵

With respect to university graduates we find a similar pattern. Females again have a higher job offer arrival rate in unemployment than males. They also have a slightly higher job offer arrival rate while employed. Their job destruction rate is lower, but they have a high exit rate to non-participation. Together these two exit

¹⁴For all four groups we found that 5 firm types were needed to fit the wage data well.

¹⁵In fact the implied value of b is negative.

avenues result in a higher exit rate to non-employment for females. They have a significantly higher value of κ_0 than males, but a slightly lower value of κ_1 . Again the lower value of κ_1 indicates a role for search friction in explaining the gender wage differential. The female reservation wage is lower than the male and here again the model has to give the females a low value of b to explain the low value for r .

These results differ somewhat from those found for the US by Bowlus (1997). Bowlus found that high school and college graduate males faced lower levels of search friction (higher values of κ_0 and κ_1) while unemployed and employed than females. The lower level while unemployed thus contributed to the explanation of a lower reservation wage for females. For high school graduates this difference occurs because, unlike in the UK, both λ_0 and λ_1 are lower for females than males. It occurs for the college graduates, because even though λ_0 and λ_1 are higher for females, as in the UK, the difference is not large enough to counter the females' higher exit rate to non-employment. Thus, the two countries display similar orderings across males and females with respect to the different arrival rates, but the ratios lead to different conclusions regarding search friction.

The most surprising result is perhaps that the UK education groups generally display lower levels of search friction than their US counterparts. The only exception to this finding is the level in unemployment for O-level males. These cross-country differences reflect the low rates of job destruction found in the early 1990's in the UK. Other differences can be found comparing each group across the two countries. For example, lower educated males in the US display higher job offer arrival rates while unemployed and employed than lower educated males in the UK. However, lower educated US workers also display much higher job destruction rates - on the order of four times greater for males and even more for females. With respect to the groups with higher education, we find that the US job offer arrival rate while unemployed is higher than that for the UK for males, but not for females. The rate while employed is slightly higher in the UK for both sexes. The job destruction rate is also much higher in the US.

The lower estimated job destruction rates for the UK may be partly explicable by stronger employment protection legislation (EPL) in the UK relative to the US. According to the Heckman and Pages (2000) 'job security' index of the wage costs of worker dismissals, the UK exhibited more employment protection than Norway, Germany, France, and Switzerland in the 1990's. In the present search model, no explicit account is made for the effects of EPL on firms' hiring or firing choices. However, in matching models such as that of Mortensen and Pissarides (1998), strong EPL will generally result in lower equilibrium labour demand and lower job destruction rates. Thus, the low job destruction rates, which here suggest less search friction in the UK than in the US, may in fact be capturing a rigidity that is not accounted for in our model.

With respect to non-participation, it is interesting to note that despite more generous maternity leave benefits females from both education groups in the UK are more likely to enter non-participation than their respective US counterparts. For

example, female college graduates in the US exit to non-participation at half the rate of female university graduates in the UK. And, even though females in the UK have longer expected employment spells, due to their much lower job destruction rates, conditional on exiting to non-employment they are 4 times more likely to have gone into non-participation.

Following Bowlus (1997), we decompose the gender wage differential into components attributable to (i) reservation wage differences, (ii) differences in κ_1 values, and (iii) differences in the productivity profiles facing the sexes. For university graduates, we find that 10.3% of the difference between male mean earnings and female mean earnings (as calculated by the model) is due to differences in reservation wages. A further 19.2% is due to differences in the κ_1 values, while 70.5% is due to differences in the productivity distributions facing each group. For O-level educated workers the contribution of differences in search frictions to gender wage differentials is larger. We find that 13.0% of gender wage differentials can be attributed to reservation wage differences, 21.9% is attributable to differences in the κ_1 values, and 65.1% attributable to differences in the productivity distribution.

The components of the wage differentials due to κ_1 values and those due to reservation wages may together be considered the search contribution to wage differentials. About 29.5% of the gender wage differential amongst those with a university degree in the UK can be attributed to differences in search behaviour. Amongst O-level individuals about 34.9% of wage differences between the sexes is attributable to search behaviour. For the US Bowlus (1997) finds that the gender wage differential in earnings attributed to search differences is 25-40% for high school graduates and 20-27% percent for college graduates. Thus, in both countries it appears that search differentials are a more important component of wage differentials amongst less educated individuals.

4.1 Fit of the Model

We turn now to examining how well the model fits the UK data. Table 3 shows averages predicted by the model that can be compared to those from the data shown in Table 1. The first row shows the mean wage from the earnings distribution. These predicted means should be compared to the mean wages of the stock of employed workers in 1991. The model is able to capture mean earnings levels fairly well with a slight tendency for the predicted mean to be too high. With respect to mean wage offers, the model in general predicts lower means than those observed in the data. Of course, we expect the predicted mean wage offer to be lower than the predicted earnings mean. However, the large gap in the means predicted by the model stems from the relatively high values of κ_1 . That is, with small levels of search friction the model predicts that agents are able to move up the wage distribution quickly over time. Therefore, the distance between the wage offer distribution and the earnings distribution is relatively large. This is, however, not consistent with the differences between these two distributions observed in the data, and reflects the model's inability to match both wage distributions at the same time. We do note that some of these

problems may be related to the small sample sizes for the wage offer distribution.

An examination of the unemployment and non-participation rates and the predicted durations of non-employment reveals several patterns. For O-level males the predicted unemployment rate is quite close to that in the data, while the predicted mean duration of unemployment is larger. For university males the predicted duration of unemployment is similar to that in the data, while the predicted unemployment rate is too high. For females the predicted mean durations of unemployment are smaller than those found in the data while the predicted unemployment rates are too large. The fit is relatively better for O-level females than for university females. The predicted non-participation rate matches that found in the data for O-level females, while it is too high for university females. Finally, for both groups the predicted duration of non-participation spells is much longer than the mean of those spells in the data, almost two times longer. In all cases the model is trying to balance several features including the length of the unemployment and non-participation durations, the high censoring rates, the low unemployment rates, and the high non-participation rates. In addition it is also trying to match the low exit rates out of jobs to unemployment.

With respect to the job-to-job transition rates the model is not able to match the high rates observed in the data for any of the four groups. The observed data generate job-to-job transition rates on the order of 0.6 to 0.85 while the model is only in the range of 0.5. This is because the model can not reconcile the observed job-to-job transition rates with the other observed features of the data. In order to generate job-to-job transition rates that are as high as those in the data the model needs very high κ_1 values. Even for O-level females, who have the lowest job-to-job transition rate, the value of κ_1 needed is 19. For the other groups the value increases exponentially ranging from 55 for university females to 2972 for university males. Such levels of search friction essentially imply a competitive labour market and thus predict that most workers should be earning the highest productivity level and that job durations at lower wages should be relatively short. As this is not the case in the data, the model is forced to move away from such high values of κ_1 and therefore is unable to reproduce the observed job-to-job transition rates.

Given these rates are so much higher than those in the US, we investigated an alternative UK data source to determine whether this was a feature particular to the UK labour market or an artefact of the BHPS data. Using the UK Labour Force Survey we find a job-to-job transition rate that is close to that of our BHPS sample at 0.72 and higher than that found in the NLSY.¹⁶ Thus, our data do not appear to overstate the degree of job mobility in the UK. Still, we recognise some possible deficiencies in the BHPS coding. Given the lack of an employer-specific code in the BHPS, it is not possible to merge multiple spells with a single employer unless the

¹⁶The spring 1989 UK LFS contains information on labour market status in the year prior to the interview, current labour market status, and the elapsed duration of current labour market spell. Although we do not have a full retrospective work history for the 1988-1989 period, we are able to calculate an upper bound on the job-to-job transition rate.

individual reports a promotion. Thus, we may be overstating the number of job-to-job transitions and the model may be doing a better job of fitting the main features of the UK labour market than is evidenced here.

The above analysis shows that despite some trouble areas the model appears to be able to match most of the averages found in the data. We now turn to examining the model's ability to fit the full distributions of wages and durations. Figure 1 shows the cumulative distribution functions (cdf) of earnings in 1991 for all four groups. Here we find that the model is able to reproduce the earnings distributions fairly well.

Figure 2 presents the empirical and predicted survivor functions for the unemployment durations for all four groups. As noted above the model predicts that the unemployment durations should be exponentially distributed. The graphs show that the data depart from the model the most in the right tails. Estimates of a Weibull model indicate that the exponential distribution assumption is rejected for all groups as there is some negative duration dependence found in the data. However, the fit of the unemployment data is also affected by the tension on the parameters in trying to fit other moments. For all groups the low unemployment rates put upward pressure on the arrival rate parameter λ_0 causing the mean duration to be lower.

Figure 3 presents the empirical and predicted survivor functions for the female non-participation durations for both education groups. Here the exponential assumption can not be rejected when estimating a Weibull model. However, the high rates of censoring preclude any tail analysis. For O-level females the graph shows that the predicted probability of surviving is too high leading to a predicted mean duration that is also too high. Here the model is likely putting pressure on these parameters because it is trying to match the relatively high non-participation rates found in the data. The opposite is true for university females. Here the predicted probability of surviving is too low so that the predicted mean is also too low. The pressure in this case is to keep the non-participation rate low as the predicted rate from the model is too high.

Job durations are assumed in the model to be distributed exponentially with respective intensity parameters for females and males of $\delta + \eta_1 + \lambda_1(1 - F(w))$ and $\delta + \lambda_1(1 - F(w))$. Following Bowlus et al. (2001) we account for the fact that observations depend on the current wage by making the transformation $z_i = (\delta + \eta_1 + \lambda_1(1 - F(w_i)))D_{Ji}$ for females and $z_i = (\delta + \lambda_1(1 - F(w_i)))D_{Ji}$ for males where D_{Ji} is the duration of the job. Under the model these observations are independently and identically distributed (*i.i.d.*) as unit exponential. Therefore, if the model is correct, we would expect the integrated hazard function to be a straight line with slope one when plotted against the transformed durations. Figure 4 presents the plots for conducting such a test. Surprisingly, while one would reject using an eye test, the relationship is closer than one might expect given the strong distributional implications of the model.¹⁷

¹⁷The test performs particularly bad in the right tails. These are censored observations and, given the transformation is not valid for them, this poor performance is to be expected.

4.2 Thought experiments

Finally we repeat the thought experiments in Bowlus (1997) and ask what the mean wages for females are under various different scenarios. Before discussing the results we note that, when resolving the model under different parameter settings, one should take into account the change in the reservation wage. However, as we have already noted, the model has a difficult time reconciling the reservation wage for females with the observed wage and transition data. We find that the reservation wage is sensitive to changes in the search friction parameters for both education groups. Because reservation wages often take on implausible values as a result, we report the results holding the reservation wage fixed in the tables but discuss what would happen if the reservation wage were allowed to vary.

Table 4 contains the thought experiment results for O-level educated workers. The first thought experiment asks what happens to females if they are faced with the same search friction levels (the same κ 's) as males. Here the mean wage offer increases slightly due to the increase in κ_1 , while mean earnings increases more causing the gender wage differential to drop from 21% to 16%. Interestingly if the reservation wage was allowed to respond to the change in κ 's, there would have been no improvement in the average wage of females. This is because the drop in κ_0 would have induced a very large drop in the reservation wage to 37 pounds, and this effect would outweigh the positive benefits of a higher κ_1 .

The second thought experiment sets η_1 to zero. This increases both κ_0 and κ_1 dramatically, causing a large increase in the mean wage offer and an even larger increase in mean earnings. Under this scenario mean earnings for females are actually higher than that for males. With respect to changing r the large increase in κ_0 would result in a substantial increase in the reservation wage to 287 pounds. This reservation wage value is so high that only one productivity type remains profitable and again female mean earnings surpass male mean earnings.

The final case considers changing the firm productivity distribution of the females to that of the males. This results in a larger increase in mean wage offers but a smaller increase in mean earnings than in the $\eta_1 = 0$ case. Allowing the reservation wage to adjust in this last case would result in an upward adjustment to 104 pounds. In conjunction with the increase in productivity values the gender wage differential would be close to being eliminated under this latter scenario.

Table 5 contains the same scenarios for higher educated workers. Because of the similarity in κ_1 values across the sexes, changing to the male levels does not increase mean wage offers or earnings significantly if r is unchanged. If the reservation wage were allowed to change, the drop in κ_0 would result in a large, implausible decrease to 46 pounds, and a drop in both means. Setting η_1 to zero increases both κ_0 and κ_1 substantially. The increase in κ_1 increases the mean wage offer slightly and has a larger effect on mean earnings. In fact mean earnings for females is now higher than that for males when r is unchanged. As in the case for O-level females, if one were to allow the reservation wage to respond, a large increase in κ_0 would result in a large increase in the reservation wage value to 190 pounds. In this case, the mean wage

offer and mean earnings for females would be higher than those for males. The final scenario indicates the importance of productivity differences in explaining the gender wage differentials. If females are given male productivity parameters, the mean wage offer and mean earnings increase substantially, both with and without allowance for changes in r , reducing the gender wage gap. Allowing the reservation wage to increase to 114 pounds in this case would essentially eliminate the wage offer gap and much of the mean earnings gap.

5 Conclusions

In this paper the role played by job search behaviour in governing gender wage differentials in the UK is investigated. Search behaviour differences across males and females are found to account for a significant portion of the gender wage differential for both education levels. This effect is larger for lower educated workers. For both education groups productivity differences play the largest role in determining the gender wage gap. The model does predict that eliminating exits to non-participation for females results in a substantial increase in their mean earnings levels.

While overall gender wage differentials are similar in the US and the UK, there are substantial differences in the search components of contributing to these differentials. Surprisingly, we find that search friction is lower in the UK than in the US. This is despite higher job offer arrival rates in the US, and results from the much lower job destruction rate in the UK. The lower job destruction rate may be partially attributable to the higher level of employment protection in the UK than in the US, which is not considered in the model. Still, our measures of search friction for the UK are similar to what Ridder and van den Berg (1999) find using aggregate data, and our participation patterns are similar to those reported for the UK in Waldfogel et al. (1999) and Booth et al. (1999). Thus, while some of observed the cross-country differences may be related to the age differences in the two samples and to the difficulties we encountered in the BHPS data with respect to merging job spells over time, they are substantial enough to suggest real behavioural differences.

Given the more generous maternity leave policies pursued in the UK than in the US, it might have been expected that exit rates to non-participation would be lower in the UK. In fact, we observe that exit rates for family concerns are higher in the UK than in the US. In part this result is likely attributable to the fact that our sample is older than the US sample of Bowlus (1997), but it may also be partly attributable to our lack of consideration of the part-time/ full-time work distinction. If women know *ex ante* that they will not be able to reduce their hours following the birth of a child, they may opt to leave their current jobs rather than commit to returning. Especially given the high job offer arrival rates and low unemployment durations found in the 1990's, exiting to family care may have been the less constraining option in the UK. The limited information available in the BHPS does suggest that new mothers tend to work substantially fewer hours after returning to the labour force than prior to their exit.

There are two areas of interest that we wish to pursue in the future. The first is dealing more satisfactorily with part-time work. Blau and Kahn (1995) find that far more women work part-time in the UK than in the US (45% versus 24%). Here we have treated all job spells as the same regardless of hours worked per week. Bowlus (1997) also ignored the part-time issue by treating spells with low hours as periods of non-employment. Given the large fraction of females employed part-time, this appears to be an important factor that may contribute to some of the observed differences between the US and UK reported here. The second issue is related to the decision to exit to non-participation. Females in the UK are more likely to enter non-participation than females in the US. As well in the UK higher educated females stay in non-participation for a shorter period of time than lower educated women, while the reverse is true in the US. We anticipate that extending the model to endogenise exits to non-participation at the time of childbirth may help to better explain the labour market behaviour of females and these cross-country differences.

References

- [1] Blau, F. and L. Kahn (1997) "Swimming Upstream: Trends in the Gender Wage Differential in the 1980s," *Journal of Labor Economics* 15, 1-42.
- [2] Blau, F. D. and L. Kahn (1995) "Wage Structure and Gender Earnings Differentials: An International Comparison," *Economica* 63, 29-62.
- [3] Booth, A., S. Jenkins, and C. G. Serrano (1999) "New Men and New Women: A Comparison of Paid Work Propensities From a Panel Data Perspective," *Oxford Bulletin of Economics and Statistics* 61(2).
- [4] Bowlus, A. J. (1997) "A Search Interpretation of Male-Female Wage Differentials," *Journal of Labour Economics* 15(4), 625-657.
- [5] Bowlus, A. J. (1998) "U.S.-Canadian Unemployment Rate and Wage Differences Among Young Low-Skilled Males in the 1980's," *Canadian Journal of Economics* 31(1), 437-464.
- [6] Bowlus, A. J., N. Kiefer, and G. Neumann (2001) "Equilibrium Job Search and the Transition From School to Work," *International Economic Review* 42, 317-343.
- [7] Bowlus, A. J., N. Kiefer, and G. Neumann (1995) "Estimation of Equilibrium Wage Distributions with Heterogeneity," *Journal of Applied Econometrics* 10, S119-131.
- [8] Dex, S., H. Joshi, and S. Macran (1996) "A Widening Gulf Among Britain's Mothers," *Oxford Review of Economic Policy* 12(1), 65-75.

- [9] Elias, P. and M. Gregory (1994) "The Changing Structure of Occupations and Earnings in Great Britain, 1975-1990: An Analysis Based on a New Earnings Survey Panel," Research Series, Sheffield Employment Department 27.
- [10] England, P. (1982) "The Failure of Human Capital Theory to Explain Occupational Sex Segregation," *Journal of Human Resources* 17, 358-370.
- [11] Heckman, J.J. and Carmen Pages (2000) "The Cost of Job Security Regulation: Evidence from Latin American Labor Markets," NBER Working Paper 7773.
- [12] Mortensen, D. (1990) "Empirical Wage Distributions: A Synthesis," in J. Hartog, G. Ridder, and J. Theeuwes (eds.), *Panel Data and Labour Market Studies*, New York: North Holland, pp. 279-297.
- [13] Mortensen, D. and G. R. Neumann (1988) "Estimating Structural Models of Unemployment and Job Duration in Dynamic Econometric Modelling," in W. Barnett, E. Berndt, and H. White (eds.), *Proceedings of the Third International Symposium in Economic Theory and Econometrics*, Cambridge: Cambridge University Press, pp. 335-355.
- [14] Mortensen, D. and C. Pissarides (1998) "Unemployment Responses to Skill-Biased Technological Shocks: The Role of Labour Market Policy," *The Economic Journal* 109, 242-265.
- [15] Ridder, G. and G. J. van den Berg (1999) "Estimating Measures of Labor Market Imperfection for Five OECD Countries Using Aggregate Data in an Equilibrium Search Framework," mimeo.
- [16] Rösén, M. and M. Sundström (1996) "Maternal Employment in Scandinavia: A Comparison of the After-Birth Employment Activity of Norwegian and Swedish Women," *Journal of Population Economics* 9(3), 267-286.
- [17] Van den Berg, G. (1990a) "Nonstationarity in Job Search Theory," *Review of Economic Studies* 57, 255-227.
- [18] Van den Berg G. (1990b) "Search Behaviour, Transitions to Nonparticipation and the Duration of Unemployment," *Economic Journal* 100, 842-865.
- [19] Waldfogel, J., Y. Higuchi, and M. Abe (1999) "Family Leave Policies and Women's Retention After Childbirth: Evidence from the United States, Britain, and Japan," *Journal of Population Economics* 12, 523-545.
- [20] Wellington, A. (1993) "Changes in the Male/Female Wage Gap, 1976-1985," *Journal of Human Resources* 28, 383-410.
- [21] Wright, R. E. and J. Ermisch (1991) "Gender Discrimination in the British Labour Market: A Reassessment," *Economic Journal* 101, 508-522.

Figure 1: Cumulative Distribution Function of Earnings, Empirical Versus Predicted







Figure 2: Unemployment Spells, Empirical Versus Predicted Survivor Functions




Figure 3: Non-participation Spells, Empirical Versus Predicted Survivor Functions




Figure 4: Integrated Hazard Function Test of Job Duration Distribution







Table 1: Means of the BHPS Stock Sample from September, 1991

	O-Level		University	
	Males	Females	Males	Females
Fraction employed in 1991	.901 (.016)	.676 (.018)	.938 (.012)	.832 (.019)
Fraction unemployed in 1991	.0937 (.015)	.0310 (.007)	.0623 (.012)	.0190 (.007)
Fraction non-participating in 1991	.006 (.004)	.293 (.019)	0	.149 (.018)
Mean residual job duration of individuals employed in 1991 (including censored)	156.6 (5.5)	143.03 (5.0)	145.38 (5.2)	128.08 (5.7)
Mean duration of jobs following a non-employment spell (including censored)	172.19 (30.4)	98.56 (23.4)	83.61 (29.2)	134.97 (35.6)
Fraction of censored job spells in stock	.532 (.030)	.398 (.025)	.417 (.027)	.290 (.028)
Fraction of completed job spells that end in a quit	.772 (.036)	.597 (.032)	.848 (.026)	.687 (.033)
Fraction of transitions to non-employment that begin in non-participation	-	.803 (.40)	-	.787 (.41)
Mean non-employment duration beginning in un-employment (including censored)	60.53 (69.6)	47.60 (65.9)	52.02 (66.3)	47.64 (66.3)
Mean non-employment duration beginning in non-participation (including censored)	-	149.22 (95.6)	-	136.02 (90.5)
Fraction of censored non-employment spells (unemployment)	.218 (.42)	.149 (.36)	.190 (.40)	.167 (.38)
Fraction of censored non-employment spells (family care)	-	.458 (.50)	-	.495 (.50)
Mean wages of individuals employed in 1991 (pounds)	178.94 (3.03)	142.54 (1.96)	234.17 (4.1)	199.20 (3.6)
Mean wage following 1991 non-employment spells (pounds)	159.09 (14.4)	138.64 (8.2)	190.01 (26.3)	221.29 (28.8)
No. of individuals in 1991	363	646	401	369

Notes: Standard errors are in parentheses. All durations and wages are expressed in terms of weeks. Wages following unemployment, non-participation and other job spells are imputed in some cases.

Table 2: Parameter Estimates for Arrival Rates Under Three State Model

	O-Level		University	
	Males	Females	Males	Females
λ_0	.0133 (.0136)	.0312 (.0032)	.01897 (.0020)	.0319 (.0045)
λ_1	.00488 (.00069)	.00605 (.00071)	.00784 (.00090)	.00874 (.00109)
δ	.00156 (.00014)	.000877 (.00013)	.00188 (.00015)	.00098 (.00019)
η_1	-	.0015 (.00013)	-	.00142 (.00017)
η_2	-	.00389 (.00031)	-	.00594 (.00071)
κ_0	8.515 (.98)	12.997 (1.43)	10.06 (1.16)	13.30 (1.97)
κ_1	3.121 (.498)	2.519 (.331)	4.161 (.538)	3.638 (.517)
Mean firm productivity	202.84	169.75	253.38	219.54
Mean worker productivity	286.20	223.71	388.94	307.64
r	103.60	83.26	114.79	94.35
Log Likelihood	-3073.29	-5405.75	-3760.51	-3626.50

Notes: Asymptotic standard errors are in parentheses.

Table 3: Averages Predicted by the Model

	O-Level		University	
	Males	Females	Males	Females
Weekly wage, earnings distribution	182.83	144.72	240.72	202.32
Weekly wage, offer distribution	146.64	120.14	179.91	154.91
Non-employment length, unemployment at start (weeks)	75.18	44.60	52.72	38.82
Non-employment length, non-participation at start (weeks)	-	301.60	-	207.19
Unemployment rate	0.105	0.051	0.090	0.056
Non-participation rate	0	0.281	0	0.193
Job-to-job transition rate	0.465	0.431	0.509	0.489

Table 4: Thought Experiments for Individuals with O-level Education

	Males	Females	Females With Male κ 's	Females With $\eta_1 = 0$	Females With Male P's
κ_0	8.5147	12.9965	8.5147	35.5758	12.9965
κ_1	3.1211	2.5188	3.1211	6.8985	2.5188
r	103.60	83.26	83.26	83.26	83.26
Mean wage offer	146.64	120.14	122.42	129.63	134.76
Mean earnings	182.83	144.72	153.08	193.15	167.28
Mean worker productivity	286.20	223.71	232.53	272.04	274.22
Mean firm productivity	202.84	169.75	169.75	169.75	202.84
Monopsony power	.2372	.2608	.2501	.2183	.3063

Table 5: Thought Experiments for Individuals with University Degrees

	Males	Females	Females With Male κ 's	Females With $\eta_1 = 0$	Females With Male P's
κ_0	10.0630	13.2992	10.0630	35.5510	13.2992
κ_1	4.1606	3.6379	4.1606	8.9184	3.6379
r	114.79	94.35	94.35	94.35	94.35
Mean wage offer	179.91	154.91	156.64	164.68	169.35
Mean earnings	240.72	202.32	209.71	259.99	227.31
Mean worker productivity	388.94	307.64	315.87	372.40	376.62
Mean firm productivity	253.38	219.54	219.54	219.54	253.38
Monopsony power	.2526	.2736	.2672	.2390	.3029