Estimating the stock-flow matching model using micro data

Martyn Andrews
University of Manchester

Steve Bradley
Lancaster University

Dave Stott
Lancaster University

Richard Upward
University of Nottingham*

For submission to *Journal of the European Economic Association*
July 2009

*The authors thank The Economic and Social Research Council (under grant R000239142) for financial assistance. The data were kindly supplied by Lancashire Careers Service and are also available from the UK Data Archive (`http://www.data-archive.ac.uk`). The comments from participants at various presentations are gratefully acknowledged. These include the Economics departments at Dundee, Erlangen, Leicester, Manchester, Nottingham, as well as EALE 2002 (Paris), ESEM 2003 (Stockholm) and RES 2003 (Warwick). We are grateful to Barbara Petrongolo for helpful clarifying conversations.

Abstract

We estimate the stock-flow matching model of Coles and collaborators using micro-level data of a well-defined labour market. Using a dataset of complete labour-market histories for both sides of the market, we estimate hazard functions for both job-seekers and vacancies. We find that the stock of new vacancies has a significant positive impact on the job-seeker hazard, over and above that of the total stock of vacancies. There is an equivalent robust result for vacancy hazards. Thus we find evidence in favour of stock-flow matching. Replacing old/new stocks by good/bad stocks does not provide a competing explanation of the data. [100 words]

Keywords: two-sided search, matched job-seeker/vacancy data, stock-flow matching, hazards
New JEL Classification: C41, E24, J41, J63, J64

Address for Correspondence:

Dr. M.J. Andrews
School of Social Sciences
University of Manchester
Manchester, M13 9PL

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

1 Introduction

This paper is an empirical investigation into how job-seekers and employers meet and match each other. The dominant model in the literature is one of friction and congestion: agents on both sides of the market take time to find a suitable partner. Pissarides' (2000) text (first published in 1990) is the original two-sided search model applied to the labour market. This model, and others like them, (see, in particular Burdett & Wright (1998)), incorporate many of the same basic structures and assumptions, as surveyed by Burdett & Coles (1999). Because the process by which agents meet each other is random, these classical two-sided models of search are referred to as random matching models.

A recent alternative view is that matching occurs via a marketplace. In the marketplace, agents can search the other side of the market in a short period of time. This will be relevant whenever there are technologies such as newspapers, employment agencies or the internet which allow simultaneous posting of vacancies or job-seekers. If an agent, say a newly unemployed job-seeker, searches the market and fails to find a match, she enters the stock of unemployed job-seekers and can then only match with the flow of new vacancies entering the marketplace. Symmetrically, employers enter the marketplace with vacancies, which they either fill, or the vacancy increases the stock. Thus, most matches in this model occur between the stock on one side of the market and the inflow on the other, which is why this alternative model is known as the stock-flow matching model (and might also be thought of as a specific form of a non-random matching model). It was originally developed by Melvyn Coles and Eric Smith (Coles & Smith 1998), although there are other papers that do not use the stock-flow terminology (Taylor 1995, Lagos 2000).

Does the stock-flow model provide a more realistic or more useful description of how labour markets operate? There are three related aspects to this question. First, does the stock-flow matching model actually provide a better description of observed behaviour of job-seekers and employers? In particular, are search hazards better characterised by the stock-flow matching model? Second, does the stock-flow matching model provide a better description of the co-movements of aggregate labour-market variables (the stocks and flows in and out of employment, unemployment and vacancies)? For example, Ebrahimy & Shimer (2009) calibrate the stock-flow matching model to the U.S. labour market and conclude that it explains much of the co-movement of aggregate labour-market variables. This turns out to be very similar to the co-movements predicted by Shimer's (2007) matching model, but are more volatile than those predicted by Pissarides's (1985) random matching model.[1] This is clearly important because the random matching model describes the labour market in many macroeconomic models.

Third, what are the policy implications of the stock-flow matching model? Coles & Petrongolo (2008) argue that these are important, particularly in the context of the large optimal unemployment insurance literature which argues that, with search frictions and unobserved search effort, unemployment benefit payments should be reduced with unemployment duration to encourage greater search effort. This is because search effort, in the random matching model, is a productive investment. In the stock-flow matching model, a job seeker who fails to find a match immediately must chase the flow of new vacancies when they come onto the market. This is pure rent-seeking behaviour because the vacancy will eventually get filled. Now, increasing search intensity has no effect on equilibrium unemployment, and a policy of reducing unemployment benefits to shorten unemployment durations is welfare reducing. Coles & Petrongolo conclude that the policy issue in the stock-flow matching model is to determine the optimal level of benefits.

The stock-flow matching model is more consistent with frictions that arise from market failure in occupational or regional segments of markets, which suggests that regional policies that move employers closer to workers, or stimulate small-firm formation to absorb the pool of unemployed, might be appropriate. The stock-flow matching model also has implications for firms who face skill shortages, a perennial problem in many economies, including the U.K.. Here, policies which lead to better or more suitable training for workers to reduce occupational mismatches would be appropriate.

There is no previous evidence on the stock-flow matching model using micro-level data; the only evidence comes from aggregate time-series data. Coles & Smith (1998) estimate job-seeker hazards using monthly aggregate time-series Job Centre data between 1987 and 1995 for the U.K.; their findings are strongly supportive of the theory. Gregg & Petrongolo (2005) use similar data and come to similar conclusions. Using similar data again, Coles & Petrongolo (2008) find evidence of one-sided stock-flow matching, whereby the stock of unemployed match with the inflow of vacancies, but not *vice versa*.

Our main contribution is to specify and implement a test of random matching as a special case of stock-flow matching using data on job-seekers and vacancies from the same market. We observe matches between job-seekers and vacancies and we observe how long each agent has been in the market when they match.

[1] See also Pissarides's (2008) interview in the Newsletter of *The Review of Economic Dynamics*, when asked to defend the Mortensen-Pissarides matching function (Mortensen & Pissarides 1999) given Shimer's recent findings.

We also observe who matches with whom. These high frequency agent-level data are superior to those previously used. We estimate the hazards of exit from the marketplace for both job-seekers and employers. Thus we can control for observed and unobserved heterogeneity and we can control for aggregation bias, a well-known potential problem with studies that use aggregate data (Burdett, Coles & van Ours 1994). With aggregate data, one cannot model the essential feature of this type of search model, that of individual agents changing behaviour in response to changing aggregate labour market conditions during the agents' period of search.

The paper is organised as follows. We first present stylised versions of both the random matching model and the stock-flow matching model and then describe how we can test between them. In Section 3, we describe the data and show how they are used to construct the key variables in the stock-flow matching model. Section 4 sets out the econometric methodology. We discuss our results and report robustness checks of our preferred model in Section 5. This includes seeing whether replacing our stocks of old and new agents by good and bad agents alters our conclusions, because it might be that the stock-flow model is picking up some heterogeneity whereby certain agents exit the marketplace quicker than others. Section 6 concludes.

2 The Stock-flow Matching Model

2.1 Stylised Stock-flow and Random Matching Models

The dominant model in the literature is one of friction and congestion: agents on both sides of the market take time to find a suitable partner.[2] The cornerstone of the model is the matching function $m = m(U, V)$, where m is the per-period matching rate between the beginning-of-period stocks of vacancies V and job-seekers U (all of whom are assumed unemployed). $m(\ )$ is assumed concave and increasing in both arguments, and is often specified as exhibiting constant returns to scale.[3] On average, an unemployed job seeker matches with a vacancy with probability of $h^w = m(U, V)/U$ per period, where h^w is the job-seeker hazard, or average number of matches per job-seeker. An increase in the stock of unemployed U has a negative effect on $h^w(U, V)$ because there are more unemployed with which job-seekers have to compete, whereas an increase in the stock of vacancies V has a positive effect because there are more vacancies on the market to search over.

There is much empirical support for the random matching model, often using

[2]This section draws heavily on Petrongolo & Pissarides (2001).

[3]Other restrictions on the matching function are $m(0, V) = m(U, 0) = 0$ and $m(U, V) \leq \min(U, V)$.

a Cobb-Douglas specification $m(U,V) = U^{\alpha}V^{\beta}$, because of the significant effect of the stocks of U and V on the exit probability from unemployment. Petrongolo & Pissarides (2001) provide a comprehensive survey of the evidence; most papers either report estimates of the matching function that use aggregate flows data or report estimates of the hazard from unemployment that use individual-level duration data.

The key feature of the random matching model is frictions, which arise from information imperfections, heterogeneities, absence of perfect insurance markets, congestion from large numbers, and so on. Further, job-seekers are assumed to randomly apply for vacant jobs, which ignores any systematic element to their searching. There are many attempts to provide some microfoundations for what is essentially a black-box relationship; most of these involve the idea that some contacts results in matches and others don't.

By contrast, there are no search frictions in the stock-flow matching model because job-seekers and employers are assumed to search the whole market in a short period of time. Unemployment and vacancies persist because suitable partners were not available on this first search of the market, and so job-seekers and employers have to wait for new opportunities to flow into the market at a later date.

Stock-flow models are usually specified in continuous time. Suppose that agents arrive on the market at flow rates u and v. The initial stock of unemployed U matches with the inflow v only, whereas the inflow u matches with both V and v. However, because the initial period is infinitesimally small, flow-flow matches between u and v cannot happen. The probability that a new job-seeker matches on entry is $1-(1-\mu)^V$, where μ is the probability that a random pairing is acceptable, and hence the hazard between the flow of job-seekers and the stock of vacancies, defined as $h^w_{12}(V)$, is written $1-(1-\mu)^V$.[4] This is because a new job-seeker has to be rejected by all V vacancies on the market for it not to form a match. The number of new matches is $u[1-(1-\mu)^V]$. By symmetry, the number of new matches due to the entry of vacancies is $v[1-(1-\mu)^U]$, and so the hazard for workers unemployed at the beginning of the period is $h^w_{21}(U,v) = v[1-(1-\mu)^U]/U$, because there are $v[1-(1-\mu)^U]/U$ matches per unemployed job-seeker. The hazard for a job-seeker drops when he joins the stock because v is likely to be smaller than U.

In discrete time, the model is modified as follows. First, because the initial period is of finite length, the new unemployed can match with the new vacancies, and so the hazard between the two can be written $h^w_{11}(u,v)$, where the effects of u and v are negative and positive respectively, as in the random matching model. Second,

[4]This notation is used throughout the paper. The first subscript is 1 if the job-seeker belongs to the flow and is 2 if job-seeker belongs to the stock. The second subscript refers to vacancies.

the stock of vacancies with which new job-seekers can match is $\bar{V} \equiv V - v$, the stock of old vacancies, rather than V as in the continuous time model. In short, we write this hazard as $h^w_{12}(u, \bar{V})$. By symmetry, the hazard for old job-seekers matching with new vacancies is $h^w_{21}(\bar{U}, v)$, where $\bar{U} \equiv U - u$.

Note that the model, specified in either continuous or discrete time, predicts that old job-seekers should not match with old vacancies, because, if there were gains to trade, they would have matched in an earlier period. In reality, it is possible that we observe such matches, and so we write the hazard for old job-seekers matching with old vacancies as $h^w_{22}(\bar{U}, \bar{V})$, where again the effects of $\bar{U}$ and $\bar{V}$ are negative and positive respectively, as in the random matching model. This is because, for example, job-seekers, having entered the old stock, might revise down their reservation utilities, and so re-examining the stock might then reveal potential matches. If there are large numbers of old job-seekers and old vacancies in the market, stock-stock matches might still happen, even though these matches occur with a low probability.

All of the above discussion about the stock-flow model can be repeated from the employer's point of view, based on the four hazards: $h^e_{11}(u, v)$, $h^e_{12}(u, \bar{V})$, $h^e_{21}(\bar{U}, v)$ and $h^e_{22}(\bar{U}, \bar{V})$.

2.2 An Estimable Stock-flow Matching Model

The stock-flow matching model just described is highly stylised. In this subsection, we develop a statistical model of stock-flow matching, that also nests random matching as a special case. Because we estimate this model, it is specified in discrete time, with the period being a week.

In general, it would be natural to model the total number of contacts between the stock of job-seekers and vacancies as a Poisson process

$$C \sim \text{Poisson}[\lambda(U, V)],$$

where the Poisson parameter λ is specified as a Cobb-Douglas: $\lambda(U, V) = AU^{\alpha}V^{\beta}$. $\lambda(U, V)$ is the average number of contacts per week.

To capture the idea that contacts might be non-random, as in the stock-flow model, we allow the parameter A to vary according to the age of the job-seeker and the vacancy. For example, the average number of contacts per week between an old job-seeker and a new vacancy is given by

$$\lambda_{21} = \left(\frac{\bar{U}v}{UV}\right) A_{21} U^{\alpha} V^{\beta}.$$

Multiplying by $\bar{U}v/UV$ gives the mean number of contacts between old job-seekers and new vacancies because this is the proportion of this type of contact out of all contacts. Similar expressions apply for λ_{11}, λ_{12} and λ_{22}. Under random matching contacts are equally likely to occur regardless of the age of the job-seeker or vacancy, and we have $A_{11} = A_{12} = A_{21} = A_{22} = A$. Under stock-flow matching, old job-seekers only contact new vacancies and so $A_{21} > A_{22} = 0$. Similarly old vacancies only contact new job-seekers and so $A_{12} > A_{22} = 0$.

We also assume that the matching probabilities, conditional on contacting, vary across pairs of agents. Suppose that an old job seeker contacts a new vacancy. The probability that they match is given by μ_{21}, and the probability that they carry on searching is $1 - \mu_{21}$. The same applies to the other three types of match, with matching probabilities given by μ_{11}, μ_{12} and μ_{22}.

Because there are four types of match, there are four different hazard rates out of unemployment. These are given by:

$$h^w_{11} \equiv \mu_{11}\lambda_{11}/u = a_{11}vU^{\alpha-1}V^{\beta-1} \tag{1}$$
$$h^w_{12} \equiv \mu_{12}\lambda_{12}/u = a_{12}\bar{V}U^{\alpha-1}V^{\beta-1} \tag{2}$$
$$h^w_{21} \equiv \mu_{21}\lambda_{21}/\bar{U} = a_{21}vU^{\alpha-1}V^{\beta-1} \tag{3}$$
$$h^w_{22} \equiv \mu_{22}\lambda_{22}/\bar{U} = a_{22}\bar{V}U^{\alpha-1}V^{\beta-1}, \tag{4}$$

where $a_{ij} \equiv A_{ij}\mu_{ij}$.[5] There are another set of hazards for vacancies, labeled h^e_{11}, h^e_{12}, h^e_{21} and h^e_{22}.

Consider the expression for h^w_{21}, the exit rate for an old job-seeker who matches with a new vacancy. $\lambda_{21}(U, V)$ gives the mean total number of contacts per period between old job-seekers and new vacancies. Multiplying by μ_{21} gives the mean total number of matches per period of this type; dividing by $\bar{U}$ gives the hazard rate h^w_{21}, or the mean number of matches per old job-seeker.

The key prediction of the stock-flow matching model is that $a_{12} > a_{22}$ and $a_{21} > a_{22}$. The pure stock-flow model predicts that $a_{22} = 0$. However, it is reasonable to acknowledge the possibility that old-old contacts could happen simply because there are large numbers of old stocks $\bar{U}$ and $\bar{V}$ in the market. It is because the parameter a_{22} is relatively small that makes old-old matches infrequent. If the matching probability μ_{22} is zero (there are contacts, but none are acceptable) or A_{22} is zero (there are no contacts), then we have the pure stock-flow matching model. The stock-flow matching model is silent about how a_{22}, a_{12}, and a_{21}, compare with

[5] Only with data on contacts and matches would be able to separately identify A_{ij} and μ_{ij}. It is also possible to specify μ_{ij} as a Cobb-Douglas function of U and V, but its parameters are not identified for the same reason.

a_{11}, but we expect a_{11} to be positive with discrete data.

The hazard for a new job-seeker, $h^w_{1.} = h^w_{11} + h^w_{12}$, comes from summing Equations (1) and (2). The hazard for an old job-seeker $h^w_{2.} = h^w_{21} + h^w_{22}$, comes from summing Equations (3) and (4).

Random matching is a special case when

$$H_0 : a_{11} = a_{12} = a_{21} = a_{22} \quad (= a, \text{ say}), \tag{5}$$

is true. Under H_0, the log-hazard for new and old job-seekers is simply

$$\log h^w = \log(a) + (\alpha - 1)U + \beta V + \epsilon^w,$$

where ϵ^w captures unobserved job-seeker heterogeneity. This is the standard hazard function for the random matching model with unobserved heterogeneity ϵ^w, and shows that the hazard is constant over the job-seeker's spell of unemployment. Furthermore, the hazard is only dependent on the total stock of job-seekers U and vacancies V; it does not depend on the proportion of new and old in the stock.

In contrast, under the stock-flow matching model, the log-hazard for new job-seekers is:

$$\begin{aligned}\log h^w_{1.}&(U, u, V, v, \epsilon^w) \\ &= \log[a_{11}v + a_{12}(V - v)] + (\alpha - 1)\log U + (\beta - 1)\log V + \epsilon^w,\end{aligned} \tag{6}$$

and the log-hazard for old job-seekers is:

$$\begin{aligned}\log h^w_{2.}&(U, u, V, v, \epsilon^w) \\ &= \log[a_{21}v + a_{22}(V - v)] + (\alpha - 1)\log U + (\beta - 1)\log V + \epsilon^w.\end{aligned} \tag{7}$$

This shows that if a_{22} is very small, as predicted by the stock-flow matching model, then the hazard for a job-seeker will fall once they have searched the market, because the flow of new vacancies (v) is likely to be much smaller than the stock of old vacancies ($\bar{V}$). Furthermore, the hazard for new and old job-seekers should respond differently to changes in the proportion of new and old vacancies on the other side of the market. If $a_{12} > a_{22}$, an increase in $\bar{V}$ should have a stronger effect on the hazard for new job-seekers than for old job-seekers.

Rather than estimate the non-linear models in Equations (6) and (7), it is much easier to linearise the model. Consider the model for $\log h^w_{2.}$, from which we can uniquely identify α, β and a_{22}/a_{21}:

$$\begin{aligned} \frac{\partial \log h^w_{2.}}{\partial \log U} &= \alpha - 1 \qquad \frac{\partial \log h^w_{2.}}{\partial \log V} = \frac{a_{22}V}{a_{21}v + a_{22}\bar{V}} + \beta - 1 \equiv \pi_1 \\ \frac{\partial \log h^w_{2.}}{\partial \log u} &= 0 \qquad \frac{\partial \log h^w_{2.}}{\partial \log v} = \frac{(a_{21} - a_{22})v}{a_{21}v + a_{22}\bar{V}} \equiv \pi_2. \end{aligned} \tag{8}$$

The estimates from the linear model can be interpreted as follows. First, an increase in the stock of unemployed job-seekers U has the familiar effect of $\alpha - 1$, and it does not matter whether the congestion comes from old or new job-seekers, which is why the extra effect from new job-seekers u is zero. Second, to obtain an estimate of β, one adds together the estimates on $\log V$ and $\log v$ (ie $\pi_1 + \pi_2 = \beta$). Third, the coefficient on v should be zero if the random matching model is true, because $a_{21} = a_{22}$ implies $\pi_2 = 0$. This is a one-sided test because, under the alternative, $\pi_2 > 0$. It is important to understand what is happening in the stock-flow model when $\pi_2 > 0$. Suppose that the stock of new vacancies v goes up whilst the stock of all vacancies V remains fixed, which means that the stock of old vacancies $\bar{V}$ falls. Under random matching, this switch between old and new has no effect on the hazard. Under stock-flow matching, v going up leads to more stock-flow matches but $\bar{V}$ going down means fewer stock-stock matches. The net effect is positive if $a_{22} < a_{21}$. The same can be seen from the estimate of a_{22}/a_{21}, obtained directly from the expression for π_2, which is given by

$$\frac{a_{22}}{a_{21}} = \frac{v}{V(1 - \pi_2)^{-1} - \bar{V}}. \tag{9}$$

If $\pi_2 > 0$, ie the effect of v is significant and positive, then $a_{22}/a_{21} < 1$.[6]

Analogous discussions apply when we use data on vacancy spells. Here we estimate $\log h^e_{.2}(U, u, V, v)$ and see whether u is significant. This tests whether $a_{12} = a_{22}$ and an estimate of a_{22}/a_{12} is obtained from:

$$\frac{a_{22}}{a_{12}} = \frac{u}{U(1 - \pi_2)^{-1} - \bar{U}}, \tag{10}$$

where π_2 now refers to the estimate on u.

Before we describe our econometric methods for estimating job-seeker and vacancy hazards, and discuss associated identification issues, we must first describe the data we use.

[6]Testing whether v is significant and positive makes a simple intuitive test. This arises because we sum the hazards h^w_{21} and h^w_{22}. If we did not do this, we would have to estimate the hazards separately, and test (across equations) whether the baseline hazards were the same.

3 The Data

3.1 The Lancashire Careers Service data

The data we use are the computerised records of the Lancashire Careers Service (LCS) over the period March 1988 to June 1992. The Careers Service was a Government-funded network which operated a free matching service for employers and youths.

The data comprise a longitudinal record of all youths in Lancashire aged 15–18, including those in education, employment, training and unemployment. For each job-seeker, we observe the start date of every labour market spell over the sample period. Some spells are right-censored. The data also include a record of all vacancies notified to the Careers Service over the sample period: again, we observe completed and right-censored spells of vacancies placed on the market. Thus our job-seeker spells and vacancy spells comprise standard flow samples. Approximately 20% of all job spells observed in the data resulted from a match with a vacancy posted with the Careers Service. Vacancies for which the Careers Service were not the method of search are not included in the data. However, in the youth labour market (in contrast to the adult labour market) vacancies posted with the Careers Service were generally representative of all vacancies available for this age group.[7]

Job-seekers are observed in one of four labour market states: unemployment, employment, government-sponsored training or education. Vacancies are either posted with the Careers Service or not. Each job-seeker therefore has four possible outcomes: they can match with a vacancy posted with the Careers Service, they can match with a vacancy not recorded in our data, they can withdraw from the labour market, or their unemployment spell can be censored by the end of the sample period. A vacancy has six possible outcomes: it can match with a job-seeker from one of the four possible labour market states, it can be withdrawn from the market, or it can be censored.

We analyse matches between job vacancies and unemployed job-seekers. Matches involving school-leavers and those on training programmes are less relevant for the purpose of testing theories of labour market matching. A spell which ends in a different kind of match is treated as censored. For example, a job-seeker finding a match from outside the Careers Service data, or a vacancy which matches a job-seeker who is still in education. The important point to note here is that we do not ignore these spells: they are included in the risk set up to the point where they are

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

censored, as is standard in all competing risks models.

Also note that we need to consider other types of job-seeker when specifying the arguments of the matching function, because it might be the case that the stock of those engaged in on-the-job search affects the probability of a match between unemployed job-seekers and vacancies because they are competing for the same vacancies. We therefore use two definitions of job-seekers. The first, narrow definition refers only to unemployed job-seekers. The second, wide definition includes those who are on training programmes and those who are in jobs, and who are registered as actively searching with the Careers Service. The narrow definition corresponds more closely to the existing literature.

In our data we are able to distinguish between the 'search duration' and the 'spell duration' for each agent. Search duration is ended when a successful contact between a job-seeker and a vacancy is recorded in the data. Spell duration is ended when a job-seeker actually starts working in a new job, which is typically some time after a successful contact. These durations form the dependent variables for estimating job-seeker hazards h^w and vacancy hazards h^e. Our preferred specification focuses on the duration of search, because this corresponds more closely to the theory. However, since almost all existing estimates of the matching function are forced to use spell durations, we also examine what happens when we use spell duration.

The point at which a job-seeker or a vacancy, when in the marketplace, changes from being new to old, is defined as k^w for job-seekers and k^e for vacancies. We refer to the first k^w and k^e weeks of a spell of search as the matching 'window'. This represents the period over which the current stock of potential partners on the other side of the market can be searched.

3.2 The Dependent Variable

The data are organised into sequential binary response form. This means that we pool over all the job-seekers in the data, and generate an unbalanced panel of job-seeker spells with t_i^w observations for each spell i. Each row in this panel corresponds to a job-seeker-week, of which there are 477,868. This defines the risk set for job-seekers. The total number of job-seeker spells is 34,657. Some job-seekers have multiple spells; the total number of job-seekers is 26,113.[8]

The dummy variable y_{is}^w indicates whether the i-th job-seeker spell ends with a match in week s. In other words, we have a sequence of observations $y_{is}^w, s =$

[8] 19,520 job-seekers have 1 spell accounting for 1,313 matches; 5,063 job-seekers have 2 spells (848 matches); 1,210 job-seekers have 3 spells (377 matches); and the rest have four or more spells (223 matches).

$1, \ldots, t_i^w$, all of which are zero except the last. For the last observation ($s = t_i^w$), if the job-seeker matches with a vacancy, $y_{is}^w = 1$ and if the job-seeker spell is censored, $y_{is}^w = 0$.

Analogous considerations apply to the j-th vacancy spell. Pooling over all the vacancies in the data, we generate an unbalanced panel of vacancies with t_j^e observations for each vacancy spell j. The risk set for this side of the market is 137,223 vacancy weeks, corresponding to 14,154 vacancy spells. Some employers have multiple vacancy spells. This is because vacancies are often bundled together in so-called vacancy orders and some employers are observed more than once. The total number of vacancy orders is 9,556 and the total number of employers is 4,121.[9]

Summing over y_{is}^w in the job-seeker panel and over y_{js}^e in the vacancy panel gives the total number of matches in the data, m, which is 2,761.

Table 1 summarises the raw data for both panels, using the 'search' definition of duration, under the assumption that $k^w = k^e = 4$ weeks. There are 477,868 job-seeker-weeks at risk, of which there are 2,761 matches and the rest where there are no matches. The 2,761 matches are disaggregated by whether the job-seeker spell is old (unemployed $\geq$ 4 weeks) or new (unemployed $<$ 4 weeks) and by whether or not the job-seeker exits to an old or new vacancy. Thus, for example, there are $m_{21} = 1,497$ cases where a job-seeker who has been unemployed for more than 4 weeks matches with a vacancy which has been open for less than 4 weeks. The four types of match add up to the total $2,761$.

In the lower panel, the 137,223 vacancy-weeks are disaggregated in the same way. There are more unemployment job-seeker-weeks at risk than there are vacancy-weeks because the youth labour market in the UK in the early 1990s was particularly slack: compare the 34,657 job-seeker spells with the 14,154 vacancy spells.

Using the cross-tabulations in Table 1, we compute the raw hazard for job-seeker spells which match with old vacancies:

$$h_{12}^w = 278/124,148 = 0.002239$$
$$h_{22}^w = 454/353,720 = 0.001284.$$

Notice that the drop in the raw hazard is $h_{22}^w / h_{12}^w = 0.573$, which is perfectly consistent with stock-flow matching.

[9] 2,123 employers have 1 spell accounting for 287 matches; 790 employers have 2 spells (332 matches); 406 employers have 3 spells (267 matches); and the rest have four or more spells (1,775 matches).

Table 1: Who matches whom? Search duration; $k^w = k^e =$ 4 weeks

	new	old	total
Job-seekers[a]			
zeros	123,338	351,769	475,107
exits to new vacancy	(m_{11}) 532	(m_{21}) 1,497	2,029
exits to old vacancy	(m_{12}) 278	(m_{22}) 454	732
Total	124,148	353,720	477,868
Vacancies[b]			
zeros	38,547	95,915	134,462
exits to new job-seeker	(m_{11}) 532	(m_{12}) 278	810
exits to old job-seeker	(m_{21}) 1,497	(m_{22}) 454	1,951
Total	40,576	96,647	137,223

[a] 477,868 job-seeker weeks correspond to 34,657 job-seeker spells and 26,113 job-seekers.

[b] 137,223 vacancy weeks correspond to 14,154 vacancy spells, 9,556 vacancy orders, and 4,121 employers.

A similar analysis applies to vacancies. It turns out that the drop in the hazard for vacancy spells matching with old job-seekers is $h^e_{22}/h^e_{21} = 0.127$. This is also consistent with stock-flow matching, and is more pronounced on this side of the market.

Table 1 can be recomputed for any values of k^e and k^w. In Figure 2 we plot the numbers of stock-stock, stock-flow, and flow-flow matches against window size, but keeping $k^w = k^e$. It is obvious that the number of flow-flow matches *must* increase with $k^w = k^e$ and that the number of stock-stock matches *must* decrease. But the number of stock-stock matches is never zero, and so a pure form of the theory does not occur in these data. The number of stock-flow matches $m_{12} + m_{21}$ increases with window size, and then decreases. Notice that the number of stock-flow matches is largest when the window size is $k^w = k^e = 4$ weeks. This is where the number stock-stock and flow-flow matches added together is minimised.

It is also possible that $k^w \neq k^e$. For example, employers, who have been in the market in previous years, know exactly what kind of job-seeker they are looking for. Young job-seekers, on the other hand, are relatively inexperienced and might only have a vague idea about what they want, and thus searching the market takes longer, on average. In other words, $k^w > k^e$. Our strategy when we come to estimation is therefore to choose a small number of (k^w, k^e) pairs, such that $k^w \geq k^e$, to see whether it makes any difference to the results. Finally, we rule out the possibility that k^w varies over job-seekers and that k^e varies over vacancies; this is because job-seekers and vacancies are using the same matching technology (the Careers Service).

3.3 The Explanatory Variables

The explanatory variables in the stock-flow model are measures of the 'stocks' and 'flows' of job-seekers and vacancies. In practice, the 'flows' are actually stocks of new job-seekers and vacancies, while the 'stocks' are stocks of old job-seekers and vacancies. As with the dependent variable, the sizes of these new and old stocks will depend on the values of k^w and k^e.

During a given week $t-1$ in calendar time, there is an inflow of job-seekers u_{t-1}^{+} into the stock of job-seekers U_{t-1}, and an outflow u_{t-1}^{-}, such that

$$U_t = U_{t-1} + (u_{t-1}^{+} - u_{t-1}^{-}). \tag{11}$$

Unfortunately, these job-seeker data are a flow sample, which means that U_t is not observed. However, we observe data for thirty weeks before the sample period, and so U_t is built up recursively from the net inflow into unemployment $u_t^{+} - u_t^{-}$ each period. In other words, U_{-30} is set to zero. We have checked that our imputed measure essentially coincides with the equivalent measure of U_t from official unemployment stock data from October to April from 1989 onwards. These are from the National Online Management Information Service (NOMIS), but cannot be disaggregated into old and new stocks.[10] The vacancy stock data are from a stock sample and so we observe V_t for all t.

If $k^w = k^e = 1$ week, then both U_t and V_t are disaggregated into old and new stocks as follows (with an analogous expression for V_t):

$$U_t = [u_{t-1}^{+} - u_{t-1}^{-} | u_{t-1}^{+}] + [U_{t-1} - u_{t-1}^{-} | U_{t-1}] \equiv u_t + \bar{U}_t. \tag{12}$$

The new stock u_t of unemployed is defined as the inflow of unemployed during the week less those who also exit during the week from the inflow, namely $u_{t-1}^{+} - u_{t-1}^{-} | u_{t-1}^{+}$. Similarly, the old stock $\bar{U}_t$ is defined as the stock of unemployed at the end of the previous week less those who also exit during the current week from the stock at the beginning of the week, namely $U_{t-1} - u_{t-1}^{-} | U_{t-1}$. The above expression generalises for any window size k:

$$U_t = \left[\sum_{i=1}^{k} u_{t-i}^{+} - \sum_{i=1}^{k} u_{t-i}^{-} | \sum_{i=1}^{k} u_{t-i}^{+}\right] + \left[U_{t-k} - \sum_{i=1}^{k} u_{t-i}^{-} | U_{t-k}\right] \equiv u_t^k + \bar{U}_t^k.$$

The data cover the whole of Lancashire, a county in the United Kingdom that

[10] Official unemployment stock data come from NOMIS (National Online Manpower Information Service). We cannot use these data as our measure of U_t because they cannot be disaggregated into old and new.

comprises 14 towns/cities. When constructing the covariates U, $\bar{U}$, u, V, $\bar{V}$, and v, in fact we group Lancashire into just three labour markets ('districts': West, Central and East), recognising that job-seekers can travel between certain towns when looking for work. 96% of all matches take place between a job-seeker and vacancy from the same labour market. This number drops to 75% when Lancashire is treated as 14 towns/cities. There are very large peaks in both new and old unemployed stocks, arising from young people leaving school between May and August each year, which, of course is when employers post their vacancies. See Figure 1. There is a similar annual variation in the data for new vacancy stocks, but less pronounced.

3.4 Temporal Aggregation Bias

Temporal aggregation bias is an important issue in this literature, and is discussed at length by Burdett et al. (1994), Gregg & Petrongolo (2005) and Coles & Petrongolo (2008). In the context of monthly data, the problem arises in not observing the instantaneous hiring rate, but rather flows over a discrete period (a month). The assumptions one needs to adjust the stock measures depend on how quickly agents are matching, which itself is being modeled, and so there is a simultaneity bias. Coles & Petrongolo (2008) estimate matching functions using a maximum likelihood technique to deal with this problem. In our data this will not be a problem as we observe weekly flows together with stocks that also vary weekly; had we used daily stocks, the issue would completely disappear. We have checked that using daily data has very little impact on our results below. What we are able to do, specifically, is assess the extent to which using monthly stocks data biases the estimates. Using the *same* flows data, we use two sets of the stocks data: (a) stocks measured weekly, ie the value observed on the Monday of each week and (b) stocks measured monthly, ie the value observed on the first week of the month. This one might label 'pure' aggregation bias. The alternative would be to collapse the flow data into months as well, thereby having both stocks and flows measured monthly. This is not 'pure' aggregation bias as there is additional measurement error in the durations.

4 Econometric Methodology

In this section, we describe how we estimate the hazard to matching on both sides of the market. The econometric framework we use is a reduced-form mixed proportional hazards (MPH) model for spells of search. The MPH framework is widely used in the estimation of reduced-form hazard models, for example models of unemployment duration. We estimate discrete-time versions of these models, using weekly data, because it allows us to estimate the baseline hazard non-parametrically.

Consider the hazard for a new job-seeker spell i, where the job-seeker matches with a new vacancy. Equation (6) describes how the observed covariates and the unobserved heterogeneity affect the hazard. We write this hazard as $h^w_{1.}(\mathbf{x}_{is}, \epsilon^w_i)$, which is defined as the probability that this job-seeker matches at some point between elapsed duration $s-1$ and s, conditional on having survived to $s-1$:

$$h^w_{1.}(\mathbf{x}_{is}, \epsilon_i) = \Pr\{T_i \in [s-1, s) | T_i \geq s\} \quad s = 1, 2, \ldots, t_i \leq k^w. \tag{13}$$

Here T_i is the latent duration of spell i, t_i is the completed duration of spell i, $\mathbf{x}_{is}$ is a vector of observed covariates, usually specified as $[\log u_{is}, \log v_{is}, \log U_{is}, \log V_{is}]$, and ϵ_i is the spell-invariant unobserved heterogeneity term discussed above.[11] Notice that the completed duration does not exceed the matching window k^w, the duration at which a job-seeker becomes old. If a job-seeker does becomes old, we re-write the hazard as

$$h^w_{2.}(\mathbf{x}_{is}, \epsilon_i) = \Pr\{T_i \in [s-1, s) | T_i \geq s\} \quad s = k^w + 1, \ldots, t_i. \tag{14}$$

As noted, to model the effect of covariates on the hazard rate, we make the proportional hazards assumption. Then the precise form of the discrete hazard is given by the complementary log-log link function:

$$\begin{aligned} h_{is} \equiv h^w(\mathbf{x}_{is}, \epsilon_i) = \\ & 1 - \exp\{-\exp[d_i \mathbf{x}_{is} \boldsymbol{\beta}_1 + (1-d_i)\mathbf{x}_{is}\boldsymbol{\beta}_2 + \gamma_s + \epsilon_i]\} \quad s = 1, \ldots, t_i. \end{aligned} \tag{15}$$

We write the model like this because these two hazards are actually estimated as one regression model by pooling the data in the usual way and interacting each covariate with a dummy variable indicating whether the job-seeker is new, $d \equiv 1(s \leq k^w)$, or old, $1 - d \equiv 1(s > k^w)$. The γ_s terms are interpreted as the log of a non-parametric piecewise linear baseline hazard. Each interval corresponds to a week, but, because of data thinning, these are grouped into longer intervals at longer durations by constraining corresponding duration dummies.[12] Note that the random effect refers to a job-seeker, not a job-seeker spell.

The log-linear specification for the hazard developed in Section 2.2 means that

[11]For notational clarity, where possible we drop the superscript w. The subscript i tells the reader that this is a job-seeker hazard; the subscript j is used for a vacancy hazard.

[12]The ten intervals are: (0,1], [1,2), [2,4), [4,6), [6,8), [8,13), [13,26), [26,39), [39,52), [52,∞).

the parameters on the four covariates are easily interpreted. Write either of the two models in Equation (15) as $h^w = 1 - \exp[-\exp(\mathbf{x}\boldsymbol{\beta} + \gamma + \epsilon^w)]$. Then $\log[-\log(1 - h^w)] \approx \log h^w = \mathbf{x}\boldsymbol{\beta} + \gamma + \epsilon^w$. This means that if the model is linear in the logs of the covariates, $\mathbf{x}_{is} = [\log u_{is}, \log v_{is}, \log U_{is}, \log V_{is}]$, we can interpret the parameters as elasticities in the usual way.

van den Berg (2000, Section 5) discusses the identification of the MPH model; essentially, identification requires that the hazard is multiplicative in elapsed duration t, the covariates $\mathbf{x}$ and the heterogeneity term ϵ. However, van den Berg also notes that allowing the effects of the covariates to vary over elapsed duration is a departure from the standard requirements for identification in a single-spell framework: in our model, it is important that the effects of the stocks of job-seekers and vacancies is allowed to have a different impact on the job-seeker hazard when he becomes ‘old’. (This is why Equations (13) and (14) are written out as separate models.) More recently Brinch (2007) shows that very little variation over time in the covariates is required for models that dispense with the mixed proportional hazards assumption. Essentially, because the stocks of job-seekers and vacancies vary through calendar time and they also vary across labour markets, he shows that our model is identified. In fact, the models would still be identified if the data were ‘collapsed’ to aggregate data comprising three weekly time-series with 226 weeks in each (one for each ‘district’ in Lancashire), a dataset that would be similar to those used in many studies surveyed by Petrongolo & Pissarides (2001). As explained, our job-seeker panel has much more variation than that. Nonetheless, all models include monthly dummy variables, to address the potential criticism that the model is identified exclusively through the variation in the stocks of job-seekers over the calendar year. Note that employer and job-seeker control variables are not included; this is because the essence of the models is to see whether individual behaviour responds to aggregate labour market conditions.

As explained in Section 3.2, a standard approach for estimating this model is to expand the data so that each job-seeker spell contributes t_i rows/weeks. It can be shown that the likelihood for spell i is:

$$L_i(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\gamma}, \dots) = \int_{-\infty}^{\infty} \left[\prod_{s=1}^{t_i} h_{is}(.)^{y_{is}} [1 - h_{is}(.)]^{1-y_{is}} \right] dF(\epsilon_i), \tag{16}$$

where $h_{is} \equiv h^w(\mathbf{x}_{is}, \epsilon_i)$ is given in Equation (15).

This is the likelihood for a binary choice random effects model, with a complementary log-log link rather than the more common logit or probit links. In the final week of the spell ($s = t_i$), either $y_{is} = 1$ if the spell of unemployment ends or $y_{is} = 0$

if the spell is censored. Note that, as our data form a flow sample, we do not need to worry about left-truncated spells.

We adopt two approaches for modeling the unobserved heterogeneity. These are: (i) Gaussian mixing and (ii) discrete mixing. The standard argument for using the latter, as advocated by Heckman & Singer (1984), is that it should affect the baseline hazard less severely than if the wrong choice of parametric mixing is made.[13]

All of the above equations apply to the hazards for a vacancy spell, replacing i by j and w by e. The random effect ϵ_j is defined for a vacancy order, not an employer.

5 Results

5.1 Base model

The model whose likelihood for the job-seeker sample is given in Equations (15) and (16) is referred to as the Base Model, and is reported in Table 2. As discussed in Section 3.1, there are two variants because we have two sets of stocks for job-seekers, narrow and wide definitions.

Described simply, our test of stock-flow matching amounts to seeing whether an increase in the number of new vacancies on the market significantly increases the exit probability for old job-seekers. In the old job-seeker hazard, using wide stocks (third column), this effect is estimated as $\frac{\partial \log h^w_{2.}}{\partial \log v} = 0.483$, and is significant. This converts to a point estimate for $a_{22}/a_{21} = 0.225$. An old job-seeker is four times more likely to match with the new vacancy than an old vacancy. A very similar estimate occurs with narrow stocks (first column). On the other side of the market, for the narrow definition, almost the same effect is detected: the effect of $\log u$ in the old vacancy hazard, $\frac{\partial \log h^e_{.2}}{\partial \log u}$ is 0.389, which converts to $a_{22}/a_{12} = 0.284$ (second column). When using the wider definition, the effect of $\log u$ is stronger, with $\frac{\partial \log h^e_{.2}}{\partial \log u} = 0.613$ and $a_{22}/a_{12} = 0.062$. Here, an old vacancy is now fifteen times more likely to match with the new job-seeker than an old job-seeker, rather than just four times, which makes obvious sense given there are three times more old job-seekers in the wider definition (Table 2, Tablenotes a,b).

[13]For Gaussian mixing, the number of quadrature points is denoted Q. Its value is determined by the investigator: the model is estimated with $Q = 8, 16, 24, \ldots$ quadrature points until the likelihood stops improving. For discrete mixing, the number of mass points, denoted M, is determined by using the Akaike Information Criterion (AIC). See Gaure, Roed & Zhang (2007). The model is fitted with $M = 2$ and $M = 3$. If $M = 2$ is the preferred model, estimation stops. Otherwise we compare $M = 3$ with $M = 4$, and repeat. Inference is conducted conditional on M.

Table 2: Base Model*

	Narrow stocks[a]		*Wide stocks*[b]	
	Job-seekers, h^w	*Vacancies, h^e*	*Job-seekers, h^w*	*Vacancies, h^e*
(a) New				
$\log u$	-0.162 (0.080)	0.075 (0.080)	-0.187 (0.064)	0.226 (0.068)
$\log U$	-0.075 (0.096)	0.599 (0.090)	-0.045 (0.068)	0.460 (0.072)
$\log v$	0.410 (0.104)	-0.200 (0.087)	0.409 (0.104)	-0.225 (0.088)
$\log V$	0.205 (0.083)	0.012 (0.074)	0.213 (0.082)	-0.022 (0.073)
α	0.764 (0.063)	0.674 (0.066)	0.768 (0.065)	0.686 (0.071)
β	0.615 (0.070)	0.812 (0.066)	0.621 (0.070)	0.751 (0.066)
$\alpha + \beta$	1.379 (0.092)	1.486 (0.094)	1.389 (0.089)	1.438 (0.094)
a-ratio[c]	0.281 [0.000]	1.381 [0.173]	0.281 [0.000]	0.263 [0.000]
(b) Old				
$\log u$	-0.193 (0.068)	0.389 (0.091)	-0.323 (0.060)	0.613 (0.073)
$\log U$	-0.263 (0.072)	0.683 (0.139)	-0.074 (0.062)	0.324 (0.123)
$\log v$	0.480 (0.066)	0.030 (0.112)	0.483 (0.066)	0.039 (0.113)
$\log V$	0.020 (0.054)	0.009 (0.110)	0.046 (0.054)	-0.060 (0.108)
α	0.545 (0.052)	1.072 (0.099)	0.604 (0.058)	0.936 (0.112)
β	0.500 (0.046)	1.039 (0.097)	0.529 (0.045)	0.979 (0.095)
$\alpha + \beta$	1.044 (0.074)	2.111 (0.142)	1.133 (0.077)	1.916 (0.144)
a-ratio[d]	0.227 [0.000]	0.284 [0.000]	0.225 [0.000]	0.062 [0.000]
SE[e]	0.688 (0.072)	1.917 (0.084)	0.700 (0.073)	1.904 (0.084)
$\log L$	-16,559.4	-11,220.7	-16,568.5	-11,226.7
Obs	477,868	137,223	477,868	137,223

*The likelihood for the two job-seeker regressions given in Equations (16) and (**??**). 4-week window: $k^w = k^e = 4$. Estimates based on 2761 matches between 34,657 unemployed job-seeker spells (26,113 job-seekers) and 14,154 LCS job vacancies (4,121 employers and 9,556 orders). Standard errors in parentheses. All regressions contain monthly dummy variables.

[a,b] Narrow stocks: the weighted averages across the 3 'districts' for u and U are 191 and 755 respectively. For the wider definition, these are 207 and 1987. For v and V, they are 55 and 203.

[c] a_{12}/a_{11} for job-seekers, a_{21}/a_{11} for vacancies. The a-ratios calculated from Equations (9) and (10). We do not report standard errors, as the a-ratios are not normally distributed. By definition, p-values are the same as for underlying parameter estimates (alternative hypothesis is one-sided).

[d] a_{22}/a_{21} for job-seekers, a_{22}/a_{12} for vacancies.

[e] Standard error for Gaussian heterogeneity. $Q = 12$ for both job-seeker regressions and $Q = 24$ for both vacancy regressions.

Estimates of α and β for each regression are reported in each panel. These are estimates of the same underlying contact technology $\lambda_{ij}(U,V) = A_{ij}U^{\alpha}V^{\beta}$, recovered from separate regressions on job-seekers and vacancies. In most cases we estimate significant increasing returns to scale ($\alpha + \beta > 1$). Returns to scale are particularly large from the vacancy regressions for old vacancies.

This is not to do with stock-flow matching *per se*, as it turns out that scale effects persist when we estimate the random matching version (see Subsection 5.2 below).

Finding scale effects is contrary to what is usually found in the literature, which comes from mainly aggregate data. Most studies find constant returns, see, for example, Broersma & van Ours (1999, Table 1) and more comprehensively, Petrongolo & Pissarides (2001). There are good reasons why we might expect scale effects using agent-level data. Petrongolo & Pissarides (2006) develop and estimate a model that has increasing returns to quality of matches, with better matches occurring in larger markets. If agents respond by increasing their reservation utilities in proportion to the match quality, the hazard function should be independent of scale. Our results therefore imply that employers do not adjust their reservation utility when facing an increase in the quality of job-seekers, whereas job-seekers do when better quality vacancies arrive onto the market. This might be because employers have more market power, because they know that there is high youth unemployment and therefore set high reservation utilities to reduce search costs: they do not need to reduce reservation utilities when better job-seekers arrive. The opposite is true for job-seekers; moreover, they know much less about the labour market than do employers. One could investigate this further if we could disentangle arrival rate effects from matching probability effects. Stock-flow matching does not help explain why scale effects are stronger for old vacancies only, because the pure theory does not allow for stock-stock matches. Again, we would need to disentangle arrival rate effects from matching probability effects.

In Figure 3, we plot and compare the hazards from the Base Model with the raw hazards. We also plot the hazards controlling for the eight stock variables only (ie without controlling for unobserved heterogeneity). This is for wide stocks, but identical figures are obtained for narrow. Looking at the three vacancy hazards first, Figure 3(b) shows that the severe fall in the raw hazard over the first 8 weeks almost completely disappears in the Base Model, and that this is primarily due to controlling for unobserved heterogeneity (adding the eight covariates to the raw hazard model makes little difference).

It is important to note that the baseline hazards in themselves tell us nothing about stock-flow matching; hazard rates may fall with elapsed duration for many reasons, including duration dependence, unobserved heterogeneity and changing reservation utilities. On the other side of the market, the shape of the job-seeker hazard is unaffected by either adding covariates or controlling for unobserved heterogeneity. This is because, compared with the vacancy hazards, the standard error of the heterogeneity is about three times smaller (0.67 compared with 1.83). But again, looking at these job-seeker hazards to examine stock-flow matching would give the wrong impression (that agents, when old, are *more* likely to exit), whereas the regression-based estimate, using vacancy data, suggests that $a_{22}/a_{12} \approx 1/4$. As already noted, our view is that the hazard increases initially because job-seekers are learning how to search.

5.2 Departures from the Base Model

In this subsection, we look at various departures from the Base Model to assess the robustness of the stock-flow matching model, in particular the size and significance of the important flow variables v and u, and to assess whether the assumptions we have made are important or innocuous.

The first row of Table 3 summarises the Base Model. Row (0) of Table 3 shows what happens when dummies for the month of match are dropped from the model. We do this to counter any suggestion that stock-flow matching depends on the strong within year variations in stocks and flows that arise because of the particular (youth) labour being studied. The results are unaffected.

Row (1) shows that the results are robust to the way the unobserved heterogeneity is modeled, because here we use discrete (Heckman-Singer) mixing. In the job-seeker regression the log-likelihood is unaffected, and the number of parameters being estimated for discrete and Gaussian mixing is the same. In the vacancy regression, the log-likelihood is 7.7 log-points higher, but there are 15 more parameters to estimate. Using the Akaike Information criterion $-2 \log L + 2q$, where q is the number of parameters, this suggests that Gaussian mixing is the better way of modeling heterogeneity.

Row (2) investigates whether the heterogeneity should be defined for a job-seeker or a job-seeker spell. Job-seekers who match more than once are those who experience repeated unemployment events, and this is likely to be correlated with individual heterogeneity. Unless the probability of re-entering the unemployment pool is separable (proportional) in covariates and individual heterogeneity, the model could be misspecified. This clearly depends on the numbers of job-seekers who experience multiple spells in the dataset; this is 25.2% (see Footnote 8). If changing the specification of the unobserved heterogeneity from job-seeker to job-seeker spell doesn't affect the results, this is not an issue. Row (2) shows that the results do not change; moreover, the log-likelihood falls by 11.4 log-points. For the vacancy regressions, should the heterogeneity be defined for employers rather than vacancy orders? It doesn't make sense to define the heterogeneity for a vacancy rather than a vacancy order, because all the vacancies in an order are identical. However, we do observe employers with more than one spell; this is about half of them (see Footnote 9). However, when we define the heterogeneity for an employer, the likelihood falls by

82.8 log-points.

Row (3) shows the result of estimating the Base Model without unobserved heterogeneity. Apart from a moderate fall in α in the vacancy regression, again very little changes, even though the likelihood is a lot lower.

Row (4) reports what happens when various observed covariates are added to the Base Model. There is very little change in any of the estimates, which implies that observable characteristics of job-seekers and vacancies are not correlated with the aggregate numbers of job-seekers and vacancies in a particular market. This is not surprising, and applies to unobservables as well. This also justifies modeling the heterogeneity using random effects techniques.

Row (5) reports estimates of the classical random matching model. We find a slight degree of increasing returns in job-seeker regressions, with $\alpha + \beta$ estimated as 1.237; this is bigger for vacancies, with of $\alpha + \beta$ estimated as 1.447. As already discussed, the increasing returns is much stronger in the Base Model for old vacancies.

In Row (6) we examine the effects of aggregation bias by replacing stocks observed at weekly intervals with those observed at monthly intervals. Now, for every week in a given month, the value of the stock is the same and equal to that of the first week of the month. The results show that aggregation bias might be a problem for investigators with monthly data. First, the estimate of α is bigger in the job-seeker regression (moving from 0.604 to 0.759) and is smaller for β (moving from 0.529 to 0.443). The effect in the vacancy hazards is the other way round, with α falling from 0.936 to 0.579—a very large change—and β decreasing from 0.979 to 0.841, so that $\alpha + \beta$ falls from 1.915 to 1.420. Thus aggregation bias really does bias the estimates. More importantly, aggregation bias affects our estimates of the coefficient on $\log v$ and the a-ratios in job-seeker regressions. Now a_{22}/a_{21} is estimated as 0.394 rather than 0.225. In the vacancy regression, a_{22}/a_{12} is estimated as 0.028 rather than 0.276.

In Row (7) we replace our preferred measure of search duration with that used hitherto in the literature, which we label spell duration (see Section 3). This has a small effect on the estimates in the vacancy regressions, with a_{22}/a_{12} increasing from 0.062 to 0.079, but a bigger effect in the job-seeker regressions, with a_{22}/a_{21} increasing from 0.225 to 0.309. In other words, there are stronger stock-flow effects when we use search duration data, which might be expected.

Table 3: Summary of departures from Base Model, wide stocks*

	avge u	avge v[a]	Old job-seekers, h^w					Old vacancies, h^e				
			$\log v$	a_{22}/a_{21}	α	β	$\log L$	$\log u$	a_{22}/a_{12}	α	β	$\log L$
Base Model	207	55	0.483 (0.066)	0.225	0.604	0.529	-16568.5	0.613 (0.073)	0.062	0.936	0.979	-11226.7
Departures												
(0) Without monthly dummies	207	55	0.542 (0.062)	0.186	0.677	0.532	-16646.3	0.559 (0.049)	0.076	1.105	0.976	-11372.6
(1) Heckman-Singer[b]	207	55	0.483 (0.066)	0.225	0.603	0.529	-16568.4	0.614 (0.072)	0.061	0.915	0.970	-11219.0
(2) Job-seeker spell random effects[c]	207	55	0.489 (0.068)	0.221	0.587	0.595	-16578.9					
(2) Employer random effects[d]	207	55						0.661 (0.064)	0.051	0.548	1.236	-11309.4
(3) Without heterogeneity	207	55	0.481 (0.065)	0.226	0.623	0.513	-16581.0	0.641 (0.061)	0.055	0.735	0.968	-11619.1
(4) With covariates[e]	207	55	0.485 (0.066)	0.223	0.581	0.479	-16438.0	0.617 (0.073)	0.060	0.834	0.938	-11060.1
(5) Classical random matching[f]	207	55	0[g]	1[g]	0.846	0.391	-16618.2	0[g]	1[g]	0.606	0.841	-11274.3
(6) Monthly stocks	221	55	0.296 (0.065)	0.394	0.759	0.443	-16608.9	0.794 (0.087)	0.028	0.579	0.841	-11250.5
(7) Spell duration[h]	208	56	0.377 (0.063)	0.309	0.652	0.528	-16544.1	0.548 (0.073)	0.079	1.002	0.959	-11516.0
(8) $k^w = 5$, $k^e = 5$	254	63	0.521 (0.075)	0.222	0.687	0.552	-16570.5	0.773 (0.080)	0.036	0.763	0.939	-11222.6
(9) $k^w = 4$, $k^e = 2$	191	33	0.359 (0.052)	0.225	0.608	0.514	-16569.3	0.493 (0.066)	0.090	0.971	0.888	-11235.4
(10) $k^w = 4$, $k^e = 1$	191	18	0.317 (0.043)	0.160	0.607	0.516	-16555.4	0.442 (0.063)	0.108	0.820	0.795	-11245.8
(11) $k^w = 3$, $k^e = 3$	158	45	0.427 (0.057)	0.229	0.677	0.561	-16566.7	0.502 (0.065)	0.073	0.767	0.870	-11236.3
(12) $k^w = 3$, $k^e = 2$	158	33	0.351 (0.050)	0.231	0.676	0.546	-16565.2	0.453 (0.063)	0.088	0.967	0.889	-11236.7
(13) $k^w = 2$, $k^e = 2$	107	33	0.351 (0.048)	0.231	0.692	0.530	-16576.6	0.358 (0.054)	0.088	0.931	0.894	-11239.8
(14) $k^w = 2$, $k^e = 1$	107	18	0.331 (0.039)	0.152	0.687	0.546	-16562.3	0.314 (0.052)	0.105	0.774	0.796	-11249.5

*In each row, the Base Model is re-estimated with one dimension altered (a single departure).

[a] Average U is 1987, average V is 203 except for spell duration (1991 and 207 respectively) and monthly stocks (1965 and 201 respectively).

[b] For job-seekers regressions, $M = 2$ mass points were used; for vacancy regressions, $M = 7$ mass points were used.

[c] Random effects defined for 34,657 job-seeker spells rather than 26,113 job-seekers.

[d] Random effects defined for 4,121 employers rather than 9,556 vacancy orders.

[e] For job-seekers regressions, these are gender (1 dummy), grades at age 16–17 (so-called GCSEs) (3), ethnicity (1), disadvantaged social background (1); for vacancy regressions, these are whether the vacancy requires a skilled employee (1), a non-manual employee (1), a written method of application (1), firm size (3) and wage (4). See Footnote 15.

[f] Estimates of $h^w = A\mu U^{\alpha-1}V^{\beta}$ for job-seekers and $h^e = A\mu U^{\alpha}V^{\beta-1}$ for vacancies, pooled across old and new agents. [g] Imposed.

[h] The number of observations in the spell duration datasets is 480,423 for job-seekers and 139,505 for vacancies.

Rows (8)–(14) report what happens when we alter the window sizes away from $k^w = k^e = 4$ weeks. We choose the following (k^w, k^e) pairs: (5,5), (4,2), (4,1), (3,3), (3,2), (2,2), and (2,1). The estimate of $\log v$ in the job-seeker hazard tends to fall with smaller windows, ranging from 0.52 to 0.32. In addition, a_{22}/a_{21} depends on the size of the average stocks (see Equation 9), which change in size as k^w and k^e vary. The net effect is that a_{22}/a_{21} is robustly estimated between 0.15 and 0.23. The same happens for the vacancy hazards: the effect of $\log u$ falls with k^w and k^e, but leaves a_{22}/a_{12} robustly estimated in the range 0.04 to 0.11.

To summarise: the stock of new vacancies $\log v$ is robustly significant in the old job-seeker regression, and the stock of new job-seekers $\log u$ is robustly significant in the old vacancy regression. This implies that $a_{22}/a_{21} < 1$ and $a_{22}/a_{12} < 1$ for all these departures from the Base Model. In particular, the result appears robust to the choice of window size. The only assumptions that really matter in the Base Model are using weekly rather than monthly stocks and using search duration rather than spell duration.

5.3 Heterogeneity as an Alternative Explanation

On the basis of the models estimated this far, we find strong evidence in favour of stock-flow matching, and this is a key result of the paper. We have shown that a_{21} is approximately four times bigger than a_{22}, which is consistent with stock-flow matching. However, if the average quality of new vacancies v is higher than the average quality of old vacancies $\bar{V}$, then it is possible that quality differences might explain some of the difference between a_{22} and a_{21}. In other words, old job-seekers are more likely to match with new vacancies, not because of stock-flow matching, but because new vacancies are, on average, higher quality.[14]

The great advantage of our micro-level data is that we observe various features of job-seekers and vacancies which might be correlated with their quality. We can therefore include these observable characteristics in our regressions, by creating 'good' and 'bad' stocks of job-seekers and vacancies. These explanatory variables are constructed as follows. We define a good job-seeker as one whose exam performance at age 16–17 is 5+ high grade GCSEs and who does not have a disadvantaged social background.[15] We define a good vacancy as one that offers a skilled and non-manual

[14]Following (Becker 1973), there are few papers that model matching with heterogeneous agents in the labour market, with and without transferable utility; they analyse whether, and under what conditions, low quality job-seekers match with low quality vacancies (positive assortative matching) or otherwise (see, for example, Burdett & Coles (1999), Delacroix (2003)). Both papers analyse steady-state equilibria and say nothing about the hazards of job-seekers and employers outside equilibria.

[15]GCSEs are national examinations taken in the school year a young person is aged 16. Career's

job, according to standard definitions. Of the 26,113 job-seekers in the dataset, this means that 43.4% are good; of the 9,556 vacancy orders, 37.8% are good. From these, we can construct the following stocks:

$$U_t = u_t + \bar{U}_t = U_t^g + U_t^b = u_t^g + u_t^b + \bar{U}_t^g + \bar{U}_t^b.$$

Figure 1(a) plots time-series of U_t, u_t, U_t^g and V_t, v_t, V_t^g. It is clear that the good stocks move very closely with the total stocks through time; this is less true when comparing the new stock of vacancies with the total stock, and definitely less true when comparing the new stock of unemployed with the total stock. A similar decomposition applies to the stocks of new, old, good and bad vacancies, see Figure 1(b).

Table 4: Job-seeker regressions, good and bad stocks*

	Base model (1)		*(2)*		*(3)*	
$\log u$	−0.323	(0.060)	−0.306	(0.060)	−0.326	(0.060)
$\log U$	−0.074	(0.062)	−0.054	(0.063)	−0.076	(0.063)
$\log v$	0.483	(0.066)			0.477	(0.066)
$\log V^g$			−0.569	(0.323)	−0.377	(0.324)
$\log V$	0.046	(0.054)	0.496	(0.102)	0.147	(0.111)
$\log L$	−16,568.5		−16,542.0		−16,570.0	

*Estimates of Equation (15), old job-seekers. Estimate of new job-seekers included, but not reported.

Column (2) of Table 4 reports estimates of the base model where $\log V^g$ replaces $\log v$. The estimated effect on $\log V^g$ is wrongly signed and badly determined, with the standard error being five times bigger than the standard error on $\log v$. In addition, the maximised log-likelihood falls by 26.5 log-points. When we add $\log V^g$ to the Base Model, reported in Column (3), all the estimates on the original covariates remain virtually unchanged, but the estimated effect on $\log V^g$ changes to -0.377 and is again insignificant.

This clear rejection of the random matching model with heterogeneity in favour of the stock-flow matching model is because $\log V$ is highly correlated with $\log V^g$, being 0.93 in both job-seeker and vacancy datasets. Because the stock of good vacancies moves exactly in line with total vacancies over the recruitment cycle (see

Officers assess whether someone has a disadvantaged social background when the young person is interviewed, on the basis of whether he or she comes from a poor home background, has been involved in crime, or has been in care. This applies to roughly 10% of the sample.

Figure 1), this means that these amended models do not capture the essential feature of the stock-flow model. If they were to provide a competing explanation of the data, good job-seekers would exit at different rates to bad job-seekers over the cycle, and the stocks would not move in line. Note that the correlations between $\log V$ and $\log v$ are 0.74 and 0.62 for the job-seeker and vacancy datasets respectively; these are much lower than above, which is why the new stocks have considerable explanatory power in the stock-flow model.

On the other side of the market, $\log U$ is highly correlated with $\log U^g$, being 0.94 and 0.97 in the two datasets, whereas the correlations between $\log U$ and $\log u$ are 0.40 and 0.61. This time, we do find a role for good and bad stocks, but the basic stock-flow model remains intact. What we find is that the effect of the stock of new job-seekers on the vacancy hazard can be split into two significant components, namely that of good new job-seekers and that of bad new job-seekers. Instead of $\log u$ having an estimate of 0.616, the effect of $\log u^g$ is estimated as 0.869 and the effect of $\log u^b$ is estimated as –0.325. The log-likelihood increases by 11.9 log-points. These estimates make good intuitive sense, in that the old vacancy hazard is much higher when there are more good new job-seekers on the market and that it is lower when there are more bad new job-seekers on the market (assuming, as always, the stock of job-seekers remains constant because the old stock is assumed to fall). Hence, on this side of the market, the stock-flow matching model becomes even richer when augmented by these good/bad variables.

To conclude, the stock-flow matching model appears to provide a better explanation of the data than a matching model with heterogeneity. This is because good stocks of data cannot be seen as proxies for the new stocks because 'good' is highly correlated with 'total', not highly correlated with 'new'.

6 Conclusion

In this paper, we report estimates of Coles & Smith's stock-flow matching model by modeling job-seeker and vacancy hazards using micro-level data from both sides of a single market. Specifically, we focus on the job-seeker hazard when the job-seeker becomes old, whose covariates are the stock of market participants, namely the stock of unemployed job-seekers and the stock of vacancies. This describes a form of the classical random matching model estimated many times in the literature with aggregate data. We then add the stock of new vacancies, and see whether it has any impact on the hazard of getting a job *over and above* the effect of the stock of all vacancies. If the effect is positive and significant, this implies that employers find it harder to match to old job-seekers once their vacancies become old. Exactly the reverse applies to the other side of the market, where the test examines the effect of the stock of new job-seekers. The test does not examine whether vacancy hazards or job-seeker hazards fall at certain durations, because this can happen for other reasons.

Our results for the stock-flow model are summarised as follows. The stock of new vacancies has a significant additional impact on the exit rate for old job-seekers, as is predicted by stock-flow matching theory, and is robust across choice of window and whether or not we use a narrow or wide definition of the stock. For the wide definition (which additionally includes those searching whilst not unemployed), this implies that the hazard rate for vacancies falls by about three-quarters when vacancies become old ($a_{22}/a_{21} \approx 0.23$). There is an equivalent robust effect for the exit rate of old job-seekers on the other side of the market for the wide definition, with the hazard falling drastically when job-seekers become old, to less than one-tenth of its value ($a_{22}/a_{12} \approx 0.06$).

However, if the quality of vacancies and job-seekers is strongly correlated with their duration in the market, this might provide an alternative explanation as to why the stock of new vacancies and job-seekers is highly significant in these duration models. To test this, we used the micro data to add a measure of "good" vacancies to the old job-seeker hazard. This measure proved insignificant, and its inclusion had very little effect on the estimated effect of new vacancies. For old vacancy regressions, the hazard is higher when there are more 'good' new job-seekers on the market compared with 'bad' new job-seekers. As always, we cannot conclude that this provides convincing evidence in favour of stock-flow matching, because we cannot rule out the possibility that some unobserved heterogeneity remains. See Coles & Petrongolo (2008), who come to the same conclusion.

One final caveat is that our data come from a particular labour market institution which provides an explicit matching service for job-seekers and vacancies, rather than from the labour market as a whole. This is an environment in which stock-flow matching is likely to be particularly appropriate, and therefore to some extent it is unsurprising that we find evidence in favour of stock-flow matching. Micro data from a more decentralised labour market would be required to test stock-flow matching more generally.

To conclude, with these caveats in mind, we find evidence in favour of stock-flow matching, using a unique dataset with high quality agent-level information from both sides of the same market. We are thus able to observe both stocks and flows over intervals shorter than one month, which is the best one can do using aggregate data. All of the analysis in this paper is conducted at the level of individual matches. Most importantly, with aggregate data we would be unable to model the essential feature of search models, that of individual agents changing behaviour in response to changing aggregate labour market conditions during the agents' stay in the market.

References

Becker, G. (1973), 'A theory of marriage: part I', *Journal of Political Economy* **81**, 813–46.

Brinch, C. (2007), 'Nonparametric identification of the mixed hazards model with time-varying covariates', *Econometric Theory* **23**, 349–54.

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

Burdett, K. & Coles, M. (1999), 'Long-term partnership formation: marriage and employment', *Economic Journal* **109**, F307–34.

Burdett, K., Coles, M. & van Ours, J. (1994), Temporal aggregation bias in stock-flow models. CEPR Discussion Paper No. 967.

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

Coles, M. & Petrongolo, B. (2008), 'A test between stock-flow matching and the random matching function approach', *International Economic Review* **49**, 1113–41.

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

Delacroix, A. (2003), 'Heterogeneous matching with transferable utility: two labour market applications', *International Economic Review* **44**, 313–30.

Ebrahimy, E. & Shimer, R. (2009), Stock-flow matching. Mimeo.

Gaure, S., Roed, K. & Zhang, T. (2007), 'Time and causality: a Monte Carlo assessment of the timing-of-events approach', *Journal of Econometrics* **141**, 1159–95.

Gregg, P. & Petrongolo, B. (2005), 'Stock-flow matching and the performance of the labour market', *European Economic Review* **49**, 1987–2011.

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

Lagos, R. (2000), 'An alternative approach to search frictions', *Journal of Political Economy* **108**, 851–73.

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, Elsevier, Amsterdam, chapter 39, pp. 2567–627.

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

Petrongolo, B. & Pissarides, C. (2006), 'Scale effects in markets with search', *Economic Journal* **116**, 21–44.

Pissarides, C. (1985), 'Short-run equilibrium dynamics of unemployment vacancies, and real wages', *American Economic Review* **75**, 676–90.

Pissarides, C. (2000), *Equilibrium Unemployment Theory*, second edn, Basil Blackwell, Oxford.

Pissarides, C. A. (2008), EconomicDynamics interviews Christopher Pissarides on the matching function. EconomicDynamics Newsletter and Mailing List.

Shimer, R. (2007), 'Mismatch', *American Economic Review* **27**, 1074–101.

Taylor, C. (1995), 'The long side of the market and the short end of the stick: bargaining power and price formation in buyers', sellers' and balanced markets', *Quarterly Journal of Economics* **110**, 837–55.

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

Figures

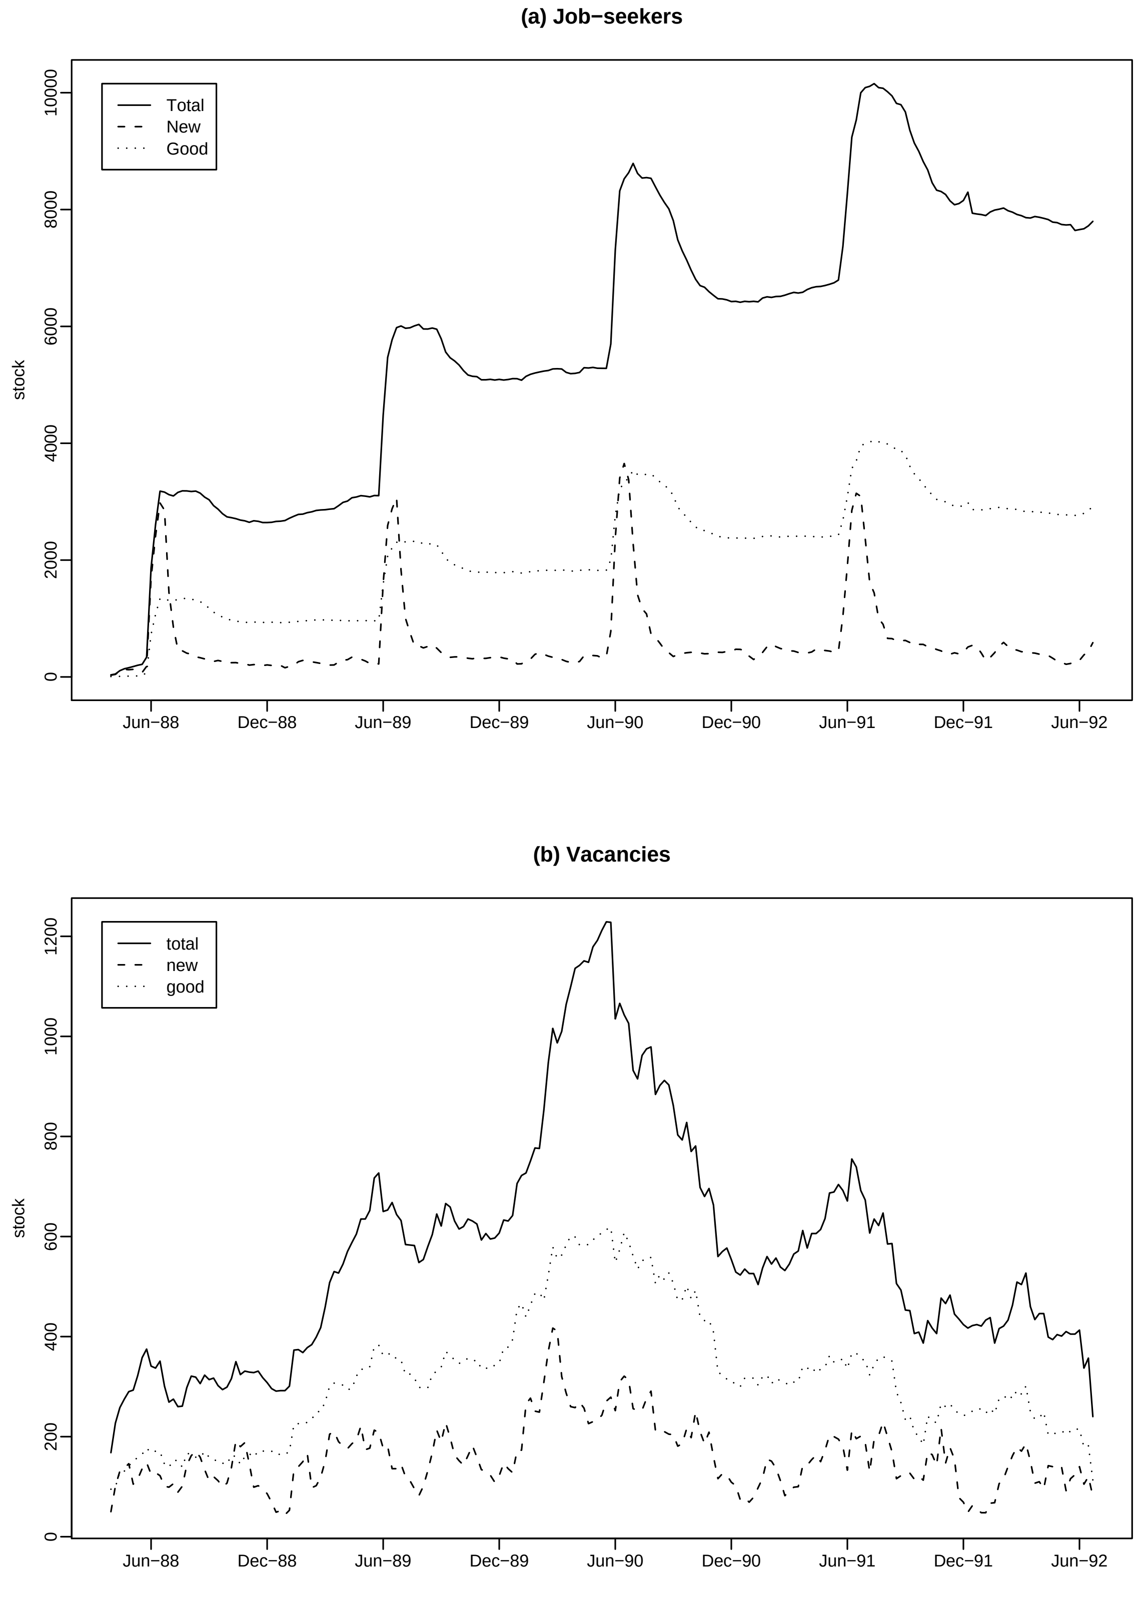

Figure 1: Unemployment and vacancy stocks: total, new and good

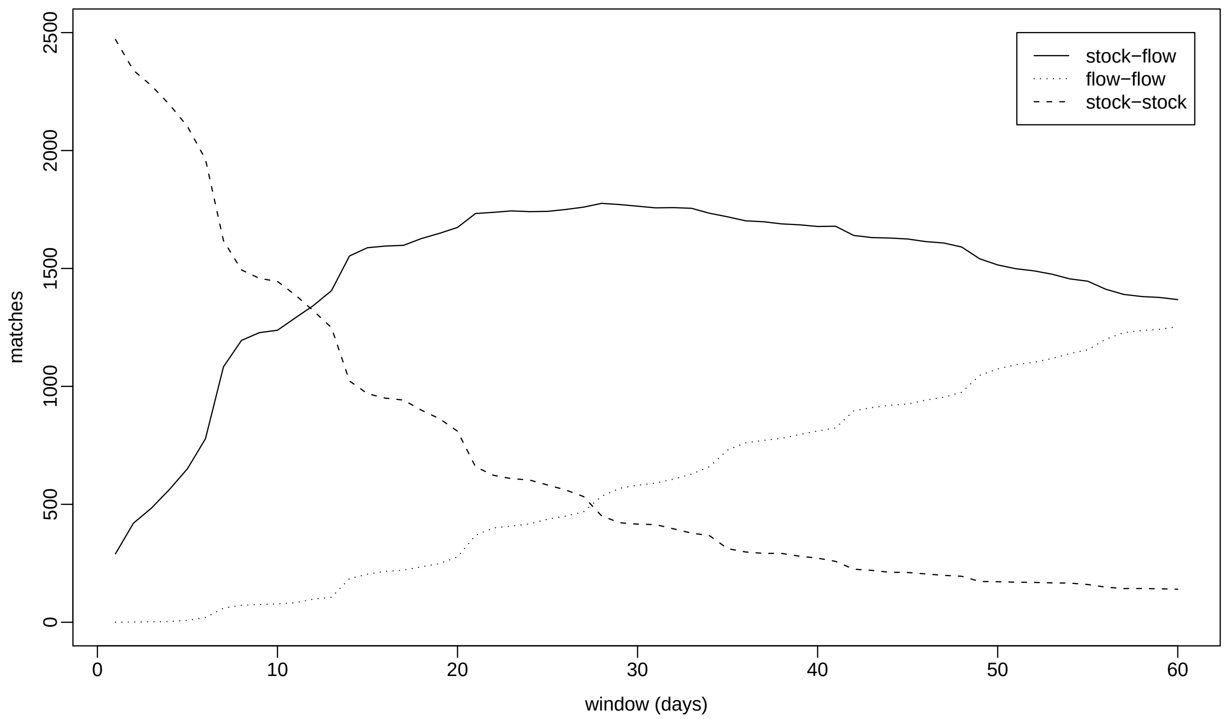

Figure 2: Stock-flow counts by window size, $k^w = k^e$

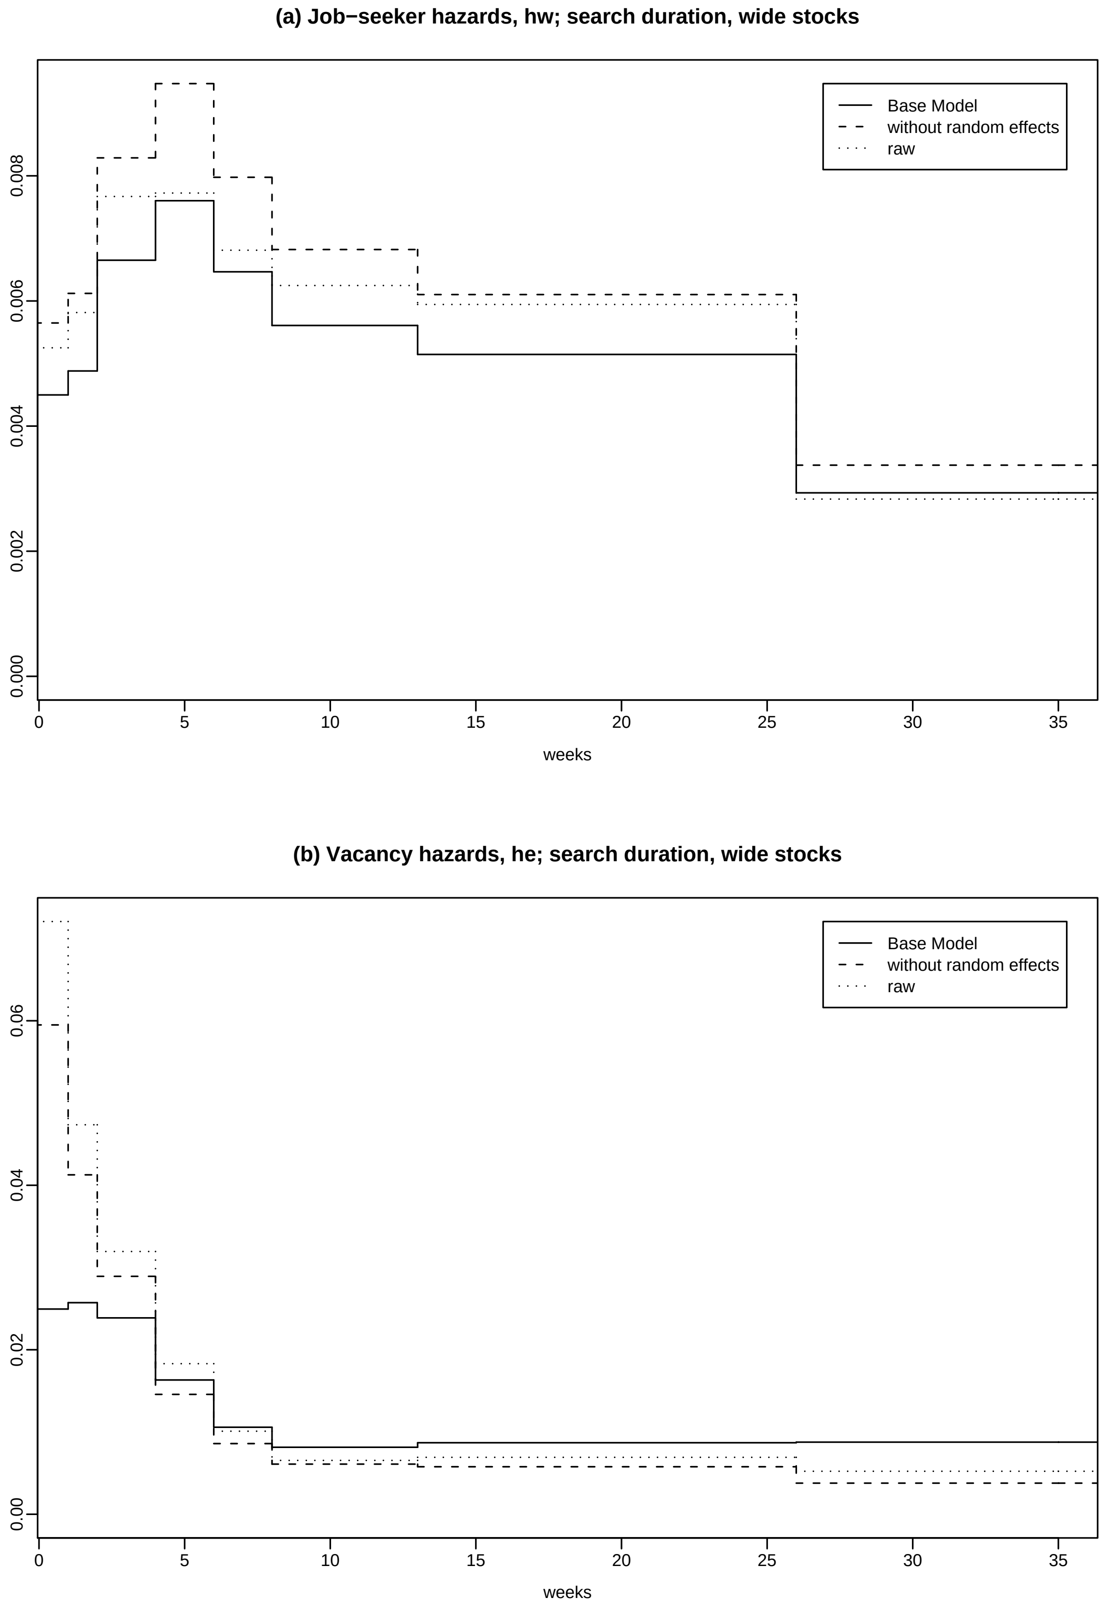

Figure 3: Conditional and raw baseline hazards