UNIVERSITY OF NOTTINGHAM

Discussion Papers in Economics

Discussion Paper
No. 01/07

EMPLOYER SEARCH, VACANCY DURATION AND SKILL SHORTAGES

by M J Andrews, S Bradley and R Upward

July 2001

DP 01/07

ISSN 1360-2438

UNIVERSITY OF NOTTINGHAM

Discussion Papers in Economics

Discussion Paper
No. 01/07

EMPLOYER SEARCH, VACANCY DURATION AND SKILL SHORTAGES

by M J Andrews, S Bradley and R Upward

Martyn Andrews is Senior Lecturer, School of Economics, University of Manchester, Steve Bradley is Senior Lecturer, School of Economics, University of Lancaster and Richard Upward is Research Fellow, School of Economics, University of Nottingham

July 2001

Employer search, vacancy duration and skill shortages*

M J Andrews
University of Manchester

S Bradley
Lancaster University

R Upward
University of Nottingham

June 2001

*The authors thank The Leverhulme Trust (under grant F/120/AS) for financial assistance and Dave Stott for research assistance. The data were kindly supplied by Lancashire Careers Service. The comments of Len Gill and Jonathan Wadsworth are gratefully acknowledged, as are those from participants at various presentations. These include the Manchester Universities' Labour Workshop, the 1997 EEEG Overnight Workshop (Royal Holloway), the Departments of Economics at Leeds and Loughborough, and the Institute of Careers Guidance Workshops in Glasgow and Newcastle. The data used in this analysis are available on request. Subsequent revisions to this paper will be available at `http://les1.man.ac.uk/ses/staff/ma/`.

Abstract

This paper provides the first analysis of employer search using duration methods for the UK. We model both the duration of employer search and whether employers succeed in filling vacancies. We present the appropriate econometric techniques for dealing with groups of identical vacancies posted simultaneously, and we examine the robustness of our results to the flexibility of the baseline hazard and unobserved heterogeneity. We compare results across two quite different markets (jobs and training places). Our results show that employers search longer for high quality vacancies; that there are skill shortages in so far as jobs requiring more qualified applicants are more likely to be withdrawn from the labour market; and that vacancy duration varies pro-cyclically with labour market tightness.

Keywords: vacancy duration, skill shortages, competing risks

New JEL Classification: C41, J41, J63, J64

Address for Correspondence:
Dr. M.J. Andrews
School of Economic Studies
University of Manchester
Manchester, M13 9PL

Email: martyn.andrews@man.ac.uk
Phone: +44-(0)161-275-4874

1 Introduction

Search theory is becoming one of the dominant paradigms used to explain both micro and macro labour-market phenomena, especially the dynamics of unemployment — see Mortensen & Pissarides' recent (1998, 1999) surveys. Two-sided search models such as Pissarides (1990) and Burdett & Wright (1998) are particularly useful since they emphasise the role of employer search as well as worker search. This is important, given that empirical work has concentrated almost entirely on workers' search, in spite of evidence which suggests that, in many labour markets, workers rarely refuse job offers.[1] If the worker's acceptance probability is close to unity, it follows that employer search is very important in understanding what factors determine transitions between unemployment, employment and non-employment. There is a large microeconometric literature that has estimated the hazard out of unemployment using unemployment duration data, but there is far less equivalent evidence for vacancies.[2] Thus, modelling employer search remains a very-much under-researched area, and is of interest in its own right.

One particular issue that has received little attention is the fact that employer search is not always successful, resulting in unfilled vacancies. Unfilled vacancies may be a result of skill shortages, and it is argued that the resulting macroeconomic implications are potentially lower productivity growth and higher wage growth (Haskel & Martin 1996). To understand properly the determinants of skill shortages, it is necessary to analyse the type of vacancies which employers find difficult to fill and therefore the type of vacancies that are eventually withdrawn from the labour market.

In this paper we estimate both the determinants of vacancy duration and the probability that the employer successfully fills the vacancy. We use far more detailed vacancy data than has previously been available. The data contain information on some 18,000 job vacancies and 30,000 training vacancies submitted by employers in Lancashire between 1985 and 1992. The data measure vacancy duration recorded to the nearest day, and include a wealth of detail about the type of vacancy being offered. We allow for the simultaneous advertising of groups of identical vacancies, and we model the fact that many vacancies remain unfilled. Our econometric methodology allows for unobserved heterogeneity, and we test the importance of modelling the underlying hazard non-parametrically.

All studies of vacancy duration are reduced-form estimates of a two-sided search process by employers and job-seekers. Our analysis is also reduced form, but we interpret our results within the theoretical framework provided by two-sided search models, in particular, that

[1] See Barron, Black & Loewenstein (1987), Holzer (1988), van den Berg (1990), Barron, Berger & Black (1997*a*) Manning (2000) and Andrews, Bradley & Upward (2001).

[2] See van den Berg (1999) for a recent list of contributions and surveys.

of Burdett & Wright (1998). If we assume that job-seekers rarely refuse job offers, this framework allows us to make predictions about the effect of measured covariates on the duration of employer search, including the effects of aggregate labour market conditions.

The two-side search model does not make any predictions about the shape of the underlying vacancy hazard. This is also an important issue, as it sheds light on the search and selection methods used by employers and job-seekers. We test whether declining vacancy hazards are a result of unobserved heterogeneity, or the result of genuine duration dependence.

Finally, we examine the issue of skill shortages from an entirely new angle, by examining which types of jobs take longer to fill, and those which are eventually removed from the market.

The paper is structured as follows. Section 2 briefly covers the relevant literature, including the issue of skill shortages and hard-to-fill vacancies. Section 3 describes the data, and Section 4 discusses our methods. Section 5 provides a theoretical framework for interpreting our results, which are discussed in Section 6. Section 7 concludes.

2 Recent literature

There have been very few microeconometric investigations of the duration of employer search, or vacancy duration, using firm-level vacancy data. This is particularly true for the UK where there are only two studies, neither of which use duration modelling techniques (Beaumont 1978, Roper 1988). More recently, van Ours & Ridder (1991, 1992, 1993) analyse Dutch data using appropriate duration techniques. Their findings illustrate some of the factors that affect vacancy duration, as well as raising implications about the nature of employer search behaviour. They show that there is an inverse relationship between the probability of successfully filling a vacancy and the total stock of vacancies on the market. This is a so-called 'congestion effect' predicted by standard models of search. Similarly, Russo, Rietveld, Nijkamp & Gorter (1996) show that when unemployment is high, the vacancy hazard is higher. Vacancy duration varies by occupation, educational requirements, the length of training and the size of the firm. Gorter & van Ommeren (1999), for instance, show that vacancies requiring a high level of education take longer to fill, especially where the employer advertises the vacancy. Similar US evidence is provided by Barron, Bishop & Dunkelberg (1985). Barron et al. (1997*a*), Barron, Berger & Black (1997*b*) and Burdett & Cunningham (1998) use US data to show that vacancy duration is increased where the training period is longer.

van Ours & Ridder (1993) investigate the timing of the applicant arrival rate and conclude that vacancy duration represents a 'selection' period rather than a pure 'search' period. This implies that employers do not search sequentially for workers and Abbring & van Ours offer further supporting evidence. Burdett & Cunningham (1998) add a further dimension to this debate by investigating the effect of 'advanced notice' on vacancy duration, which suggests that employers begin their search *before* a worker has quit, increasing vacancy duration.

Although the analysis of employer search using individual vacancy data has improved, previous work still suffers several limitations. Some of the earlier studies are narrowly focused on certain types of vacancy, or have imprecise measures of vacancy duration. All previous work excludes multiple vacancies advertised simultaneously, vacancies which are withdrawn from the market before they are filled, training vacancies and have few control variables with which to model vacancy duration. Few studies explore the issue of unobserved heterogeneity, or more specifically whether this should be modelled parametrically or non-parametrically. Our own work overcomes all of these shortcomings.

It is often claimed that the UK economy is particularly prone to periodic skill shortages, which move pro-cyclically and threaten productivity, competitiveness and growth. A recent *Skill Needs in Britain Survey* shows that skill shortages have increased as the economy has pulled out of recession, with hard-to-fill vacancies reported by employers rising from 5% in 1992 to 17% in 1996 (Walsh 1997). Such shortages appear to be particularly acute in associate professional and technical, personal and protective services, professional and craft occupations. Changes in production processes, technology and work practices requiring multiple skills were cited as the most frequent causes of an increase in skill needs.[3]

Hard-to-fill vacancies will have lower hazards and higher probabilities of being withdrawn from the search process. However, vacancies which take longer to fill may not represent skill shortages at all. In fact, the opposite may be true. They could be harder to fill because they are 'bad' jobs, in the sense that they offer low wages or little training, for instance. Haskel & Martin (1993) attempt to overcome this problem by analysing reported shortages for skilled jobs using establishment level data. They find that higher wages have no effect on skill shortages, whereas firms located in areas with more skilled workers are less likely to suffer shortages. What is not made clear from previous work is whether

[3]Machin (1996) argues that the shift in labour demand in favour of non-manual or skilled workers can be attributed to various indicators of technological change—R&D or innovation activity and the introduction of computers. Most of this shift, it is argued, can be accounted for by within-industry and within-establishment changes in the demand for labour. Our analysis of vacancies may also shed light on this issue, albeit from a completely different angle.

firms succeed in filling vacancies for skilled workers. Our data allow us to distinguish between the duration of vacancies and the outcome of the search process: are vacancies for skilled workers eventually filled or are they withdrawn from the labour market? The consequences for the skills gap are clearly very different in each case.[4]

3 Data and institutional background

The data we use are the computerised records of the Lancashire Careers Service over the period 1985–1992. The Careers Service fulfills a similar role for the youth labour market as Employment Offices and Job Centres provide for adults. Its main responsibilities are to provide vocational guidance for youths and to act as an employment service to employers and youths. The latter includes a free pre-selection service for employers. Use of the Careers Service is voluntary for employers with job vacancies, whereas notification of training vacancies is compulsory, so that the government offer of a guaranteed training place for all 16-17 year old youths can be monitored.[5]

The Careers Service holds records on *all* youths aged between 15 and 18, including those who are seeking employment. We observe every vacancy notified by employers to the Careers Service between March 1985 and June 1992. All training vacancies and about 30% of job vacancies are notified to the Careers Service. Unlike vacancies posted at Job Centres in the adult labour market, the job vacancies in the Careers Service data require both high- and low-quality job-seekers, and are representative of all entry-level jobs in the youth labour market.

Although our data only cover one method of search, it is an important method. 19% of all jobs and 87% of training schemes for those aged 16–18 are filled by the Careers Service. In addition, a further 18% of jobs follow directly on from training schemes. The Careers Service is therefore involved, directly or indirectly, in 37% of job placements for young people in Lancashire. (See Upward (1998, ch. 4) for fuller details.)

Employers notify the Careers Service of the type of vacancy, including detailed information about the occupation, the wage, a closing date for applications and selection criteria. Job-seekers are then selected for interview and a contact is made. Either a match occurs or the pair each continue their search.

[4]See the collection of papers in Booth & Snower (1996) for a discussion of the nature of the skills gap in Britain, and its causes and consequences.

[5]Training vacancies are limited duration subsidised placements with firms under the Youth Training Scheme, which constituted a major part of the youth labour market over this time period. The time period of the data cover three different incarnations of the Youth Training Scheme: the one-year programme from 1983–1986, the two-year programme from 1986–1990, and the variable length programme since 1990.

A vacancy has one of two possible outcomes. Either the employer successfully fills the vacancy with applicants submitted by the Careers Service, or the use of this search method is abandoned before the vacancy is filled. In this case the vacancy is described as *lapsed*. This occurs either if the Careers Service stops sending applicants for interview (perhaps because there are no suitable applicants) or because the employer decides to adopt another search strategy.

In this paper we argue that vacancies which lapse are genuinely 'hard to fill'. For this to be true, we need to be sure that vacancies which are lapsed by the Careers Service are not subsequently filled by some other method of search. We check this by searching the career histories of all school-leavers between 1988 and 1992. We find that only a tiny proportion (about 1%) of unfilled vacancies previously notified to the Careers Service were subsequently filled by other means. This still leaves the possibility that these vacancies were subsequently filled by older job-seekers. We cannot rule this out, since the Careers Service data only covers those aged 15–18. However, these vacancies are almost all specifically aimed at those who have recently entered the labour market. They offer low wages and many have some element of basic training. Indeed, training vacancies are intended only for those aged under 18. Further evidence that lapsed vacancies are not subsequently filled by older job-seekers is provided by the fact that a high proportion of lapsed vacancies subsequently re-appear at a later date.[6] Of 6054 lapsed job vacancies, 40% re-appear. We can therefore be confident that these other search channels are indeed for different vacancies than those recorded in the Careers Service database.

The dependent variable in this study measures the length of time the employer is engaged in search. This is the duration from the point when the employer posts the vacancy at the Careers Service to the point when the vacancy is either filled or lapsed by the Careers Service. A very small proportion of vacancies are right-censored at the time data collection stopped (June 1992). Left-censoring does not occur as we have a flow sample.

One further feature of these data is that employers may advertise several identical vacancies simultaneously. This is particularly true of training vacancies, although a proportion of job vacancies also have this feature. For example, a firm may want to take on 10 identical apprentice welders at the same time. These vacancies are called *multiple vacancy orders*. In principle, it is vacancies within an order that are the unit of observation, not the order itself. Unfortunately, the duration of individual vacancies within an order is not recorded, and needs to be inferred from the total duration of the whole vacancy order. Our methods for dealing with this problem are described in the following section.

[6] A vacancy is counted as re-appearing if the same firm subsequently notifies a vacancy with identical characteristics to the initial vacancy.

4 Methodology

In this section we outline our econometric methods for estimating the determinants of the vacancy hazard. Because a vacancy can exit into one of two states, namely filled or lapsed, the appropriate methodology is a standard competing risks model. This also incorporates a third state for censored vacancies. Standard methods can be applied to those vacancies which are advertised individually (i.e. not in multiple orders), where we assume proportional hazards and incorporate unobserved heterogeneity. We then amend the likelihood function to cater for the problems caused by vacancy orders which contain multiple vacancies.

4.1 A standard competing risks model

Most vacancies exit to one of two states: *filled* or *lapsed*. Some (a small number) exit to a third state, namely *censored*. The trichotomous random variable $R = 1, 2, 0$ describes these three states. The three resulting sub-samples of data are $\mathcal{E}_1$, the set of filled vacancies; $\mathcal{E}_2$, the set of lapsed vacancies; and $\mathcal{C}$, the set of censored vacancies. Each vacancy, subscripted i, belongs to one and only one set. The random variables T, $\bar{T}$, and C represent the time it takes a vacancy to be filled, lapsed, or censored respectively. The corresponding survivor functions are $S_1(t)$, $S_2(\bar{t})$, and $S_0(c)$, with density functions $f_1(t)$, $f_2(\bar{t})$, and $f_0(c)$ respectively. Clearly t, $\bar{t}$, and c are the realisations from T, $\bar{T}$, and C, but for any vacancy we only observe only one of these three outcomes. This is denoted z, given by

$$z = \min(t, \bar{t}, c).$$

Data are observed in unit intervals (days):

$$[0, 1), [1, 2), \ldots$$

For each vacancy a duration z is recorded if it is observed either filling, lapsing or censoring in the interval $[z - 1, z)$.

We assume that the three underlying stochastic processes describing time to fill, lapse and censor are mutually independent.[7] Then the likelihood of observing a filled vacancy in the interval $[z - 1, z)$ is

$$\Pr[T \in [z - 1, z)] \Pr(\bar{T} \geq z) \Pr(C \geq z) = [S_1(z - 1) - S_1(z)] S_2(z) S_0(z),$$

[7] Of these, the assumption that filling and lapsing are independent is the most contentious. However, we have an unusually large number of covariates which reduces the likelihood of correlated outcomes.

whereas the likelihood of observing a lapsed vacancy between $[z-1,z)$ is

$$\Pr(T \geq z)\Pr[\bar{T} \in [z-1,z)]\Pr(C \geq z) = S_1(z)[S_2(z-1) - S_2(z)]S_0(z).$$

Finally, the likelihood of observing a censored vacancy in the interval $[z-1,z)$ is

$$\Pr(T \geq z)\Pr(\bar{T} \geq z)\Pr[C \in [z-1,z)] = S_1(z)S_2(z)[S_0(z-1) - S_0(z)].$$

It follows then that the two (independent) likelihoods for the parameters describing the durations to filling and lapsing are respectively:

$$\prod_{i\in\mathcal{E}_1}[S_1(z_i-1) - S_1(z_i)]\prod_{i\in\mathcal{E}_2}S_1(z_i)\prod_{i\in\mathcal{C}}S_1(z_i) \tag{1}$$

and

$$\prod_{i\in\mathcal{E}_2}[S_2(z_i-1) - S_2(z_i)]\prod_{i\in\mathcal{E}_1}S_2(z_i)\prod_{i\in\mathcal{C}}S_2(z_i). \tag{2}$$

This is because the likelihood can be partitioned into two terms, that is the parameters for filled vacancies can be estimated by single-risk methods, as can the parameters for lapsed vacancies. It is for this reason why the terms describing censoring (subscripted 0) have been dropped, given we are not interested in estimating the parameters of these distributions. These likelihoods are exactly equivalent to the continuous time case, except that the density $f(z)$ is replaced by the discrete change in the survival function $S(z-1) - S(z)$.

Our methods now depend on whether we analyse the whole sample or just the sub-sample of vacancy orders that comprise a single vacancy. It is much easier to deal with unobserved heterogeneity in the latter case where numerical, as well as analytical, methods can be used.

4.2 Single vacancy orders

The standard way to estimate discrete-time duration data is to form a panel of vacancies with the i-th vacancy contributing $j = 1, 2, \ldots, z_i$ observations. This is the 'sequential binary response' form (Prentice & Gloeckler 1978, Han & Hausman 1990).[8] For both exit states $r = 1, 2$, all observations y_{ij} are zero except the last, and only if the vacancy exits to state r (that is, is not censored) is unity recorded. For example, for $r = 1$, this is the z_i-th period for the vacancy that corresponds to $[S_1(z_i-1) - S_1(z_i)]$ in (1) above.

[8]This subsection draws heavily on Stewart (1996) and references within.

Censored observations and those exiting to the other exit state are grouped together to form the zeros. The likelihood for the i-th vacancy is now written

$$L_i = \prod_{j=1}^{t_i} h_j(\mathbf{x}_i')^{y_{it}}[1 - h_j(\mathbf{x}_i')]^{1-y_{it}} \tag{3}$$

where h_j is the hazard of exit and $\mathbf{x}_i'$ is a vector of observable covariates. If one makes a proportional hazards assumption,

$$h_j(\mathbf{x}_i') = \bar{h}_j \exp(\mathbf{x}_i'\boldsymbol{\beta}), \tag{4}$$

where $\bar{h}_j$ is the baseline hazard, it can be shown that the covariates affect the hazard via the complementary log-log link:

$$h_j(\mathbf{x}_i') = 1 - \exp(-\exp(\mathbf{x}_i'\boldsymbol{\beta} + \gamma_j)). \tag{5}$$

The γ_js are interpreted as the log of a non-parametric piecewise linear baseline hazard, as $\gamma_j \approx \log \bar{h}_j$ when $\mathbf{x}_i'\boldsymbol{\beta} = 0$. We have no priors about the shape of the baseline hazard and because there are a large number of vacancies in the data, this flexible non-parametric approach is feasible. Each interval corresponds to a day, but, because of data thinning, these are grouped into longer intervals at longer durations (by constraining the appropriate γ_js). Equation (5) is substituted into Equation (3) to define a likelihood $L_i(\boldsymbol{\beta}, \boldsymbol{\gamma})$ for each vacancy with observed covariates $\mathbf{x}_i'$, where the γ_j are collected into a vector $\boldsymbol{\gamma}$. This equation is estimated separately for each exit state ($r_i = 1, 2$). These two likelihoods are reparameterisations of Equations (1, 2) above, and are useful for a number of reasons. For example, introducing time-varying covariates in this framework is straightforward.

A possible restriction on the shape of the baseline hazard is provided by the Weibull hazard, $\bar{h}_j = \gamma\alpha j^{\alpha-1}$. In this case the γ_j in (5) are replaced by $\log \alpha\gamma + (\alpha - 1)\log j$, greatly reducing the number of parameters to be estimated. These two 'homogeneous proportional hazards' models are referred to as Models A (non-parametric) and A' (Weibull).

It is well established that the failure to control for unobserved heterogeneity may induce severe bias on the shape of the baseline hazard (particularly if it is parametrically estimated). Following standard practice, the positive-valued random variable (or mixture) v is added to the hazard function given in (4) as follows:

$$h_j(\mathbf{x}_i', v_i) = \bar{h}_j v_i \exp(\mathbf{x}_i'\boldsymbol{\beta}) \tag{6}$$

$$= \bar{h}_j \exp(\mathbf{x}_i'\boldsymbol{\beta} + u_i), \tag{7}$$

where $u \equiv \log v$ and has density $f_u(u)$. The likelihood for a vacancy in this 'mixed proportional hazards' model is

$$L_i(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \int_{-\infty}^{\infty} \left[\prod_{j=1}^{t_i} h_j(\mathbf{x}_i', u_i)^{y_{it}} [1 - h_j(\mathbf{x}_i', u_i)]^{1-y_{it}} \right] f_u(u_i) du_i, \tag{8}$$

$$h_j(\mathbf{x}_i', u_i) = 1 - \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta} + \gamma_j + u_i)] \tag{9}$$

replacing Equations (3) and (5) above. For single vacancy orders, we adopt three approaches for modelling the unobserved heterogeneity.

Gamma mixing

Here $V \sim \text{Gamma}(1, \sigma_v^2)$. See Meyer (1990) for the algebraic detail for integrating out v to obtain the likelihood $L_i(\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma_v)$ in terms of observables. Note that for comparisons with distributions for U rather than V, $\mathrm{E}(\log V) = \log \sigma_v^2 + \psi(1/\sigma_v^2)$ and $\mathrm{var}(\log V) = \psi'(1/\sigma_v^2)$, where ψ is the Digamma function (Lancaster 1990, Eqn 16).

Gaussian mixing

Here $U \sim \mathrm{N}(0, \sigma_u^2)$. This can be justified by a combination of a vast number of unobserved characteristics. These enter the hazard in the same way as the observables. It is possible, but less plausible, to justify a Gamma mixture in a similar way (Stewart 1996, van den Berg 2000). Replacing $f_u(u_i)$ in Equation (8) by the Normal density $(2\pi\sigma_u^2)^{-1/2} \exp(-u_i^2/2\sigma_u^2)$, the integral can be approximated using Gauss-Hermite quadrature of the form

$$\int_{-\infty}^{\infty} \exp(-s^2) g(s) ds \approx \sum_{q=1}^{Q} w_q g(s_q)$$

where w_q denotes the quadrature weights and s_q the quadrature abscissas, both of which can be computed using standard iterative algorithms. Using the transformation $s = u/(\sigma_u\sqrt{2})$, the likelihood for the i-th vacancy becomes

$$L_i(\boldsymbol{\beta}, \boldsymbol{\gamma}, \sigma_u) = \frac{1}{\sqrt{\pi}} \sum_{q=1}^{Q} w_q \left[\prod_{j=1}^{t_i} h_j(\mathbf{x}_i')^{y_{it}} [1 - h_j(\mathbf{x}_i')]^{1-y_{it}} \right] \tag{10}$$

$$h_j(\mathbf{x}_i') = 1 - \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta} + \gamma_j + s_q\sigma_u\sqrt{2})]. \tag{11}$$

The number of quadrature points Q is chosen by the investigator, which we set at 12, after experimentation.

Non-parametric (Heckman & Singer) mixing

The standard argument for not using either of the two parametric densities above is the lack of justification for either choice. Heckman & Singer (1984) advocate the use of non-parametric mixing, arguing the effect on the baseline hazard should be less severe. Here the u_i in (8), and associated densities $f_u(u_i)$, are replaced by a discrete mass point approximation: $\bar{u}_1, \ldots, \bar{u}_M, \pi_1, \ldots, \pi_M$, collected in a vector $\boldsymbol{\theta}$. All are parameters to be estimated, and must satisfy

$$\sum_{m=1}^{M} \pi_m = 1, \quad \pi_m \geq 0 \; \forall \, m, \quad \sum_{m=1}^{M} \pi_m \bar{u}_m = 0.$$

We follow recommended practice 'that the likelihood is maximised for $M = 2$ and add mass points singly until the optimisation algorithm collapses back to the optimum for the previous value of M combined with a value of zero for one of the π parameters' (Stewart 1996). We find that M is always 2. The likelihood for the i-th vacancy is

$$L_i(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\theta}) = \sum_{m=1}^{M} \pi_m \left[\prod_{j=1}^{t_i} h_j(\mathbf{x}_i')^{y_{it}} [1 - h_j(\mathbf{x}_i')]^{1-y_{it}} \right] \tag{12}$$

$$h_j(\mathbf{x}_i') = 1 - \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta} + \gamma_j + \bar{u}_m)]. \tag{13}$$

Notice that we assume that the two u_i are uncorrelated across the risks. These three 'mixed proportional hazards' models are referred to Models $B(g)$, $B(n)$ and $B(h)$ respectively (for Gamma, Normal, and Heckman & Singer), where again a transpose indicates that a Weibull baseline hazard has been imposed.

4.3 Multiple vacancy orders

As noted in Section 3, in practice many vacancies are grouped together in 'orders'. Each order contains V_i vacancies, where the orders are numbered $i = 1, \ldots, N$. Within an order, any number of individual vacancies may be filled before the whole order is lapsed by the Careers Service. If the duration of every filled vacancy within an order were recorded, then the fact that vacancies are grouped into orders would be of no consequence. Unfortunately, this is not the case. If all vacancies are filled before the order is lapsed, we only observe the duration of the vacancy filled *last*. Further, if any vacancies within an order remain unfilled when the order is lapsed, we only observe the time of lapsing. For each order we know how many vacancies are filled, denoted W_i. We need to infer the parameters describing the distribution of a single vacancy.

Four types of order are possible, suppressing the i subscript for clarity:

1. All are filled before any are lapsed or censored ($W = V$):

$$t_1 < \bar{t};\ t_2 < \bar{t};\ t_3 < \bar{t}; \ldots;\ t_V < \bar{t} \quad \text{or}$$
$$t_1 < c;\ t_2 < c;\ t_3 < c; \ldots;\ t_V < c,$$

   but we only observe $y = \max(t_1, \ldots, t_V)$. The likelihood of observing this type of order is

$$V[1 - S_1(y)]^{V-1}[S_1(y-1) - S_1(y)]S_2(y)^V S_0(y)^V.$$

2. All are lapsed before any are filled or censored ($W = 0$):

$$\bar{t} < t_1;\ \bar{t} < t_2;\ \bar{t} < t_3; \ldots;\ \bar{t} < t_V \quad \text{or} \quad \bar{t} < c$$

   Here we only observe $\bar{t}$. The likelihood of observing this type of order is

$$[S_2(\bar{t}-1) - S_2(\bar{t})]^V S_1(\bar{t})^V S_0(\bar{t})^V.$$

3. All are censored before any are filled or lapsed:

$$c < t_1;\ c < t_2;\ c < t_3; \ldots;\ c < t_V \quad \text{or} \quad c < \bar{t}.$$

   Here we only observe c. The likelihood of observing this type of order is

$$S_1(c)^V S_2(c)^V [S_0(c-1) - S_0(c)]^V.$$

4. W are filled and $V - W$ are lapsed before any are censored:

$$t_1 < \bar{t}; \ldots;\ t_W < \bar{t};\ \bar{t} < t_{W+1}; \ldots;\ \bar{t} < t_V$$

   we only observe $\bar{t}$ and W. The $V - W$ lapsed vacancies have a likelihood

$$[S_2(\bar{t}-1) - S_2(\bar{t})]^{V-W} S_1(\bar{t})^{V-W} S_0(\bar{t})^{V-W},$$

   and the W filled vacancies have a likelihood

$$[1 - S_1(\bar{t})]^W [S_2(\bar{t}-1) - S_2(\bar{t})]^W S_0(\bar{t})^{V-W}.$$

In each expression above, contributions to the likelihood from the censored distributions are suppressed. The likelihood for the whole sample is (now explicitly indexing each vacancy order i and replacing $\bar{t}_i$, y_i and c_i by z_i):

$$\begin{aligned}
&\prod_{i \in \{W_i = 0\}} [S_2(z_i - 1) - S_2(z_i)]^{V_i} S_1(z_i)^{V_i} \times \\
&\qquad \prod_{i \in \{0 < W_i < V_i\}} [S_2(z_i - 1) - S_2(z_i)]^{V_i} S_1(z_i)^{V_i - W_i} [1 - S_1(z_i)]^{W_i} \times \\
&\qquad \prod_{i \in \{W_i = V_i\}} V_i [1 - S_1(z_i)]^{V_i - 1} [S_1(z_i - 1) - S_1(z_i)] S_2(z_i)^{V_i} \times \\
&\qquad\qquad \prod_{i \in \{C_i = 1\}} S_1(z_i)^{V_i} S_2(z_i)^{V_i}. \qquad (14)
\end{aligned}$$

This generalises Equations (1, 2) above, seen by setting $V_i = 1$. The important difference between this likelihood and the one for single vacancy orders is that the data cannot be organised into sequential binary response form. The likelihood can only be written in terms of survivor, not hazard, functions, and therefore to complete the econometric specification of the model, the functions S_1 and S_2 need deriving from the hazard function given in (6). Stewart shows that each survivor function is given by

$$S_r(z_i) = \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \delta_{rj})] \qquad r = 1, 2 \tag{15}$$

where $\delta_{rj} \equiv \log \bar{H}(z)$ is the integrated baseline hazard over the interval $[z-1, z)$. Notice that for single vacancy orders, we are able to estimate the baseline hazard directly, not its integral. To recover the γ_{rj} from the δ_{rj}, use

$$\gamma_{rj} = \log[\exp(\delta_{rj}) - \exp(\delta_{r,j-1})].$$

(See Stewart 1996, Equations 14 and 16). To examine whether the discrete-time Weibull $\bar{h}(z) = \gamma\alpha z^{\alpha-1}$ is an appropriate special case, then

$$S_r(z_i) = \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \log\gamma_r + \alpha_r \log z_i)] \qquad r = 1, 2. \tag{16}$$

replaces (15) above. In both cases, to obtain the likelihood, (15) or (16) are substituted directly into Equation (14). These two 'homogeneous proportional hazards' models for multiple vacancy orders are labelled Models C and C'.

In principle, unobserved heterogeneity is easily incorporated into this framework. Equation (15) generalises to

$$S_r(z_i, v_{ir}) = \int_0^\infty \exp[-v_{ir}\exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \delta_{rj})] f_v(v_{ir}) dv_{ir}, \qquad r = 1, 2.$$

If v_{ir} is also assumed to be Gamma distributed, with variance σ_r^2, then integrating out v_{ir} gives

$$S_r(z_i) = [1 + \sigma_r^2 \exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \delta_{rj})]^{-1/\sigma_r^2} \qquad r = 1, 2$$

Given a similar expression for a Weibull baseline hazard, these two 'mixed proportional hazards' models for multiple vacancy orders are labelled Models $D(g)$ and $D'(g)$. Because the data cannot be organised into sequential binary response form, corresponding models for Gaussian mixing or Heckman-Singer cannot be estimated. This is because the numerical methods referred to in Section 4.2 cannot be implemented.

4.4 Interpreting the parameter estimates

The vectors $\boldsymbol{\beta}_r$, $r = 1, 2$, convey no information about the effect of a single covariate x on either the likelihood of exit via risk r (Π_r), or the expected waiting time until exit via risk

r (E_r) (Lancaster 1990, Thomas 1996). This is because Π_r (and therefore E_r) depend on both h_{1j} and h_{2j} via the overall survivor function

$$\Pi_r = \sum_{j=1}^{\infty} h_{rj} S_{j-1}, \qquad E_r = \frac{1}{\Pi_r} \sum_{j=1}^{\infty} j h_{rj} S_{j-1}, \qquad S_j = \prod_{s=1}^{j} (1 - h_{1s} - h_{2s}) \,. \tag{17}$$

However, a result provided by Thomas (1996) is particularly useful when proportional hazards are assumed. Instead of examining the effects of x on the unconditional probability of exit, it is computationally much easier to focus on the probability of filling conditional on exiting during the interval j, denoted P_{1j}:

$$P_{1j} = \frac{h_{1j}}{h_{1j} + h_{2j}}. \tag{18}$$

In order to interpret our results, we compute both E_r and P_{rj}. The baseline hazards used to compute these are

$$\hat{h}_{rj} = 1 - \exp(-\exp(\bar{\mathbf{x}}'\hat{\boldsymbol{\beta}}_r + \hat{\gamma}_{rj})), \qquad r = 1, 2, \tag{19}$$

where $\bar{\mathbf{x}}$ is set at the sample mean. We report the marginal effect of a covariate x on the conditional exit probability, given by

$$\frac{\partial P_{1j}}{\partial x} = \frac{h_{1j} h_{2j} (\beta_1 - \beta_2)}{(h_{1j} + h_{2j})^2} \equiv -\frac{\partial P_{2j}}{\partial x}. \tag{20}$$

Marginal effects will vary across j, and so we evaluate them at the expected waiting time until exit via risk 1 (filling).

We also report the 'simulated' marginal effect of a covariate on the expected waiting time. E_r is initially computed at the sample mean. For dummy variables, we then re-evaluate at $x = 0$ and $x = 1$. For continuous variables, we evaluate at $x + 5\%$ and $x - 5\%$.[9] We denote these simulated marginal effects by $\Delta E_r / \Delta x$ and can be viewed as close approximations to marginal effects.

In our results we report (a) marginal effects $\partial h_r / \partial x \approx \hat{\boldsymbol{\beta}}_r$, (b) marginal effects $\partial P_{1j} / \partial x$ and (c) the simulated marginal effects $\Delta E_r / \Delta x$. $\partial P_{1j} / \partial x$ reveals how a particular covariate affects the probability of employer search being successful, while $\Delta E_r / \Delta x$ reveals how a covariate affects the duration of that search, conditional on exit to risk r.

[9]For all continuous variables, shifts of this size are well within the range of the data.

5 A theoretical framework for interpreting the regression results

In this section we outline stylised one-sided and two-sided search models in order to establish a framework for analysing the impact of our covariates on the employer's hazard and vacancy duration. This is taken from Burdett & Wright (1998). We then describe various predictions about the likely shape of the baseline hazard. Finally, we comment on the likely impact of skill shortages on the probability of vacancies being removed from the market and on the shape of the hazard to lapsing.

5.1 The impact of covariates on the duration of employer search

The flow utility or "payoff" to the employer for a particular applicant is given by

$$Z = Z(p - w, E), \tag{21}$$

where p is the revenue flow if the match occurs, w is the known wage rate for the type of labour to be supplied (defined by occupation, location etc); and E is a random variable capturing the idiosyncratic utility unique to the match, with variance s. E comprises considerations about the applicant that the employer finds relevant, such as personality or ability. p can be thought of as comprising the average quality of the vacancy and the average quality of the applicant for the type of labour to be supplied. An increase in $p-w$ shifts the mean of the distribution of Z to the right. The variance of Z is denoted s.

We assume that w is exogenous; this is a reasonable assumption (discussed in more detail below) in the youth labour market, as the training allowance is institutionally determined and the wage for jobs is closely related, although may also respond to competitive forces. The employer (denoted by a superscript e) gets a draw z from $F(z)$. Its reservation payoff, r^e, is given by the standard first order condition for continuing to search:

$$r^e = -c^e + \frac{\alpha^e}{\delta + r}\rho(r^e) \quad \rho(r^e) = \int_{r^e}^{\infty} (z - r^e) dF(z)$$

$\rho(r^e)$ is the surplus function and is decreasing in r^e. c^e is the cost to the employer of keeping the vacancy open (search costs, opportunity cost of keeping capital idle). α^e is the arrival rate of applicants to the employer, i.e. job-seekers who have already found the vacancy acceptable. r is the discount rate and δ is the separation rate at which workers leave the employer. As neither appear in the empirical analysis below, they are ignored henceforth.

The solution to this first order condition is

$$r^e = r^e(c^e, \alpha^e, p - w, s^e) \quad r_1^e < 0,\ r_2^e > 0,\ 0 < r_3^e < 1,\ r_4^e > 0.$$

The comparative statics are absolutely standard (e.g. Mortensen 1986). Employers increase their reservation utility if the costs of search reduce, if the arrival rate of applicants increases, if the wage falls, if the revenue flow from the match increases or if there is a mean-preserving increase in the variance of E.[10]

The probability that an employer will find an applicant acceptable, μ^e, is then

$$\begin{aligned} \mu^e &= 1 - F(r^e) \\ &= \mu^e(c^e, \alpha^e, p - w, s^e) \quad \mu_1^e > 0,\ \mu_2^e < 0,\ \mu_3^e > 0,\ \mu_4^e \lesseqgtr 0. \end{aligned} \tag{22}$$

Thus employers become more selective (in that μ^e falls) as search costs decrease, as the arrival rate of applicants increases, if the wage increases or if the revenue flow decreases. Note that although an increase in the wage causes the employer to decrease its reservation utility, it does so by less than the increase in the wage because $0 < r_3^e < 1$, and so the employer becomes more selective.[11] Also note that the effect of s is ambiguous. These initial 'one-sided' impacts on r^e and μ^e are employer selection effects.

The hazard for a vacancy is therefore

$$h^e = \alpha^e \mu^e = \lambda^e \mu^w \mu^e \equiv \lambda^e \mu$$

Here the arrival rate of job-seekers who are willing to take the job, α^e, has been decomposed into λ^e, the average arrival rate of job-seekers to a given vacancy and μ^w, their acceptance probability.[12] $\mu \equiv \mu^w \mu^e$ is the matching probability, the probability that a vacancy is filled for a given job-seeker.

We denote the average contact rate between job-seekers and vacancies as $\lambda(U, V)$, where U is the stock of job-seekers and V is the stock of vacancies, and assume that $\lambda(\ )$ has usual properties, in particular constant returns to scale. Most evidence supports this (Broersma & van Ours 1999, Table 1). Then the average rate at which job-seekers contact a given vacancy is given by

$$\lambda^e(V/U) = \lambda(U, V)/V = \lambda(U/V, 1) \quad \lambda_{V/U}^e < 0.$$

[10] Mortensen shows that $\partial r^e/\partial(p - w) = h^e/(r + h^e)$, where h^e is the vacancy hazard, defined below. It lies in the (0,1) interval.

[11] Mortensen shows that $\partial \mu^e/\partial(p - w) = \alpha^e F' r/(h^e + r) > 0$.

[12] Strictly speaking, the job-seeker's acceptance probability for an employer who has already made an offer.

which is decreasing in labour-market tightness $\theta \equiv V/U$. Note that in contrast, the average arrival rate of applicants refers to those job-seekers who have already found the employer acceptable, $\alpha^e = \mu^w \lambda^e(\theta)$.

Making all the substitutions, the vacancy hazard is given by

$$h^e = \lambda^e(\theta).\mu^w.\mu^e(c^e, \mu^w \lambda^e(\theta), p - w, s^e) \tag{23}$$

The job-seeker's acceptance probability, μ^w, depends on the same set of arguments as the employer's acceptance probability, but superscripted w rather than e:

$$\mu^w = \mu^w(c^w, \mu^e \lambda^w(\theta), w, s^w) \quad \mu_1^w > 0,\ \mu_2^w < 0,\ \mu_3^w > 0,\ \mu_4^w \lesseqgtr 0. \tag{24}$$

The important point about Burdett & Wright's two-sided search model is that the reservation utility of the worker affects the reservation utility of the employer, and *vice versa*. If, for some reason, r^w goes up and workers become more selective, this reduces the flow of acceptable applicants to the employer, who becomes less selective. This can be seen as the negative impact of μ^w on μ^e in (23); there is a symmetrical negative impact of μ^e on μ^w in (24). From the employer's point of view, these might be labelled applicant arrival effects. Given two negative functions $r^w = f^w(r^e)$ and $r^e = f^e(r^w)$, the equilibrium solutions for (r^e, r^w) in terms of c^w, c^e, θ, p and w simply reinforce the initial one-sided responses. This means that comparative statics effects are best seen in the following expression for the vacancy hazard, where for clarity the dependence of μ^w on μ^e and *vice versa* has been suppressed, even though they will still be present in the data:

$$h^e = \lambda^e(\theta).\mu^e(c^e, \mu^w \lambda^e(\theta), p - w, s^e).\mu^w(c^w, \mu^e \lambda^w(\theta), w, s^w). \tag{25}$$

Burdett & Wright show that it is possible to have an equilibrium with job-seekers being especially 'easy' in the sense that they accept all job offers. This could happen in a very slack labour market.[13] As noted in Footnote 1, there is evidence that job-seekers rarely refuse job offers from employers. If it is the case that the job-seeker's acceptance probability is close to unity, or less stringently, that the employer's effects dominate those of the job-seeker, then the vacancy hazard can be interpreted as representing employers' search behaviour, and this is the assumption we make in interpreting our results. In Andrews et al. (2001) we confirm that this assumption is justified, using the same data that we analyse below, where we estimate models for the matching probability μ. Because we expect the employer's effect to dominate for all variables that operate through both sides of the market, we can more simply write our empirical model as:

$$h^e = \lambda^e(\theta).\mu(c^e, \theta, p - w, s^e) \quad h_\theta^e \lesseqgtr 0,\ h_{p-w}^e > 0,\ h_{c^e}^e > 0,\ h_{s^e}^e \lesseqgtr 0. \tag{26}$$

[13] In fact, Burdett & Wright normalise $V/S = 1$, but the model easily generalises.

This clearly distinguishes employer selection effects that operate via μ and from applicant arrival effects that operate via λ^e.

To summarise, our predictions about the effects of observed covariates on the vacancy hazard follow from (25), and summarised in (26), are as follows.

1. Labour-market tightness ($\theta \equiv V/U$). Employers get fewer contacts per vacancy, ($\lambda^e_\theta < 0$). They respond by becoming less selective, lowering r^e so that μ^e goes up ($\mu^e_2 \lambda^e_\theta > 0$). Job-seekers respond by increasing r^w and so μ^w falls ($\mu^w_2 \lambda^w_\theta < 0$). Because we assume that the job-seeker's effect is dominated by μ^e, then the net effect is positive, and so the matching probability $\mu \equiv \mu^e \mu^w$ is increasing in labour-market tightness (pro-cyclical). However, the direct effect on $\lambda^e(\theta)$ is negative, so in general the effect of labour market tightness on the vacancy hazard is ambiguous.

2. The wage (w). The effect of the wage is absolutely standard in search theory, affecting employers and job-seekers in opposite directions. If the mean of the offer distribution (in utility terms) increases, the optimal response of the job-seeker is to increase r^w, but by not as much as the shift in the distribution, and so μ^w increases (is less selective). By symmetry, the employer is more selective and so μ^e falls. Again, we assume that the employer's response dominates.

3. Quality of the match (p). Assuming that this increases the employer's revenue flow relative to its labour costs, the employer is better off, and so it is less selective (μ^e up). The effect on the vacancy hazard is positive.

4. Costs of search (c^e). The unambiguous prediction is that employers are less selective and so the vacancy hazard increases.

5. Variance of payoff distribution (s^e). Although the theory is ambiguous, a common-sense prediction is employers will wait longer for a "bargain", of which there are relatively more, in which case the effect on the vacancy hazard is negative.

Finally, note that in the stylised two-sided search model above, it is only labour-market tightness that influences the arrival rate of job-seekers to a given vacancy. This is because, *ex-ante*, all vacancies look identical to a job-seeker. In reality, there will be observed and unobserved (to the investigator) characteristics of vacancies that make some more attractive than others. Such variables, such as the wage, could then operate through the applicant arrival rate $\lambda^e(\ )$. For certain covariates we can infer the determinants of λ^e, using probit estimates for the matching probability μ, reported in a Andrews et al. (2001), which uses a subsample of the data where we observe both job-seeker and employer characteristics.

5.2 The shape of the baseline hazard

If the environment is stationary, then the employer's reservation utility is constant, and the vacancy hazard is flat. However, it seems likely that the hazard will decline with duration as it typically does for unemployed job-seekers, and what little evidence there is suggests that this is so (van Ours 1990). As with unemployment hazards, a declining hazard might be caused by neglected unobserved heterogeneity, meaning that 'better' vacancies fill first and 'worse' vacancies fill later.

The main reason why we might observe genuine negative duration dependence is that the average quality of the applicant falls, because the most suitable having been chosen first by the Careers Service. The employer's response is to become more selective ($\mu_3^e < 0$). For example, van Ours (1990) estimates a model of the vacancy hazard where he also observes the arrival rate of applicants λ^e, from which a model for the matching probability μ can also be estimated. He finds that λ^e declines with duration but that μ increases with duration, the latter implying that employers become less selective, with an overall effect of positive duration dependence consistent with van Ours (1989). However, casual empiricism and some evidence suggests that the initial arrival rate of applicants is high and tails off quickly, ie the hazard is not continuous, but has a sharp discrete fall at some low duration. Coles & Smith (1998) provide a convincing explanation. Once a given stock of employers and job-seekers have contacted and subsequently rejected each other, then employers will only search the flow of new arrivals of job-seekers (and *vice versa*), which necessarily lowers the rate at which they contact each other.

Other theoretical issues have been noted in the literature that might affect the shape of the hazard. First, if the employee gives 'notice' and announces that she will leave at some point in the future (Burdett & Cunningham 1998) the costs of search to the employer are no longer constant, since lost output due to search will not occur until the current employee actually leaves. This suggests that the hazard will be increasing up to the point where the employee quits, and constant thereafter. Unfortunately we do not observe whether vacancies have a future start date. Second, van Ours & Ridder (1993) suggest that there is an 'application' period when vacancies are collected and a subsequent 'selection' period at the end of which the best applicant is chosen. Both periods are optimally chosen by the employer, making up vacancy duration. By definition, the vacancy hazard is low (possibly zero) during the application period. The shape of the hazard might reveal whether there is an application period in the data.

5.3 Towards a theory of lapsing and skill shortages

As already noted, in our sample a large proportion of vacancies never fill, which we model in a competing risks framework as the vacancy exiting to another state called ‘lapsing’. This is analogous to the process by which job-seekers exit unemployment by leaving the workforce. In the search framework set out above, lapsing occurs if either the employer’s acceptance probability or the arrival rate of applicants falls to zero. This would imply that the lapsing hazard, denoted h_2^e will be a decreasing function of both μ and λ^e. The combination of both hazards should imply that the conditional probability of lapsing P_{2j} should be an increasing function, asymptoting towards unity as j increases.

What is of particular interest is identifying which types of vacancy are more likely to lapse. The argument that there are skill shortages (Section 2) suggests that ‘better’ jobs have very low applicant arrival rates, shifting the hazard to filling downwards. Note, however, that low applicant arrival rates will also make employers less selective. For skill shortages to cause lapsing, we require that at some point the applicant arrival rate is effectively zero. However, it seems likely that the relationship between skill shortages and the hazard is more complicated than this, because the arrival rate for high-skill jobs might initially be higher than for low-skill jobs. In short, it is an empirical issue as to the shapes of the two hazards and which way they are shifted by higher quality vacancies, and the resulting impact on the conditional probability of exit to filling.

6 Results

6.1 The raw data

Table 1 describes the sample, which covers the period 1985–1992. There are 14510 job vacancy orders containing a total of $\sum_{i=1}^{N} V_i = 17759$ vacancies. Most job vacancy orders (12840) therefore contain a single vacancy. In contrast, most training vacancy orders contain multiple vacancies: the 4185 training vacancy orders contain a total of 29656 training vacancies. What is possibly surprising is the number of vacancies which lapse: 34% of job vacancies and 11% of training vacancies.

[TABLE 1 ABOUT HERE]

Table 1 also summarises the dependent variable, the total time that a vacancy is open on Careers Service records. Mean duration can only be calculated for single vacancy orders. Job vacancies which fill have a mean duration of three weeks, compared with over 10 weeks for training vacancies. Job vacancies which lapse have a mean duration of six weeks compared with 16 weeks for training vacancies. The higher mean duration of training vacancies is partly due to an institutional feature: over 70% of training vacancies are notified between March and June, whereas most school-leavers who start training schemes do so between June and August. These details suggest that employer search in job and training markets is very different, and should be modelled separately. Furthermore, Andrews et al. (2001) also note that the probability of a job-seeker being offered a training vacancy is much higher than being offered a job vacancy because one objective of government sponsored training schemes was to mop up the excess supply of youth labour.

To calculate the underlying average duration of filled and lapsed vacancies in the absence of censoring from either the other exit state or genuine censoring, we compute ML estimates of the parameters γ_1 and γ_2 from an exponential distribution (i.e. Equation 5 with $\gamma_{rj} = \gamma_r$). The inverse of γ_r is an estimate of average duration for outcome j. In all four cases (lapsed/filled pairwise with jobs/training) this estimate of duration is about twice as long as the raw mean duration, which does not take account of censoring by the other outcome.

Using the likelihood for competing risks with multiple orders, equation (14), we also estimate average duration for all vacancies (second panel), again assuming an exponential distribution. These estimates produce slightly lower average durations for filling, and much higher average durations for lapsing, which is a result of a relatively small proportion of multiple orders which lapse.

Because the sample period is long relative to the average length of vacancies, the number of censored vacancies (vacancies which were still open at the end of June 1992) is small. There are just 460 (147+313) censored vacancy orders, comprising 2416 (220+2196) individual vacancies. Finally, the third panel of Table 1 gives the number of the four types of order which make up the modified likelihood.

[TABLES 2 to 5 ABOUT HERE]

Tables 2 to 5 summarise the other sample characteristics, weighted by the size of the vacancy order. Again, differences between the training market and the jobs market are pronounced. Training vacancies are highly concentrated in the 'Other Services' sector — often the public sector and personal services. Job vacancies are more dispersed, with higher proportions in manufacturing, distribution and banking.

Given that a sizable proportion of vacancies fail to find a successful match (i.e. lapse), it is interesting to note their characteristics. Vacancies which lapse are more likely to occur in small firms, in firms which are not centrally located, and in firms which provide training placements. Vacancies for skilled jobs, non-manual jobs and jobs which offer apprenticeship training are all more likely to lapse. Similarly, vacancies which have more demanding selection criteria such as requiring higher qualifications, older applicants or written applications are disproportionately lapsed. Perhaps most significantly, job vacancies which lapse tend to offer higher wages. Figure 1 shows some evidence of a counter-cyclical variation in the proportion of training vacancies which lapse: between 1986 and 1988 the proportion is less than 10%, but over 20% in 1985 and 1989–1991. The variation in job vacancies is much less pronounced.[14]

Figure 2 shows that the total number of job vacancies opened is pro-cyclical; this is less obviously so for training, suggesting a possible active (counter-cyclical) labour-market policy to mop-up youth unemployment.

In Figure 3 we plot the time it takes job vacancies to fill or lapse over the business cycle, estimated from an exponential competing risks model with only a set of year dummies as regressors. Vacancies took about half as long again to fill in the peak (1990) than in the year in the data with the highest unemployment rate (1985). The relationship between the business cycle and the time it takes to lapse job vacancies is less clear cut. Although there is a substantial increase between 1985 and 1989, there is a sharp drop in 1990.

6.2 Preferred specification

In Section 4 we described several possible specifications for modelling the data. Different specifications are required for (a) multiple vacancy orders, (b) unobserved heterogeneity and (c) the choice between parametric and non-parametric baseline hazards. In this section we briefly discuss our method for selecting a preferred specification. Table 6 summarises all the possible specifications.

[TABLE 6 ABOUT HERE]

We note at the outset that some models with non-parametric hazards and/or multiple vacancy orders with unobserved heterogeneity could not be estimated because the likelihoods are badly conditioned near potential maxima. In other words, such models are asking too much of the data. This was particularly true for training vacancies. Table 7 summarises results from those models that could be estimated.

[14] The unemployment rate in Figures 1 and 2 is the adult rate for Lancashire, source NOMIS (National Online Manpower Information System).

[TABLE 7 ABOUT HERE]

For job vacancies, Table 7 shows evidence that the mixing models are preferable to the homogeneous models. In most cases we cannot reject the hypothesis that the variance of the distribution of unobserved characteristics is zero. Figures 4 and 5 compare the estimated baseline hazard from the parametric specification $B'(n)$ with that of the non-parametric estimates from specification $B(n)$. The Weibull specification ignores the non-monotonicity of the hazard which occurs after 3 days (see below), but is generally a reasonable approximation. Estimates of the Weibull shape parameter vary little between the three mixing models; all confirm that the hazard to filling is decreasing ($\alpha_1 \approx 0.75$) (vacancies become harder to fill the longer they remain open) while the hazard to lapsing is increasing ($\alpha_1 \approx 1.4$) (vacancies are more likely to lapse the longer they remain open). This seems intuitively sensible. The variance of unobservables is greater for the whole-sample specification $D'(g)$ ($\hat{\sigma}_v^2 = 1.62$) than for single job vacancies, specification $B'(g)$ ($\hat{\sigma}_v^2 = 0.49$). Given that the Weibull specification closely follows the data, and that mixing models dominate homogeneous models, our preferred specification for job vacancies is therefore $D'(g)$.

It should be emphasised that parameter estimates $\boldsymbol{\beta}_r$ are very robust across *all* specifications. Only 35 estimated parameters (out of a total of 118) change sign across any of the specifications, and of those only *two* are significantly different from each other. Whichever specification we report, significant estimates are very similar. The only lack of robustness is for wages, which we discuss in more detail later.

For training vacancies, the parametric mixing model $B'(n)$ collapses back to its equivalent homogeneous model A'. Similarly, the non-parametric mixing model $B(n)$ is almost identical to model A. Thus for training vacancies, the homogeneous parametric specifications are satisfactory. This makes sense, in so far as training vacancies are much more similar to each other than jobs (partly because vacancies within a given order are identical). In Figures 6 and 7 we plot estimates of the baseline hazard for both the non-parametric and parametric specifications, and it again appears that the Weibull model is a reasonable approximation. Again, there are some non-monotonic regions of the hazard to filling in the first two weeks. In addition, the Weibull shape parameter does not pick up the large spike in the hazard on the first day. Nevertheless, once again parameter estimates are extremely close whether or not we use a non-parametric baseline.

However, it is important to include multiple vacancy orders in the estimation, since model C' gives quite different estimates both of the Weibull shape parameter α_1, α_2 and the parameter estimates $\boldsymbol{\beta}_1$, $\boldsymbol{\beta}_2$. This is not surprising as very few vacancies (1977 out of

29656) come as single vacancy orders, whereas for jobs it is 12840 out of 17759. For the complete sample of training vacancies, the (downward) hazard to filling and the (upward) hazard to lapsing are steeper than for single training vacancies only—compare $(\alpha_1, \alpha_2) = (0.363, 1.979)$ from specification C' with $(\alpha_1, \alpha_2) = (0.793, 1.232)$ from specification A'. Our preferred specification for training vacancies is therefore C'.

6.3 Parameter estimates

Tables 8 and 9 report parameter estimates from specification $D'(g)$ for job vacancies and C' for training vacancies. In discussing the results, we have three objectives. First, to examine the shape of the underlying vacancy hazards. To do this we plot estimates of h_j for various specifications, calculated from (5). Second, to test whether the predictions of the two-sided matching model are consistent with the effects of covariates on the duration of employer search. To do this we primarily examine the coefficient estimates $\partial h_1/\partial \mathbf{x} \approx \hat{\boldsymbol{\beta}}_1$ (rather than the associated marginal effects $\partial P_{1j}/\partial x$ and $\Delta E_1/\Delta x$). Third, to see whether there is any evidence of skill-shortages by examining which types of vacancy are more likely to be removed from the market before they are filled. This is done by examining the marginal effect of covariates on the probability of filling, $\partial P_{1j}/\partial x$.[15]

6.3.1 The shape of the baseline hazard

In Figures 8 and 9 we plot an completely unrestricted estimate of the baseline hazard for job and training vacancies. We compare this with the restricted non-parametric baseline hazard used in our estimates, which groups the γ_js into weekly intervals from 2 weeks to 8 weeks, and monthly intervals from then on. Figure 8 shows that the probability of successfully filling a job vacancy within one day of opening is nearly 0.08. This drops dramatically to about 0.02 for the next few days, with subsequent spikes in the hazard appearing at weekly intervals. After about 60 days the hazard appears extremely flat. For training vacancies (Figure 9), the hazard is lower throughout, with a similar spike on

[15] It is a moot point whether one should examine the effect of covariates on h_1 rather P_1 when examining the theory, which is the same as treating lapsed observations the same as genuinely censored observations. There are two reasons (but the reader may think differently). First, we do not have a good theory why vacancies do get lapsed, except that ultimately they do not get filled, when (we suspect) it is administratively neater to remove them for the market. It seems sensible to separate this "administrative" decision from the firm's "economic" hiring decisions. It is only when we explicitly consider *why* some vacancies never get filled (skill shortages?) do we examine P_1. Second, relatively few vacancies do lapse (especially training vacancies) and so the estimates of P_1 are "contaminated" by relatively poorer estimates for $\hat{\boldsymbol{\beta}}_2$, compared with $\hat{\boldsymbol{\beta}}_1$.

the first day. There is less evidence of spikes in the hazard at weekly intervals, and the hazard becomes flat sooner than for job vacancies.

The non-monotonicity of the hazard in the first week, and the subsequent spikes in the hazard at weekly intervals are also interesting phenomena. This may be an institutional feature of the data, perhaps if the Careers Service searches the pool of applicants for a particular vacancy on a weekly basis. It might also suggest that for some vacancies there is a period of application where the hazard is very low, followed by a period of selection, as suggested by van Ours & Ridder (1993). Clearly, the non-monotonicity of the hazard in Figures 4–7 occurs because this institutional 'seasonality' has not been smoothed away by grouping at short durations.

We have already established that the shape of the baseline hazard is the result of omitted heterogeneity. Figure 10 compares the estimated baseline for job vacancies between the homogeneous non-parametric specification A and the mixed specification $B(n)$. For training vacancies, the mixed model collapses back to the homogeneous model. The effect of the mixed model is to rotate the hazard anti-clockwise, reducing the extent to which it declines over time. The same effect can be seen in the equivalent Weibull specifications, where shape parameter is always larger in mixing models.[16] This is precisely what one would expect if (unobservably) better vacancies are exiting earlier than (unobservably) worse ones. Nevertheless, the estimated hazard from specification $B(n)$ still declines from nearly 0.05 to 0.02 within the first few weeks of a vacancy being opened.

Our data therefore provides some strong evidence for a decline in the hazard to filling after a very short period of time, consistent with Coles & Smith's theory of stock-flow matching. In this market the selection of applicants is computerised and practically instantaneous. Provided that there is a suitable applicant in the stock of job-seekers, a new vacancy can fill extremely quickly. Further information is provided by Andrews et al. (2001), who report that, for job vacancies, the matching probability falls by 0.013 percentage points after one month and by another 0.013 after 2 months, but remains flat thereafter. Given the hazard's shape is roughly similar after one month, the implication is that the applicant arrival rate λ^e is roughly constant for durations longer than one month. Unfortunately, we are unable to infer whether the sharp drop in the hazard in the first few days is due to μ or λ^e. Roughly the same occurs for training vacancies, except that the matching probability *rises* after 6 months, because one of the training programme's objectives was to mop up excess supply of youth labour.

[16] For single vacancy orders, compare the three mixing B' models $(\alpha_1, \alpha_2) \approx (0.75, 1.4)$ with model A' $(\alpha_1, \alpha_2) \approx (0.6, 1.1)$; for all vacancies compare $(\alpha_1, \alpha_2) = (0.78, 1.21)$ with $(\alpha_1, \alpha_2) \approx (0.64, 1.04)$.

6.3.2 The predictions of the search model

The predictions of the stylised search model were summarised in Section 5.1. Provided we assume that employer effects dominate job-seeker effects, the hazard to filling should decrease with the wage, the quality of the match and the variance of the payoff distribution. The hazard should increase with the costs of search. Decreasing labour market tightness (θ) should decrease the hazard provided that the direct effect on $\lambda^e(\theta)$ outweighs the indirect effect of employer selectivity on the matching probability μ.

Labour market tightness

Our measures of labour market tightness are the number of unemployed aged 18 or less and the number of vacancies in each local district for each month of the data. The vacancy data is calculated from the Careers Service data itself, while the unemployment data is taken from NOMIS.[17] For vacancies, we are able to distinguish the total stock of vacancies open in each district-month from those vacancies which 'compete' directly in the sense of being in the same occupation. We also separate job vacancies from training vacancies.

The coefficient on "Unemployed $\leq$ 18" in Table 8 shows that job vacancies in labour markets with higher youth unemployment have significantly higher hazards to filling ($\hat{\beta}_1 = 0.29$), lower hazards to lapsing ($\hat{\beta}_2 = -0.13$) and are therefore significantly more likely to fill ($\partial P_{1j}/\partial \log U = 0.090$). The effect on expected waiting time conditional on filling is however very small ($\Delta E_1/\Delta \log U = -0.3$ days).

As noted, labour market tightness operates via two opposing channels, (assuming that job-seeker selection effects are minimal). There is the direct effect on applicant arrival rates $\lambda^e(\theta)$ and the indirect effect via the matching probability μ^e. However, additional evidence on the effect of the labour market tightness is available from Andrews et al. (2001), where we established that an increase in the district-level stock of unemployment reduces the probability of a match, conditional on a contact. The exact estimate was $\partial\mu/\partial \log U = -0.17$. From (26), the hazard can be decomposed in $h^e(\theta) = \lambda^e(\theta)\mu(\theta)$, and so if we acknowledge the estimate cames from a smaller sample, we infer that $\partial \log \lambda^e/\partial \log U = 0.46$. This large effect is consistent with the theory (although only guaranteed if the offer distribution is log-concave) and so our conclusion is that, in the jobs market, an increase in unemployment increases the hazard because the increase in the number of contacts per vacancy outweighs the employers' 'more selective' response.

[17] We cannot use the Careers Service data to calculate the stock of unemployment because the data only covers entrants to the labour market from 1988 onwards.

In subsection 5.2, we conjectured that the hazard to lapsing should be a negative function of both λ^e and μ. This is consistent with 3 estimates reported immediately above—$\partial\mu/\partial\log U = -0.17$, $\partial\log\lambda^e/\partial\log U = 0.46$, $\partial\log h_2^e/\partial\log U = -0.13$—providing, of course, that the effect of unemployment on the arrival rate of applicants dominates the effect of unemployment on the matching probability. This seems extremely plausible, since this is true for the filling hazard.

As expected, the impact of the stock of vacancies is the reverse: the hazard to filling is reduced ($\hat{\beta}_1 < 0$) although the effect on the hazard to lapsing is insignificant. The net effect on the probability of employer search being successful is therefore negative, with $\partial P_{1j}/\partial x = -0.012$, but this does not hold for measures of competing job vacancies, nor for training vacancies. However, results on our other measures of vacancies are less intuitive. Increases in the stock of vacancies in the same occupation (which ought to have greater congestion effects) actually increase the hazard to filling. It would also appear that the number of training vacancies in the market has no significant effect on the hazard to filling for job vacancies, with $\hat{\beta}_1$) insignificantly different from zero. This is sensible if training vacancies are not competing with job vacancies, possibly because they are less attractive to job-seekers. In all cases the impact of vacancies on expected waiting times is extremely small.

The results for training vacancies (Table 9) suggests that this market does not respond in the same way to labour market tightness. The stock of unemployment has no significant effect on the duration of employer search or its outcome ($\hat{\boldsymbol{\beta}}_1$, $\hat{\boldsymbol{\beta}}_2$ and $\partial P_{1j}/\partial x$ are insignificantly different from zero). The number of job and training vacancies in the market does reduce the hazard to filling for training vacancies ($\hat{\boldsymbol{\beta}}_1 = -0.06, -0.07$), but the stock of competing job vacancies (those in the same occupation) actually has a positive effect on the hazard to filling.

The wage

There are three different types of wage offer in the data.[18] Over 95% of training vacancies and about 80% of job vacancies have a set pre-announced wage, where the wage is non-negotiable (see Table 3). The majority of these vacancies specify age and tenure profiles, which reflects the rigid institutional nature of wage setting in the youth labour market. A small proportion of both training and job vacancies have a set pre-announced wage offer, but are still open to negotiation. The remaining job vacancies have a negotiable wage offer and no pre-announced wage. For this third category there is no wage recorded in

[18]The term 'wage' also refers to the training allowance for training vacancies.

the data.

The important point is that both job-seekers and employers take the wage as given when they decide whether or not to form a match. Burdett & Wright's (1998) model in Section 5 assumes that the wage is not negotiable after agents meet and assume that an agent cannot transfer utility to the other party by varying the wage or by other means.[19] We argue that this is an accurate characterisation of the youth labour market, given the vast majority of job and training vacancies in the data have a non-negotiable wage.

We model these effects as follows. Define N as a dummy variable indicating whether a vacancy has a negotiable wage offer, and D as a dummy variable indicating whether the wage is pre-announced. Interacting the dummies with the log real hourly wage rate w, where it exists, allows us to include all observations, even where a wage is not observed.

The predictions of the model are that a higher wage should make employers more selective, but job seekers less selective. If the employer selection effect dominates, this will lower the hazard to filling a vacancy. But this ignores another potentially important channel by which the wage can affect the hazard, via the applicant arrival rate λ^e. In the two-sided search model, the rate at which job-seekers contacts vacancies is purely random except for the relative numbers of job seekers and vacancies in the market (θ). In reality, all vacancies are *not ex ante* identical and so the wage signals something about the quality of the vacancy irrespective of who the job seeker is (Burdett & Cunningham 1998); in other words, a higher wage will increase λ^e and increase the hazard, thereby offsetting the negative employer selection effect.

It turns out that the results on the wage variable are sensitive to the choice of specification — we find that the inclusion of multiple vacancy orders in the analysis significantly affects several of the parameter estimates, particular for training vacancies. A complete set of estimates for $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ across all specifications are reported in Table 10.

The coefficients on wDN and $wD(1-N)$ capture the relationship between the wage and the vacancy hazard for vacancies with a pre-announced wage with and without negotiation. For job vacancies neither has any significant effect on the hazard to filling. The effect of $wD(1-N)$ on the hazard to lapsing is positive for single vacancies but insignificant for the whole sample (specifications C, C' and $D'(g)$). But Andrews et al. (2001) found clear negative effects on the matching probability, in line with the theory, which implies that a higher wage must generate more applicants, thereby supporting, as just suggested, an unobserved quality effect (over and above other covariates that reflect the quality of the

[19]Note that in Pissarides (1990), utility is transferable, and the wage is determined by splitting the total surplus of both parties.

vacancy). For training vacancies, the same occurs, that is the filling hazard is positive (but barely significant), and we ignore the large estimates for all other specifications as they are based on a small number of vacancies. Similarly the large hazard to lapsing can be ignored, as so few training vacancies actually lapse.

For job vacancies, those that have a pre-announced wage but which are still open to negotiation (DN) have significantly lower hazards to filling than the base group, vacancies with a set pre-announced wage only. This is intuitive for two reasons: first because negotiation takes time and second because these jobs offer potentially higher wages than those with inflexible wages. Note that for training vacancies we get a similar and stronger result, but only for specification C where multiple vacancy orders are included. If these are not included estimates of $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ are insignificantly different from zero.

Job vacancies can also be negotiable but have no pre-announced wage ($(1-D)N$). Results are insignificant for $\hat{\boldsymbol{\beta}}_1$ and $\hat{\boldsymbol{\beta}}_2$ across all specifications unless multiple orders are included, in which case we get a similar result that negotiated wages cause a decrease in the hazard to filling, although this is only just significant at the 10% level.

The revenue flow from a match

We cannot observe the revenue flow from a match directly (p in Equation 21). We do, however, have a number of vacancy characteristics which are likely to be good proxies. These include the skill level, whether the vacancy is in a non-manual occupation, and the amount of training offered. Existing empirical evidence (Section 2) suggests that the greater the potential investment by the employer in the worker (because of a higher wage, more training and so on) the longer it takes to fill a vacancy (lower vacancy hazard). This seems to suggest some kind of investment cost which outweighs future revenue p, i.e. p is actually lower for more skilled jobs. In addition, if higher quality vacancies attract more variable applicants, the vacancy hazard will be lower if the employer waits longer for a "bargain".[20] It may also be the case that vacancies with better characteristics (more training, higher skill) have steeper wage profiles which are not picked up by the starting wage.

Table 8 shows that non-manual job vacancies have significantly lower hazards to filling ($\hat{\beta}_1 = -0.49$), but that skilled job vacancies are insignificantly different from unskilled vacancies. The –0.49 estimate implies that it takes 20 days longer to fill non-manual job vacancies.[21] Stronger evidence that employer search is longer for "better" vacancies

[20]Though note that in theory $\mu_4^e \lesseqgtr 0$ from Equation (22).

[21]Assuming that hazard is flat, –0.49 means a 39% longer duration at a sample mean of about 50 days

is provided by the training information. Vacancies offering day release ($\hat{\beta}_1 = -0.35$) or apprenticeship training ($\hat{\beta}_1 = -0.51$) have lower hazards.[22] It seems plausible that vacancies with these characteristics are unambiguously preferred by job-seekers, and will therefore have higher applicant arrival rates and worker acceptance probabilities. Only one of these four variables had any effect on the matching probability in Andrews et al. (2001)—an elasticity of –0.20 for non-manual—and so we are estimating large applicant arrival effects in all four cases. In contrast, the corresponding estimates for training vacancies (Table 9) are insignificant, presumably because training vacancies are much more homogenous.

A number of measurable characteristics refer to the selection criteria associated with a vacancy. These include the required level of educational qualification, required subjects studied at school, age and whether or not a written application is required.[23] We argue that some of these characteristics, such as qualifications, are directly related to the employer's reservation utility r^e, and are an attempt to limit the pool of potential applicants. We would therefore expect that higher criteria imply higher r^e and longer search durations. On the other hand, there may be fewer of the better qualified applicants available to such vacancies. A consistent finding across several studies is that higher educational requirements increase the duration of a vacancy (van Ours 1991).

Tables 8 and 9 show that both job and training vacancies requiring higher educational qualifications have significantly lower hazards to filling and therefore have longer search durations. Vacancies requiring 4 or more GCSEs take about 20 days longer to fill (based on the same calculation in Footnote 21). Estimates of the effect on the matching probability were not reported in Andrews et al. (2001) and so we cannot disentangle the arrival and selection effects.

The largest estimate is from the requirement of a written application. Increases in duration might be because written applications increase the application period, or because written applications increase the selection period, as in van Ours & Ridder (1993). For job vacancies, we estimate of large elasticity of –0.94, implies such vacancies take approximately one month longer to fill. In fact, most of this effect is a much lower matching probability, an elasticity of –0.72. Finally, applications requiring older job seekers have lower hazards, with an elasticity of –0.29 for job vacancies and –0.68 for training vacancies.

(Table 1).

[22]Day release involves employees studying at College for a set time each week (typically one day), while apprenticeships offer a longer-term training programme.

[23]For vacancies where no written application was required, the Careers Service would undertake the application procedure.

Efficient or costly search and matching

Generally, a number of covariates have influenced the filling hazard through applicant arrival rates. Here we see if any other variables influence the arrival rate, as well as examine covariates that proxy the cost of search for employers and job seekers. The role of the Careers service on the hazard is also considered.

In addition to labour market tightness variables above, we also consider two other characteristics of the local labour market in which the vacancy is open. In Andrews et al. (2001) we used log population density (log population minus log area) to see whether the matching probability is higher in cities compared with rural areas, and found no effect. Clearly we might also expect population density to influence the applicant arrival rate, but it is only for training vacancies that a positive effect of population is observed: districts with larger populations have higher hazards to filling ($\hat{\beta}_1 = 0.41$).

Further evidence on the effects of applicant arrival rates are provided by a measure of firm location. Firms located in town centres and which are therefore more accessible to job-seekers (lower search costs), have significantly higher hazards to filling, particularly for filling a job vacancy, where the differential is some 0.30 log-points. This variable had no effect on the matching probability in Andrews et al. (2001).

The second local labour variable considered in Andrews et al. (2001) is the number of staff in a given Careers Office, normalised on the population of each district. Although it had a negative effect on the matching probability, it has a positive effect on the hazard to filling a training vacancy in Table 9, suggesting that more staff can generate more applicants per vacancy, as would be expected.

The final variable we consider is firm size, where it is clear that the bigger the firm, the easier it is to fill a job vacancy (there is a clear gradient over size bands 1–10, 11-30, and 31+). As this variable has no effect on the matching probability, again this is an applicant arrival effect. Like larger Careers Offices, larger firms can process more applicants (perhaps because the costs of search per vacancy are lower). For training vacancies, the results differ slightly. The firm-size effect is only between very small firms and the rest, with a differential of about 10 percentage points towards large firms. This represents an arrival rate effect of about 20 percentage points, as the matching probability is about 10 percentage points higher for small firms (as is predicted by the theory).

6.3.3 Lapsing and skill shortages

An important feature of our data is that a substantial minority of vacancies are removed from the market before they are filled. In Figure 11 we plot the probability of filling a vacancy P_{1j}, calculated from Equations (18, 19). The declining hazard for filling and the increasing hazard for lapsing (see Figures 4 to 7) imply that the probability that a vacancy fills declines with duration. For job vacancies, the probability of filling falls below 0.5 after only about one month. For training vacancies, lapsing becomes the most likely outcome after about four months. As we suggested in Section 5.3, lapsing may occur if the arrival rate of applicants or the matching probability falls to zero. Once the Careers Service have searched the pool of potential applicants, and found that none are suitable, a vacancy becomes more likely to lapse. For job vacancies, this process takes only a few weeks. It is often thought that training vacancies are regarded as inferior to job vacancies by potential applicants (perhaps because they are temporary), who will contact job vacancies first. This is one reason why training vacancies take longer to fill.

In the competing risks framework we are also able to determine which characteristics of a vacancy increase the probability of lapsing. The hypothesis that skill shortages cause vacancies to lapse suggests that those vacancies requiring more skilled applicants will take longer to fill and be more at risk of lapsing. An alternative hypothesis is that lapsing is a result of low-quality jobs being refused by potential applicants, in which case it will be low-quality vacancies which are more at risk of lapsing.

Estimates of $\partial P_{1j}/\partial x$ are reported in Tables 8 and 9. Because we are in a competing risks framework, the probability of a vacancy lapsing is a function of both the hazard to filling and the hazard to lapsing (Equation 18). For example, one possible response of the Careers Service to skill shortages might be to keep high-skill vacancies open for longer. In this case we would observe high-skill vacancies having lower hazards to filling and to lapsing, and the net effect on P_{1j} would be ambiguous.

The results seem consistent with the hypothesis that "good" vacancies requiring more skilled applicants are more difficult to fill. Job and training vacancies requiring higher educational qualifications are more likely to lapse ($\partial P_{1j}/\partial x < 0$). The net effect on $\partial P_{1j}/\partial x$ is negative because the hazard to filling falls proportionately more than the hazard to lapsing. Other measures of vacancy quality also tend to have $\partial P_{1j}/\partial x < 0$, and some with quite large semi-elasticities for job vacancies: non-manual vacancies (-0.21), vacancies offering apprenticeship training (-0.09) and vacancies requiring a written application (-0.10) and older applicants (-0.12) all have lower probabilities of filling.[24] The effects

[24] P_{1j} is measured in percentage points, whereas h_{1j} and h_{2j} enter the regressions in logged form.

for training vacancies tend to be weaker, except vacancies which offer higher wages and require older applicants, and are more likely to lapse than the base group. This is to be expected as fewer training vacancies lapse compared with job vacancies.

Since job-seeker selection effects are almost certainly positive for these measures of vacancy quality (higher μ^w), this result provides strong evidence that employer search is unsuccessful because the supply of suitable applicants is insufficient. This may occur because the Careers Service filters out unsuitable applicants (low λ^e), or because employers reject applicants (lower μ^e). Notice that most of these variables have been discussed already in the context of their impact on the hazard to filling (which is only half the story here), but it is generally the case these vacancies have lower applicant arrival rates rather than lower matching probabilities. All in all, we conclude that it is skill shortages which cause increases in employer search duration rather than the unattractiveness of certain vacancies to job-seekers.

7 Conclusions

This paper provides the first analysis of vacancy duration and the outcome of employer search using duration modelling methods for the UK. Previous work in this area has been restricted by a lack of detailed information about the characteristics of vacancies and has ignored the fact that vacancies may be removed from the market before they are filled. In addition, we present the appropriate econometric techniques for dealing with groups of identical vacancies posted simultaneously, and we examine the robustness of our results to different assumptions about unobserved heterogeneity. We are also able to compare results across two quite different markets, jobs and training places.

Our results are interpreted in the framework of recent two-sided search models. These models have the feature that the impact of variables which operate on both sides of the market are expected to have opposite effects on employers and job-seekers. For example, an increase in the wage offer causes employers to become more selective but job-seekers to become less selective. We therefore start with the assumption that job-seekers rarely refuse job offers—see Andrews et al. (2001) for evidence that this is true for these data—which means that we are estimating the response of employers. These models also demonstrate that the duration of employer search can be decomposed into applicant arrival rate effects and selection effects, again using Andrews et al. (2001).

Our key results are as follows:

1. A surprisingly large proportion (one-third) of job vacancies are removed from the

market before they are filled ('lapsed'). A smaller proportion of training vacancies lapse. This is evidence either that employers find it difficult to find suitable applicants, or that this search channel is inefficient. Our results suggests the former, because these vacancies are *not* subsequently filled by other means, and many re-appear at a later date.

2. Despite some institutional 'seasonality', simple monotonic parametric (Weibull) specifications appear adequate and greatly improve the feasibility of estimation. We find no evidence that estimating flexible non-parametric baseline hazards has no effect on parameter estimates.

3. For job vacancies, evidence of duration dependence is significantly affected by the inclusion of unobserved heterogeneity. As predicted, mixture models have less sharply declining hazards to filling. For training vacancies there is no evidence of unobserved heterogeneity.

4. Despite this, there is some evidence of a sharp decline in the baseline hazard to filling after just one day, suggesting that potential stock of applicants can be searched very quickly in this market, supporting the ideas of Coles and Smith.

5. A key variable in all search models is labour market tightness. For job vacancies, an increase in unemployment increases the hazard because the increase in the number of contacts per vacancy outweighs the employers' 'more selective' response. The duration of employer search in the market for training vacancies does not respond in the same way to labour market tightness. The effect of the aggregate stock of vacancies is not well determined.

6. The wage does not generally affect the duration of employer search. Andrews et al. (2001) found clear negative effects on the matching probability, which implies that a higher wage generates more applicants, because of an unobserved quality effect for such vacancies.

7. A number of other covariates have a negative influence on the employer's hazard, including the type of vacancy (non-manual and involves training) and selection criteria (qualification, written application, older applicant). Generally, but not always, these are because the arrival rate of applicants is lower.

8. Vacancies are increasingly likely to lapse as time passes. Almost all the evidence presented here suggests that it is *good* rather than *bad* vacancies which are hard to fill, both in the sense that employer search takes longer, and in the sense that they

are more likely to be withdrawn from the market before they are filled. Generally, these are the same vacancies as listed immediately above.

9. Finally, we note that the market for training vacancies is quite distinct. We know from Andrews et al. (2001) that an application to a training vacancy is more likely to result in a hire: employers are generally less selective when filling training vacancies, presumably because they are of limited duration. In general, results are less clear cut for training vacancies, partly because they are more homogenous in nature (most coming in multiple vacancy orders) and because they are less likely to lapse.

References

Abbring, J. & van Ours, J. (1994), 'Sequential or non-sequential employers' search', *Economics Letters* **44**, 323–8.

Andrews, M., Bradley, S. & Upward, R. (2001), 'Estimating the probability of a match using micro-economic data for the youth labour market', *Labour Economics,* forthcoming.

Barron, J., Berger, M. & Black, D. (1997*a*), Employer search, training and vacancy duration. Upjohn Institute for Employment Research, Kalamazoo: Michigan.

Barron, J., Berger, M. & Black, D. (1997*b*), 'Employer search, training and vacancy duration', *Economic Inquiry* **35**, 167–192.

Barron, J., Bishop, J. & Dunkelberg, W. (1985), 'Employer search: the interviewing and hiring of new employees', *Review of Economics and Statistics* **67**, 43–52.

Barron, J., Black, D. & Loewenstein, M. (1987), 'Employer size: the implications for search, training, capital investment, starting wages, and wage growth', *Journal of Labor Economics* **5**1, 76–89.

Beaumont, P. (1978), 'The duration of registered vacancies: an exploratory exercise', *Scottish Journal of Political Economy* **25**, 75–87.

Booth, A. & Snower, D., eds (1996), *Acquiring skills: market failures, their symptoms and policy responses*, Cambridge University Press, Cambridge.

Broersma, L. & van Ours, J. (1999), 'Job searchers, job matches and the elasticity of matching', *Labour Economics* **6**, 77–93.

Burdett, K. & Cunningham, E. (1998), 'Toward a theory of vacancies', *Journal of Labor Economics* **16**, 445–78.

Burdett, K. & Wright, R. (1998), 'Two-sided search with nontransferable utility', *Review of Economic Dynamics* **1**, 220–45.

Coles, M. & Smith, E. (1998), 'Marketplaces and matching', *International Economic Review* **39**, 239–55.

Gorter, C. & van Ommeren, J. (1999), 'Sequencing, timing and filling rates of recruitment channels', *Applied Economics* **31**, 1149–60.

Han, A. & Hausman, J. (1990), 'Flexible parametric estimation of duration and competing risk models', *Journal of Applied Econometrics* **5**, 1–28.

Haskel, J. & Martin, C. (1993), 'The causes of skill shortages in Britain', *Oxford Economic Papers* **45**, 573–86.

Haskel, J. & Martin, C. (1996), Skill shortages, productivity growth and wage inflation, *in* A. Booth & D. Snower, eds, 'Acquiring skills: market failures, their symptoms and policy responses', Cambridge University Press, Cambridge, chapter 8.

Heckman, J. & Singer, B. (1984), 'Econometric duration analysis', *Journal of Econometrics* **24**, 63–132.

Holzer, H. (1988), 'Search method use of unemployed youth', *Journal of Labor Economics* **6**, 1–20.

Lancaster, T. (1990), *The Econometric Analysis of Transition Data*, Cambridge University Press, Cambridge.

Machin, S. (1996), Changes in the relative demand for skills, *in* Booth & Snower (1996), chapter 7.

Manning, A. (2000), 'Pretty vacant: recruitment in low-wage labour markets', *Oxford Bulletin of Economics and Statistics* **62 (1)**, 747–93.

Meyer, D. (1990), 'Unemployment insurance and unemployment spells', *Econometrica* **58**, 757–82.

Mortensen, D. (1986), Job search and labor market analysis, *in* O. Ashenfelter & R. Layard, eds, 'Handbook of Labor Economics', Vol. 2, North-Holland, Amsterdam, chapter 15, pp. 849–919.

Mortensen, D. & Pissarides, C. (1998), Job reallocation, employment fluctuations and unemployment differences, *in* M. Woodford & J. Taylor, eds, 'Handbook of Macroeconomics', North-Holland, Amsterdam, chapter 18, pp. xxxx–xxx.

Mortensen, D. & Pissarides, C. (1999), New developments in models of search in the labor market, *in* O. Ashenfelter & D. Card, eds, 'Handbook of Labor Economics', Vol. 3B, North-Holland, Amsterdam, chapter 39, pp. 2567–627.

Pissarides, C. (1990), *Equilibrium Unemployment Theory*, Basil Blackwell, Oxford.

Prentice, R. & Gloeckler, L. (1978), 'Regression analysis of grouped survival data with application to breast cancer data', *Biometrics* **34**, 57–67.

Roper, S. (1988), 'Recruitment methods and vacancy duration', *Scottish Journal of Political Economy* **35**, 51–64.

Russo, G., Rietveld, P., Nijkamp, P. & Gorter, C. (1996), Labour market conditions and recruitment behaviour of Dutch firms: spatial and temporal aspects. 8th EALE Annual Conference, Papers.

Stewart, M. (1996), Heterogeneity specification in unemployment duration models, mimeo, University of Warwick, September.

Thomas, J. (1996), 'On the interpretation of covariate estimates in independent competing-risks models', *Bulletin of Economic Research* **48**1, 27–39.

Upward, R. (1998), Estimating the impact of the Youth Training Scheme on individual labour market outcomes, PhD thesis, University of Manchester.

van den Berg, G. (1990), 'Search behaviour, transitions to non-participation and the duration of unemployment', *Economic Journal* **100**, 842–865.

van den Berg, G. (1999), 'Empirical inference with equilibrium search models of the labour market', *Economic Journal* **109**, F283–306.

van den Berg, G. (2000), Duration models: specification, identification, and multiple durations, *in* J. Heckman & E. Leamer, eds, 'Handbook of Econometrics', Vol. 5, North-Holland, Amsterdam.

van Ours, J. (1989), 'Durations of Dutch job vacancies', *De Economist* **137**3, 309–327.

van Ours, J. (1990), An empirical analysis of employers' search, *in* J. Hartog, G. Ridder & J. Theeuwes, eds, 'Panel Data and Labor Market Studies', North-Holland, Amsterdam, pp. 191–214.

van Ours, J. (1991), 'The efficiency of the Dutch labour market in matching unemployment and vacancies', *De Economist* **139**3, 359–78.

van Ours, J. & Ridder, G. (1991), 'Cyclical variations in vacancy durations and vacancy flows: an empirical analysis', *European Economic Review* **35**, 1143–1155.

van Ours, J. & Ridder, G. (1992), 'Vacancies and the recruitment of new employees', *Journal of Labor Economics* **10**, 138–155.

van Ours, J. & Ridder, G. (1993), 'Vacancy durations: search or selection?', *Oxford Bulletin of Economics and Statistics* **55**, 187–198.

Walsh, R. (1997), 'Skill needs in Britain: 1996 report', *Labour Market Trends* **February**, 73–76.

Figures

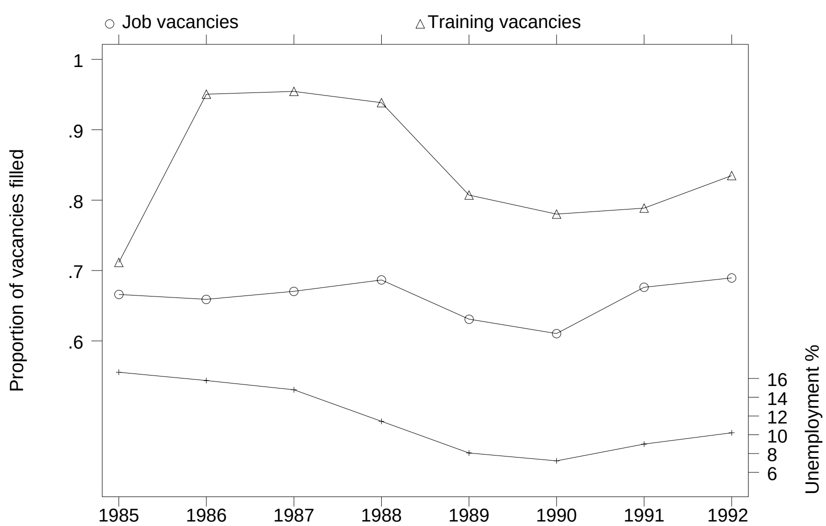

Figure 1: Proportion of vacancies filled and the business cycle

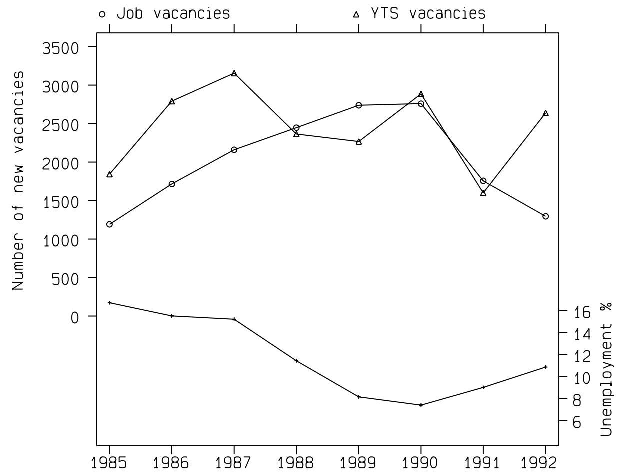

Figure 2: Number of vacancies notified and the business cycle

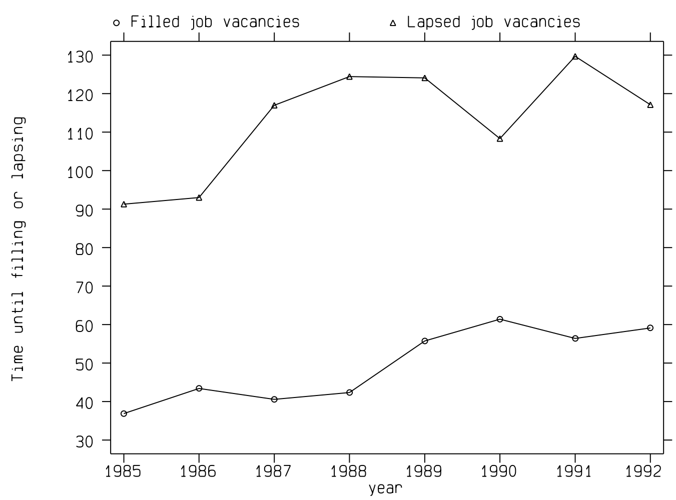

Figure 3: Estimated time until filling and lapsing, job vacancies

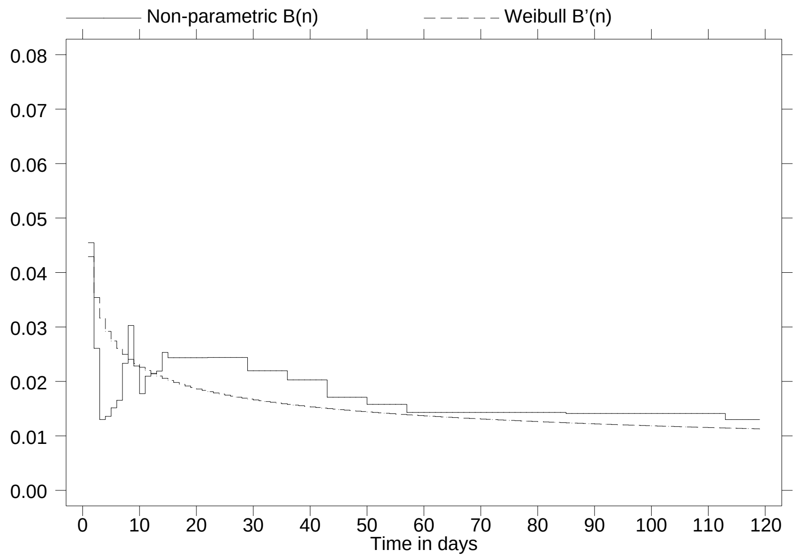

Figure 4: Baseline hazards to filling, job vacancies

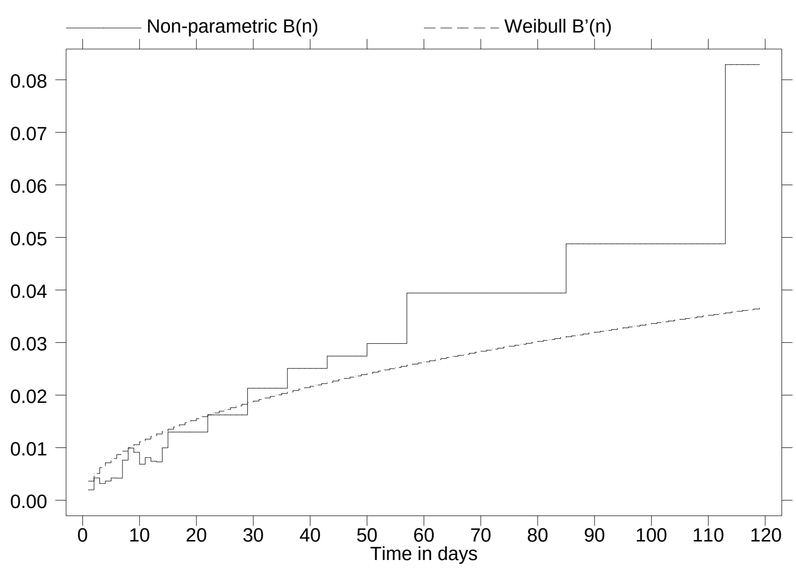

Figure 5: Baseline hazards to lapsing, job vacancies

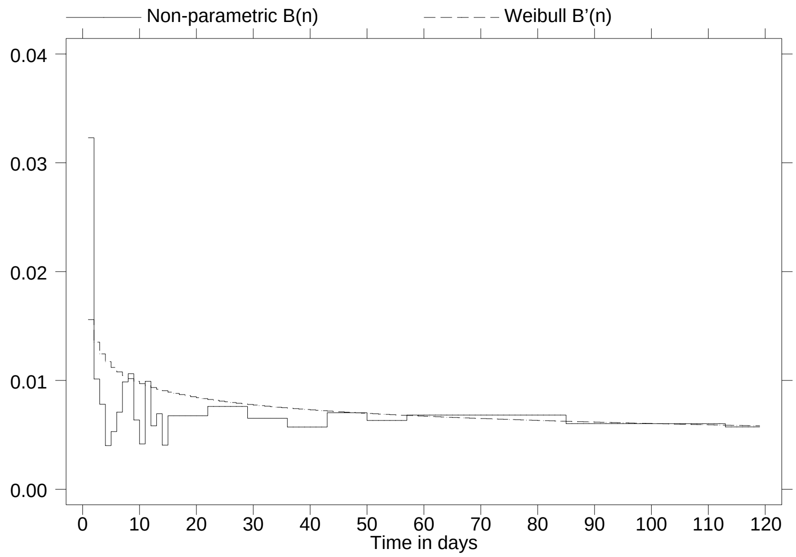

Figure 6: Baseline hazards to filling, training vacancies

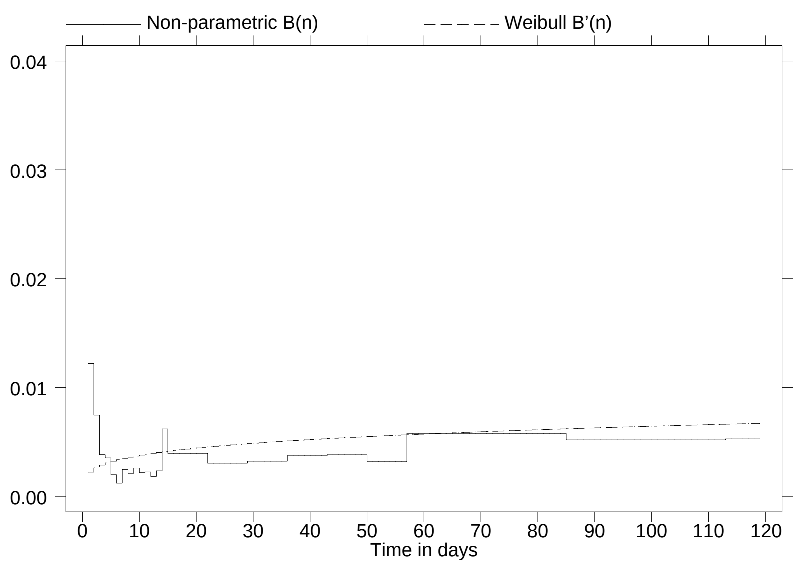

Figure 7: Baseline hazards to lapsing, training vacancies

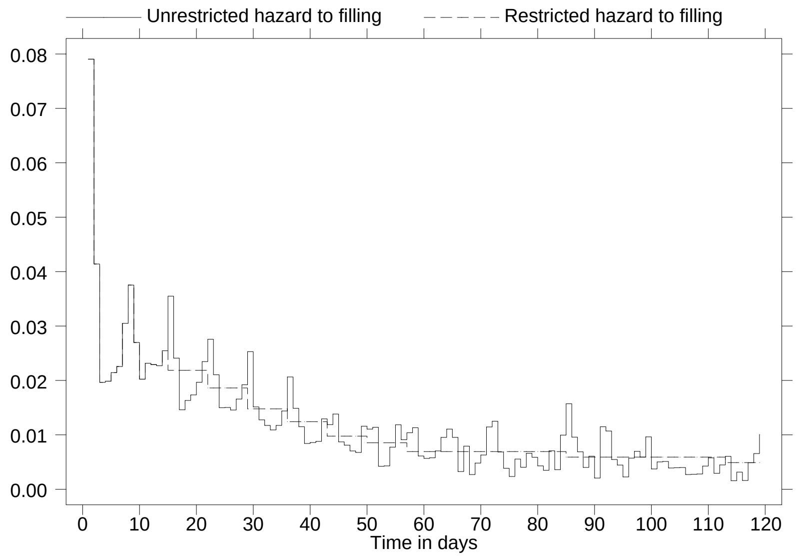

Figure 8: Unrestricted baseline hazard to filling, job vacancies, no covariates

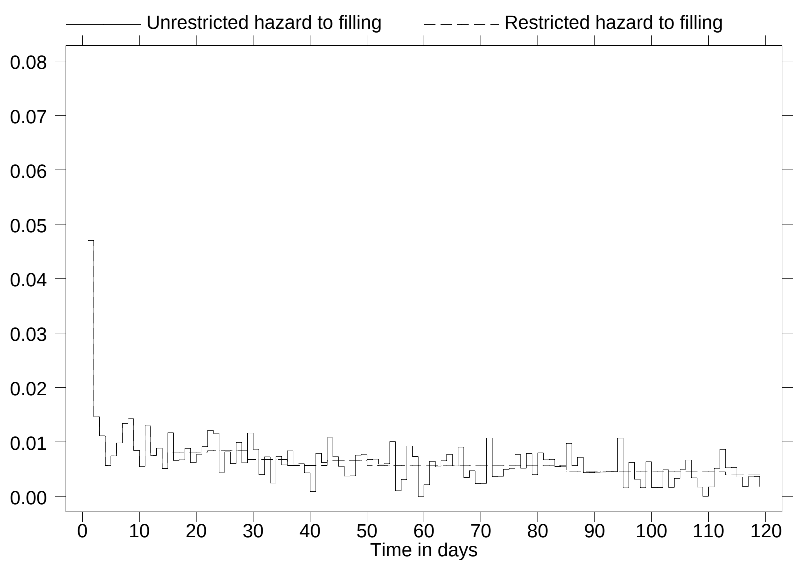

Figure 9: Unrestricted baseline hazard to filling, training vacancies, no covariates

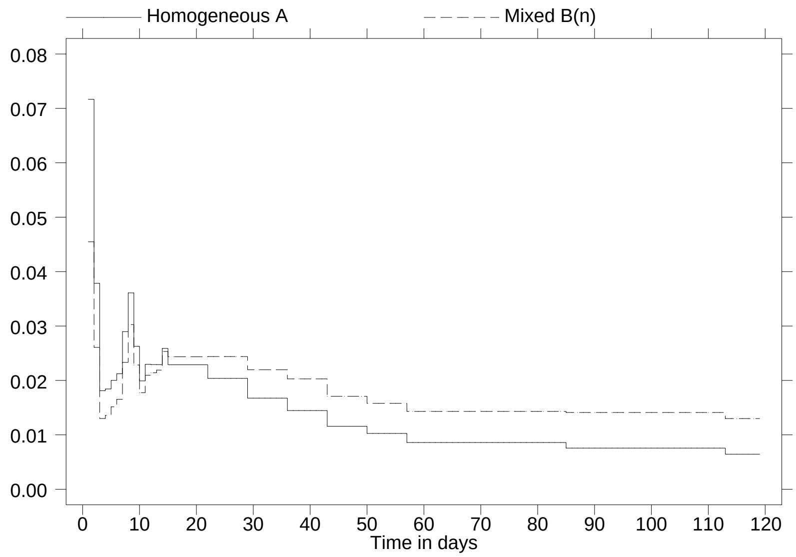

Figure 10: Impact of unobserved heterogeneity on the baseline hazard

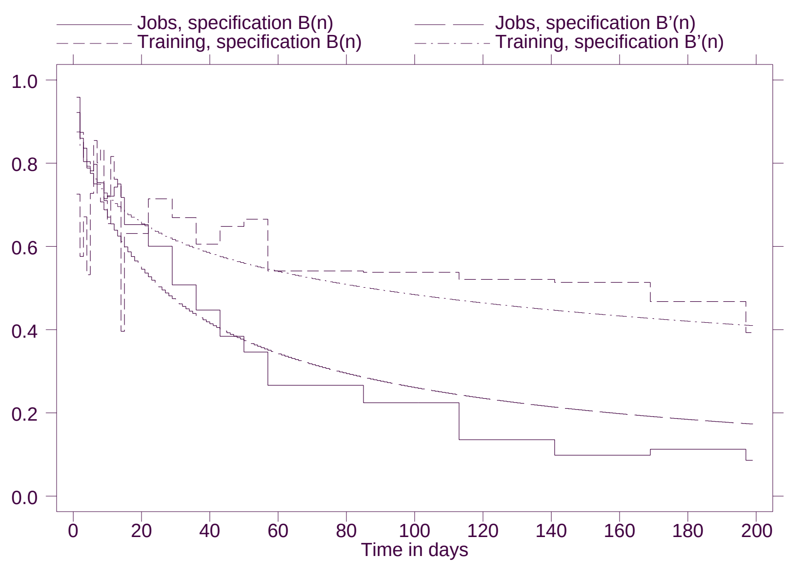

Figure 11: Probabilities of filling, job and training vacancies

Tables

Table 1: Vacancy duration

	Job vacancies			*Training vacancies*		
	Mean duration (days)	*ML estimate of duration*[a] *(days)*	*No.*	*Mean duration (days)*	*ML estimate of duration*[a] *(days)*	*No.*
Single vacancies, $V_i = 1$						
Filled ($W_i = 1$)	21.07	53.97	7234	72.50	151.45	1136
Lapsed ($W_i = 0$)	42.05	71.19	5484	111.86	232.50	740
Censored	60.69		122	68.46		101
Number of single vacancies			12840			1977
All vacancies						
Filled		50.80	11485		112.88	23928
Lapsed		111.82	6054		837.42	3532
Censored			220			2196
Total no. of vacancies ($\sum_{i=1}^{N} V_i$)			17759			29656
All vacancies, by order						
All filled ($W_i = V_i$)			8548			2529
Some filled ($0 < W_i < V_i$)			242			395
All lapsed ($W_i = 0$)			5573			948
Censored ($C_i = 1$)			147			313
Total number of orders (N)			14510			4185

[a] Assuming exponential distribution, i.e. $1/\hat{\gamma}_j$.

Table 2: Mean values of employer characteristics

Characteristic	*Job vacancies*		*Training vacancies*	
	Filled	*Lapsed*	*Filled*	*Lapsed*
Firm size				
$<$ 10 employees	0.390	0.455	0.343	0.342
11–30 employees	0.239	0.205	0.294	0.206
31–100 employees	0.183	0.152	0.150	0.211
$>$ 100 employees	0.188	0.188	0.213	0.240
Firm activity (SIC)				
Agriculture	0.011	0.015	0.003	0.003
Energy and water supplies	0.002	0.004	0.013	0.023
Extraction of minerals, metals	0.010	0.015	0.005	0.005
Metal goods, engineering	0.142	0.113	0.058	0.077
Other manufacturing	0.232	0.207	0.031	0.052
Construction	0.062	0.057	0.017	0.036
Distribution, catering and hotels	0.300	0.325	0.112	0.200
Transport and communication	0.015	0.014	0.024	0.021
Banking, finance	0.102	0.102	0.036	0.038
Other services	0.123	0.148	0.257	0.285
Training agent	0.000	0.000	0.443	0.260
Location				
Firm located in town centre	0.468	0.380	0.397	0.393
Firm located elsewhere	0.532	0.620	0.603	0.607
Involvement in training scheme				
Firm provides training vacancies	0.351	0.378	1.000	1.000
Firm does not provide training vacancies	0.649	0.622	—	
Sample size				
Individual vacancies ($\sum_{i=1}^{N} V_i$)	11485	6054	23928	3532
Censored observations		220		2196

Table 3: Mean values of vacancy characteristics[a]

Characteristic	*Job vacancies*		*Training vacancies*	
	Filled	*Lapsed*	*Filled*	*Lapsed*
Wage offer[b]				
Wage announced and non-negotiable ($D = 1$, $N = 0$)	0.801	0.759	0.990	0.969
Wage announced and negotiable ($D = 1$, $N = 1$)	0.037	0.052	0.010	0.031
Wage not announced and negotiable ($D = 0$, $N = 1$)	0.162	0.189	—	
Hourly wage rate (1987 prices) if $D = 1$, $N = 0$	£1.39	£1.43	£0.70	£0.72
Hourly wage rate $D = 1$, $N = 1$	£1.32	£1.38	£0.67	£0.71
Hourly wage rate $D = 0$, $N = 1$	—		—	
Occupational type				
Unskilled	0.542	0.452	0.430	0.437
Skilled	0.458	0.548	0.570	0.563
Manual	0.552	0.422	0.508	0.485
Non-manual	0.448	0.578	0.492	0.515
Training				
No or little training provided	0.720	0.695	—	
'In-house'	0.070	0.058	—	
Day release	0.081	0.102	0.673	0.694
Apprenticeship	0.129	0.144	0.327	0.306
Sample size				
Individual vacancies ($\sum_{i=1}^{N} V_i$)	11485	6054	23928	3532
Censored observations		220		2196

[a] Vacancy characteristics also include dummies for the month and year in which the vacancy opened (not shown).

[b] Unless specified otherwise, training vacancies are assumed to offer the standard training allowance.

Table 4: Mean values of selection criteria

Characteristic	*Job vacancies*		*Training vacancies*	
	Filled	*Lapsed*	*Filled*	*Lapsed*
Qualifications required				
Non-exam or low ability	0.328	0.243	0.359	0.334
Average GCSE or just below	0.427	0.451	0.491	0.462
High GCSE	0.164	0.201	0.129	0.165
4 or more 'O' levels	0.082	0.104	0.022	0.040
Subjects required				
No subjects required	0.651	0.635	0.644	0.668
English required	0.051	0.042	0.029	0.027
Maths required	0.040	0.037	0.017	0.013
English and Maths	0.176	0.183	0.263	0.236
Science plus English or Maths	0.048	0.054	0.045	0.050
Other subjects	0.034	0.049	0.001	0.005
Age required				
16 year-old applicants accepted	0.857	0.784	1.000	0.998
Older applicants (over 16)	0.143	0.216	0.000	0.002
Application method required				
Careers Service makes application	0.839	0.760	0.607	0.583
Written application required	0.161	0.240	0.393	0.417
Sample size				
Individual vacancies ($\sum_{i=1}^{N} V_i$)	11485	6054	23928	3532
Censored observations		220		2196

Table 5: Local labour market characteristics

Characteristic	*Job vacancies*		*Training vacancies*	
	Filled	*Lapsed*	*Filled*	*Lapsed*
Unemployed ≤ 18	343.057	310.673	354.109	254.321
Job vacancies	60.790	57.122	57.931	67.039
Job vacancies in same occupation	14.982	13.405	12.587	15.389
Training vacancies	580.207	509.117	503.603	503.647
Training vacancies in same occupation	112.315	96.749	115.363	104.774
Population (000's)	110.413	108.474	106.847	105.511
Area (000's hectares)	20.822	19.226	17.896	19.937
Careers Service staff	11.609	11.350	11.530	10.454
Sample size				
Individual vacancies ($\sum_{i=1}^{N} V_i$)	11485	6054	23928	3532
Censored observations		220		2196

Table 6: Definitions of hazard specifications

	Homogeneous	*Mixed*		
		Gamma	Normal	Non-parametric
Single vacancies				
Non-parametric baseline	A	$B(g)^{\mathrm{a}}$	$B(n)$	$B(h)^{\mathrm{a}}$
Parametric (Weibull) baseline	A'	$B'(g)^{\mathrm{c}}$	$B'(n)$	$B'(h)^{\mathrm{c}}$
All vacancies including "multiple orders"				
Non-parametric baseline	C^{c}	$D(g)^{\mathrm{a}}$	—[b]	—[b]
Parametric (Weibull) baseline	C'	$D'(g)^{\mathrm{c}}$	—[b]	—[b]

[a] Data cannot support non-parametric baseline.

[b] Not estimated because data cannot be organised into sequential binary form.

[c] Data cannot support parametric baseline for training vacancies.

Table 7: Comparison of hazard specifications

		Job vacancies				*Training vacancies*			
		Filled		*Lapsed*		*Filled*		*Lapsed*	
Non-parametric									
A	Log L	−33463.545		−28767.670		−6474.495		−5009.060	
$B(n)$	Log L	−33421.291		−28702.285		−6472.912		−5009.058	
	σ_u^2	0.908	[0.632]	2.081	[0.000]	0.264	[0.058]		
C	Log L^{a}	−56921.763				—			
Parametric (Weibull)									
A'	Log L	−33727.522		−28865.275		−6540.280		−5102.170	
	α	0.603	[0.000]	1.120	[0.000]	0.793	[0.000]	1.231	[0.000]
$B'(g)$	Log L	−33687.397		−28791.212		—		—	
	σ_v^2	0.491	[0.000]	0.493	[0.000]				
	α	0.765	[0.000]	1.373	[0.000]				
$B'(n)$	Log L	−33703.020		−28783.112		−6540.280		−5102.170	
	σ_u^2	0.375	[0.000]	0.941	[0.681]	0.000		0.000	
	α	0.717	[0.000]	1.486	[0.000]	0.793	[0.000]	1.231	[0.000]
$B'(h)$	Log L	−33644.351		−28773.651		—		—	
	$\bar{u}_1$, π_1	−1.922	0.189	−0.933	0.434				
	$\bar{u}_2$, π_2	0.448	0.811	0.714	0.567				
	σ_u^2	0.861		0.666					
	α	0.758	[0.000]	1.390	[0.000]				
C'	Log L^{a}	−77891.805				−75404.464			
	α	0.635	[0.000]	1.042	[0.000]	0.363	[0.000]	1.979	[0.000]
$D'(g)$	Log L^{a}	−77720.668				—			
	σ_v^2	1.624	[0.000]	1.613	[0.000]				
	α	0.779	[0.000]	1.206	[0.000]				

[a] Likelihood defined jointly for filling and lapsing.

Table 8: Specification $D'(g)$, job vacancies*

	Filled		*Lapsed*				
	$\hat{\boldsymbol{\beta}}_1$	p-value	$\hat{\boldsymbol{\beta}}_2$	p-value	$\partial P_{1j}/\partial x^{\mathrm{a}}$	p-value	$\Delta E_1/\Delta x^{\mathrm{b}}$
11–30 employees	0.1296	[0.000]	−0.1769	[0.000]	0.066	[0.000]	−0.2
31–100 employees	0.2040	[0.000]	−0.4404	[0.000]	0.138	[0.000]	0.6
$>$ 100 employees	0.2099	[0.000]	−0.5669	[0.000]	0.167	[0.000]	1.3
Firm located in town centre	0.2963	[0.000]	−0.0733	[0.043]	0.079	[0.000]	−3.4
Firm provides training vacancies	−0.1239	[0.000]	−0.0177	[0.608]	−0.023	[0.017]	1.8
Wage announced and negotiable (DN)	−0.2164	[0.006]	−0.0452	[0.626]	−0.037	[0.160]	3.3
Wage not announced and negotiable ($(1-D)N$)	−0.0059	[0.873]	0.0259	[0.574]	−0.007	[0.591]	−0.1
Log wage if $D(1-N)$	0.0356	[0.531]	−0.0811	[0.229]	0.025	[0.186]	0.0
Log wage if DN	−0.0409	[0.840]	0.2898	[0.185]	−0.071	[0.268]	−0.2
Skilled	0.0607	[0.215]	0.1189	[0.037]	−0.013	[0.438]	−1.7
Non-manual	−0.4890	[0.000]	0.4841	[0.000]	−0.209	[0.000]	2.9
In house training	0.0820	[0.103]	−0.1059	[0.102]	0.040	[0.022]	−0.4
Day release training	−0.3480	[0.000]	−0.4928	[0.000]	0.031	[0.042]	9.6
Apprenticeship training	−0.5148	[0.000]	−0.1026	[0.170]	−0.089	[0.000]	7.8
Average GCSE or just below	−0.1513	[0.000]	0.0798	[0.063]	−0.050	[0.000]	1.4
High GCSE	−0.3608	[0.000]	−0.0503	[0.410]	−0.067	[0.000]	5.3
4 or more GCSEs	−0.4359	[0.000]	−0.3226	[0.000]	−0.024	[0.266]	9.3
English required	0.0028	[0.967]	−0.2297	[0.003]	0.050	[0.023]	1.7
Maths required	0.0533	[0.538]	−0.6465	[0.000]	0.150	[0.000]	4.0
English and Maths required	0.0068	[0.939]	−0.3623	[0.001]	0.079	[0.008]	2.7
Science required	−0.0534	[0.446]	−0.2809	[0.001]	0.049	[0.035]	2.9
Other subject required	0.0766	[0.386]	−0.3774	[0.000]	0.098	[0.001]	1.8
Older applicants required (over 16)	−0.2930	[0.000]	0.2825	[0.000]	−0.124	[0.000]	1.2
Written application required	−0.9437	[0.000]	−0.4863	[0.000]	−0.098	[0.000]	20.1
Unemployed ≤ 18	0.2933	[0.000]	−0.1260	[0.022]	0.090	[0.000]	−0.3
Job vacancies	−0.0664	[0.001]	−0.0107	[0.655]	−0.012	[0.073]	0.1
Job vacancies in same occ.	0.0732	[0.000]	−0.1159	[0.000]	0.041	[0.000]	0.0
Training vacancies	0.0226	[0.152]	−0.0858	[0.000]	0.023	[0.000]	0.0
Training vacancies in same occ.	0.0092	[0.343]	0.0104	[0.396]	0.000	[0.937]	0.0
Population (000's)	−0.0421	[0.719]	0.4624	[0.001]	−0.108	[0.006]	−0.3
Area (000's hectares)	0.0398	[0.051]	−0.2433	[0.000]	0.061	[0.000]	0.1
Careers Service staff	−0.0070	[0.774]	−0.1500	[0.000]	0.031	[0.000]	0.1

*Also includes dummies for SIC (9), year (7) and month (11).
[a]Marginal effect on the probability of filling, evaluated at predicted duration. See Equation (20).
[b]Change in predicted duration conditional on exit to risk r, in days.

Table 9: Specification C', training vacancies*

	Filled		*Lapsed*				
	$\hat{\boldsymbol{\beta}}_1$	p-value	$\hat{\boldsymbol{\beta}}_2$	p-value	$\partial P_{1j}/\partial x^{\mathrm{a}}$	p-value	$\Delta E_1/\Delta x^{\mathrm{b}}$
11–30 employees	0.1322	[0.000]	−0.1020	[0.005]	0.019	[0.000]	−2.9
31–100 employees	0.0310	[0.248]	0.0217	[0.599]	0.001	[0.850]	−1.0
> 100 employees	0.0978	[0.000]	0.1753	[0.000]	−0.006	[0.123]	−4.1
Firm located in town centre	0.0675	[0.001]	−0.4077	[0.000]	0.038	[0.000]	1.3
Wage announced and negotiable (DN)	−0.6005	[0.000]	0.5140	[0.002]	−0.089	[0.000]	−0.2
Log Wage if $D(1-N)$	0.0815	[0.062]	1.2951	[0.000]	−0.097	[0.000]	−1.2
Log Wage if DN	−0.9216	[0.008]	1.5721	[0.000]	−0.198	[0.000]	−1.1
Skilled	0.0065	[0.767]	0.1328	[0.000]	−0.01	[0.002]	−1.2
Non-manual	0.0107	[0.652]	−0.1421	[0.000]	0.012	[0.000]	0.8
Apprenticeship training	0.0732	[0.007]	0.0534	[0.178]	0.002	[0.680]	−2.4
Average GCSE or just below	−0.1141	[0.000]	0.0704	[0.029]	−0.015	[0.000]	2.8
High GCSE	−0.2605	[0.000]	0.2221	[0.000]	−0.038	[0.000]	5.3
4 or more GCSEs	−0.4893	[0.000]	0.2342	[0.006]	−0.058	[0.000]	10.5
English required	0.0842	[0.639]	−0.9861	[0.000]	0.085	[0.000]	8.7
Maths required	0.1931	[0.298]	−1.0896	[0.000]	0.102	[0.000]	6.5
English and Maths required	0.3397	[0.078]	−0.7736	[0.004]	0.089	[0.001]	0.0
Science required	0.1341	[0.455]	−1.1494	[0.000]	0.102	[0.000]	8.6
Other subject required	0.2207	[0.228]	−1.0194	[0.000]	0.099	[0.000]	5.2
Older applicants required (over 16)	−0.6808	[0.077]	0.8010	[0.028]	−0.118	[0.005]	4.1
Written application required	−0.0826	[0.000]	−0.1100	[0.000]	0.002	[0.443]	3.1
Unemployed ≤ 18	−0.0251	[0.454]	0.0602	[0.255]	−0.007	[0.173]	0.0
Job vacancies	−0.0569	[0.000]	−0.0406	[0.065]	−0.001	[0.528]	0.2
Job vacancies in same occ.	0.0779	[0.000]	0.0563	[0.000]	0.002	[0.200]	−0.3
Training vacancies	−0.0729	[0.000]	0.1198	[0.000]	−0.015	[0.000]	0.1
Training vacancies in same occ.	−0.0023	[0.732]	−0.0295	[0.004]	0.002	[0.025]	0.0
Population (000's)	0.4121	[0.000]	−0.7246	[0.000]	0.090	[0.000]	−0.6
Area (000's hectares)	−0.0184	[0.210]	0.1701	[0.000]	−0.015	[0.000]	−0.1
Careers Service staff	0.2021	[0.000]	−0.2619	[0.000]	0.037	[0.000]	−0.3

*See footnotes to Table 8.

Table 10: Differences in wage estimates across specifications

		Job vacancies				Training vacancies			
		Filled		Lapsed		Filled		Lapsed	
		$\hat{\boldsymbol{\beta}}_1$	p-value	$\hat{\boldsymbol{\beta}}_2$	p-value	$\hat{\boldsymbol{\beta}}_1$	p-value	$\hat{\boldsymbol{\beta}}_2$	p-value
DN	A	−0.1948	[0.009]	−0.0992	[0.221]	0.3065	[0.110]	0.2880	[0.189]
	$B(n)$	−0.2752	[0.006]	−0.2042	[0.133]	0.3391	[0.111]	0.2880	[0.190]
	C	−0.1763	[0.006]	−0.0379	[0.621]				
	A'	−0.1912	[0.010]	−0.0977	[0.228]	0.2977	[0.119]	0.2993	[0.169]
	$B'(g)$	−0.2433	[0.007]	−0.1544	[0.131]				
	$B'(n)$	−0.2298	[0.007]	−0.1642	[0.138]	0.2977	[0.119]	0.2993	[0.169]
	$B'(h)$	−0.2158	[0.013]	−0.1181	[0.252]				
	C'	−0.1596	[0.013]	−0.0040	[0.958]	−0.6005	[0.000]	0.5140	[0.002]
	$D'(g)$	−0.2164	[0.006]	−0.0452	[0.626]				
$(1-D)N$	A	−0.0251	[0.459]	−0.0007	[0.985]				
	$B(n)$	0.0375	[0.416]	−0.0046	[0.943]				
	C	−0.0498	[0.091]	0.0291	[0.446]				
	A'	−0.0237	[0.484]	−0.0051	[0.896]				
	$B'(g)$	0.0389	[0.360]	0.0023	[0.962]				
	$B'(n)$	0.0128	[0.749]	0.0002	[0.996]				
	$B'(h)$	0.0583	[0.157]	0.0020	[0.968]				
	C'	−0.0481	[0.101]	0.0178	[0.640]				
	$D'(g)$	−0.0059	[0.873]	0.0259	[0.574]				
$wD(1-N)$	A	−0.0185	[0.735]	0.1593	[0.006]	0.9319	[0.000]	0.5034	[0.000]
	$B(n)$	0.0176	[0.811]	0.2656	[0.008]	0.9726	[0.000]	0.5034	[0.000]
	C	−0.0184	[0.676]	−0.0640	[0.249]				
	A'	−0.0155	[0.775]	0.1612	[0.006]	0.8683	[0.000]	0.4299	[0.001]
	$B'(g)$	0.0204	[0.762]	0.2213	[0.003]				
	$B'(n)$	0.0051	[0.936]	0.2264	[0.005]	0.8683	[0.000]	0.4299	[0.001]
	$B'(h)$	0.0383	[0.559]	0.2030	[0.007]				
	C'	−0.0099	[0.822]	−0.0846	[0.126]	0.0815	[0.062]	1.2951	[0.000]
	$D'(g)$	0.0356	[0.531]	−0.0811	[0.229]				
wDN	A	0.0242	[0.897]	0.2217	[0.232]	0.4922	[0.239]	0.3810	[0.393]
	$B(n)$	0.0383	[0.877]	0.3980	[0.210]	0.4638	[0.312]	0.3810	[0.393]
	C	−0.0435	[0.794]	0.2319	[0.192]				
	A'	0.0235	[0.900]	0.2170	[0.243]	0.4346	[0.297]	0.3478	[0.434]
	$B'(g)$	0.0368	[0.871]	0.3197	[0.175]				
	$B'(n)$	0.0340	[0.874]	0.3364	[0.190]	0.4346	[0.297]	0.3478	[0.434]
	$B'(h)$	0.0187	[0.933]	0.2746	[0.267]				
	C'	−0.0473	[0.776]	0.2041	[0.248]	−0.9216	[0.008]	1.5721	[0.000]
	$D'(g)$	−0.0409	[0.840]	0.2898	[0.185]				