SUCCESSFUL EMPLOYER SEARCH? AN EMPIRICAL ANALYSIS OF VACANCY DURATION USING MICRO DATA

M J Andrews
University of Manchester

S Bradley
Lancaster University

D Stott
Lancaster University

R Upward
University of Nottingham

December 2006

Abstract

This paper provides the first estimates of the determinants of the duration of employer search in the UK. We model duration until either a vacancy is successfully filled or withdrawn from the market. The econometric techniques deal with multiple vacancies and unobserved heterogeneity (dependent risks), using flexible and parametric baseline hazards. The hazards to filling and withdrawing exhibit negative and positive duration dependence respectively, implying that the conditional probability of successful employer search decreases with duration. We also find that non-manual vacancies are less likely to fill, consistent with there being skill shortages in the sample period.

Keywords: vacancy duration, lapsed vacancies, competing risks

New JEL Classification: C41, J41, J63, J64

Address for Correspondence:
Dr. R. Upward
School of Economics
University of Nottingham
Nottingham,
United Kingdom, NG7 2RD
Email: richard.upward@nottingham.ac.uk
Phone: +44-(0)115-951-4735

Search theory is becoming one of the dominant models used to explain both micro and macro labour-market phenomena, especially the dynamics of unemployment — see Mortensen & Pissarides' recent (1998, 1999) surveys. However, empirical work has concentrated far more on workers' than on employers' search behaviour. This is in spite of evidence which suggests that, in many labour markets, workers rarely refuse job offers — see Barron, Black and Loewenstein (1987), Holzer (1988), van den Berg (1990), Barron, Berger and Black (1997) and Manning (2000). If, in equilibrium, the worker's acceptance probability is close to unity, it follows that employer search is 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 evidence for vacancies. Employer search remains an under-researched area.

One particular issue that has received little attention is the fact that employer search is not always successful, resulting in vacancies which are withdrawn from the market without being filled. Genuinely unfilled vacancies may be a result of skill shortages, with associated macroeconomic implications such as potentially lower productivity growth and higher wage growth (Haskel and Martin, 1996). To understand the determinants of skill shortages, it is necessary to analyse both those vacancies that tend to have longer durations, and those vacancies that are eventually withdrawn from the labour market.

Although there is some evidence on vacancy characteristics which lead to longer search durations (see Section I), there is almost no evidence on vacancies which are withdrawn from the market. This may be because economists do not believe that employers post vacancies and then subsequently withdraw them from the market, perhaps because it suggests that employers are 'irrational' in their search strategies. But in many ways the process of withdrawing a vacancy is analogous to the process by which job-seekers leave the labour market, a process which has been studied in the literature on labour market transitions (van den Berg, 1990; Frijters and van der Klaauw, 2006).

A second reason may be that there is very little information on the eventual fate of vacancies. In fact, in the UK, 'cancelled' vacancies are common. Vacancy data from the Office of National Statistics show that the proportion of all vacancies notified to the public employment service that are subsequently cancelled is between 20% and 30% over the period of our study (1985–2001). Machin (2003) notes that a substantial proportion of these cancelled vacancies are regarded by employers as 'no longer existing'.

Are these cancelled vacancies actually withdrawn from the market, or do they merely represent employers filling their vacancies using other search strategies? We provide some evidence that many cancelled vacancies are indeed withdrawn from the market. We do this using a large sample of vacancies notified to a particular labour market in the UK. We estimate both the determinants of vacancy duration and the probability that an employer is ultimately unsuccessful in filling the vacancy (hereafter we refer to such a vacancy as *lapsed*). We are able to check that vacancies are genuinely withdrawn because we observe all job-seekers in the same market. Another check is to examine whether lapsed vacancies re-appear in the market in the following year.

We use far more detailed vacancy data than has previously been available. The data measure vacancy duration recorded to the nearest day, and provide detailed information on vacancy characteristics. We allow for the simultaneous advertising of groups of identical vacancies. Our econometric methodology allows for unobserved heterogeneity, which might not be independent between the filling and lapsing risks. We also model the underlying hazards non-parametrically and parametrically.

We show that it is non-manual vacancies that employers find hard to fill, which we interpret as evidence of skill shortages. We also shed light on employers' search strategies, given that a substantial number of vacancies fill within the first week. Finally, we provide new evidence that the hazard to filling is downward sloping and that the hazard to lapsing is upward sloping.

The paper is structured as follows. Section I briefly covers the relevant literature, and Section II provides the theoretical framework. Section III describes the data, and Section IV discusses our econometric methods. Results are presented in Section V, and Section VI concludes.

I RECENT LITERATURE

There are few microeconometric investigations of the duration of employer search, or vacancy duration, using firm-level data. This is particularly true for the UK where there are only three studies, none of which use duration modelling techniques (Beaumont, 1978; Roper, 1988; Adams, Greig and McQuaid, 2002). van Ours and Ridder (1991, 1992, 1993) analyse Dutch data using duration techniques. Their findings suggest that the vacancy hazard displays positive duration dependence, even after allowance is made for the effect of unobserved heterogeneity: employers become less choosy as vacancy duration increases. van Ours and Ridder (1993) have suggested that vacancy durations are mainly periods of selection rather than search, thereby casting doubt on the conventional sequential search model. Weber (2000) offers further supporting evidence in favour of the non-sequential search model. Gorter and van Ommeren (1999) and Gorter, Nijkamp and Rietveld (1996) show that vacancies that are advertised exhibit positive duration dependence, whereas those that use informal contacts exhibit negative duration dependence. Finally, Burdett and Cunningham (1998) estimate a non-parametric vacancy hazard which increases rapidly in the first week, and falls slowly thereafter. Burdett and Cunningham argue that the majority of firms in their sample could *not* have used non-sequential search because so many vacancies fill in a very short space of time.

Thus it would appear that the shape of the baseline hazard depends on the search technology adopted by the employer. Those studies which find positive duration dependence tend to be those that analyse advertised vacancies, where employers tend to adopt a search period followed by a selection period. On the other hand, studies whose data comprise informal search methods or data from public exchanges tend to exhibit negative duration dependence, where the applicant arrival tends to be initially high but falls thereafter.

Of the factors that affect vacancy duration, the most important are the relationships with the total stock of vacancies and the total stock of job-seekers in the market. Increasing the vacancy stock increases search duration (the so-called 'congestion effect'), while increasing the stock of job-seekers reduces search duration. Both are predicted by standard models of search. More stringent entry requirements with respect to age, education and work experience increase vacancy duration (Gorter and van Ommeren, 1999; Behrenz, 2002). Similar US evidence is provided by Barron, Bishop and Dunkelberg (1985). Barron *et al.* (1997) and Burdett and Cunningham (1998) also show that vacancy duration is increased where the training period is longer. Most studies find that the wage does not have a significant effect on vacancy duration, an exception being Adams *et al.* (2002) who find a positive and significant effect.

Vacancies with long durations are, almost by definition, 'hard-to-fill'. However, a literature also exists which analyses the determinants of hard-to-fill vacancies as reported in employer surveys. This literature shows that reports of hard-to-fill vacancies are negatively related to hourly wages, union recognition, the share of part-time vacancies in the firm, the amount of training offered, and the local unemployment rate, and positively related to firm size and employment growth (Mason and Stevens, 2003; Haskel and Martin, 2001; Green, Machin and Wilkinson, 1998; Campbell and Baldwin, 1995; Bosworth, 1993). The existence of hard-to-fill vacancies has often been interpreted as an indication of skill shortages. However, this view has been questioned. Green *et al.* (1998) and Haskel and Martin (2001) show that there is only a partial overlap between firms reporting a skill shortage and simultaneously reporting a hard-to-fill vacancy.

Evidence on the existence of lapsed vacancies is as follows. van Ours and Ridder (1992) find that 4% of vacancies are cancelled according to employers ("because the need for the new employee disappears ... [or] because of the changing economic performance or re-organisation of the firm ... [or] these vacancies were hard to fill and lasted too long.") However, the authors ignore them because they are a small percentage. Barron *et al.* (1985) also use employer survey data and find that 28% of employers did not recruit for the position. They could not analyse these 'lapsed' vacancies because of the design of the survey. This is the same data used by Burdett and Cunningham (1998), who also ignore lapsed vacancies. There is therefore some evidence that lapsing is a widespread phenomenon, albeit one which has not been formally analysed.

II A THEORETICAL FRAMEWORK

The canonical model is one of sequential search, and this has been applied to employers as well as job-seekers: see Lipmann and McCall (1976, pp.181–185) for an early description of the basic employer-search model. Burdett and Cunningham (1998) is a more recent example. However, some empirical work has suggested that employers use a non-sequential search strategy whereby a period of search is used to accumulate a number of applicants, at least one of whom is subsequently selected. Therefore, although the arrival rate of applicants is initially high because the vacancy has been advertised or posted with an employment agency, the hazard rate is initially low (possibly zero) during the search period.

A high applicant arrival rate is also consistent with so-called 'stock-flow' theories of matching as proposed by Coles and Smith (1998) and tested by Coles and Petrongolo (2003). Stock-flow matching suggests that, if there is some kind of 'marketplace' in which search frictions are low, a new vacancy to the market will have an initially high applicant arrival rate because potential matches come from the entire stock of job-seekers. Clearly, these two explanations are closely related since an advert or an employment agency serves as a marketplace. However, the non-sequential model suggests that although the applicant arrival rate will be initially high, the hazard should be low because firms wait for a pool of applicants to arrive. In contrast the stock-flow theory allows for firms to accept applicants immediately.

Is the sequential or non-sequential model consistent with the data? Figure 1 plots the raw hazard to filling for the vacancies in our sample (described more fully in Section 6). It is clear that many applicants arrive and are accepted almost immediately: the hazard is actually highest on day 1, although there are also subsequent peaks at weekly intervals.

[FIGURE 1 HERE]

This is consistent with the finding of Burdett and Cunningham (1998), who suggest that the majority of firms in their sample cannot have used non-sequential search. In what follows we therefore use a sequential search model. This does not rule out, however, the possibility that applicants arrive initially very quickly but then at a greatly reduced rate because of stock-flow matching considerations.

Define V_m to be the discounted revenue stream of a firm which successfully fills a vacancy, and V_n to be the revenue stream of the firm which has a vacancy to fill and which is searching for a suitable applicant. A standard expression (e.g. Cahuc and Zylberberg, 2004) for the discounted flow of revenue per period is

$$rV_m = z_m + q(V_n - V_m), \tag{1}$$

where r is the discount rate, z_m is the revenue over and above the wage paid and q is the per-period probability of the match separating. Over any short time period dt a match yields a flow of revenue z_m and an expected cost deriving from the break-up of a match $V_n - V_m$.

The standard result in these models is the adoption of a stopping rule. If an applicant arrives, the optimal strategy is to accept if $V_m > V_n$. If no applicant arrives, the firm continues to search. The reservation productivity level, z^*, is that productivity level which leaves a firm indifferent between accepting and rejecting the applicant i.e. where $V_m(z^*) = V_n$.

The discounted revenue stream for a firm with a vacancy is

$$rV_n = (z_n - c) + \frac{\lambda}{r+q}\int_{z_m}^{\infty}(z_m - z^*)dF(z_m), \tag{2}$$

where z_n is the revenue from producing with a vacancy. In general $z_n > 0$ because the firm can produce with a vacancy, albeit at a lower level of profit. z_n is analogous to income received by a job-seeker which is not dependent on search, such as a means-tested benefit. c is the per-period cost to the employer of keeping the vacancy open. This includes advertising and screening costs. λ is the arrival rate of applicants to the employer. The remaining term is the standard surplus function, decreasing in z^*.

Finally, consider the discounted revenue stream from producing with an unfilled vacancy but *not* searching. The simplest assumption would be that the revenue flow for a firm which chooses not to search, denoted V_l, is just

$$rV_l = z_n. \tag{3}$$

In other words, the firm's revenue is z_n whether or not it chooses to search. If $V_l > V_n$ then the firm chooses not to search. Equivalently, if $z_n > z^*$ then it is more profitable to carry on producing with an unfilled vacancy than it is to search.

In a stationary model we would not observe a firm choosing to search and then changing its behaviour (i.e. lapsing a vacancy). However, in general, we would expect some of those parameters which affect z^* (such as λ and those that characterise $f(z_m)$) to change with elapsed duration. If they cause z^* to fall, at some point the firm may choose to stop searching because $z^* < z_n$. Comparing (2) with (3) we can see that V_n will fall below V_l if the value of the surplus function falls below the cost of search, c.

The most likely explanation for this is that the arrival rate of applicants falls with elapsed duration. Theories of stock-flow matching suggest that this is the case. Employers who advertise a new vacancy initially receive a high rate of applicants because the potential pool of applicants comprises all those job-seekers currently in the market. If none of these initial applicants are acceptable, the applicant arrival rate falls. At this point the costs of keeping the vacancy open may outweigh the discounted benefits of continuing to search.

A second reason why V_n might fall below V_l is that the productivity distribution of applicants falls with elapsed duration. If $f(z_m)$ shifts to the left, the surplus function falls, and the benefit of search reduces. This seems particularly plausible in a market with a 'recruitment cycle' where there is a substantial inflow of potential applicants at particular times of the year. As time passes the better applicants leave the pool, shifting the distribution of remaining applicants to the left. In this case the firm may decide to re-advertise the vacancy at a later date.

The framework described above is essentially that analysed by Frijters and van der Klaauw (2006), who consider a non-stationary job-search model where job-seekers have the option of leaving the labour force. In our case, we have a non-stationary employer-search model where employers can choose to withdraw a vacancy from the market. In Frijters and van der Klaauw, job-seekers have perfect foresight regarding the job-offer arrival rate and the wage offer distribution. They show that, in this case, the job-seeker's optimal policy is a sequence of reservation wages and a maximum duration of search. In our case, employers would have a sequence of (presumably declining) reservation productivities z_t^* and would decide how long to keep a vacancy open. This duration would be determined by the point at which z_t^* falls below z_n.[1]

The probability that an employer will find an applicant acceptable is

$$\mu = 1 - F(z^*) = \mu(c, \lambda, z_m, \sigma_{z_m}) \quad \mu_c > 0,\ \mu_\lambda < 0,\ \mu_z > 0,\ \mu_\sigma \lesseqgtr 0. \tag{4}$$

The comparative statics are standard (e.g. Mortensen, 1986). Employers become less selective (in that μ increases) as search costs increase, as the arrival rate of applicants decreases, as the wage decreases, or as the revenue flow increases. An increase in the variance of $f(z_m)$, denoted σ_{z_m}, has an ambiguous effect on μ in theory, although a commonsense prediction is that employers will wait longer for a 'bargain'. An increase in σ_{z_m} increases the number of bargains and so $\mu_\sigma < 0$.

The hazard for a vacancy which fills, denoted h_1, is the product of the applicant arrival rate and the probability of acceptance. Therefore

$$h_1 = \lambda \cdot \mu(c, \lambda, z_m, \sigma_{z_m}). \tag{5}$$

The arrival rate of applicants therefore has two influences on the vacancy hazard. There is the direct positive effect, but also an indirect negative effect. A fall in λ causes employers to become less selective, lowering z^* so that μ goes up. van den Berg (1994) shows that the net effect is positive for all reasonable wage offer distributions.

If employers have perfect foresight then the hazard to lapsing is strictly not defined where $z_t^* > z_n$: vacancies will never be withdrawn while the benefit of searching is greater than the benefit of not searching. However, one can still estimate a reduced form lapsing hazard from the observed distribution of lapsing times. This distribution will depend on the distribution of z_n and z_t^* amongst employers. Thus we would expect the hazard for a vacancy which is lapsed, denoted h_2, to depend on the same variables as h_1:

$$h_2 = h_2(c, \lambda, z_m, \sigma_{z_m}).$$

Here one should interpret the effect of observed variables on h_2 as affecting the distribution of lapsing times, and not the individual hazard of a vacancy. For example, vacancies in markets with lower applicant arrival rates λ will have lower benefits to search and will tend to lapse more quickly.

The proportion of vacancies which lapse at each duration is given by

$$P = \frac{h_2}{h_1 + h_2}. \tag{6}$$

Those vacancies which are more likely to lapse are those where z_t^* falls below z_n more quickly. For example, if "high skill" vacancies have a very small pool of potential applicants which is quickly exhausted, we would expect these vacancies to be withdrawn from the market more quickly.

III DATA AND INSTITUTIONAL BACKGROUND

The data we use are the computerised records of the Lancashire Careers Service over the period 1985–1992. During this period, the Careers Service fulfilled a similar role for the youth labour market as Employment Offices and Job Centres currently 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 vacancies.

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. Vacancies in the 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 by employers, it is an important method. 19% of all jobs for those aged 16–18 are filled by the Careers Service. In addition, a further 18% of jobs follow directly on from the Youth Training Scheme. The Careers Service is therefore involved, directly or indirectly, in 37% of job placements for young people in Lancashire.[2]

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. 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.

We argue that vacancies which lapse are genuinely unfilled. For this to be true, we need to be sure that vacancies which lapse are not subsequently filled by some other method of search. Fortunately the database of school-leavers includes information on *all* successful matches. A variable on the database records whether the match was made via the Careers Service or not. For each match we have a date, a location, an occupation as well as the age and the qualifications of the school-leaver. We then search the vacancy database for any lapsed vacancies which match on all these characteristics. We find that only a tiny proportion (about 1%) of lapsed vacancies previously notified to the Careers Service are filled by school-leavers using other search methods.

This still leaves the possibility that these lapsed vacancies were filled by older job-seekers. This is unlikely for three reasons. First, it is implausible to imagine that lots of older job-seekers are matching with these vacancies via some other search method, while almost none of the 15–18 year olds do so (recall that we observe the population of 15–18 year-olds). Second, 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. Third, a high proportion of firms with lapsed vacancies subsequently post vacancies with identical characteristics the following year. We cannot tell, however, whether these vacancies are the same vacancies which were lapsed the previous year, or new vacancies with the same characteristics.

Thus we have some evidence consistent with the notion that lapsed vacancies are not being filled by other means. But because we do not observe all methods of search, and because we do not observe all possible job-seekers, it remains possible that some lapsed vacancies are in fact filled. Our results, of course, also shed some light on this issue.

Our data comprise a standard flow sample: we observe completed durations of vacancies placed on the market and subsequently filled or lapsed between March 1985 and June 1992. 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. Left-censoring does not occur as we have a flow sample.

One further feature of these data is that employers may advertise several vacancies simultaneously. For example, a firm may want to hire 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. In other words, the unit of observation in our analysis is a vacancy order. The precise details of what is observed, and what is not, is deferred to when we derive a non-standard likelihood function to deal with multiple vacancy orders in Section IV.

IV ECONOMETRIC METHODS

The econometric framework we use is a reduced-form mixed proportional hazards (MPH) model with multiple destinations for discrete date, and with dependent risks. These are often referred to as *correlated competing risks models*. The MPH framework is widely used in the estimation of reduced-form hazard models, for example models of unemployment duration with multiple outcomes. If we could observe all determinants of search duration without measurement error, the mixed part of the MPH assumption would be unnecessary, but this rarely happens. The proportional part of the assumption is, in fact, hard to justify in terms of the standard search model (see, for example, van den Berg) and this is also the case for the model discussed in Section II (Frijters and van der Klaauw, 2006). In our treatment of dependent risks, we adapt Lancaster (1990), using van den Berg (2000, Section 8.2.1). We first discuss standard methods for modelling single vacancy orders, which we then amend for multiple vacancy orders.

Single orders

Most vacancies exit to one of two destinations, filled (denoted $r = 1$), or lapsed ($r = 2$), and a small number are censored ($r = 0$). Each vacancy, subscripted i, exits to one and only one destination. The random variables T_1, T_2, and T_0 represent the time it takes a vacancy to be filled, lapsed, or censored respectively. There are three heterogeneity terms v_0, v_1 and v_2, and we write the joint density of $\mathbf{v} \equiv (v_1, v_2)$ as $g(\mathbf{v})$. In general, we write the conditional (on v_r) hazards as $h_1(T_1 \mid \mathbf{x}', v_1)$ and $h_2(T_2 \mid \mathbf{x}', v_2)$, where $\mathbf{x}'$ is a vector of observed covariates that is the same for both destinations. Both v_1 and v_2 are independent of $\mathbf{x}'$. Conditional on ($\mathbf{x}'$, v_0, v_1, v_2), the latent durations T_1 and T_2 are assumed to be independent. The way any dependence between T_1 and T_2 is modelled is by allowing v_1 and v_2 to be dependent, but with both independent of v_0. The latter means that we do not have to model the parameters of the lapsing process. If v_1 and v_2 are also independent, the model reduces to two unrelated mixed proportional hazards models for T_1 and T_2.

It is well-known that competing risks models are not identified in that a general competing risks model with an arbitrary joint distribution for T_1 and T_2, but without covariates, is observationally equivalent to a model with independent T_1 and T_2. The role played by the observed covariates is important (Heckman and Honoré, 1989). The key assumption is proportional hazards; see Abbring and van den Berg (2003) for a comprehensive discussion of the assumptions needed for identification in such models. Loosely speaking, there is identification if there are two continuous covariates which affect both the filling and lapsing hazards, and with different parameters.

Our data are discrete and are observed in unit intervals (days): $[0, 1), [1, 2), \ldots$. For each vacancy i a duration t_i is recorded if it is observed either filling, lapsing or censoring in the interval $[t - 1, t)$. In terms of the underlying durations, two of which are latent, this means that $\min(T_1, T_2, T_0)$ for vacancy i falls in the interval $[t - 1, t)$.

The standard way to estimate discrete-time duration models is to form a panel of vacancies with the i-th vacancy contributing $j = 1, 2, \ldots, t_i$ observations. This is the 'sequential binary response' form (Prentice and Gloeckler, 1978; Han and Hausman, 1990; Stewart, 1996; Wooldridge, 2002). To do this, one defines a discrete hazard. Because we assume that, conditional on $\mathbf{v}$, the three underlying stochastic processes describing time to filling, lapsing and censoring are mutually independent, the discrete hazard is:

$$h_{rt}(v_r) = \Pr(t-1 \leq T_r < t \mid T_r \geq t-1, \mathbf{v}).$$

In words, this is the probability of exiting to destination r in the interval $[t-1, t)$, given that an exit to destination r has not already occurred.

The proportional hazards assumption is modelled as

$$h_{rt}(\mathbf{x}_i' \mid v_{ri}) = \exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \delta_{rt} + u_{ri}), \tag{7}$$

where $u \equiv \log v$, u_1 and u_2 have a joint distribution function $F(u_1, u_2)$, and δ_{rt} is the logarithm of the integrated baseline hazard at duration t for destination r. As with all discrete-time competing-risks models, we ignore the possibility that the latent durations could coincide in the same observational interval (a day). For example, a vacancy could have lapsed and filled on the same day, but it was filled in the morning before it was due to be lapsed in the afternoon. However, because the interval is a day, this is an unlikely event.

It can be shown that the likelihood for vacancy i in this mixed proportional hazards model is (see Wooldridge (2002), for the single risks case):

$$\begin{aligned} &L_i(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2, \ldots) = \\ &\quad \int_{-\infty}^{\infty} \left[\prod_{j=1}^{t_i} h_{1ij}(.)^{y_{1ij}} [1 - h_{1ij}(.)]^{1-y_{1ij}} h_{2ij}(.)^{y_{2ij}} [1 - h_{2ij}(.)]^{1-y_{2ij}} \right] dF(u_{1i}, u_{2i}), \end{aligned} \tag{8}$$

where

$$h_{rij}(.) = 1 - \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \gamma_{rj} + u_{ri})]. \tag{8'}$$

The dummy variable y_{rij} indicates whether vacancy i exits to destination r in the interval $[j-1, j)$. In other words, for both filling and lapsing, we have a sequence of observations $y_{rij}, j = 1, \ldots, t_i$, all of which are zero except the last. If the vacancy is filled the last observation y_{1it_i} is recorded as unity, if it is lapsed $y_{2it_i} = 1$, and if the vacancy is censored, both $y_{1it_i} = y_{2it_i} = 0$. The proportional hazards assumption means that the covariates affect the hazard via the complementary log-log link. Were v_{1i}, v_{2i} independent, filling and lapsing could be modelled separately, and we would have a binary choice random effects model, with a complementary log-log link rather than the more common logit or probit links.

The parameters to be estimated are $\boldsymbol{\beta}_1$, $\boldsymbol{\beta}_2$, γ_{1j} and γ_{2j}, where the γ_{rj} are collected into the vectors $\boldsymbol{\gamma}_1$ and $\boldsymbol{\gamma}_2$. The γ_{rj}s are interpreted as the log of a non-parametric piecewise linear baseline hazard. Because there are a large number of vacancies in the data, a flexible non-parametric approach is feasible. A possible restriction on the shape of the baseline hazard is provided by the Weibull hazard. In this case, the γ_{rj} in (8$'$) are replaced by $\log \alpha_r \gamma_r + (\alpha_r - 1) \log j$, greatly reducing the number of parameters to be estimated.

We adopt two approaches for modelling the bivariate unobserved heterogeneity. These are: (i) *Gaussian mixing* and (ii) *discrete mixing.* The standard argument for using the latter, as advocated by Heckman and Singer (1984), is that it should affect the baseline hazard less severely than if the wrong choice of parametric mixing is made.

For bivariate Gaussian mixing, u_{1i} and u_{2i}, which have variances σ_1^2, σ_2^2 and correlation ρ, are reparameterised to independent standard Normal variates ϵ_{1i} and ϵ_{2i}, which are then approximated by bivariate Gauss-Hermite quadrature. Thus the likelihood for each observation is written:

$$L_i(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2, \sigma_1, \sigma_2, \rho) = \\ \sum_{k=1}^{Q} \sum_{l=1}^{Q} \left[\prod_{j=1}^{t_i} h_{1ij}(.)^{y_{1ij}} [1 - h_{1ij}(.)]^{1-y_{1ij}} h_{2ij}(.)^{y_{2ij}} [1 - h_{2ij}(.)]^{1-y_{2ij}} \right] \omega_k \omega_l, \quad (9)$$

where

$$h_{1ij}(.) = 1 - \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_1 + \gamma_{1j} + \sigma_1\epsilon_{1ik})]$$
$$h_{2ij}(.) = 1 - \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_2 + \gamma_{2j} + \rho\sigma_2\epsilon_{1ik} + \sqrt{1-\rho^2}\sigma_2\epsilon_{2il})], \tag{9'}$$

and where $\epsilon_{1ik}, k = 1, \ldots, Q$ denotes the Q known quadrature points for ϵ_{1i}, and $\epsilon_{2il}, l = 1, \ldots, Q$, is defined analogously. The corresponding quadrature weights are denoted ω_k and ω_l. The value of Q is determined by the investigator.

For bivariate discrete mixing, u_{1i}, u_{2i} and associated joint density $f(u_{1i}, u_{2i})$ in (8) are replaced by a bivariate discrete mass point approximation $(\bar{u}_{1m}, \bar{u}_{2m}, \pi_m), m = 1, \ldots, M$. Three constraints are imposed on these $3M$ parameters to be estimated:

$$\sum_{m=1}^{M} \bar{u}_{1m}\pi_m = 0, \quad \sum_{m=1}^{M} \bar{u}_{2m}\pi_m = 0, \quad \sum_{m=1}^{M} \pi_m = 1.$$

Unlike Gaussian quadrature, σ_1^2, σ_2^2 and ρ are computed afterwards, rather then being parameters to be estimated. Collecting $(\bar{u}_{1m}, \bar{u}_{2m}, \pi_m), m = 1, \ldots, M$, into vectors $\bar{\mathbf{u}}_1$, $\bar{\mathbf{u}}_2$, $\boldsymbol{\pi}$, the likelihood is written:

$$L_i(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2, \bar{\mathbf{u}}_1, \bar{\mathbf{u}}_2, \boldsymbol{\pi}) = \sum_{m=1}^{M} \pi_m \left[\prod_{j=1}^{t_i} h_{1ijm}(.)^{y_{1ij}} [1 - h_{1ijm}(.)]^{1-y_{1ij}} h_{2ijm}(.)^{y_{2ij}} [1 - h_{2ijm}(.)]^{1-y_{2ij}} \right], \tag{10}$$

where

$$h_{rijm}(.) = 1 - \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \gamma_{rj} + \bar{u}_{rm})]. \tag{10'}$$

The number of mass points M is determined by experimentation; it is usually obvious when to stop adding mass points as the new mass point might have a very low π, or an existing mass point is just split into two, with little improvement in the likelihood. Inference is conducted conditional on M.

Multiple orders

As noted in Section III, many vacancies are grouped together into multiple vacancy orders. Each order contains V_i vacancies, where the orders are numbered $i = 1, \ldots, N$. It is entire orders, rather than an order's individual vacancies, that lapse or are censored. Within an order, any number of individual vacancies may be filled before all the remaining vacancies are lapsed. 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 or censored, we only observe the time of lapsing/censoring, but not identity, of the order's vacancies filled before that time. For each order we know how many vacancies are filled, denoted W_i. Although the unit of observation is a vacancy order, we need to infer the parameters describing the distribution of a single vacancy.

There are five possible outcomes for a vacancy order: (i) the vacancy order is filled. We only observe $\max(t_1, \ldots, t_V)$, the duration of the vacancy order; (ii) the vacancy order is lapsed; (iii) the vacancy order is censored; or (iv) the vacancy order is partially filled before it is lapsed; (v) the vacancy order is partially filled before it is censored. The likelihood for the whole sample is (each product corresponding to (i)–(v)):

$$\int_{-\infty}^{\infty} L \; dG(v_{1i}, v_{2i})$$

where

$$\begin{aligned}
L = & \prod_{i \in \{W_i = V_i\}} V_i [1 - S_1(t_i \mid v_{1i})]^{V_i - 1} [S_1(t_i - 1 \mid v_{1i}) - S_1(t_i \mid v_{1i})] S_2(t_i \mid v_{2i}) \times \\
& \prod_{i \in \{W_i = 0\}} [S_2(t_i - 1 \mid v_{2i}) - S_2(t_i \mid v_{2i})] S_1(t_i \mid v_{1i})^{V_i} \times \prod_{i \in \{C_i = 1\}} S_1(t_i \mid v_{1i})^{V_i} S_2(t_i \mid v_{2i}) \times \\
& \prod_{i \in \{0 < W_i < V_i\}} {}_V\mathcal{C}_W [S_2(t_i - 1 \mid v_{2i}) - S_2(t_i \mid v_{2i})] S_1(t_i \mid v_{1i})^{V_i - W_i} [1 - S_1(t_i \mid v_{1i})]^{W_i}. \quad (11)
\end{aligned}$$

(See the Appendix.) Here $S_1(t_i \,|\, v_{1i})$ and $S_2(t_i \,|\, v_{2i})$ denote the survivor functions for filling and lapsing respectively; these denote the probability of survival to t, given departure to destination r. The fourth product term captures all partially filled vacancy orders, whether (iv) eventually lapsed or (v) censored.[3]

The survivor functions S_1 and S_2 that correspond to the hazard functions given in (7) are:

$$S_r(t_i \mid v_{ri}) = \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \delta_{rt} + u_{ri})], \tag{12}$$

where, recall, δ_{rt} is the integrated baseline hazard over the interval $[t-1, t)$. In other words, the duration dummies in Equation (12) are not the same as those in Equation (8′). Note that $\exp(\gamma_{rt}) = \exp(\delta_{rt}) - \exp(\delta_{r,t-1})$. To examine whether the discrete-time Weibull $\bar{h}_{rt} = \gamma_r \alpha_r t^{\alpha_r - 1}$ is an appropriate special case, then

$$S_r(t_i \mid v_{ri}) = \exp[-\exp(\mathbf{x}_i'\boldsymbol{\beta}_r + \log\gamma_r + \alpha_r \log t_i + u_{ri})] \tag{13}$$

replaces (12) above.

We integrate out the bivariate unobserved heterogeneity in the same way as in the single-order case, using dependent discrete mixing or dependent Gaussian mixing. Thus Equations (12) and (13) need amending in analogous ways, but, to save space, we do not present the formulae here.

Interpreting the parameters

Although we estimate a separate vector of coefficients $\boldsymbol{\beta}_r$, $r = 1, 2$ for filled and lapsed, each vector conveys 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

$$\Pi_r = \sum_{j=1}^{\infty} h_{rj} S_{j-1}, \qquad E_r = \frac{1}{\Pi_r} \sum_{j=1}^{\infty} j h_{rj} S_{j-1}, \qquad S_j = \prod_{s=1}^{j} (1 - h_{1s} - h_{2s}). \tag{14}$$

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 or lapsing conditional on exiting during the interval j. The conditional probability of lapsing is defined as:

$$P_j = \frac{h_{2j}}{h_{1j} + h_{2j}}. \tag{15}$$

This is the empirical equivalent of Equation (6) in Section II. Thus, in addition to estimates of β_1 and β_2, we report the marginal effect of a covariate x on the conditional lapsing probability, given by

$$\frac{\partial P_j}{\partial x} = \frac{h_{1j} h_{2j} (\beta_2 - \beta_1)}{(h_{1j} + h_{2j})^2}. \tag{16}$$

This formula applies to discrete variables as well as continuous ones, and applies whether or not the heterogeneity terms are present, let alone dependent. Standard errors can be obtained using the Delta Method.

V RESULTS

The raw data

Table 1 describes the sample, which covers the period 1985–1992. There are 14,510 vacancy orders containing a total of $\sum_{i=1}^{N} V_i = 17,759$ vacancies. Most vacancy orders (12,840) therefore contain a single vacancy. As already noted, a substantial proportion (34%) of vacancies lapse. Table 1 also summarises the dependent variable, the total time that a vacancy is open on Careers Service records. To calculate the underlying average duration of filled and lapsed vacancies that allows for genuine censoring or for the vacancy exiting to the other destination, we compute ML estimates of the parameters from an Exponential distribution using Equation (11). Vacancies which fill have a mean duration of eight weeks, whereas vacancies which lapse have a mean duration of ten weeks. For multiple orders, we observe very similar mean durations. There are just 147 censored vacancy orders, comprising 220 individual vacancies. Finally, the third panel of Table 1 gives the number of the four types of order which make up the likelihood function Equation (11). Note that none of the censored vacancy orders contain any filled vacancies.

[TABLE 1 HERE]

Preferred specification and baseline hazard

In Section IV we described several possible specifications for modelling the data. Different specifications are required for (a) the choice between parametric and non-parametric baseline hazards, (b) the type of unobserved heterogeneity (Gaussian or discrete mixing) and (c) when modelling multiple, rather than single, vacancy orders. In this subsection we explain how we select our preferred specification. The way we do this is to focus on single orders first of all, decide on the appropriate specification, and then use the same specification for multiple orders. Table 2 summarises the following six specifications for single orders, labelled A to F.

[TABLE 2 HERE]

Specifications A, B, and C all have non-parametric baseline hazards.[4] Specification A is discrete mixing, dependent risks, whereas Specification B is Gaussian mixing, dependent risks. Specification C is Gaussian mixing, but with independent risks. See Equations (9) and (10) respectively. Specifications D, E, and F repeat these three, but with Weibull hazards.

In choosing between Gaussian and discrete mixing (Specifications A and B), the discrete mixing model has a log-likelihood that is 150.5 log-points higher. It also has 12 more parameters, having $3(M-1) = 15$ parameters for $M = 6$ discrete mass points, whereas the Gaussian mixing model estimates two variances and a covariance. Any of the common model selection criteria that take into account these 12 extra parameters favour discrete mixing given the large difference in log-likelihoods. The two main differences between the models are that the negative correlation between u_1 and u_2 is stronger for discrete mixing (–0.649 rather than –0.313), as is the variance of u_1 (unobserved propensity to fill). In terms of parameter estimates, it does not matter which is preferred. The negative correlation between u_1 and u_2 suggests that vacancies with a high unobserved propensity to fill have a low unobserved propensity to lapse. Suppose one is unable to observe potential revenue flow, defined as z_m in the theory above. A "good" vacancy with a high z_m and a high λ is more likely to fill, but also likely to have a longer optimal lapsing time, thereby creating a negative correlation in the unobservables.

In Specification C we re-estimate B, but with independent risks. We just reject the hypothesis of a zero correlation (the likelihood decreases by 2.21), giving a p-value of 0.036 on the likelihood ratio test. Very little else alters, which does suggest that it is not particularly important to estimate this correlation, except to show that it is indeed negative. Specification D is the Weibull equivalent of A. The Weibull estimate implies decreasing duration dependence in the filling hazard and increasing duration dependence in the lapsing hazard. These mimic the non-parametric hazards estimated in Specification A (not plotted). In the raw data, Figure 1, there are spikes ('seasonality') in both hazards which occur at regular intervals. These spikes reflect the institutional nature of recruitment in the youth labour market. Because the Weibull cannot capture this feature of the data at all, there is a clear rejection of the likelihood ratio test compared with A. (Specification A estimates 24 extra parameters per risk.) Also note that Specification E has a zero estimated correlation between the risks, unlike Specification B.

To conclude, discrete mixing is preferred to Gaussian mixing, and in the latter we clearly reject independent risks. Moreover, non-parametric hazards are preferred to the Weibull. In other words, Specification A is our preferred model for single orders. We now report, in Table 3, the estimates from applying discrete mixing to the likelihood given in Equation (11), which allows for multiple orders in the data. This is labelled Specification G. Adding multiple orders to the sample means that the likelihood is harder to maximise (that is, becomes "flatter" and with more local maxima). We were therefore forced to specify a more parsimonious baseline hazard, with 10 rather than 26 parameters per risk.[5] But this is preferable to fitting Weibull hazards.

[FIGURES 2 AND 3 HERE]

In Figures 2 and 3 we plot the estimated hazards and compare them with the raw data. In both cases, estimated hazards are flatter than the raw hazards, and lie below the raw hazards. If it is the heterogeneity that partly causes duration dependence, then this is what one expects. (Actually, there are 5 hazards per risk, one for each of $m = 1, \ldots, 5$, but the figure plots the average.) The extreme case is when the hazards become completely flat after controlling for unobserved heterogeneity, which means that it would be unnecessary to model firms in a non-stationary environment in the theory above. There is no evidence that this is true. Thus the shapes of the filling and lapsing hazards (and therefore the conditional probability of lapsing a vacancy) are entirely consistent with Section II, providing the applicant arrival rate falls over time. As search duration increases, the benefit of searching falls until the desired lapsing time is reached. The distribution of lapsing times across the sample of vacancies is such that the lapsing hazard increases with elapsed duration. However, because we cannot be certain that lapsed vacancies have not been filled elsewhere, it is possible that the upward sloping hazard to lapsing h_2 is also consistent with non-sequential search via another search channel; for example, employers might fill vacancies via the LCS quickly, but fill vacancies via newspapers only after some delay.[6]

[FIGURE 4 HERE]

In Figure 4 we plot the conditional probability of lapsing a vacancy P_j, calculated from Equation (15). Recall that the theory suggests that this should be interpreted as the proportion of vacancies which lapse at each duration, rather than the probability of lapsing for a given vacancy. For both the raw data and Specification G, the declining hazard for filling and the increasing hazard for lapsing (see Figures 2 and 3) means that the probability that a vacancy lapses must increase with duration. Because both hazards are flatter, the curve in Figure 4 should be flatter than the raw one. Indeed, at long durations where all but a handful of vacancies have either lapsed or filled, P_j is close to the sample proportion of vacancies that lapse (roughly one-third).

In comparing Specifications A and G, the estimated correlation between the unobserved filling and lapsing terms is the same, changing from –0.650 to –0.606. Similarly, the estimates change do not move much; out of 106 covariates (excluding those that model the hazards) 22 lie outside the 95% confidence for Specification A. Of these, the most significant change is the estimate on the stock of job vacancies in the filling hazard, which now has a correct negative effect. It should be noted that the estimated effects for almost all the observed covariates are robust across all seven specifications reported in this paper.[7]

Determinants of vacancy duration

There are two distinct objectives in reporting our parameter estimates. First, in this subsection, we report the effects of covariates on the duration of employer search, using (primarily) the coefficient estimates $\partial h_1/\partial \mathbf{x} \approx \widehat{\boldsymbol{\beta}}_1$, and in particular we examine whether they are consistent with the predictions of the employer-search model outlined in Section II. We report evidence from Andrews, Bradley and Upward (2001) [hereafter ABU], who use the same data to report estimates of $\partial \log \mu/\partial x$, so that we infer the effect of a covariate on λ as well as h_1. Second, in Subsection V, we examine the marginal effect of covariates on the conditional probability that an employer lapses a vacancy, $\partial P/\partial \mathbf{x}$, given in Equation (15). In particular, we assess 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.

[TABLE 3 HERE]

Labour market tightness and applicant arrival effects, λ

In almost all models of search in the labour market, the arrival rate of applicants is a decreasing function of labour market tightness, $\lambda(V/U)$. 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 coefficient on 'Unemployed $\leq$ 18' in Table 3 shows that vacancies in labour markets with higher youth unemployment have significantly higher hazards to filling ($\widehat{\beta}_1 = 0.17$). As noted, labour market tightness operates via two opposing channels. There is the direct effect on applicant arrival rates λ and the indirect effect via the matching probability μ (see Equation (5)). Additional evidence on the effect of the labour market tightness is available from ABU, who established that $\partial \log \mu/\partial \log U = -0.17$. Given $h_1 = \mu\lambda$, we can infer that $\partial \log \lambda/\partial \log U = 0.34$. This large effect is consistent with the theory (in fact, only guaranteed if the offer distribution is log-concave) and so our conclusion is that an increase in unemployment increases the hazard to filling because the increase in the number of contacts per vacancy outweighs the employer's 'more selective' response. This is, of course, the same reason why vacancy hazards are downward-sloping as the applicant arrival rate declines over time.

The estimates for the stock of vacancies show a corresponding negative effect on the filling hazard, with a small elasticity of -0.11. We are only just unable to impose the homogeneity restriction that allows the covariate to be labour market tightness V/U (p-value=0.026); in fact, finding increasing returns to scale is the norm in this literature (Petrongolo and Pissarides, 2001).

In addition to the labour-market-tightness variables above, we also consider two other characteristics of the local labour market in which the firm is located. The first is log population density to capture whether the matching probability or the arrival rate of applicants is higher in cities compared with rural areas; there is a weak but significant effect of 0.05. A better variable with which to capture applicant arrival effects is our measure of firm location. Firms located in town centres, and which are therefore more accessible to job-seekers (who have lower search costs), have significantly higher hazards to filling, with a differential of some 0.16 log-points. This offsets possible effects from higher competition, with more potential employers in town centres.

The third local labour variable we consider is the number of staff in a given Careers Office, normalised on the population of each district. It has a negative elasticity of –0.13 on the hazard to filling a vacancy, suggesting that more staff generate fewer applicants per vacancy, which is somewhat unexpected.

The next variable we consider is firm size, where it is clear that the bigger the firm, the easier it is to fill a 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 in ABU, again this is an applicant arrival effect λ. (The negative effect of λ on μ cancels out with the positive effect of lower search costs c for larger firms.) Unlike larger Careers Offices, larger firms can process more applicants. If larger firms have higher applicant arrival rates the theory predicts that h_2

will fall, which is what we find. The net effect on the conditional probability of exit $\partial P/\partial \mathbf{x}$ is that larger firms tend to fill vacancies more successfully than smaller ones.

The revenue flow from a match, z_m

The most important observable component of z_m is the wage. There are three different types of wage offer in the data. About 75% of vacancies have a set pre-announced wage, where the wage is non-negotiable. 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 vacancies have a set pre-announced wage offer, but are still open to negotiation. The remaining vacancies have a negotiable wage offer and no pre-announced wage. For this third category there is no wage recorded in 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; in other words, simultaneity bias is not an issue. We argue that this is an accurate characterisation of the youth labour market, given the vast majority of vacancies in the data have a non-negotiable wage.

We model these effects as follows. We 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 prediction of the simple model above is that a higher wage should make employers more selective. Moreover, the wage can also affect the hazard via the applicant arrival rate λ. A higher wage will increase λ, which has an ambiguous effect on the hazard. Added to the employer selection effect above, the overall prediction is ambiguous.

The coefficient on the 'Log wage' is $\log wD$, that is, captures the relationship between the wage and the vacancy hazard for vacancies with a pre-announced wage. The effect is positive and insignificant. (It is generally insignificant in the single-order specifications as well [check when Dave finishes].) Most of the literature does not find any significant effect on the hazard to filling. ABU found clear negative effects on the matching probability, which implies that a higher wage must generate more applicants in both cases.

Those 23% of vacancies that have a negotiable wage have the same hazard to filling as those without, suggesting that negotiation does not take long at all. This is probably because, in this particular market, employers make take-it-or-leave-it offers. On the other hand, those vacancies which do lapse are open longer if the wage is negotiable.

Apart from the wage, we do not observe the revenue flow from a match directly. 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 I) suggests that the greater the potential investment by the employer in the worker (through more training and so on) the longer it takes to fill a vacancy (lower vacancy hazard). Vacancies which require more investment by the firm may take longer to fill because the revenue flow is initially low during the training period. Once training has finished, the revenue flow is higher than for low-skill jobs, but this is discounted because of the probability of a future separation: the newly trained employee may leave the firm. Employers therefore become more selective (μ falls) and the hazard is lower. In addition, if higher quality vacancies attract more variable applicants, the vacancy hazard will be lower if the employer waits longer for a 'bargain'. 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 3 shows that non-manual vacancies have sizeable and significantly lower hazards to filling ($\widehat{\beta}_1 = -0.32$) than their manual counterparts, but that skilled vacancies do not have a significantly higher hazard to filling compared with unskilled vacancies. More evidence that employers search longer for 'better' vacancies is provided by the training information. Vacancies offering day release to College ($\widehat{\beta}_1 = -0.36$) or apprenticeship training ($\widehat{\beta}_1 = -0.43$) have much lower hazards. Only one of these three variables had any effect on the matching probability in ABU—an elasticity of -0.20 for non-manual—and so we are estimating large negative applicant arrival effects in the other two cases.

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. (For vacancies where no written application was required, the Careers Service would undertake the application procedure.) Some of these characteristics, such as qualifications, are directly related to the employer's reservation productivity z^*, and are an attempt to limit the pool of potential applicants. We would therefore expect that higher criteria imply higher z^* and longer search durations. In addition, 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). Table 3 shows that vacancies requiring higher educational qualifications have significantly lower hazards to filling and therefore have longer search durations, taking longer the more educated the applicant.

The largest estimated effect is the requirement that a written application be submitted by the job-seeker ($\widehat{\beta}_1 = -0.68$). Increases in duration might be because written applications increase the application period, or because written applications increase the selection period, as in van Ours and Ridder (1993). In fact, most of this effect is because of a much lower matching probability, an elasticity of -0.72 (ABU), suggesting that the written application reveals the non-suitability of the candidate before they get to interview. Also notice that the hazard to lapsing is also much lower, with no significant effect on $\partial P/\partial \mathbf{x}$.

Skill shortages — which vacancies lapse?

An important feature of our data is that a substantial number of vacancies are removed from the market before they are filled. As noted in the theory section, in a stationary world one would not observe lapsing. We predict that lapsing occurs because the applicant arrival rate falls, pushing the discounted revenue stream rV_n below its reservation level rV_l. When the applicants dry up completely, this is when the employer learns that there is a skill shortage and hence lapses the vacancy. In this subsection 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.

Alternative descriptive evidence comes from the ONS data referred to in the Introduction, namely the proportion of all vacancies notified to the public employment service which are subsequently cancelled. For July 1993, not long after the end of our sample period, this proportion is highest for skilled manual workers (35%), 30% for semi-skilled non-manual workers, 23% for skilled manual workers and 20% for unskilled manual workers.

In the competing risks framework that we adopt, we are able to determine which characteristics of a vacancy reduce the time to lapsing. The hypothesis that skill shortages cause vacancies to lapse suggests that those vacancies requiring more skilled applicants will take longer to fill and, at the same time, take longer to lapse. Conditional on exiting, such vacancies will have a higher probability of lapsing. An alternative hypothesis is that lapsing is a result of low-quality jobs, in that they offer low wages or little training, being refused by potential applicants, in which case it will be low-quality vacancies which are more likely to lapse.

Estimates of $\partial P_j/\partial x$ are reported in Table 3. We examine $\partial P/\partial \mathbf{x}$ rather than $\partial h_2/\partial \mathbf{x}$, because some vacancies might take longer to lapse than others, but also might take longer to fill, therefore having no net effect on the conditional probability of lapsing P. It is the eventual fate of the vacancy that we are interested in.

We find that non-manual vacancies have significantly higher lapsing hazards and significantly lower filling hazards, implying that the probability that these vacancies lapse is significantly higher. If we believe that employers correctly predict the applicant arrival rate, this suggests that employers decide to lapse non-manual vacancies earlier because they know that the applicant arrival rate will dry up earlier.

In contrast, most other measures of vacancy quality (such as training and academic requirements) are associated with a lower lapsing hazard even though such vacancies, as already discussed, have lower hazards to filling. The net effect is that $\partial P/\partial \mathbf{x}$ is insignificant for these vacancies.

The only evidence we have for skill shortages is for the non-manual dummy. The lack of significant effects for training and academic requirements is because non-manual is highly correlated with these variables. On the other hand, there is no evidence at all that the less-skilled vacancies are more likely to lapse. 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. However, we cannot be certain that lapsed vacancies are actually withdrawn from the market. As noted in Section III, it is possible, although unlikely, that these vacancies are subsequently filled by older applicants via a different search channel. Even if this is the case, these results still demonstrate that employers cannot fill these vacancies from the youth labour market; the conclusion that there are skill shortages still holds up.

VI CONCLUSIONS

This paper provides the first analysis of vacancy duration and the outcome of employer search using duration modelling methods for the UK, using a large sample of vacancies for a particular market. Our results are of interest for two reasons. First, understanding employers' search behaviour is an important, and yet under-researched, area. Second, we provide plausible evidence that a sizeable proportion of vacancies are removed from the market before they are filled, or are 'lapsed'. Evidence on 'hard-to-fill' vacancies (those with long durations) is thin; apart from Beaumont (1978), our paper provides the only other analysis of lapsed vacancies.

Our key results are as follows:

1. For all specifications with non-parametric hazards, after controlling for unoberved heterogeneity, the hazard to filling is downward-sloping over most of its range and the hazard to lapsing is upward-sloping, implying that the conditional lapsing probability increases with duration. This is consistent with the claim that lapsed vacancies are being withdrawn from the market (if they were being filled elsewhere, they would probably have downward-sloping hazards), but we cannot be certain that this is so. It is also consistent with a fall in the arrival rate and quality of applicants; the employer's response is to increase the probability of lapsing a vacancy and decrease the probability of a filling a vacancy.

2. For all models with dependent risks, we always estimate a negative correlation. It makes good sense that a vacancy that has a high unobserved propensity to fill (a

'good' vacancy) will also have a low unobserved propensity to lapse.

3. A key variable in all search models is labour market tightness. An increase in unemployment increases the hazard because the increase in the number of applicants per vacancy outweighs the employers' more selective response. The smaller negative effect of the aggregate stock of vacancies implies increasing returns.

4. In most specifications, the wage does not affect the duration of employer search. However, in earlier work, ABU found clear negative effects on the matching probability, which implies that a higher wage generates more applicants.

5. 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.

6. Finally, we find that employers find it difficult to find suitable applicants for non-manual vacancies, because search takes longer and because these vacancies are more likely to be withdrawn from the market before they are filled. Thus it is *good* rather than *bad* vacancies which are hard to fill.

In short, we find that lapsed vacancies are an important aspect of labour markets that have hitherto been ignored. In our data, we have evidence that lapsing is caused by skill shortages insofar as it is non-manual vacancies that are more likely to lapse. However, without more information on the choice of search channel used by employers it remains a possibility that higher-skill vacancies are filled by other search methods, and that those search channels have a different hazard to filling. A priority for future research is to see whether lapsing is a more general phenomenon, for which we would require data on multiple search channels.

APPENDIX: LIKELIHOOD FOR MULTIPLE VACANCY ORDERS

There are five possible outcomes for a vacancy order, suppressing the i subscript for clarity:

1. The vacancy order is filled. This occurs if all V individual vacancies are filled before any are lapsed or censored:

$$T_1 < \bar{T}, T_2 < \bar{T}, \ldots, T_V < \bar{T} \quad \text{or}$$
$$T_1 < C, T_2 < C, \ldots, T_V < C.$$

   We only observe y, the duration of the vacancy order, when $\max(T_1, \ldots, T_V)$ falls in the interval $[t-1, t)$. The likelihood of observing this outcome is

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

   Note that this expression is slightly different in the first period, because we know that V individual vacancies fill in the interval $[0, 1)$, and that $\bar{T}$ or C occurs later. Thus the likelihood for $y = 1$ is written $[1 - S_1(1)]^V S_2(1)S_0(1)$.

2. The vacancy order is lapsed. This occurs if none of the V individual vacancies are filled and the vacancy order is not censored before the vacancy order is lapsed:

$$\bar{T} < T_1, \bar{T} < T_2, \ldots, \bar{T} < T_V, \bar{T} < C.$$

   Denoting $\bar{t}$ as the duration of the lapsed vacancy order, observed if $\bar{T}$ falls in the interval $[t-1, t)$, the likelihood of observing this outcome is

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

   In the first period, $S_2(0) = 1$.

3. The vacancy order is censored. This occurs if none of the V individual vacancies are filled and the vacancy order is not lapsed before the vacancy order is censored

$$C < T_1, C < T_2, \ldots, C < T_V, C < \bar{T}.$$

   Denoting c as the duration of the lapsed vacancy order, observed if C falls in the interval $[t-1, t)$, the likelihood of observing this outcome is

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

   In the first period, $S_0(0) = 1$.

4. The vacancy order is partially filled before it is lapsed. Suppose that W individual vacancies are filled before lapsing:

$$T_1 < \bar{T}, \ldots, T_W < \bar{T}, \bar{T} < T_{W+1}, \ldots, \bar{T} < T_V, \bar{T} < C.$$

   We only observe $\bar{t}$ and W. We observe $\bar{t}$ if $\bar{T}$ falls in the interval $[t-1, t)$. The likelihood of observing this outcome is:

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

Notice that this encompasses Case 2 above when setting $W = 0$. In the first period, $S_2(0) = 1$.

5. The vacancy order is partially filled before it is censored. Suppose that W individual vacancies are filled before censoring:

$$T_1 < C, \ldots, T_W < C, C < T_{W+1}, \ldots, C < T_V, C < \bar{T}.$$

We only observe c and W. We observe c if C falls in the interval $[t - 1, t)$. The likelihood of observing this outcome is:

$$_V\mathcal{C}_W[1 - S_1(c)]^W[S_0(c - 1) - S_0(c)]S_1(c)^{V-W}S_2(c).$$

Notice that this encompasses Case 3 above when setting $W = 0$. In the first period, $S_0(0) = 1$.

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 t_i) given by Equation (11) of the main text. Notice that we suppress contributions to the likelihood from the censored distributions.

NOTES

[1]Relaxing the assumption of perfect foresight, van den Berg (2000, Section 3.1.2) considers the case of nonstationarity without anticipation. Such a model is estimated by Narendranathan (1993). In this case employers have a single reservation productivity z^*. "Shocks" to λ or $f(z_m)$ could then cause the employer to revise z^*, which might cause it to fall below z_n.

[2]See Upward (1998, ch. 4) for fuller details, especially Section 4.3 on the representativeness of our data.

[3]Substituting $V_i = 1$ into Equation (11) gives the standard survivor form likelihood that corresponds to Equation (8). The important difference between this likelihood for single vacancies and Equation (11) is that the data cannot be organised into sequential binary response form with multiple vacancy orders.

[4]There are 26 pieces per risk. The intervals, in days, are: [0,1), ..., [13,14), [14,21), [21,28), [28,35), [35,42), [42,49), [49,56), [56,84), [84,112), [112,140), [140,168), [168,196), [196,∞).

[5]The intervals, in days, are: [0,7), [7,14), [14,21), [21,28), [28,56), [56,84), [84,112), [112,140), [140,168), [168,∞).

[6]We are grateful to two referees for pointing this out.

[7]The problem of identification noted in Section IV does not arise here; all of the models have a large number of covariates, six of which are continuous, and with different estimates between the filling and lapsing hazards.

REFERENCES

Abbring, J. and van den Berg, G. (2003), "The identifiability of the mixed proportional hazards competing risks model", *Journal of the Royal Statistical Society: Series B* **65**, 701–10.

Adams, J., Greig, M. and McQuaid, R. (2002), "Mismatch in local labour markets in central Scotland: the neglected role of demand", *Urban Studies* **39**, 1399–416.

Andrews, M., Bradley, S. and Upward, R. (2001), "Estimating the probability of a match using microeconomic data for the youth labour market", *Labour Economics* **8**, 335–57.

Barron, J., Berger, M. and Black, D. (1997), "Employer search, training and vacancy duration", *Economic Inquiry* **35**, 167–92.

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

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

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

Behrenz, L. (2002), "The employment service and vacancy durations", *Evaluation Review* **26**, 602–17.

Bosworth, D. (1993), "Skill shortages in Britain", *Scottish Journal of Political Economy* **40**, 241–71.

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

Cahuc, P. and Zylberberg, A. (2004), *Labor Economics*, MIT Press.

Campbell, M. and Baldwin, S. (1995), "Recruitment difficulties and skill shortages: an analysis of labour market information in Yorkshire and Humberside", *Regional Studies* **27**, 271–80.

Coles, M. and Petrongolo, B. (2003), A test between unemployment theories using matching data, mimeo.

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

Frijters, P. and van der Klaauw, B. (2006), "Job search with non-participation", *Economic Journal* **116**, 45–83.

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

Gorter, R., Nijkamp, P. and Rietveld, P. (1996), "Employers' recruitment behaviour and vacancy durations: an empirical analysis for the Dutch labour market", *Applied Economics* **28**, 1463–74.

Green, F., Machin, S. and Wilkinson, D. (1998), "The meaning and determinants of skill shortages", *Oxford Bulletin of Economics and Statistics* **60**, 165–87.

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

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

Haskel, J. and Martin, C. (2001), "Technology, wages and skill shortages: evidence from UK micro data", *Oxford Economic Papers* **62**, 642–58.

Heckman, J. and Honoré, B. (1989), "The identifiability of competing risks models", *Biometrika* **76**, 325–30.

Heckman, J. and 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.

Lipmann, S. and McCall, J. (1976), "The economics of job search, a survey (part I)", *Economic Inquiry* **XIV**, 155–189.

Machin, A. (2003), "The Vacancy Survey: a new series of National Statistics", *Labour Market Trends* pp. 349–62.

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

Mason, G. and Stevens, P. (2003), The determinants of hard-to-fill vacancies and skill shortage vacancies in key occupational groups, mimeo, National Institute of Economic and Social Research.

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

Mortensen, D. and Pissarides, C. (1998), "Job reallocation, employment fluctuations and unemployment differences", *in* M. Woodford and J. Taylor, eds, *Handbook of Macroeconomics*, North-Holland, Amsterdam, chapter 18, pp. 1171–228.

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

Narendranathan, S. (1993), "Job search in a dynamic environment: an empirical analysis", *Oxford Economic Papers* **45**, 1–22.

Petrongolo, B. and Pissarides, C. (2001), "Looking into the black box: a survey of the matching function", *Journal of Economic Literature* **39**, 390–431.

Prentice, R. and 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.

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

Thomas, J. (1996), "On the interpretation of covariate estimates in independent competing-risks models", *Bulletin of Economic Research* **48**, 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**, 842–65.

van den Berg, G. (1994), "The effects of changes of the job offer arrival rate on the duration of unemployment", *Journal of Labor Economics* **12**, 478–98.

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

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

van Ours, J. and Ridder, G. (1991), "Cyclical variations in vacancy durations and vacancy flows: an empirical analysis", *European Economic Review* **35**, 1143–55.

van Ours, J. and Ridder, G. (1992), "Vacancies and the recruitment of new employees", *Journal of Labor Economics* **10**, 138–55.

van Ours, J. and Ridder, G. (1993), "Vacancy durations: search or selection?", *Oxford Bulletin of Economics and Statistics* **55**, 187–98.

Weber, A. (2000), "Vacancy durations: a model of employer search", *Applied Economics* **32**, 1069–75.

Wooldridge, J. (2002), *Econometric Analysis of Cross Section and Panel Data*, MIT Press.

FIGURES

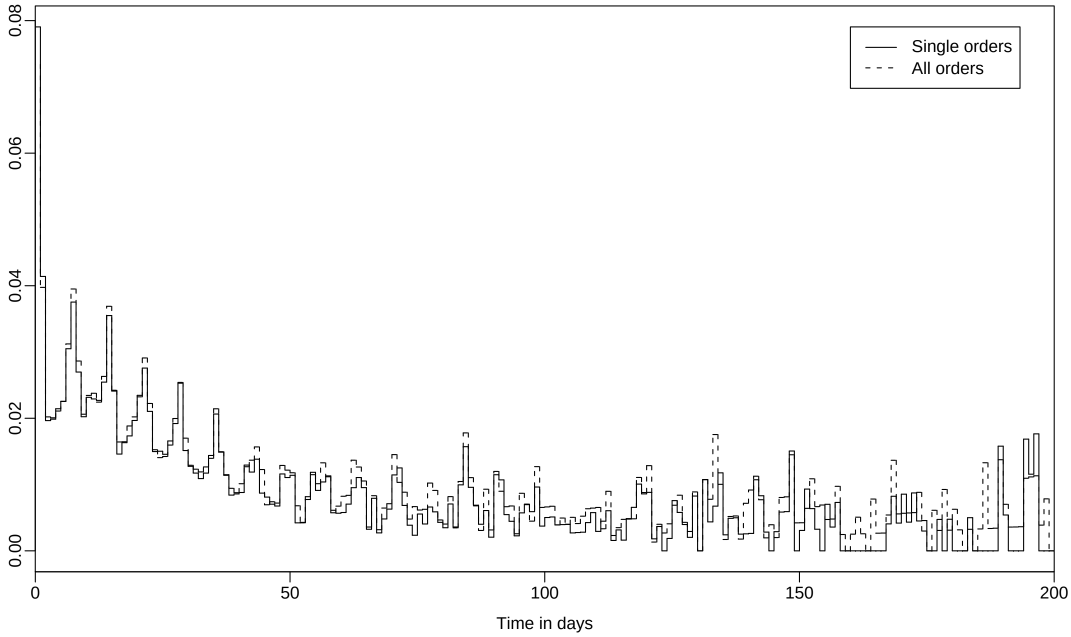

Figure 1: Raw vacancy (filling) hazards

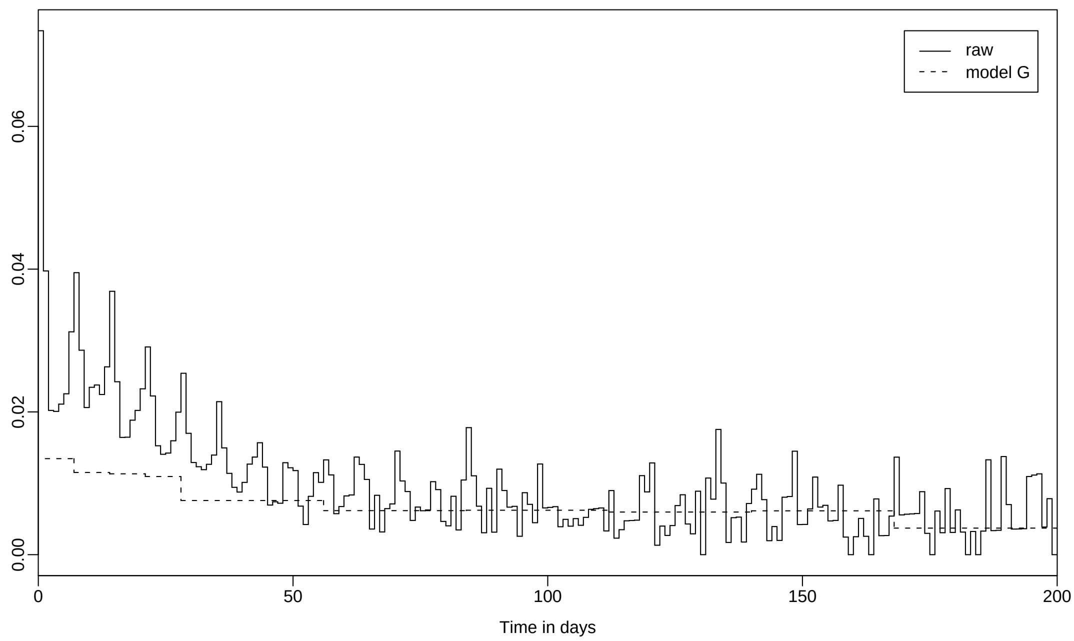

Figure 2: Baseline hazards to filling, multiple orders

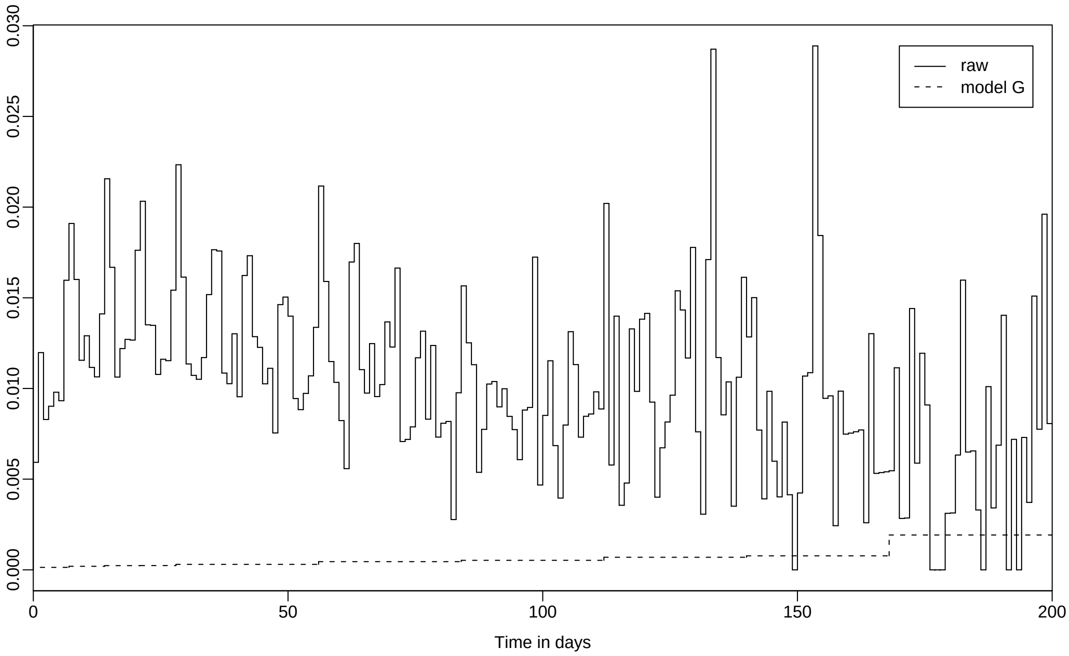

Figure 3: Baseline hazards to lapsing, multiple orders

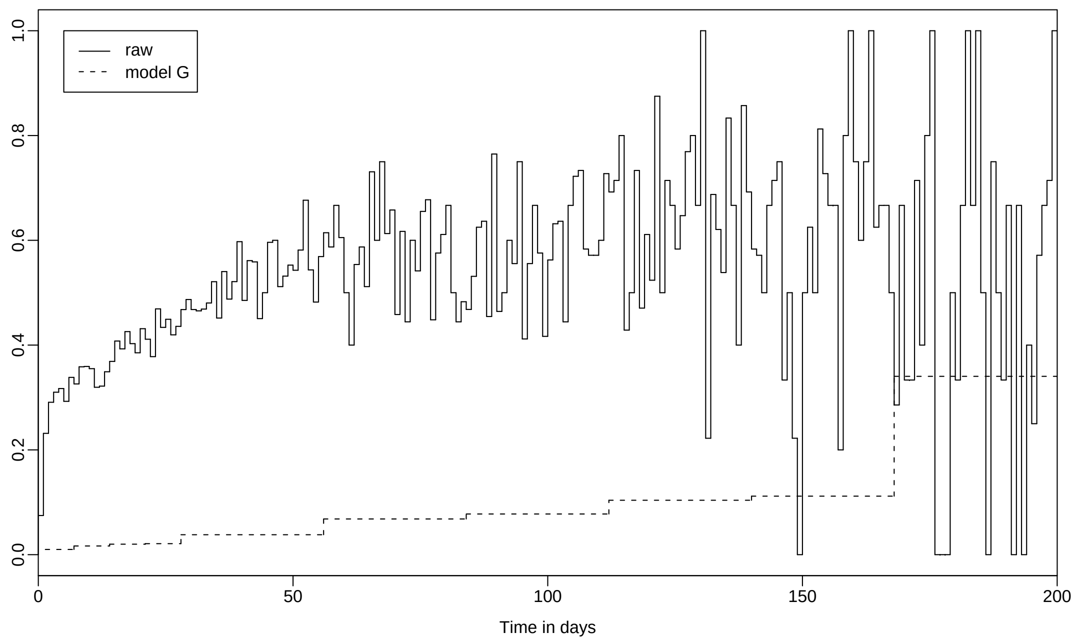

Figure 4: Probability of lapsing, multiple orders

TABLES

Table 1: Vacancy duration

	Mean duration (days)	*ML estimate of duration*[a] *(days)*	*No.*
Single vacancies, $V_i = 1$			
Filled ($W_i = 1$)	21.07	53.97	7234
Lapsed ($W_i = 0$)	42.05	71.19	5484
Censored	60.69		122
Number of single vacancies			12840
All vacancies			
Filled		50.33	11485
Lapsed		81.53	6054
Censored			220
Total no. of vacancies ($\sum_{i=1}^{N} V_i$)			17759
All vacancies, by order			
All filled ($W_i = V_i$)			8548
Partially filled ($0 < W_i < V_i$)			242
All lapsed ($W_i = 0$)			5573
Censored ($C_i = 1$)			147
Total number of orders (N)			14510

[a] Assuming an Exponential distribution: substitute $S_r(t_i) = \exp(-\gamma_r t_i)$ into Equation (11), without mixing. Mean duration is $1/\gamma_r$.

Table 2: Summary of specifications, single orders[*]

	Non-parametric hazard											
	A				B				C			
	Discrete, dep risks				Gaussian, dep risks				Gaussian, indep risks			
	Filled		Lapsed		Filled		Lapsed		Filled		Lapsed	
Unemployed ≤ 18 ($\log U$)	0.2647	(0.0298)	-0.0741	(0.0456)	0.3017	(0.0343)	-0.0854	(0.0468)	0.2873	(0.0321)	-0.0533	(0.0399)
Job vacancies ($\log V$)	0.0193	(0.0232)	-0.1690	(0.0342)	0.0495	(0.0250)	-0.1536	(0.0340)	0.0474	(0.0238)	-0.1475	(0.0305)
Log wage if $D = 1$	-0.0243	(0.0661)	0.2137	(0.0984)	-0.0243	(0.0725)	0.2323	(0.1010)	-0.0232	(0.0684)	0.2210	(0.0881)
σ_1, σ_2	4.5498	(n/a)	1.3814	(n/a)	1.0386	(0.1126)	1.5906	(0.1966)	0.9137	(0.0949)	1.2983	(0.0966)
ρ	-0.6448 (n/a)				-0.3128 (0.0859)				0[a]			
$\log L$	-62013.77 ($M = 6$)				-62164.23 ($Q = 8$)				-62166.44 ($Q = 8$)			

	Weibull hazard											
	D				E				F			
	Discrete, dep risks				Gaussian, dep risks				Gaussian, indep risks			
	Filled		Lapsed		Filled		Lapsed		Filled		Lapsed	
Unemployed ≤ 18 ($\log U$)	0.2452	(0.0285)	-0.0701	(0.0387)	0.2566	(0.0281)	-0.0484	(0.0367)	0.2560	(0.0275)	-0.0475	(0.0355)
Job vacancies ($\log V$)	0.0360	(0.0218)	-0.1423	(0.0296)	0.0489	(0.0211)	-0.1303	(0.0273)	0.0489	(0.0211)	-0.1303	(0.0272)
Log wage if $D = 1$	-0.0209	(0.0617)	0.1830	(0.0830)	-0.0243	(0.0595)	0.2020	(0.0766)	-0.0243	(0.0594)	0.2020	(0.0764)
$\alpha_1 - 1, \alpha_2 - 1$	-0.2841	(0.0149)	0.5837	(0.0417)	-0.2904	(0.0218)	0.4836	(0.0430)	-0.2907	(0.0216)	0.4834	(0.0429)
σ_1, σ_2	4.4991	(n/a)	1.0026	(n/a)	0.5835	(0.0794)	0.9814	(0.0806)	0.5781	(0.0622)	0.9771	(0.0700)
ρ	-0.7382 (n/a)				-0.0158 (0.1447)				0[a]			
$\log L$	-62400.76 ($M = 6$)				-62525.31 ($Q = 8$)				-62525.31 ($Q = 8$)			

[*] See Table 3 for the other covariates that are included.

[a] Imposed.

[b] For consistency, both discrete mixing models have same number of mass points. Adding a seventh made little improvement to the log-likelihoods.

Table 3: Multiple orders, discrete mixing, non-parametric hazard (Spec'n G)*

	Filled		*Lapsed*		*Pr(lapsing)*[a]		*Sample means*
	$\widehat{\boldsymbol{\beta}}_1$	st. err.	$\widehat{\boldsymbol{\beta}}_2$	st. err.	$\partial P/\partial x$	st. err.	
Unemployed ≤ 18 ($\log U$)	0.1681	(0.0390)	0.0695	(0.0582)	-0.0063	(0.0057)	334.327
Job vacancies ($\log V$)	-0.1053	(0.0261)	-0.1193	(0.0340)	-0.0009	(0.0032)	52.297
Whether wage negotiable (N)	0.0054	(0.0328)	-0.1382	(0.0527)	-0.0091	(0.0042)	0.235
Log wage if $D = 1$	0.0894	(0.0566)	0.0410	(0.0844)	-0.0031	(0.0069)	£1.383
Log population density	0.0497	(0.0232)	0.1691	(0.0296)	0.0076	(0.0027)	10.398
Firm located in town centre	0.1610	(0.0288)	-0.0470	(0.0474)	-0.0072	(0.0037)	0.426
Log careers Service staff per person	-0.1308	(0.0308)	0.0000	(0.0339)	0.0083	(0.0036)	106.3
11–30 employees	0.0424	(0.0321)	-0.0699	(0.0550)	-0.0071	(0.0043)	0.225
31–100 employees	0.0676	(0.0378)	-0.1022	(0.0650)	-0.0108	(0.0051)	0.159
>100 employees	0.2303	(0.0465)	-0.2848	(0.0698)	-0.0327	(0.0057)	0.161
Firm provides training vacancies	-0.0781	(0.0285)	-0.0749	(0.0460)	0.0002	(0.0036)	0.357
Skilled	0.0802	(0.0506)	0.0570	(0.0693)	-0.0015	(0.0058)	0.527
Non-manual	-0.3164	(0.0511)	0.1105	(0.0708)	0.0271	(0.0059)	0.530
In house training	0.0030	(0.0516)	-0.0365	(0.0905)	-0.0025	(0.0071)	0.061
Day release training	-0.3648	(0.0444)	-0.3993	(0.0707)	-0.0022	(0.0056)	0.094
Apprenticeship training	-0.4270	(0.0640)	-0.3045	(0.0953)	0.0078	(0.0077)	0.146
Average GCSE or just below	-0.1305	(0.0339)	-0.1218	(0.0574)	0.0006	(0.0045)	0.451
High GCSE	-0.3188	(0.0492)	-0.3351	(0.0797)	-0.0010	(0.0063)	0.188
4 or more GCSEs	-0.4880	(0.0647)	-0.6507	(0.1043)	-0.0103	(0.0082)	0.092
English required	0.0915	(0.0634)	-0.1778	(0.0938)	-0.0171	(0.0075)	0.630
Maths required	0.0452	(0.0835)	-0.2775	(0.1303)	-0.0205	(0.0104)	0.044
English and Maths required	0.0351	(0.0858)	-0.2399	(0.1379)	-0.0175	(0.0108)	0.039
Science required	-0.0534	(0.0671)	-0.1738	(0.1034)	-0.0077	(0.0082)	0.189
Other subject required	0.1200	(0.0852)	-0.3337	(0.1329)	-0.0288	(0.0106)	0.053
Older applicants required (over 16)	-0.0274	(0.0365)	-0.0591	(0.0530)	-0.0020	(0.0043)	0.183
Written application required	-0.6846	(0.0423)	-0.7857	(0.0650)	-0.0064	(0.0052)	0.178
$\bar{u}_{11}$, $\bar{u}_{21}$, π_1	1.6254	(0.0778)	-2.6240	(0.1504)	0.5220	(0.0089)	
$\bar{u}_{12}$, $\bar{u}_{22}$, π_2	-2.3232	(0.2725)	5.3038	(0.1382)	0.2912	(0.0133)	
$\bar{u}_{13}$, $\bar{u}_{23}$, π_3	-2.3884	(0.2038)	3.1386	(0.0560)	0.0840	(0.0076)	
$\bar{u}_{14}$, $\bar{u}_{24}$, π_4	0.2264	(0.1337)	-27.5917	(1.1230)	0.0263	(0.0037)	
$\bar{u}_{15}$, $\bar{u}_{25}$, π_5	0.2959	(n/a)	3.7332	(n/a)	0.0767	(n/a)	
Log L		-37345.89					
σ	1.8540	(n/a)	5.8024	(n/a)			
ρ		-0.6063					
Vacancies	11485		6274				
Orders	8548		5962				

* Specification G. Discrete mixing ($M = 5$), dependent risks, non-parametric baseline hazard (10 pieces). Also includes dummies for SIC (9), year (7) and month (11). Mean duration to filling E_1 is estimated as 73 days, mean duration to lapsing E_2 as 143 days, and, P, the conditional-on-exit probability that a vacancy fills, as 0.870. See Equation (14). Hence mean duration to exit is 82 days.

[a] Marginal effect on the probability of lapsing, evaluated at mean duration of 82 days. See Equation (16). p-value assumes that h_1 and h_2 are non-random. It can be shown that, because $h_1 \approx h_2$ at mean duration, this assumption is innocuous.

ACKNOWLEDGMENTS

The authors thank The Leverhulme Trust (under grant F/120/AS) for financial assistance. We used Sabre 4.0 for the estimation of the models with dependent risks, which was developed by the Centre for e_Science, Lancaster University and funded under grant ESRC award RES-149-25-0010. The data were kindly supplied by Lancashire Careers Service. The comments of Len Gill, Jonathan Wadsworth, Alan Manning and especially Chris Orme 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.