Equilibrium Job Search and Gender Wage Differentials in the UK

Audra J. Bowlus *
Louise Grogan †

August 30, 2000

Abstract

Gender wage differentials in the UK are examined using the general equilibrium job search framework in Bowlus (1997). The role of differences in labour market transition patterns across males and females is examined using longitudinal data from the British Household Panel Study. We find that search behavioural differences play an important role in explaining gender wage differentials in the UK for lower educated workers, but only a minor role for higher educated workers. For both groups, however, productivity differences are found to be the greatest factor determining the gender wage gap. We also find that the level of search friction is lower in the UK than in the US. The lower search friction level in the UK is primarily attributable to very low job destruction rates.

Keywords: labour force participation, search models, gender wage differentials
JEL codes: J64, J62, J16

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

The authors are grateful to the Brendan Halpin at the University of Essex, the staff of CEPS/INSTEAD and IRISS facility (which sponsored the initial stages of the data preparation), and to the Social Sciences and Humanities Research Council of Canada.

*Department of Economics, University of Western Ontario
†CILN, Dept. of Economics, McMaster University

1 Introduction

Gender wage differentials are pervasive across countries, ages, and skill groups. The UK is no exception. Our sample from the British Household Panel Survey (BHPS) reveals a female-male hourly earnings ratio of approximately 75% for the early 1990's. This ratio is similar in magnitude to those found for the US and other developed countries. The gender wage differential in the UK has been declining in recent years, while the labour force participation of women has continued to increase.

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. The remainder is often attributed to discrimination against women in the labour market. Another possible source of gender wage differentials is differences in job search behaviour between males and females. Search behaviour has long been noted as a potential source of wage differentials. 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. (2000) find that a large portion of the black-white wage differential in the US can be traced back to differences in job search behaviour. With respect to male-female wage differentials, Bowlus (1997) finds that one quarter of the US male-female wage differential can be explained by search behaviour differences across males and females. Our study of the UK closely follows the frameworks laid out in these two studies.

While there is evidence that search friction plays a role in determining wage differentials in the US, there is also evidence that search friction levels vary widely across countries. Thus one can not say a priori whether or not search friction plays a greater or lesser role in the UK labour market. In their cross-country study, Ridder and van den Berg (1999) find that, while the level of search friction in the UK is more similar to that in the US than to other European countries, the UK still exhibits more search friction than the US. Thus one might expect that search friction may play a larger role in determining wage differentials in the UK than in the US. However, their study reveals nothing about differences between men and women with regard to search behaviour.

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 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 that the US, which until recently had no national policy on family leave. However, it is difficult to know how such policy differences may affect behaviour and hence gender wage differentials. The US system is often 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 the model we present the latter would result in higher wages for women. In support of this notion Rönsen 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 return to work after childbirth. We attempt to examine differences in maternity leave policies between the two countries and their effect on behaviour. However, our results find little evidence that maternity leave policies are a significant factor in determining the differences we do find across the two countries.

Gender differences in labour market behaviour, especially with respect to non-participation and child rearing, have long been studied in labour economics. The literature documenting such differences for the UK is more recent due to available panel data. Booth et al. (1999) use waves 1 to 5 of the BHPS and panel probit models to examine the participation rate of males and females in a longitudinal context. For individuals aged 18 to 55 at the first BHPS interview they find that the year-on-year persistence in paid work propensities is higher for males than for females. This evidence indicates that females have less labour force attachment than males in the UK. Booth et al. (1999) also find that 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 study establishes for the UK the existence of differences across men and women in their labour market behaviour. 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 such a 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.

In addition to the large existing literature in both the US and the UK on gender wage differentials (for recent examples see Wright and Ermisch (1991), Elias and Gregory (1994) for

the UK; Wellington (1993) and Blau and Kahn (1997) for the US), there is also a growing literature in both countries on the so-called "family gap". This gap refers to differences in wages between women who do and those who do not have children. In both the US and the UK, gender pay gaps are larger for mothers than non-mothers (Waldfogel (1998)).

Joshi et al. (1999) examine the family gap for the UK, and look at how continuity of employment around childbearing affects the future pay of mothers. Using cohort studies of individuals born in a week in March 1946 and a week in March 1958, Joshi et al. (1999) are able to look at what happened to the "family gap" as gender equality advanced in the UK and the labour market was liberalised. A key result of this study is that women who exited the labour force for childbearing were paid less than childless women upon re-entry, whereas mothers who maintained their employment while having children were as well paid as childless women. Joshi et al. (1999) find that, once the presence of children is controlled for in the model specification, marriage or partnership does not appear to have a significant impact on women's wages. These results, and similar findings by Waldfogel (1998) for the US, suggest strongly that the labour force behaviour of women around childbirth may be important in explaining gender wage differentials.

One aspect of the model used in Bowlus (1997) that we find unsatisfactory is the exogeneity of the decision to exit the labour force. However, in order to facilitate the comparison across countries, we do not augment the model and leave this for further research. Some features of the BHPS data do enable us to extend the analysis further for the UK than Bowlus was able to do for the US. For example the BHPS contains a representative sample of the population. Thus it is possible to focus on different age cohorts. In contrast, Bowlus used the National Longitudinal Survey of Youth (NLSY) and was therefore restricted to a sample of young individuals just entering the labour force after finishing their education. We can thus present estimates for a wider and more representative age range for the UK. We also have the potential to study the return of women to the labour force after an extended period away. Further, the estimates stem from a more recent time period. The BHPS starts in 1991, whereas the NLSY began in 1979. Finally, we are able to use the BHPS data to provide some evidence on wage outcomes related to various labour market transition patterns.

Our results show that there appear to be similar behavioural differences between males and females in the UK and the US. For example, in both countries females are more likely to enter non-participation at the time of a family concern, i.e. birth of a child, than males. However, there are some differences across the countries. The UK exhibits a much higher job-to-job transition rate than the US. The rate is so high that the model has difficulty in reconciling it with the observed earnings distributions and in the end can not match both. In essence with such a rate,

the model would predict that firms would have very little monopsony power and that there should be very little variation in wages. Workers would climb the wage distribution very quickly and earn their marginal product. We suspect that to some degree the UK job-to-job transition rate is overstated in the BHPS. However, the difference is so great that this is likely a significant difference between the two countries.

In addition the unemployment rate for women in the UK appears to be very low. So, while women have a greater tendency to exit to non-participation, this is offset by a lower tendency to exit to unemployment. Thus the relative level of search friction is similar between men and women in the UK especially for higher educated workers. According to the model differences in search behaviour explain one-third of the gender wage differential for lower educated workers, but only 14% for the differential for higher educated workers.

The paper is organised as follows. Section 2 highlights relevant aspects of the British labour market and presents evidence from the BHPS on important male-female differences in labour market patterns. Section 3 gives a brief overview of the model and estimation methodology taken from Bowlus (1997). 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 The UK labour market

Our study focuses on the early 1990's. The first year of the BHPS is 1991. During this period, the UK experienced moderate unemployment rates, with a high fraction of long-term unemployed amongst the non-employed. The OECD (1990) reports an unemployment rate of 6.5% for the UK in 1990. More than 45% of individuals unemployed in 1990 were long-term unemployed (that is, they had been unemployed for more than one year at the time of interview), despite the fact that unemployment had dropped slightly since the late 1980's.

Unemployment rates and durations differ substantially in the UK across demographic and skill groups (Layard et al. (1991)). The female unemployment rate at 5.1% was slightly lower than that of males at 6.8% in the spring 1989 Labour Force Survey (LFS). Overall labour force participation rates for males and females were 95.6% and 73.3%, respectively. Participation rates for women in the early 1990's were very similar to those in the US. As in other industrialised countries, female labour force participation rates in the UK are higher amongst more educated groups of women.

The participation rate for women continued to grow during the 1990's. During this period wage inequality increased substantially in the UK. The situation for women had consistently improved over time, although in the early 1990's a substantial gender wage gap still remained.

In contrast with other European countries (see for example van den Brink (1994) for the Netherlands), this gap was larger for part-time than full-time workers. The female to male full-time equivalent weekly earnings ratio in the 1991 BHPS sample was about 23% larger for full-time workers than for part-time earners. Several studies for the UK have looked at wage differences between part-time and full-time workers (for recent examples see Makepeace et al. (1999) and Joshi and Paci (1998)).

Employment rates of young women in the United Kingdom compared favorably to those of women in other Western European countries at the beginning of the 1990's. Employment rates of more educated women were similar to those observed in the US in the 1990's, while they were substantially higher in the UK for O-level (highschool) educated women. On the other hand, non-participation rates of women in the 20-40 age range were far higher in the UK than in Eastern European countries. The impact of factors such as taxation structure, child care provisions, and the prevalence of flexible working hours on female labour supply in Western Europe has been extensively analysed and modeled (see for example Heckman (1974), Hausman (1980), Hartog and Theeuwes (1985), Hagenaars (1989), Groot and Pott-Butter (1992), Gustafsson and Bruyn-Hundt (1991), Gustafsson (1992), and Gustafsson and Stafford (1992)). This literature points to the large effects of government policies toward families on the nature and extent of women's labour force insertion.

Like many other countries, the UK has several government programs and policies that directly affect the labour market. For example, the UK has an unemployment insurance program, equal pay and affirmative action laws, and mandatory maternity leave policies. In most respects the programs in the UK are more generous than those in the US. Statutory limits on the time in which women have rights to return to work following maternity leave have been found by Gustafsson et al. (1996) to have effects on the length and nature of labour market interruptions around childbearing. The UK's maternity leave policy changed over the sample period of the BHPS. In 1991 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 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. Women in the UK receive maternity leave benefits from national (social) insurance funds. Thus the potential for longer maternity leave spells is greater in the UK than the US where maternity leaves depend on firm policies and often consist of 6 weeks of sick leave. Maternity leave *perse* is not covered

by social insurance funds in the US[1].

In both the UK and the US, firms are able to make private agreements with their employees regarding maternity leave and other benefits. Prior to 1993, an estimated 40% of US women had explicit maternity leave provisions due to state laws, unions, or voluntary employer provisions (Waldfogel (1998)). In the UK, Dex et al. (1996) have documented the fact that employers have increasingly provided for working mothers additional career-break schemes, top-ups to maternity benefits, workplace nurseries, and flexible hours. Given the prevalence of private arrangements in both countries, laws regarding maternity leave entitlements tell only part of the story about the incentives and constraints facing women when they decide to have children. In general, it is also to be expected that the replacement rate (the fraction of a woman's salary she receives while on maternity benefits) will also have an effect on her behaviour around childbearing.

2.1 BHPS data

To examine the relationships between gender wage differentials, labour force participation and unemployment rate gaps, and the labour market policies in the UK, we use the British Household Panel Survey (BHPS). The BHPS is the primary panel data source for the UK, and 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. We briefly describe our sample and its construction here. Further construction details can be found in the data appendix.

Our sample contains the stock of individuals present in the first wave of the BHPS. Labour market status in the BHPS is based on self-definition. To be in our sample individuals must have recorded being out of the labour force, unemployed or working at the time of the 1991 interview. Thus students, those on government training programs, those who do not report a labour market status, and those who have retired are excluded. As well, individuals who are observed to transit to retirement, training schemes, or higher education are dropped from the sample.

We restrict the age range of the sample to 20-40 years of age in 1991. This draws in a considerably larger and more diverse group of workers than that used in Bowlus (1997) for the US. Here we have the advantage of being able to observe older women returning to the workforce after an extended period of leave. Unfortunately, because of small sample sizes in the BHPS, we are not able to restrict the age range further to conduct a more direct comparison. We do

[1]Women in the US can now take up to 3 months of unpaid maternity leave without fear of losing their jobs.

divide our sample into two education groups: those educated at the O-level and those with higher education. We consider individuals to have achieved a higher education 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.[2] 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. A full discussion of the latter problem can be found in the data appendix. 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 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. This appears to be possible in the BHPS data, but in actuality individuals under-record such transitions.

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

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

lay-offs, no matter how short, as unemployment spells. Bowlus (1997) 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 we need to be concerned with the changes in the UK maternity leave legislation in 1994. Given that "old style" benefits were maintained, and more coverage was introduced, we would expect to observe more women taking maternity leave as a fraction of total exits to care (non-participation or maternity) after October 1994. In particular, we would expect that individuals who had not attained two years of consecutive job tenure would be more likely to take maternity leave in the new system. Under the European rules adopted by the UK, all working women are entitled to fourteen weeks of maternity leave, and are required to take a minimum of two. Given that 58% of women in our sample did not have elapsed job tenures of two years or more at the time of exiting for caring activities, we would expect the new rules to have a large effect. However, the BHPS data provide only limited evidence that the changes in the rules regarding maternity leave benefits in October 1994 had an effect on the likelihood that working women choose maternity leave over non-participation around childbirth.

Table 1: Percentages of women in maternity leave, before and after changes in maternity leave provisions. Individuals aged 20-40 in 1991

	O-level		higher	
	maternity leave	famiy care	maternity leave	famiy care
autumn 1993	1.2	27.6	2.09	10.9
	(.4)	(1.6)	(.6)	(1.3)
autumn 1995	1.0	24.2	3.3	12.3
	(.4)	(1.6)	(.7)	(1.3)

Source: authors' calculations using the British Household Panel Survey 1991-1998.
Note: Standard errors are in parentheses.

Table 1 shows that there was no evident jump in the take-up rate of maternity benefits in the year immediately after the implementation of the new maternity leave provisions. While the fraction of O-level educated women on maternity leave fell between autumn 1993 and autumn 1995, the fraction of higher-educated women on such leave rose. However, the proportion of O-

level women in non-participation for family reasons also dropped over the period, while amongst higher educated women the level of non-participation rose. It appears that higher-educated women are postponing the advent of children.

Table 2: Maternity leave durations before and after 1994 changes, working women aged 20-40, fractions of individuals with spells in duration category

	Before Oct. 16th 1994	After Oct. 16th 1994
2 weeks or less	.01 (.01)	0
2-14 weeks	.04 (.02)	.1 (.06)
14 to 40 weeks	.24 (.04)	.07 (.05)
Longer than 40 weeks	.71 (.05)	.83 (.07)

Source: authors' calculations using the British Household Panel Survey 1991-1998.
Notes: Above contains both censored and uncensored spells. Standard errors are in parentheses.

Table 2 shows that, both before and after the 1994 changes, maternity leaves were generally longer than forty weeks. Though the sample sizes are small, and a majority of maternity leave spells are censored in our data, we find that most women take longer than the "old style" provisions when they go on leave. This suggests that women who begin care spells in maternity leave generally run out of financial support before returning to enterprises. As well, enterprises are not under obligation to take back employees for whom leave time has expired. Surprisingly, there is no tendency for women to take either exactly two (mandatory) or exactly forty weeks of maternity leave, either before or after October 1994.

In the estimation of the job search model, we adopt a similar procedure to that in Bowlus (1997) with regard to maternity leave spells. Maternity leave spells that are shorter than 14 weeks are subsumed into job spells, while those that are longer are treated as non-participation spells.

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

Table 3: Fraction of completed job spells ending in a quit (job-to-job transition) reported under different methods of data coding

coding criteria	O-level		higher	
(M=males, F=females)	M	F	M	F
reason for leaving	.683	.819	.756	.875
	(.027)	(.027)	(.027)	(.020)
subsequent job status report	.829	.828	.780	.880
	(.022)	(.027)	(.030)	(.024)

Source: authors' calculations using the British Household Panel Survey 1991-1998.

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 job loss, 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. As Table 3 indicates this choice results in slightly lower job-to-job transition rates, although only for lower educated males is the difference large.

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. 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 wages without correcting for hours of work, with the implicit assumption that firms offer workers monthly wages that they must accept or reject. Estimation using uncorrected 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.[3]

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

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

interviews. Surprisingly this affects a large number of spells. Thus, while we have wage data for most working individuals at the start of the survey in 1991, we are able to collect only limited wage information for many later job spells. We do try to impute wages for some of these spells. This procedure is described in the data appendix.

2.2 Descriptive statistics

Before discussing the model and how the above data are used to estimate it, 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 4 provides various sample statistics of interest.

In 1991 the sample contains individuals aged 20-40. Across the four groups the mean age ranges from 29 to 31 years of age. As expected we find a higher employment rate amongst men than women, and a higher rate amongst higher 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. Maternity leave status is also self-defined.

The employment and unemployment rates are similar to those found in national unemployment surveys. Using the spring, 1989 UK labour force survey (LFS), the distribution of individuals across labour market states is close for all four groups to what we find with the BHPS.[4] Small differences may be attributed to the fact that the LFS uses an ILO-based (job search) distinction between unemployment and non-participation, while in the BHPS individuals self-describe their labour market state.

The fifth row in Table 4 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 accounting for censored spells. 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 less for the higher education group than for the O-level educated. The next row shows the job spell means following an initial spell of non-employment. These are shorter than those at the

[4]Given that we have excluded the self-employed from our BHPS sample, we also exclude them from our LFS calculations.

Table 4: Means of the BHPS stock sample from September, 1991

	O-level		higher	
(M=males, F=females)	F	M	F	M
mean age of individuals	29.3	28.8	30.2	31.1
	(5.3)	(5.4)	(5.2)	(5.1)
fraction employed in 1991	.693	.888	.831	.929
	(.01)	(.01)	(.02)	(.01)
fraction unemployed in 1991	.044	.105	.040	.071
	(.01)	(.01)	(.01)	(.01)
fraction non-participating in 1991	.263	.007	.129	0
	(.01)	(.004)	(.01)	
fraction of censored job spells	.406	.524	.300	.373
	(.02)	(.11)	(.02)	(.02)
fraction of completed job spells that end in a quit	.683	.819	.747	.872
	(.03)	(.03)	(.03)	(.02)
fraction of transitions to non-emp.	.615	.036	.705	0
that begin in non-part.	(.04)	(.03)	(.05)	–
mean residual job duration of individuals employed in	150.22	167.08	142.77	149.96
1991 (including censored)	(4.59)	(5.10)	(5.42)	(4.81)
mean duration of jobs following a non-employment	84.34	120.64	114.07	105.27
spell (including censored)	(11.89)	(18.59)	(15.19)	(26.62)
mean non-emp. duration beginning in unemp.	60.98	92.24	57.70	71.76
(including censored)	(8.43)	(8.56)	(11.02)	(8.02)
mean non-emp. duration beginning in non-part.	146.06	137.82	147.35	
(includes censored)	(5.15)	(52.38)	(8.16)	
fraction of censored non-employment spells	.37	.35	.24	.29
(unemployment)	(.05)	(.05)	(.06)	(.05)
fraction of censored non-employment spells	.55	.67	.56	
(family care)	(.03)	(.21)	(.04)	
mean wages of individuals employed	137.09	175.91	198.12	236.49
in 1991(pounds)	(1.61)	(2.46)	(3.22)	(3.55)
mean wage following 1991 non-employment	143.54	171.16	190.70	163.87
spells (pounds)	(5.40)	(8.69)	(15.34)	(14.52)
modal number of consecutive job spells prior to	1	1	1	1
spell in 1991				
no. of individuals in 1991	948	562	550	594

Source: authors' calculations using the British Household Panel Survey 1991-1998. 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.

start because of the fixed end date of the panel and are composed of fewer observations. They 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.

We find 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 68% to 88%. 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 higher educated workers. The greater tendency of higher educated females to exit to family care is not consistent with the stock levels of non-participants found in 1991, but may indicate a greater tendency of higher educated women to exit for family care reasons when they are older. 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% (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.[5] On average women spend a little over one year in unemployment while men spend close to 1.5 years. With respect to non-participation we find little difference across the education groups for women with a mean duration of approximately 3 years. The sample sizes for men at less than 5 individuals are just too small to make any comment. Here again the censoring rates are quite high with higher rates for lower educated workers than higher educated workers.

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. Higher educated women earn on average 84% of the salary of higher educated men, while women with an O-level education earn 78% of the salary O-level educated men. Thus women appear to fare better in our sample

[5]Here 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.

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 4 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 O-level educated females. The means in Table 4 do not distinguish between individuals with and without children. For women in particular, the presence of children in the household can be expected to have an important impact on participation decisions.

Table 5: Fraction of individuals with given demographic characteristics

	O-level		higher	
(M=males, F=females)	M	F	M	F
previous exit for caring duties	.0014	.593	.0017	.375
previous maternity leave spell	0	.189	0	.22
married throughout job	.343	.371	.414	.364
single throughout job	.317	.207	.279	.264
married, no children	.340	.094	.414	.140
married, responsibility for children	.004	.277	0	.224
single, no children	.317	.164	.271	.244
single, responsibility for children	0	.043	0	.02

Source: authors' calculations using the British Household Panel Survey 1991-1998.

For women in particular, the presence of young children in the household is known to be an important determinant of their labour market attachment (Booth et al. (1999)). At the macroeconomic level, changes in the participation patterns of women with young children are a major factor determining changes in UK unemployment figures over time. Evans (1998) finds that the fall in the unemployment inflow rate of women with young children over the 1984-1993 period is the main reason for the fall in unemployment rates for females over time. Evans argues that improvements in provisions for mothers returning to work after childbirth have reduced labour market frictions associated with women having young children. While the model we estimate is stationary (thus implying that the history underlying the job transition process is not relevant), it is nevertheless of interest to describe differences between the four groups in the extent of their labour force attachment. In Table 5 we look at aspects of individuals' labour market histories that may be expected to have an influence on their current labour market behaviour.

It is evident that the less educated group of women generally has a history of stable less labour force attachment than the more educated group. Of higher-educated females in our sample, 37% has had at least one spell of non-participation since completing full-time education. Amongst females with O-level education, 59% has had at least one spell. The mean reported length of care spells in the working life history data is slightly longer for O-level educated women than for their higher educated counterparts. As well, this less educated group of women is more likely to have children under the age of 16 at the beginning of our panel.[6]

Table 6: Mean wages of females by previous work history, 1991 stock

mean wages	O-level	higher
with previous caring spell	134.67	197.32
	(2.47)	(6.60)
without previous caring spell	138.36	196.82
	(2.00)	(3.37)
with children under age 16	133.90	209.35
	(2.33)	(5.55)
without children under age 16	138.85	190.86
	(2.02)	(3.497)
married individuals	137.27	201.87
	(2.73)	(4.93)
single individuals	127.49	190.75
	(2.93)	(5.23)

Source: authors' calculations using the British Household Panel Survey 1991-1998. Notes: Standard errors are in parentheses. Wages have been trimmed 5% at top and bottom of the wage distributions. Wages are expressed in terms of pounds per week at full time equivalent.

Previous labour market trajectories of individuals can be expected to influence the wages they receive at the September 1991 BHPS interview. In Table 6 we look at the mean wages of women by previous work histories. We find that mean wages are greater for higher educated women who have previous spells of non-participation for caring activities, while they are lower for O-level educated women who have experienced such spells. For higher educated women in our sample, then, there does not appear to be a wage penalty for previous labour force exit, although these women may simply be older (have more work experience). For the O-level educated group of women, mean wages are found to be higher for those without children. The opposite is true

[6]The BHPS collects information about the age of all household members. However, it is not known whether or not these are biological children of one or both parents.

for the higher educated women. Given the result that higher educated women with previous labour force exits are generally earning more than those without, there does not appear to be a clear pattern across education groups.

It is also of interest to compare the work experience of individuals in our sample by education and gender group. Looking at months of full-time work experience since the completion of full-time education, it is clear that men have much more experience than women, and that O-level educated women have more full-time work experience than higher educated women. O-level women have an average of 78 months of full-time work experience at the start date while O-level men have an average of 98. For the higher educated groups, the figures are 71 and 99 months, respectively. As well, O-level educated females have, on average, significantly more part-time work experience than their higher educated counterparts, at 34 and 39 months on average, respectively. This finding is consistent with the fact that O-level educated women generally stop full-time education at a younger age than those with higher education. Thus the less stable labour force attachment for O-level women (noted in Table 5) is tempered by their greater experience.

Far more women than men of either education group report having part-time work experience, and the mean cumulative part-time work experience (for those with any) is in both cases larger for women. In our sample 42% of higher educated females reports previous part-time work experience, as do 53% of O-level educated females. Amongst men, less than 10% of men of either education group has any part-time work experience. While we are unable to distinguish between part-time and full-time employment in the present study (the BHPS does not contain information on hours worked for spells that occur entirely between interviews), we are aware that the propensity to undertake part-time work is a prominent aspect of behavioural differences between the sexes.

Finally, we use the working life history spell files constructed by Brendan Halpin to summarise work experience prior to the spell ongoing at the first BHPS interview. These files show that higher educated females in our sample are more likely than O-level educated females to have had a previous maternity leave spell. Of O-level educated females 19% has had previous maternity leave spells, while 22% of higher educated females has experienced such spells. Given that more O-level educated women in the sample have children, it appears that this result stems from the higher labour force participation rate of more educated women.

3 Model and estimation

Given the many behavioural differences across males and females it is important to know if they have an effect on the observed gender wage differentials. To study such effects 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. 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 transition from non-participation to employment is not allowed.

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. The decision to exit is effectively suppressed as well. In our study of the UK labour market, such exits consist solely of caring for family members, primarily children. It is likely that such 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 (1997) and leave this important extension for further research. We do provide some evidence on the validity of this assumption in section 5.4.

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)):

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

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 parameters κ_0 and κ_1 can be thought of as measures of search friction in the labour market. A greater presence of search friction results in lower levels of κ_0 and κ_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 $r = b$.

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)):

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

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 then given by the following formula:

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

Gender wage differentials can be generated in this model easily by allowing firms to post gender-specific wage offers. Then, if the arrival rate of family concerns is higher for women than for men, women will earn, on average, less than men. This occurs because the reservation wage for women will be less than that for men, so that the wage offer distribution for women is shifted to the left of that for men. Also 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), and Bowlus et al. (2000) and estimate the model assuming discrete productivity heterogeneity. A full description of the estimation methodology for the three–state model is given in Bowlus (1997). Here we only briefly point out how the structure of the UK data fits into this framework. Bowlus essentially deals with a flow sample of unemployment, non-participation and job spells. In contrast, the BHPS starts with a representative stock sample. Thus the data we collect differ slightly from that collected by Bowlus for the NLSY. 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,

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

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

and

$$\Pr(E) = 1 - \Pr(N) - \Pr(U). \tag{6}$$

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.[7] 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 start in unemployment are exponential with

[7] This result is particular to the exponential distribution.

parameter $\lambda_0\eta_2/(\eta_1+\eta_2)$ while those that start 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. As shown in Table 4 this holds true for the UK data.

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

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

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

and

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

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) and Bowlus et al. (2000) and use simulated annealing[8] to handle the discontinuties 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 each male education group. As shown in Table 4 so few males are in the family care state or exit to the family care state that this simplification has little effects our results. We estimate the model separately for males and females for both education groups - higher education and O-level education. For all four groups

[8]This is a procedure for global optimisation which distinguishes between local optima. In our estimation, the routine picks wage cuts for each firm type, then considers the implied fit of the data to the model. Rather than searching over all possible combinations, this subroutine stops when a set of wage cuts is found which satisfies the model's specification of the relationship between wage cuts and other firm-specific parameters, and which satisfies an optimal stopping rule (see Szu and Hartley (1987)). Recall that each firm type offers a unique set of wages, and that the wage distribution is discontinous at the wage cuts.

we find that we need five firms types to fit the wage data well. We then compare the estimated parameter values across population subgroups, and to those obtained by Bowlus (1997) using the NLSY. We also examine the fit of the model and conduct several 'thought' experiments by analysing the effect on the gender wage differential of changing various parameters.

4 Estimation results

The parameter estimates for the four subgroups are shown in Table 7. For the O-level education group we find the following relationship across males and females. First, females have a much higher job offer arrival rate while unemployed. Thus 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 similar offer arrival rates while employed, while males have a much greater chance of having their jobs being destroyed. 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 even larger than 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. The fact that females exit to non-participation lowers their reservation wage and hinders their movement up the wage distribution.

Search friction, however, is not the whole story. We see that 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. However, this finding is only consistent with the model if females have an implausibly low value of non-market time. Since females have a much higher level of κ_0 than males and κ_0 raises the reservation wage , one would expect (given similar productivity distributions) females to have a higher reservation wage than males. Offers are arriving so quickly while unemployed that they should be pickier. The value of κ_0 is so high that the only way the model can explain the reservation wage strategy of the females is to give them a negative value for b.

With respect to higher educated workers we find a similar pattern. Females again have a

Table 7: Parameter estimates for arrival rates under three state model

	O-level		higher	
(M=male, F=females)	M	F	M	F
λ_0	.0082	.0195	.0137	.0212
	(.0007)	(.0017)	(.0011)	(.0024)
λ_1	.0052	.0058	.0080	.0084
	(.0006)	(.0006)	(.0008)	(.0009)
δ	.0014	.0006	.0018	.0008
	(.0001)	(.00001)	(.0001)	(.0001)
η_1	–	.0015	–	.0012
	–	(.0001)	–	(.0001)
η_2	–	.0040	–	.0057
	–	(.0003)	–	(.0007)
κ_0	6.0502	9.5215	7.7090	10.5801
	(.5461)	(.8664)	(.6956)	(1.2668)
κ_1	3.8442	2.8210	4.5341	4.2000
	(.5160)	(.2960)	(.5069)	(.5272)
mean firm prod.	187.08	157.65	246.62	208.63
mean worker prod.	265.40	210.09	361.32	299.73
r	100.83	83.00	114.79	97.87
LL	-4373.47	-7009.64	-5162.34	-4815.55

Source: authors' calculations using the British Household Panel Survey 1991-1998. Notes: Asymptotic standard errors are in parentheses.

higher job offer arrival rate in unemployment than males. They also now have a slightly higher job offer arrival rate while employed. Their job destruction rate is lower, but they have a very high exit rate to non-participation. Together these two exit avenues result in a higher exit rate to non-employment for females. They have a much higher value of κ_0 than males, but a slightly lower value of κ_1. The female reservation wage is lower than the males and here again the model has to give the females a low value of b to explain the low value for r. As the values of κ_1 are very similar, search friction is expected to explain little of the gender wage differential for this education group. The model must therefore attribute the differential to differences in productivity, giving females a much lower average productivity level.

These results differ somewhat from those found for the US by Bowlus (1997). The most surprising result is perhaps that the UK education groups display lower levels of search friction than their US counterparts. This may reflect age differences but appears to be connected to the very low rates of job destruction found in the early 1990's in the UK. With respect to gender differences there are other differences as well. In particular, Bowlus (1997) found that males in both education groups, high school and college graduates, faced lower levels of search friction (higher values of κ_0 and κ_1) while unemployed and employed. 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 levels of the rates also vary across the countries. For example, lower educated workers in the US (high school graduates) display much higher job offer arrival rates while unemployed and employed than lower educated workers in the UK (O-level educated). For males the difference is two times. However, the US workers also display much higher job destruction rates - on the order of four times greater for males and even more for females. Interestingly the exit rate to non-participation for females is similar for this education level across the countries. Some of the differences found between the US and Canada may reflect the difference in ages across the two subpopulations. It may be difficult for older individuals to locate new jobs, while younger individuals may face more job uncertainty and therefore higher exit rates. With respect to the groups with higher education, we find that again the US job offer arrival rate while unemployed is higher than that for the UK. The rate while employed is slightly higher in the UK. The job

destruction rate is also much higher in the US while for females the exit rate to non-participation is lower in the US. Female college graduates in the US exit to non-participation at half the rate of higher educated females in the UK.

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 higher educated workers, we find that 6.2% 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 7.9% is due to differences in the κ_1 values, while 85.9% 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 much larger. We find that 9.1% of gender wage differentials can be attributed to reservation wage differences, 26.6% is attributable to differences in the κ_1 values, and 64.3% 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 14% of the gender wage differential amongst those with higher education in the UK can be attributed to differences in search behaviour. Amongst O-level individuals, about 36% of wage differences between the sexes is attributable to search behaviour. These results contrast with the US results in that we find large differences between education classes in the importance of search in explaining gender wage differentials. In contrast, Bowlus (1997) finds that, for both high-school and college-educated Americans, 20-25% of the gender wage differential can be attributed to search differences.

We turn now to examining how well the model fits the UK data. Table 8 shows averages predicted by the model that can be compared to those from the data shown in Table 4. 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 except for the case of higher educated males. For that group the model overestimates the mean by over 10 pounds. With respect to mean wage offers, the model in general predicts lower means than those observed in the data. Of course, we expect the mean wage offer to be lower than the 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 entirely consistent with the differences

Table 8: Averages predicted by the model

	O-level		higher	
(M=male, F=females)	M	F	M	F
weekly wage –earnings distribution	179.00	138.50	249.04	199.58
weekly wage –offer distribution	141.07	115.43	181.08	152.35
non-employment duration – unemployed at start (weeks)	122.04	70.24	73.18	56.91
non-employment duration – nonparticipation at start (weeks)	-	320.82	-	232.36
unemployment rate	.142	.069	.115	.071
nonparticipation rate	-	.271	-	.172
job-to-job transition rate	.497	.449	.521	.510

Source: authors' calculations using the British Household Panel Survey 1991-1998. Note that wages are trimmed at both extremes by 5%.

between these two distributions observed in the data, and reflects the model's inability to match both wage distributions.

An examination of the predicted durations of non-employment and the unemployment and non-participation rates reveals several patterns. For males the predicted mean unemployment durations are similar to those observed in the data (if you combine the different means given for non-employment spells in Table 4). The model is trying to balance several features including the length of the unemployment durations, the high censoring rates, the low unemployment rates, and the very low exit rate out of jobs to unemployment. As it is, the predicted unemployment rate is higher than the observed rate for both groups of males. Amongst females the predicted mean durations of spells starting in unemployment are also close to those observed in the data albeit somewhat high for O-level females. For both groups the predicted means for spells starting in non-participation are higher. The model is trying to match the very low female unemployment rate observed in the data, and the high non-participation rate. The model is able to match the high non-participation rate but predicts an unemployment rate that is too high.

The model is not able to match the very high job-to-job transition rates observed in the data for any of the four groups. The observed data generates job-to-job transition rates on the order of 0.7 to 0.9 while the model is only in the range of 0.5. This is because the model can not reconcile the observed high job-to-job transition rates with the other observed features of the data. In order to generate job-to-job transition rates that are this high 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 higher educated females to 2972 for higher educated males. Such levels of search friction essentially imply a competitive labour market and thus predict that all workers should be earning the highest productivity level. 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 LFS we find a job-to-job transition rate that is slightly lower than in our BHPS sample, but higher than that found in the NLSY. [9] This implies that our data overstate the degree of job mobility in the UK. Given the lack of an employer-specific code in the BHPS, it is not possible to merge multiple spells with a single employer unless an individual reports a promotion. Thus the model may be doing a better job of fitting the main features of the UK labour market than is evidenced here. Our λ_1 and κ_1 parameters are below those estimated by Ridder and van den Berg (1999) for the UK using aggregate data indicating our estimates are still plausible.

Next we repeat the exercises in Bowlus (1997) and ask what are the mean wages for females under various different scenarios. Before discussing the results we note that technically 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 women with the observed wage and transition data. We find that the reservation wage is very sensitive to changes in the parameters, and often takes on implausible values. Therefore the results we present in tables 5.9 and 5.10 hold the reservation wage constant at its current level. For each case we do note what the model would predict regarding the reservation wage change and its effect on the mean wage.

Table 9 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 while the mean earnings level of females increases a lot. This reduces the gap between average earnings of males and females by 27%, due to the increase in κ_1. However, if we would have allowed the reservation

[9] The 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 of .72.

Table 9: Thought experiments for individuals with O-level (high school) education

	males	females	females with male κ's	females with $\eta_1 =$ 0	females with male P's
κ_0	6.0502	9.5215	6.0502	34.5465	9.5215
κ_1	3.8442	2.8210	3.8442	10.2352	2.8210
R	100.83	83.00	83.00	83.00	83.00
mean wage offer	141.07	115.43	118.07	124.71	130.05
mean earnings	179.00	138.50	149.26	197.31	161.54
mean worker prod.	265.40	210.09	222.96	275.84	250.33
mean firm prod.	187.08	157.65	157.65	157.65	187.08
monopsony power	.2157	.2345	.2212	.1905	.2821

Source: authors' calculations using the British Household Panel Survey 1991-1998.

wage to adjust, the mean wage offer would actually have fallen and mean earnings only increased slightly. The large drop in κ_0 would have resulted in a large decrease in the reservation wage, a decrease to an unrealistically low level of 34 pounds sterling.

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 is 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 135 pounds. This reservation wage value is so high that the lowest productivity firm type would no longer be able to afford to hire any workers. This results in even larger increases in mean offers and earnings.

The final case considers changing the firm productivity distribution of the females to that of the males. This results in the largest increase in the mean wage offer of all the thought experiments. Mean earnings also increase substantially. This indicates the importance of productivity differences in determining the gender wage differential for this group. Allowing the reservation wage to adjust in this last case would result in an upward adjustment to 101 pounds. Under this solution the mean wage offer for females would be 138 pounds and the mean earnings level 166.

Table 10 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 the reservation wage were allowed to change, the drop in κ_0 would result in a large, implausible decrease to 54 pounds, and a drop in both means. Setting η_1 to

Table 10: Thought experiments for individuals with higher education

	males	females	females with male κ's	females with $\eta_1 = 0$	females with male P's
κ_0	7.7090	10.5801	7.7090	25.8506	10.5801
κ_1	4.5341	4.2000	4.5341	10.2619	4.2000
R	114.79	97.87	97.87	97.87	97.87
mean wage offer	181.07	152.35	153.18	160.40	173.04
mean earnings	249.04	199.58	203.49	254.04	240.14
mean worker prod.	361.32	299.73	304.51	366.99	356.18
mean firm prod.	246.62	208.63	208.63	208.63	246.62
monopsony power	.2464	.2475	.2443	.2179	.2857

Source: authors' calculations using the British Household Panel Survey 1991-1998.

zero increases both κ_0 and κ_1 substantially. For the case presented in the Table 10, the increase in κ_1 increases the mean wage offer slightly and has even more of an effect on mean earnings. In fact mean earnings for females is now higher than that for males. As in the case for O-level females, if one were to allow the reservation wage to respond, the large increase in κ_0 would result in a reservation wage value of 204 pounds such that the lowest productivity type could not afford to pay. In this case, the reservation wage would increase so much that both the mean wage offer and mean earnings levels 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. This reduces the gender wage gap. Allowing the reservation wage to increase to 122 pounds in this case would essentially eliminate the gap.

The model estimated above does not take into account the possible relationship between a woman's wage and her labour market behaviour around childbirth. Whereas the probability of exit for a personal concern is considered to be exogenous to the current wage, microeconomic theories of labour supply would predict otherwise. In particular, women with higher wages should be less likely to exit to non-participation, and more likely to take maternity leave , when a personal concern arises. There is some empirical evidence suggesting that this is the case in both the US and UK. Ermisch (1989) estimates the elasticity of female participation with respect to wages at .69-.87 in the UK, and Hotz and Miller (1988) estimates an elasticity of 1.23 for the US. Table 11 suggests that, for higher educated women, those who have higher wages are less

likely to exit the labour force with the advent of a child. However, for O-level educated women this does not appear to be the case at the mean. Comparing the fit of models with and without endogenous exits to non-participation would allow clarification of the relative importance of this phenomena.

Table 11: Mean wages of females by response to a personal concern

	O-level	higher
no exit for personal concern	136.4115	195.369
no maternity leave	(1.69)	(3.843)
maternity leave	133.0611	247.681
no exit for personal concern	(15.30)	(31.574)
exit for personal concern	151.3056	189.1614
no maternity leave	(7.73)	(21.11)

Source: authors' calculations using the British Household Panel Survey 1991-1998. Notes: Standard errors are in parentheses. Wages are trimmed 5% at top and bottom of the wage distributions. Wages are expressed in terms of weeks, and come from both the 1991 stock of jobs and the flow into jobs over the sample period.

As well as the issue of the endogeneity of exits to non-participation around childbirth, the present model does not take account of the relationship between fertility decisions and a woman's wage. In our model, the advent of a child is modeled as a stochastic process. The landmark theoretical work of Becker (1960) on the issue of economic decisions surrounding fertility, spawned a generation of empirical work on female labour supply and fertility (see for example Browning (1992), Schultz (1978), as well as Hotz and Miller (1988), and Ermisch (1989)). Tommaso (1999) develops a structural model of labour force participation, fertility, and wages, and estimates it for Italy. A general result of existing empirical investigations is that estimates of the elasticity of fertility with respect to the wages of females and their partners tend to vary much more widely than estimates of participation elasticities. For the UK in 1989, Ermisch (1989) estimate elasticities of fertility with respect to the female wage that vary between -0.26 and -1.01. As Becker (1960)'s work predicts, Ermisch (1989) finds that there is a positive relationship between the wages of a woman's partner and her fertility decisions. Despite the fact that theory and evidence points to the endogeneity of fertility decisions, we are not optimistic about being able to endogenise fertility decisions within a general equilibrium job search framework. Whereas the present model does not allow for a relationship between a woman's wage and her fertility decisions, it is doubtful that modeling these decisions as entirely dependent on her wage would provide a better fit to the data. As well, the Burdett-Mortensen style job search model cannot easily be adapted

to take into account factors such as accumulated wealth, or the labour market status of other household members. While we are optimistic about the feasibility and usefulness of extending the model to endogenise exits to non-participation, we are less so about endogenising fertility decisions.

5 Conclusions

In this paper, the role played by job search behaviour in governing gender wage differentials in the UK was investigated. Search behaviour differences across males and females were found to account for a significant portion of the gender wage differential for O-level educated workers, but not for higher educated workers. For both groups, productivity differences play the largest role in determining the gender wage gap. Differences in reservation wages also play a somewhat minor role, although these differences are difficult to reconcile within the context of the model. We also find that eliminating exits to non-participation for females results in a substantial increase in their mean earnings levels.

Our comparison of the UK results with those for the US from Bowlus (1997) indicate cross-country differences as well. Surprisingly, the measure of search friction is lower for the UK than the US. This is despite higher job offer arrival rates in the US, and results from a much lower job destruction rate in the UK. As well, the levels of search friction vary more across education groups in the UK than they do in the US. We also observe a much higher exit rate to non-participation for higher educated women in the UK. We surmise that some of these 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.

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. A related concern is the finding by Blau and Kahn (1995) that a larger proportion of part-time workers are married in the UK than in the US. This highlights the important interaction between marital status and labour market behaviour. The second issue is related to the decision to exit to non-participation. In the UK higher educated females stay in non-participation for a shorter period of time than lower educated women. The reverse is true in the US. In contrast, higher educated women in the UK

are more likely to enter non-participation than higher educated women 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 women and these cross-country differences.

6 Appendix A: Data description

The British Household Panel Survey (BHPS) was designed to answer socio-economic questions on topics including the relationship between health changes and unemployment, the effects of life events on socio-economic values, life cycle variation in income, and returns in the labour market to training and education. The initial sample for Wave One of the BHPS consisted of 8167 issued addresses drawn from the Postcode Address File. Interviews were attempted with all private households found at these addresses. The waves that follow include all adults in all households containing at least one member who was resident in a household interviewed in Wave One. The BHPS is representative at the national level of private households in the UK.

Individual-level questionnaires (administered to all adults) include information on current labour market status, remuneration from work, transitions made between interview periods, and the reasons for such transitions. The BHPS makes use of the practice of "feed forward", in which interviewers are given information about the labour market status of respondents in previous years for the purpose of clarifying apparent contradictions in responses across waves.

6.1 Sample restrictions

Individuals must be either non-participants, unemployed or working in the 1991 BHPS interview in order to be included in the sample. Thus students, those on government training schemes, those who do not report a labour market status, and those who have retired are excluded. As well, individuals who are observed to transit to retirement, to training schemes, or to higher education before being censored or completing a job cycle are dropped from the sample. Finally, we select individuals who have at least one work spell since completing full time education.

In the NLSY sample used in Bowlus (1997) the mean age of high school graduates is 18.7 years, and of college graduates 23.9 years. The BHPS sample is older, with the mean age of all four groups about 30 years in September 1991. Whereas 10% of the NLSY sample of high school educated individuals is married at the start of their jobs, about one third of our sample is continuously married throughout the 1991-1996 period.

An employment spell is defined for our purposes as a continuous period of work for a single employer. Although promotions within enterprises are considered as separate spells according to

the question scheme of the BHPS, we do not consider internal enterprise mobility here. Direct job-to-job transitions in the BHPS spell file were merged if no employer transition took place. We considered only direct transitions between jobs at different employers as separate spells.

Whereas the NLSY interviewing procedure allows individuals to discuss multiple jobs in which they may be engaged, the BHPS collects job transition information for only one primary job. In practise this should not cause comparability problems between the NLSY and BHPS results, since Bowlus (1997) only focuses on primary jobs.

6.2 Creating spell files

In September, 1992 BHPS interviewers collected information from individuals on all employment status spells from leaving full time education to the date of interview. In the following year's BHPS interview, information was collected on all jobs held since leaving full-time education. To learn more about the work history of the individuals in our sample, we make use of a composite file of these two information sources. This file is one of many constructed by Brendan Halpin at the University of Essex using information from the BHPS.

Although job and non-employment spells that ended prior to September 1991 are not used in our estimation of the three-state Burdett-Mortensen model, the work history information is used to test the validity of the assumptions behind the model. From the composite life job spell and work status files, we extract variables relating to the total part-time and full-time months worked since completing full-time education, total time spent in full-time caring activities, and previous maternity leave spells. For more on the reliability of the working life history information in the BHPS, see Dex and McCulloch (1997).

It is not possible to compute full-time equivalent weekly wages for all individuals who complete an employment spell in the data. Due to the fact that hours of work are not reported for individuals who complete job spells entirely between spells, it is not possible to use the wages reported for these spells in the analysis. In cases where we know the wage received at the interview immediately following such a spell, we use the interview-date wage to impute the missing wage.

It is well known that individuals make errors in reporting the timing of changes in labour market status (and in the reporting substantive transitions), that lead to apparent inconsistencies in reporting over years of interviews. For example, an individual being unemployed for one month in one interview, and twelve months later that the current job has been held for 15 months. Despite the practise of "feed forward" in the BHPS (mentioned previously), there are still some substantial differences in year-on-year records of the timing of labour market transitions.

In our construction of the spell files, we treat such apparent recall errors (referred to in the literature as a "seam problem") according to a rule, and assume that the report of labour market status changes closest to date of the transition is the correct one. If the reports of the start date of a given spell vary by more than three months across subsequent interviews, the individual is censored at the date of the earlier interview, and only spells observed up to this date contribute to our likelihood. If the inconsistency in reports of labour market status is less than three months, the individual is not censored, and the date of the labour market transition is on or about the start date of the earlier interview.

References

Becker, G. (1960). An economic analysis of fertility. *Demographic and Economic Change in Developed Countries*.

Blau, F. and L. Kahn (1997). Swimming upstream: Trends in the gender wage differential in the 1980's. *Journal of Labor Economics 15*, 1–42.

Blau, F. D. and L. Kahn (1995). Wage structure and gender earnings differentials: An international comparison. *Economica 63*, 29–62.

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).

Bowlus, A. J. (1997). A search interpretation of male-female wage differentials. *Journal of Labour Economics 15*(4), 625–657.

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.

Bowlus, A. J., N. Kiefer, and G. Neumann (1995). Estimation of equilibrium wage distributions with herterogeneity. *Journal of Applied Econometrics 10*, 119–132.

Bowlus, A. J., N. Kiefer, and G. Neumann (2000). Equilibrium job search and the transition from school to work. *International Economic Review (forthcoming)*.

Browning, M. (1992). Children and household economic behaviour. *Journal of Economic Literature 30*, 1434–75.

Dex, S., H. Joshi, and S. Macran (1996). A widening gulf among Britain's mothers. *Oxford Review of Economic Policy 12*.

Dex, S. and A. McCulloch (1997). The reliability of retrospective unemployment history data. *Institute For Social and Economic Research Working Paper 17*.

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*.

Ermisch, J. (1989). Purchased child care, optimal family size and mother's employment. *Journal of Population Economics 2*, 79–102.

Evans, P. (1998). Why has the female unemployment rate fallen so much in Britain? *Bank of England, ISSN 1368-5562*, 76–108.

Groot, W. and H. Pott-Butter (1992). The timing of maternity in the Netherlands. *Journal of Population Economics 5*, 61–85.

Gustafsson, S. (1992). Separate taxation and married women's labour supply, a comparison of West Germany and Sweden. *Journal of Population Economics 5*, 61–85.

Gustafsson, S. and M. Bruyn-Hundt (1991). Incentives for women to work: A comparison between the Netherlands, Sweden, and West Germany. *Journal of Economic Studies 18*(56), 30–65.

Gustafsson, S. and F. Stafford (1992). Child care subsidies and labor supply in Sweden. *Journal of Human Resources 27*(1), 204–230.

Gustafsson, S., C. Wetzels, D. Vlasblom, and S. Dex (1996). Women's labour force transitions in connection with childbirth: A panel data comparison between Germany, Sweden, and Great Britain. *Journal of Population Economics 9*.

Hagenaars, A. (1989). Female labor supply in micro-simulation models. In H. Petersen (Ed.), *Micro-simulation models: Prospects and limitations*. University of Gessen.

Hartog, J. and J. Theeuwes (1985). The emergence of the working wife in Holland. *Journal of Labor Economics 3*, 235–355.

Hausman, J. (1980). The effects of wages, taxes and fixed costs on women's labor force participation. *Journal of Public Economics 14*(2), 161–194.

Heckman, J. (1974). The effects of child-care programs on women's work effort. *Journal of Political Economy 82*(2), 679–694.

Hotz, J. and R. Miller (1988). An empirical analysis of life-cycle fertility and female labor supply. *Econometrica 56*(1).

Joshi, H. and P. Paci (1998). *Unequal Pay for Women and Men*. MIT Press: Cambridge, MA.

Joshi, H., P. Paci, and J. Waldfogel (1999). The wages of motherhood: better or worse? *Cambridge Journal of Economics 23*, 543–564.

Layard, R., S. Nickell, and R. Jackman (1991). *Unemployment: Macroeconomic Performance and the Labour Market*. Oxford: Oxford University Press.

Makepeace, G., P. Paci, H. Joshi, and P. Dolton (1999). Has equal treatment in pay progressed since the 1970's? a study across men and women in two cohorts. *Journal of Human Resources 34*(3), 534–556.

Mortensen, D. (1990). Empirical wage distributions: A synthesis. In J. Hartog, G. Ridder, and J. Theeuwes (Eds.), *Panel Data and Labour Market Studies.*, pp. 279–297. New York: North Holland.

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*, pp. 335–355. Cambridge: Cambridge University Press.

OECD (1990). *Employment Outlook*. OECD: Paris.

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)*.

Rönsen, 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.

Schultz, P. (1978). The influence of fertility on labour supply of married women: Simultaneous equations estimates. *Research in Labour Economics 2*.

Szu, H. and R. Hartley (1987). Fast simulated annealing. *Physics Letters 12*, 157–162.

Tommaso, M. L. D. (1999). A trivariate model of participation, fertility, and wages: the Italian case. *Cambridge Journal of Economics 23*, 623–640.

van den Brink, H. M. (1994). *Female Labor Supply, Child Care and Marital Conflict*. Amsterdam University Press: Amsterdam.

Waldfogel, J. (1998). The family gap for young women in the United States and Britain: Can maternity leave make a difference? *Journal of Labour Economics 16*, 505–545.

Wellington, A. (1993). Changes in the male/female wage gap, 1976-1985. *Journal of Human Resources 28*, 383–410.

Wright, R. E. and J. Ermisch (1991). Gender discrimination in the British labour market: A reassessment. *Economic Journal 101*, 508–522.