INTERNATIONAL ECONOMIC REVIEW

November 2008
Vol. 49, No. 4

A TEST BETWEEN STOCK-FLOW MATCHING AND THE RANDOM MATCHING FUNCTION APPROACH*

BY MELVYN COLES AND BARBARA PETRONGOLO[1]

University of Essex, U.K.; Institució Catalana de Recerca i Estudis Avançats and Institut d'Análisi Económica, Spain, London School of Economics and CEP (LSE), U.K.

This article tests between the standard "random matching function" approach and "stock-flow" matching while controlling for temporal aggregation bias. Consistent with previous empirical work, the random matching function fits the matching data reasonably well. But match flows are more highly correlated with vacancy inflows than is consistent with the random matching approach. Instead the data support stock-flow matching, where unemployed workers match directly with suitable new vacancies as such vacancies come on to the market.

1. INTRODUCTION

The random matching approach has provided an important framework for analyzing labor market policy (Pissarides, 2000). But the empirical literature, when estimating the random matching function, rarely tests the matching function against a meaningful alternative (see, for example, Blanchard and Diamond, 1989; Petrongolo and Pissarides, 2001, for a recent survey). This article uses matching data to test between the random matching hypothesis and stock-flow matching.

Stock-flow matching assumes that when laid off, a worker contacts friends, consults situations vacant columns in newspapers, registers with job agencies, and so observes the stock of vacancies currently on the market. If the worker is lucky and a suitable vacancy already exists, the worker can quickly exit unemployment. If a suitable vacancy does not exist, the worker then has to wait for something suitable to come onto the market. This worker does not then match with the stock of vacancies—the stock has already been sampled and no match exists. The worker instead matches with the inflow of newvacancies coming onto the market. The

* Manuscript received June 2005; revised August 2006.

[1] Please address correspondence to: Melvyn Coles, Department of Economics, University of Essex, Wivenhoe Park, Colchester, Essex, CO4 3SQ, U.K. Phone: (01206)-87-4875. Fax: 00 44 1206 872724. E-mail: *mcole@essex.ac.uk*.

problem is symmetric for vacancies. If a firm has a vacancy, it might first ask employees if they know someone suitable or advertise the post in a situations vacant column, job agencies, etc. If the firm is lucky, a suitable worker already exists in the stock of unemployed workers and the post is quickly filled. If not, the firm has to wait for someone suitable to come onto the market. This implies "stock-flow" matching as the stock of unmatched agents on one side of the market matches with the inflow of new agents on the other side. Papers in this literature include Taylor (1995), Coles and Smith (1998), Coles and Muthoo (1998), Coles (1999), Lagos (2000), and Gregg and Petrongolo (2005), but also see Jones and Riddell (1999) who consider "wait" unemployment. More recent work includes Ebrahimy and Shimer (2006), who calibrate the stock-flow matching model to the U.S. labor market.

Lagos (2000) perhaps provides the most useful perspective for understanding the results obtained here. Using a taxi market analogy, he supposes that cabs meet potential customers at taxi ranks. At the microlevel there is stock-flow matching, i.e., at any given taxi rank there is either a stock of customers waiting for cabs or a stock of cabs waiting for customers. Aggregation over all taxi ranks and the restriction to *steady state* imply that at the macrolevel, the flow number of cab rides depends only on the total stock of taxi cabs and on the total stock of potential customers in the market. Casual introspection also suggests there will be constant returns to matching: doubling the total number of participants should double the (steady state flow) number of taxi rides. Aggregate matching seemingly has the properties of a standard random matching function, even with stock-flow matching at the microlevel.

This view of matching is consistent with the fact that around 25–30% of new vacancies posted in Jobcentres in Britain are filled on the first day (Coles and Smith, 1998). Burdett and Cunningham (1998)also report for the United States that most vacancies (55%) are filled within a week. This suggests that for many vacancies matching frictions may not be a significant factor. Instead such vacancies are snapped up by workers who have been waiting for something suitable to come onto the market.

The equivalence between random matching and stock-flow matching, however, only holds in steady state. Outside of a steady state, stock-flow matching at the microlevel implies that a higher inflow of new vacancies, say into an unemployment black spot, will yield an immediate increase in matches. Consider for example, Figures 1 and 2 below, which describe aggregate matching time series in Great Britain. A critical feature of these time series is that the unemployment outflow or the vacancy outflow are much more volatile than both the unemployment and the vacancy stock. Matching data for the United States provide a very similar picture; see Blanchard and Diamond (1989, figure 5). This data pattern can only happen when the number of matches is mostly driven by the inflow of new agents into the market. Otherwise, if most matches originate in the stocks, a large increase in the number of matches must result in a large fall in the stock of vacancies and unemployment, which is not observed in the data. These data therefore imply that a spike in the number of matches coincides with spikes in the inflows of new agents and *not* with spikes in the existing stocks. The standard random matching approach is inconsistent with this facet of the data.

The underlying point is that outside of a steady state, random matching and stock-flow matching imply quite different equilibrium hazard rates of re-employment. In Shapiro and Stiglitz (1984), for example, the re-employment hazard rate of an unemployed worker, denoted λ, is simply $\lambda = v/U$, where v is the flow of new vacancies into the market, which match immediately (and randomly) with one unemployed worker, and U describes the current stock of unemployed workers. We shall refer to this case as job queueing. In contrast, the random matching approach assumes the re-employment hazard rate depends on the vacancy/unemployment ratio, where frictions imply it is the stock of vacancies V that is the relevant state variable (throughout uppercase refers to stock variables, lowercase to flow variables). The stock-flow explanation instead implies there are two parameters of interest. Proportion p of laid-off workers are on the short side of their market and are re-employed immediately (or at least within a very short period of time). Proportion $(1-p)$ are on the long side of their market and so chase new vacancies. Their subsequent hazard rate λ depends on the flow of vacancies into their particular specialization. Clearly this latter hazard rate λ has a similar structure to the job queueing hazard, depending on v instead of V.

Consistent with Lagos (2000), we find that the random matching function does indeed provide a good fit of the United Kingdom aggregate matching data and the hypothesis of constant returns is not rejected (see column 1, Table 2, below). The critical overidentifying test for the random matching approach is that the inflow of new vacancies should not have any additional explanatory power for the matching rates of workers (given we control for temporal aggregation of the data). We find instead that the vacancy inflow coefficient is highly significant. The principal difficulty for the random matching approach, however, is that once vacancy inflow is included as a conditioning variable, the estimated vacancy stock term becomes wrong-signed. The results instead show that a stock-flow matching specification provides a much better fit of the data. The central finding is that the longer-term unemployed match with the inflow of new vacancies, which implies that these workers are waiting for suitable vacancies instead of facing high matching frictions. As pointed out in the conclusion, this has important implications for the design of optimal unemployment insurance schemes.

There are two closely related papers. Coles and Smith (1998) introduces the notion of stock-flow matching and shows that unemployment outflows of workers with durations longer than one month are highly correlated with vacancy inflows, whereas vacancy stocks typically yield no significant correlations.[2] Consistent with stock-flow matching, Coles and Smith (1998) also find that unemployment outflows for workers with durations shorter than one month are significantly correlated with vacancy stock measures. There is no formal testing in that paper. Gregg and Petrongolo (2005) were the first to distinguish between stock-flow matching and the random matching function on aggregate match data. But their estimates do not properly control for temporal aggregation bias. In particular, they only use the vacancy stock measure at the start of each quarter as a determinant of the job-finding hazard for the unemployment stock (see Equations (8) and (9), p. 1998). As will be made clear below, even with random matching, the total

[2] Except at one year durations where worker unemployment benefits expire.

unemployment outflow over the period depends on the inflow of new vacancies within the period. As unemployment outflows are highly correlated with vacancy inflows, a goodness-of-fit test in Gregg and Petrongolo (2005) is biased against random matching. A central contribution here is that we show how to construct "at-risk" measures of vacancies and unemployed workers within each month and so fully control for the temporal aggregation problem. A formal test between the two matching frameworks is then possible using maximum likelihood techniques.

The article is structured as follows. Section 2 discusses identification. It argues that with unobserved search effort, a distinguishing test between the two matching frameworks is not possible using microdata, but that identification is possible on macrodata. Section 3 shows how to distinguish between competing matching frameworks, while controlling econometrically for temporal aggregation bias. Section 4 describes the data used. Section 5 presents our results. Section 6 concludes.

2. IDENTIFICATION ON AGGREGATE DATA

The hazard function literature establishes there is substantial variation in re-employment rates across unemployed workers (see Machin and Manning, 1999, for a recent survey). To understand how worker heterogeneity may affect identification, suppose unemployed worker i at time t makes search effort k_{it} and let

$$K_t = \sum_{i=1}^{U_t} k_{it}$$

denote aggregate search effort across the U_t unemployed. The standard random matching approach (e.g., Pissarides, 2000) assumes that the re-employment rate of this worker, denoted λ_{it}, is

$$\lambda_{it} = \frac{k_{it}}{K_t} M(K_t, V_t),$$

where V_t is the stock of vacancies at time t and $M(.)$ is a standard matching function.

Consider instead a job queueing framework. In this world the stock of unemployed workers matches with new vacancies as they are created. Using the taxi rank analogy, suppose unemployed worker i makes effort k_{it} to catch the next job. Then this worker's re-employment rate is

$$\lambda_{it} = \frac{k_{it}}{K_t} v_t,$$

where v_t is the inflow of new vacancies and k_{it}/K_t is the probability that worker i gets the next job to come onto the market.

The aim of the article is to identify which better explains the re-employment rates of unemployed workers: Is it the stock of vacancies (as implied by random matching) or the inflow of new vacancies (as implied by job waiting)?

Unfortunately it is not possible to identify either of these processes on microdata. For most individuals we have only one data point—that individual's observed unemployment spell (a starting and end date)—and we cannot control for the fixed effect k_{it}. We might potentially control for the fixed effect by focusing on individuals with repeat unemployment spells, but such individuals are a biased sample of the market.

Note also that this matching structure generates congestion effects, where greater job search effort by one worker reduces the matching probability of other workers (via an increase in aggregate K). Unobserved search heterogeneity implies a negative correlation in individual worker re-employment rates—if one worker gets the job, the others do not. Aggregating across workers, however, nets out these congestion effects. For example, job queueing implies individual hazard rate $\lambda_{it} = k_{it}v_t/K_t$, but aggregating over i implies average re-employment rate

$$\bar{\lambda}_t = \frac{1}{U_t}\sum_{i=1}^{U_t}\lambda_{it} = \frac{1}{U_t}\sum_{i=1}^{U_t}\frac{k_{it}}{K_t}v_t = \frac{v_t}{U_t}.$$

Note that the average re-employment rate is driven by the inflow of new vacancies; there is crowding out by the unemployment stock, but pure crowding out implies the average re-employment rate is independent of the distribution of search efforts $\{k_{it}\}$.[3]

The random matching approach instead implies $\lambda_{it} = \frac{k_{it}}{K_t}M(K_t, V_t)$ and aggregating over i implies average re-employment rate

$$\bar{\lambda}_t \equiv \frac{1}{U_t}\sum_{i=1}^{U_t}\lambda_{it} = \frac{1}{U_t}\sum_{i=1}^{U_t}\frac{k_{it}}{K_t}M(K_t, V_t) = \frac{1}{U_t}M(K_t, V_t).$$

If there are constant returns to matching then the average re-employment rate simplifies to

$$\bar{\lambda}_t = M\left(\bar{k}_t, \frac{V_t}{U_t}\right),$$

where $\bar{k}_t = K_t/U_t$ is average search effort. As there is only partial crowding out, the average re-employment rate depends on average search effort but is otherwise independent of the distribution of search efforts $\{k_{it}\}$.

As $\bar{k}_t$ is unobserved, a critical identifying assumption is that $\bar{k}_t$ changes slowly and systematically over time. This seems reasonable as the unemployment and vacancy stocks change slowly over time (see Figures 1 and 2, below).[4] In that case

[3] This is not true for the distribution of uncompleted unemployment spells. As $1/\lambda_{it}$ is a convex function of k_{it}, a mean preserving spread in k_{it} yields an increase in the average uncompleted unemployment spell. Indeed if $k_{it} = 0$ for one individual, the average spell is infinity.

[4] In a job search framework this implies job offer arrival rates change slowly over time and so aggregate job search effort will also change slowly over time.

we can control for unobserved changes in aggregate job search effort by using time dummies. In particular, using month dummies to capture seasonal variations in job search effort (for example, the job market is quiet in August) and year dummies to capture business cycle variations, random matching implies that observed short-run variations in the average re-employment rates of unemployed workers are driven by variations in the vacancy-unemployment ratio. Job queueing instead implies that those variations are driven by variations in vacancy inflow. A second approach, discussed further in the data section, is to use an HP filter that takes out all trends.

3. THE EMPIRICAL FRAMEWORK

Our data do not record unemployment and vacancy stocks over the month. Instead we have information on the stocks available at the start of each month and the gross inflows during that month. From now on we adopt the following time notation. U_n denotes the stock of unemployed workers at the beginning of month $n \in \mathcal{N}$ and V_n denotes the stock of vacancies. u_n denotes the total inflow of newly unemployed workers within the month and v_n denotes the inflow of new vacancies.

We suppose that at time $t \in [n, n+1]$ in month n, the average re-employment probabilities of unemployed workers are denoted by a pair $(p(t), \lambda(t))$. $p(t)$ is the proportion of workers laid off at date t who find immediate re-employment, whereas $\lambda(t)$ is the average re-employment rate of workers who have been unemployed for some (strictly positive) period of time. The competing theories suggest alternative functional forms for p and λ.

As previously described, the random matching approach implies $p = 0$—it takes time to find work—and the average re-employment rate, now simply denoted λ, is $\lambda = M(kU, V)/U$, where k is average search effort. Constant returns to matching implies a functional form $\lambda = \lambda^M(V/U, k)$.

Pure job queueing as described by Shapiro and Stiglitz (1984) also implies $p = 0$—it takes time to find work—but in this case the average re-employment rate is $\lambda = v/U$, where the stock of unemployed workers U matches with the inflow of new vacancies v. For econometric purposes we consider a more general specification of the form $\lambda = \lambda^Q(v, U)$.

Stock-flow matching is a generalization of the job queueing approach. Shapiro and Stiglitz (1984) imply all unemployed workers are on the long side of the market. In reality, some suitably skilled workers may find themselves on the short side of the market and can quickly find re-employment. We suppose with probability $p = p^{\mathrm{SF}}$, a newly unemployed worker finds there is a suitable vacancy already on the market and so becomes (very quickly) re-employed. With probability $1 - p^{\mathrm{SF}}$ there is no suitable vacancy and the worker must then wait for a suitable vacancy to come onto the market. Job queueing implies average re-employment rate $\lambda = \lambda^{\mathrm{SF}}(v, U)$. The same argument implies that each new vacancy may be either on the short or long side of the market. Vacancy queueing—waiting for a suitably skilled worker to come onto the market—implies average matching rate $\mu = \mu^{\mathrm{SF}}(u, V)$ where u is the inflow of newly unemployed workers onto the market. Symmetry suggests a functional form $p^{\mathrm{SF}} = p^{\mathrm{SF}}(V, u)$ where laid-off workers on the short side of the market match with vacancies on the long side. Pure job queueing (e.g., Shapiro and Stiglitz, 1984) implies one-sided stock flow matching and the testable restriction $p^{\mathrm{SF}} = 0$.

We now show how to identify (p, λ) using monthly time series. As the identifying equations depend on the assumed theory, we consider each case separately, starting with random matching.

3.1. *Temporal Aggregation with Random Matching.* To obtain a discrete time representation of the underlying continuous time matching process, we assume the inflows of new agents, u_n and v_n, are constant within a month. We construct at-risk measures for the stock of vacancies and unemployed workers by considering a representative worker who matches at average rate λ_n. Also suppose this representative unemployed worker withdraws into nonparticipation at rate δ_n^U. Then, if all the unemployed match at the same average rate λ_n over month n, total matches are

$$M_n = \int_0^1 U_n e^{-(\lambda_n+\delta_n^U)t} \lambda_n \, dt + \int_{x=0}^1 u_n \, dx \left[\int_{t=x}^1 e^{-(\lambda_n+\delta_n^U)(t-x)} \lambda_n \, dt \right].$$

The first term describes those workers in the initial stock who successfully match in month n. The second term describes those workers in the inflow who match, where—assuming entry is uniform and given entry at date $x \in [0, 1]$—each matches with probability given by the bracketed integral. Calculating these integrals, temporal aggregation implies the expected total number of matches in month n is

$$M_n = \lambda_n U_n \frac{\left[1 - e^{-(\lambda_n+\delta_n^U)}\right]}{\lambda_n + \delta_n^U} + \lambda_n u_n \left[\frac{e^{-(\lambda_n+\delta_n^U)} - 1 + (\lambda_n + \delta_n^U)}{(\lambda_n + \delta_n^U)^2} \right].$$

Define now the "at-risk" measure of unemployment:

$$\bar{U}_n = U_n \left[\frac{1 - e^{-(\lambda_n+\delta_n^U)}}{\lambda_n + \delta_n^U} \right] + u_n \left[\frac{e^{-(\lambda_n+\delta_n^U)} - 1 + (\lambda_n + \delta_n^U)}{(\lambda_n + \delta_n^U)^2} \right]. \tag{1}$$

The expected number of matches in month n is thus

$$M_n = \lambda_n \bar{U}_n,$$

with $\bar{U}_n$ given by Equation (1) and λ_n is the matching rate of the representative worker. To see why $\bar{U}_n$ is an "at-risk" measure, note that letting $\delta_n^U \to \infty$ implies $\bar{U}_n \to 0$: Arbitrarily high withdrawal rates into nonparticipation implies each unemployed worker is at no risk of finding a job. Conversely with no withdrawal into nonparticipation, $\delta_n^U = 0$, the unemployment at-risk measure becomes

$$\bar{U}_n = U_n \left[\frac{1 - e^{-\lambda_n}}{\lambda_n} \right] + u_n \left[\frac{e^{-\lambda_n} - 1 + \lambda_n}{{\lambda_n}^2} \right].$$

In addition, letting $\lambda_n \to 0$ then implies $\bar{U}_n \to U_n + \frac{1}{2} u_n$. Given that nobody finds work ($\lambda_n = 0$), each newly unemployed worker in month n is, on average, unemployed in that month for exactly half of it (given the identifying assumption that newly unemployed workers enter the market at a uniform rate). Hence $U_n + 0.5 u_n$ measures the average number of unemployed workers at risk over the whole month. In contrast, the limit $\lambda_n \to \infty$ instead implies each worker is only unemployed for an instant and this limit yields $\lambda_n \bar{U}_n \to U_n + u_n$: Given all workers immediately find work, total matches equal the total number of workers unemployed at any stage in the month. For any arbitrary matching rate $\lambda_n \geq 0$ and withdrawal rate δ_n^U, (1) computes $\bar{U}_n$, and $M_n = \lambda_n \bar{U}_n$ then describes the predicted total number of matches in month n. This formulation is usefully compared to the standard continuous time case where match flows at time t are given by $M_t = \lambda_t U_t$. With temporally aggregated data, $\bar{U}_n$ is the appropriate aggregated measure of unemployment.

The above temporal aggregation argument also applies to vacancies. If in month n vacancies match at average rate μ_n, are withdrawn at rate δ_n^V, and if v_n describes the total inflow of new vacancies within the month, then the relevant at-risk measure for vacancies is

$$\bar{V}_n = V_n \left[\frac{1 - e^{-(\mu_n + \delta_n^V)}}{\mu_n + \delta_n^V} \right] + v_n \left[\frac{e^{-(\mu_n + \delta_n^V)} - 1 + (\mu_n + \delta_n^V)}{(\mu_n + \delta_n^V)^2} \right], \tag{2}$$

and expected matches are given by

$$M_n = \mu_n \bar{V}_n.$$

This in turn implies the identifying restriction

$$\lambda_n \bar{U}_n = \mu_n \bar{V}_n, \tag{3}$$

as the expected number of workers who match must equal the expected number of vacancies that match.

For random matching, we adopt the functional form

$$\lambda_n = \lambda^M(\bar{U}_n, \bar{V}_n; \theta). \tag{4}$$

This specification says that the average re-employment rate λ_n in month n does not simply depend on the initial stocks U_n, V_n, but on the measures of unemployed workers and vacancies who try to match over the entire month. As in much of the matching literature (see Petrongolo and Pissarides, 2001), we will express λ_n as a log-linear function of its determinants.

Our working paper, Coles and Petrongolo (2002),estimates the equations above but with $\delta^U, \delta^V = 0$. Those results are qualitatively identical to the ones reported below, except the predicted average completed spell of a vacancy was severely overestimated. This overestimation occurred as the withdrawal rate of vacancies is surprisingly high; around one third of all posted vacancies are eventually withdrawn. As a steady state implies the average completed vacancy spell is $1/(\mu + \delta^V)$, ignoring δ^V (which is large) results in predicted vacancy spells being way too long.

In this version we control for vacancy withdrawal as part of the estimating equations. Specifically the "at-risk" structure above implies withdrawn vacancies W_n in month n are

$$W_n = \delta_n^V \bar{V}_n. \tag{5}$$

Using data on vacancy withdrawals in each month allows us to identify δ_n^V period by period. Unlike withdrawn vacancies, however, we do not have direct information on worker withdrawals into nonparticipation (especially as search effort is unobserved). Fortunately, with $\delta^U = 0$ the results, both here and in Coles and Petrongolo (2002), fit the average completed spell of unemployment well, and so setting $\delta^U = 0$ seems a reasonable approximation.

Thus given a functional form $\lambda^M(.)$, parameters θ and period n data $\{U_n, u_n, V_n, v_n, W_n\}$, Equations (1)–(5) can be solved numerically for the five unknowns $\bar{U}_n, \bar{V}_n, \lambda_n, \mu_n, \delta_n^V$. The predicted number of matches is then

$$M_n(\theta) = \lambda_n \bar{U}_n, \tag{6}$$

where the identifying restriction (3) ensures the matching rates of workers are consistent with the matching rates of vacancies.[5] Further predictions of interest are unemployment and vacancy durations, $1/\lambda_n$ and $1/(\mu_n + \delta_n^V)$, respectively, and the proportion of vacancies eventually withdrawn $\delta_n^V/(\mu_n + \delta_n^V)$.

Given data on actual matches, Section 5 provides maximum likelihood estimates of θ. The next section, however, derives the identifying equations for stock-flow matching.

3.2. *Temporal Aggregation with Stock-Flow Matching.* Suppose now that during month n, proportion p_n of newly unemployed workers are on the short side of their markets and re-match immediately. All other unemployed workers are on the long side and again consider a representative worker who matches at average rate λ_n and who withdraws into nonparticipation at rate δ_n^U. If all match according to these parameters, then total expected matches over month n are

[5] In contrast, Gregg and Petrongolo (2005) do not compute these at-risk measures and instead estimate (3) assuming $\lambda = \lambda^M(U_n, V_n; \theta)$, which ignores the matching effects due to the inflow of new vacancies.

$$M_n = \int_0^1 U_n e^{-(\lambda_n+\delta_U)t}\lambda_n\,dt + p_n u_n + \int_{x=0}^1 \left[\int_{t=x}^1 e^{-(\lambda_n+\delta_U)(t-x)}\lambda_n\,dt\right](1-p_n)u_n\,dx.$$

The first term again describes the number of workers in the original stock who match within the month, the second now describes those in the unemployment inflow who are on the short side of their markets and so immediately re-match, and the third describes those in the unemployment inflow who are on the long side of their markets but are sufficiently fortunate to re-match before the end of the month. Integration implies the temporally aggregated matching function

$$M_n = \lambda_n U_n \frac{\left[1-e^{-(\lambda_n+\delta_U)}\right]}{\lambda_n+\delta_U} + p_n u_n + \frac{\lambda_n(1-p_n)u_n}{\lambda_n+\delta_U}\left[1 - \frac{1}{\lambda_n+\delta_U} + \frac{e^{-(\lambda_n+\delta_U)}}{\lambda_n+\delta_U}\right].$$

In this case the appropriate "at-risk" measure of (long side) unemployment is

$$\bar{U}_n = U_n\left[\frac{1-e^{-(\lambda_n+\delta_U)}}{\lambda_n+\delta_U}\right] + (1-p_n)u_n\left[\frac{e^{-(\lambda_n+\delta_U)} - 1 + (\lambda_n+\delta_U)}{(\lambda_n+\delta_U)^2}\right] \tag{7}$$

and expected matches in month n are given by

$$M_n = \lambda_n\bar{U}_n + p_n u_n.$$

Note, stock-flow matching implies matches can be decomposed into those workers who are on the long side of the market and who match slowly (at average rate λ_n) and those newly unemployed who are on the short side and re-match very quickly (proportion p_n).

The above temporal aggregation argument also applies to vacancies. If, in month n, proportion q_n of new vacancies are on the short side and match immediately, whereas vacancies on the long-side match at average rate μ_n and are withdrawn at rate δ_n^V, then the appropriate "at-risk" measure of (long-side) vacancies is

$$\bar{V}_n = V_n\left[\frac{1-e^{-(\mu_n+\delta_V)}}{\mu_n+\delta_V}\right] + (1-q_n)v_n\left[\frac{e^{-(\mu_n+\delta_V)} - 1 + (\mu_n+\delta_V)}{(\mu_n+\delta_V)^2}\right] \tag{8}$$

and expected matches are

$$M_n = \mu_n\bar{V}_n + q_n v_n.$$

As in the random matching case, the expected number of unemployed workers who match must equal the expected number of vacancies that match. But stock-flow matching implies two identifying restrictions. First as the stock of (long side) unemployed workers matches with the inflow of new vacancies on the short side, we have

$$q_n v_n = \lambda_n\bar{U}_n. \tag{9}$$

Second, newly unemployed workers on the short-side match with the stock of vacancies on the long side and so

$$p_n u_n = \mu_n \bar{V}_n. \tag{10}$$

A closed form econometric structure is then obtained by specifying

$$\lambda_n = \lambda^{\mathrm{SF}}(v_n, \bar{U}_n; \theta) \tag{11}$$

$$p_n = p^{\mathrm{SF}}(u_n, \bar{V}_n; \theta) \tag{12}$$

which imply that the stock of (long side) unemployed workers $\bar{U}_n$ matches with the inflow of new vacancies v_n, whereas proportion p_n of (short side) newly laid-off workers u_n re-match with current vacancies $\bar{V}_n$.

We again use Equation (5), i.e., $W_n = \delta_n^V \bar{V}_n$, and data on vacancy withdrawals to identify δ_n^V, and set $\delta_n^U = 0$. Given functional forms for $\lambda^{\mathrm{SF}}(.)$, $p^{\mathrm{SF}}(.)$, parameters θ, and period n data, Equations (5) and (7)–(12) jointly determine $(\bar{U}_n, \bar{V}_n, \lambda_n, \mu_n, p_n, q_n, \delta_n^V)$. Expected matches are then

$$M_n(\theta) = \lambda_n \bar{U}_n + p_n u_n \tag{13}$$

and we use maximum likelihood techniques to estimate θ.

Predicted unemployment and vacancy durations are now $(1 - p_n)/\lambda_n$ and $\mu_n(1 - q_n)/[(\mu_n + q_n \delta_n^V)(\mu_n + \delta_n^V)]$,[6] and the proportion of vacancies eventually withdrawn is given by $(1 - q_n)\delta_n^V/(\mu_n + \delta_n^V)$.

4. THE DATA

Construction of the "at-risk" measures $\bar{U}_n$, $\bar{V}_n$ requires data that distinguish between flows (u, v) and stocks (U, V). Using inches of help-wanted advertisements to measure vacancies, as is the general procedure for the United States,[7] is not sufficient as there is no information on whether a particular job advertisement is new or is a re-advertisement. Such information is however available for the U.K. labor market, as registered at *Jobcentres*.

The U.K. Jobcentre system is a network of government funded employment agencies, where each town or city typically has at least one Jobcentre. A Jobcentre's services are free of charge to all users, both to job seekers and to firms advertising

[6] Conditional on being filled, the expected completed spell is

$$\frac{q[0] + (1-q)\displaystyle\int_0^\infty t e^{-(\mu+\delta^V)t}\mu\, dt}{\dfrac{\mu + q\delta^V}{\mu + \delta^V}} = \frac{\mu(1-q)}{(\mu + q\delta^V)(\mu + \delta^V)}.$$

[7] See Abraham (1987) for a description of U.S. vacancy data.

vacancies. To be entitled to receive welfare payments, an unemployed benefit claimant in the United Kingdom is required to register at a Jobcentre.[8]

The vast majority of Jobcentre vacancy advertisements are for unskilled and semi-skilled workers. Certainly the professionally trained are unlikely to find suitable jobs there. Nevertheless, as the bulk of unemployment is experienced by unskilled and semi-skilled workers, instead of by professionals, it seems reasonable that understanding the determinants of re-employment hazard rates at this level of matching provides useful differentiating information between competing theories of unemployment.

The data we use are monthly time series running from September 1985 to April 2001 (188 observations). The choice of the sample period is driven by data availability. Flow data are not available on a monthly basis before 1985, and vacancy data were discontinued in April 2001.

The unemployment and vacancy counts are usually made on the second Thursday and on the first Friday of each month, respectively. The monthly timing throughout the article thus does not strictly refer to calendar months. The data record the number of unemployed workers and the number of unfilled vacancies at each count date, U_n and V_n, respectively, and the number of new job seekers and vacancies that register between two consecutive count dates, u_n and v_n, respectively. The data also record the number of workers who leave unemployment, M_n^U, the number of vacancies that leave the register, M_n^V, and the number of those that are filled at a Jobcentre, $\tilde{M}_n^V$. The difference between these two measures is accounted for by vacancies that are withdrawn, either because they are cancelled, or because they are filled through another search channel. Any matching measure M_n^U, M_n^V or $\tilde{M}_n^V$ could in principle be used as our dependent variable in Equations (6) and (13). Finally, since months have different lengths, all flow variables are standardized to a 4.33 week accounting month (see also Berman, 1997, pp. S270–1). In our notation, all stocks indexed by n are measured at the beginning of period (month) n, whereas all flows indexed by n are measured between the start of period n and the start of period $n + 1$.

All data used are extracted from the Nomis databank (http://www.nomisweb.co.uk/) and not seasonally adjusted. The main unemployment and vacancy series are plotted in Figures 1 and 2, respectively. The unemployment claimant count declines substantially during our sample period, from about three to one million individuals in a 16-year span. Although it is possible to detect a long cycle between 1986 and 1994, the second half of our sample period is characterized by a secular fall in unemployment figures. The flows experience a more modest decline than the stock and tend to be much more volatile. The vacancy series is clearly negatively correlated to the unemployment series, although its level is about one order of magnitude smaller. The end of our sample period is characterized by both record low unemployment and record high vacancy rates. The vacancy flows are very volatile, display strong seasonality, and are roughly untrended.

[8] Gregg and Wadsworth (1996) report that Jobcentres are used by roughly 80–90% of the claimant unemployed, 25–30% of employed job seekers, and 50% of employers.

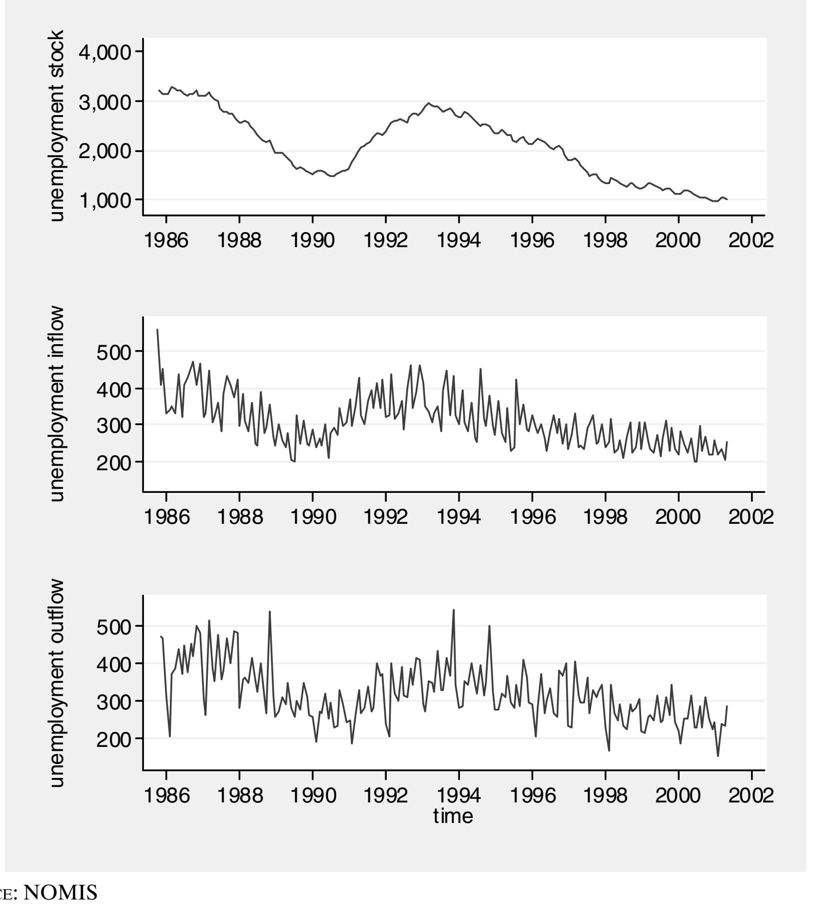

SOURCE: NOMIS

FIGURE 1

THE UNEMPLOYMENT STOCK, INFLOW AND OUTFLOW IN BRITAIN, 1985:09–2001:04

The main sample statistics of these series are summarized in Table 1. As one would expect from Figures 1 and 2, stock variables display a much higher degree of persistence than flow variables, as shown by the monthly autocorrelation figures. The data also show a much higher turnover rate for vacancies than for the unemployed: the relevant monthly inflow/stock ratio being 0.158 for the unemployed and 1.081 for vacancies. Although information on duration, as obtained from the NOMIS, is relatively limited[9], we report duration data in the same table. The average duration of unemployment is around 6.5 months and for filled vacancies it is about three quarters of a month, or 22 days. This latter figure is the duration

[9] Unemployment duration is obtained from information on the number of workers who leave the register, disaggregated by 16 duration classes. We assign to each duration class its middle value, having closed the last open class of 6+ years at 7 years. The duration of filled vacancies is only available quarterly from 1986Q2 to 2000Q2.

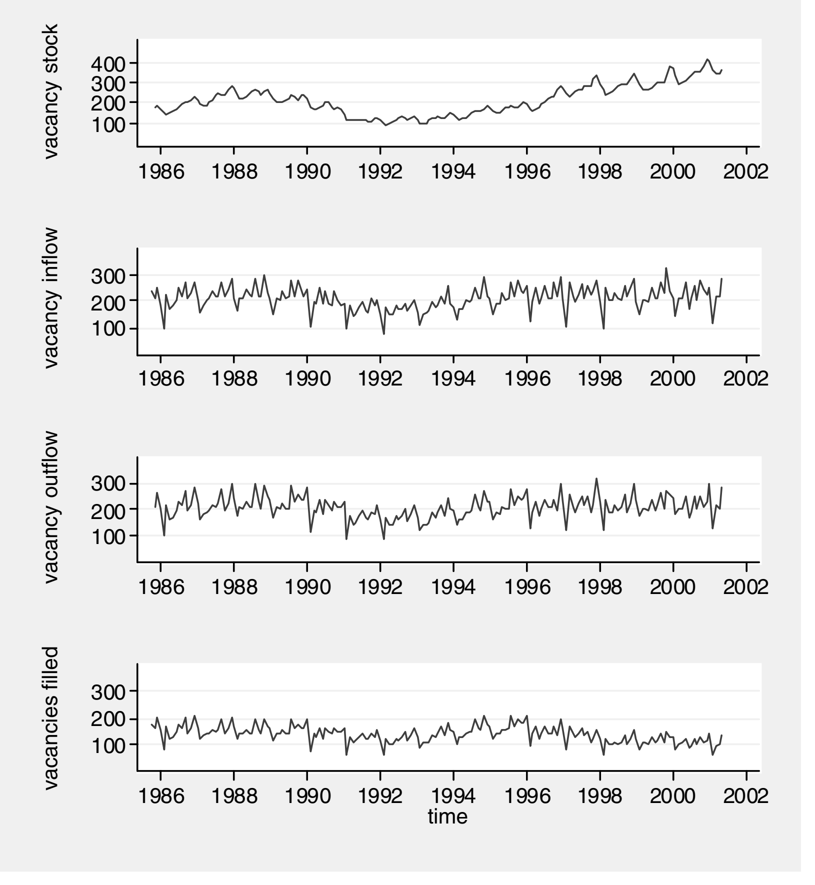

SOURCE: NOMIS

FIGURE 2

THE VACANCY STOCK, INFLOW AND OUTFLOW, AND FILLED VACANCIES IN BRITAIN, 1985:09–2001:04

between the time when a vacancy is first posted and when it is filled by a worker and does not include vacancies that are not filled at Jobcentres.

As also suggested by Figures 1 and 2, the bottom part of Table 1 shows that the monthly vacancy outflow is very highly correlated with the inflow of new vacancies and more weakly correlated with the vacancy stock. Correlation coefficients are 0.93 and 0.46, respectively. When only including vacancies that are filled, the correlation between vacancies filled and the vacancy stock becomes negative (although very small), whereas the one between filled and new vacancies stays high at 0.72. For the unemployed, the correlation coefficient between the inflow and the outflow is 0.70, and the one between the outflow and the stock is 0.60. Both persistence and correlation statistics from this table, as well as visual inspection of the time series in Figures 1 and 2, suggest the bulk of matches are driven by the inflow of new agents.

TABLE 1

MAIN STOCHASTIC PROPERTIES OF VARIABLES

	Mean	Standard Deviation	Monthly Autocorr.	Min	Max	ADF(4)	ADF(12)
Stocks and flows							
U_n	2074956	671559	0.996	960571	3282024	−1.341	−4.694
u_n	309168	69822	0.494	199396	559044	−2.459	−1.482
M_n^U	319880	75288	0.431	154172	545019	−2.436	−1.130
V_n	213288	76836	0.979	90118	413447	−0.852	−1.357
v_n	209186	42872	0.235	78492	322271	−4.560	−0.689
M_n^V	207979	41118	0.189	87325	314562	−3.840	−0.615
$\tilde{M}_n^V$	140115	31007	0.372	59404	212264	−4.126	−0.753
Turnover							
u_n/U_n	0.158	0.041	0.633	0.098	0.286	−1.529	0.346
v_n/V_n	1.081	0.337	0.651	0.316	2.020	−1.272	−0.040
Durations							
Unemployment	6.567	1.133	0.956	4.664	8.677	−1.089	−2.843
Vacancies	0.727	0.111	0.989	0.535	0.953	−0.931	−1.124

Correlations

$\rho\ (M_n^U, U_n) = 0.601$	$\rho\ (M_n^U, u_n) = 0.696$	$\rho\ (M_n^U, M_n^V) = 0.342$
$\rho\ (M_n^V, V_n) = 0.464$	$\rho\ (M_n^V, v_n) = 0.930$	$\rho(M_n^V, \tilde{M}_n^V) = 0.762$
$\rho(\tilde{M}_n^V, V_n) = -0.115$	$\rho(\tilde{M}_n^V, V_n) = 0.720$	$\rho(\tilde{M}_n^V, M_n^U) = 0.625$

NOTES. Sample period: 1985:09–2001:04. ADF(4) and ADF(12) represent Augmented Dickey–Fuller statistics for a unit-root in the relevant series, with 4 and 12 lags, respectively. The 5% critical value is −2.88.

There are a number of data issues to be addressed here. A simultaneity bias in matching function estimates typically arises when M_n is measured as a flow over a time period and U_n and V_n as stocks at some point during the period (or monthly stock averages). In this case measured stock variables are depleted by matches M_n, and this generates a downward bias in the coefficients of interest. See the related discussions by Burdett et al. (1994) and Berman (1997, pp. S272–3). Constructing "at-risk" measures solves this problem. Specifically the at-risk measures $\bar{U}_n$ and $\bar{V}_n$ are a weighted sum of initial stocks at the beginning of the month and inflows within the month. These at-risk measures are not affected by the realized number of matches and so there is no feedback from the unemployment outflow to right-hand side variables.[10]

Our estimates as obtained using the unemployment outflow are the dependent variable. Results using vacancy outflow as the dependent variable are very similar and, if anything, provide stronger evidence in favor of stock-flow matching. We prefer the results using unemployment outflow as the count date for unemployment is the second Thursday of each month, whereas for vacancies it is the first

[10] A related issue is whether a higher than expected matching rate induces more vacancies and unemployed workers to flow into the market within that same month. But this would rest on a degree of labor market responsiveness that is hard to imagine, as contemporaneous hiring rates are unobserved and the relevant information is only released the following month by the Department for Work and Pensions.

Friday of each month. As the unemployment stock series is counted four working days after the vacancy count, using vacancy outflow as the dependent variable would re-introduce the simultaneity problem described above—a high vacancy outflow yields a lower measured unemployment stock. This simultaneity problem does not arise if instead we use unemployment outflow as our dependent variable.

Simple OLS estimates of the random matching function, using unemployment outflow as the left-hand side variable, generates results that are reasonably consistent with the literature. In particular, estimating a log-linear matching function à la Blanchard and Diamond (1989) gives

$$\ln M_n = \underset{(3.819)}{-3.385} + \underset{(0.147)}{0.250}\ln V_n + \underset{(0.194)}{0.859}\ln U_n, + \varepsilon_n, \tag{14}$$

where constant returns are not rejected ($F = 0.15$) and $R^2 = 0.749$ (the regression includes both monthly and yearly dummies, with standard errors reported in parentheses). These results are fairly close to those obtained by Pissarides (1986) on a similar log linear specification for the United Kingdom.[11]

A second issue concerns the identifying assumption made in Section 3.1, namely, that the inflow of unemployed workers and vacancies is uniform within each month. This assumption is important as it yields a tractable empirical structure.[12] The potential problem with this assumption is the following. Suppose that the assumption is violated such that, say, most labor contracts are terminated at the end of a month, and thus most newly unemployed workers register at Jobcentres on the first day of each month. Similarly, employers might register most new vacancies at the beginning of each month. If our stock variables were also measured as those surviving from the end of the previous month, we would end up classifying as variations in inflows what are essentially variations in stocks. This would be the worst-case scenario for our identifying assumption and would bias our estimates in favor of stock-flow matching. More generally, if there is a systematic downward (upward) trend in inflow rates between two count dates, we would tend to bias our estimates in favor (against) stock-flow matching.

But it easy to argue that these scenarios cannot arise in our data. In our data set, we do not have information on daily or weekly inflows, but we can use information from other data sets to investigate time patterns of inflow rates. In particular, the Survey of Incomes In and Out of Work records the day of signing on for British unemployed workers who registered at selected unemployment benefit offices in

[11] Note that in both Equation (14) and Pissarides (1986) the vacancy elasticity is lower and the unemployment elasticity is higher than in the findings of Blanchard and Diamond (1989), who obtain estimates around 0.6 and 0.4, respectively. These differences are due to the different choice of dependent variable; see Petrongolo and Pissarides (2001, section 4.2) for a discussion.

[12] Berman (1997) also uses this uniform inflow assumption to assess temporal aggregation bias in matching function estimates. Berman simulates a daily matching process using the assumption of uniform inflows during each month, and then compares true underlying elasticities with respect to vacancies and unemployment with the estimates obtained on monthly data. He finds the size of the resulting aggregation bias to be negligible.

the four weeks starting March 16, 1987.[13] The data show a regular weekly pattern, with 50% of workers joining on Mondays, 15% on Tuesdays, and so on, down to 0.5% on Saturdays. There is also evidence of a small spike in enrollment rates on April 1. But it should be noted that neither the enrollment spikes on Mondays nor the spike at the beginning of April coincide with our count dates (which for unemployment is the second Thursday of each month and for vacancies is the first Friday of each month). To conclude, as there does not seem to be a trend in inflow rates between two consecutive count dates, the assumption of uniform inflow during each month does not systematically bias our results either in favor or against stock-flow matching.

A third issue is that vacancies advertised at Jobcentres are only a fraction of existing job openings. Unfortunately there is no readily available information on total job vacancies in the United Kingdom during our sample period, but a very crude measure of the fraction of Jobcentre vacancies can be obtained by looking at the fraction of total hires represented by vacancies filled at Jobcentres. Assuming that all employment variations happen through variations in unemployment, total hires from unemployment can be proxied by $H_q = u_q + \Delta N_q$, where ΔN_q is the net quarterly change in aggregate employment and u_q is the quarterly inflow into unemployment.[14] The ratio between the number of vacancies filled at Jobcentres and the constructed series of hires is on average 0.447 during 1985:09–2001:04, and it does not display a definite trend or a significant correlation with the business cycle, as measured by the vacancy/unemployment ratio. If one assumes that the fraction of hires that happen via Jobcentres is a good proxy of the fraction of vacancies advertised there, one can thus rescale both Jobcentre vacancy measures V_n, and V_n, by dividing through by 0.447. This is the adjustment procedure that we use in our estimates. Such rescaling, however, is largely cosmetic: As we use log-linear functional forms for our matching rates λ_n and p_n (see Equations (15)–(17) below), the scale factor only shows up in the constant term and simply allows us to better replicate the observed unemployment duration in our estimates.[15]

Independent information on the total number of job openings in the United Kingdom comes from a recent enterprise-based survey, conducted by the Office for National Statistics between April 2001 and August 2002 (see Machin and Christian, 2002, for details). An additional question that was included in the survey in May 2002 revealed that 44% of vacancies reported to the ONS had also been registered at a Jobcentre—although due to sampling error a more conservative estimate would range between a third and one half. The 44% figure is

[13] Unfortunately there are no similar data available for vacancies.

[14] This follows from the identity $\Delta N_q = H_q - u_q$, where u_q proxies job separations into unemployment. The employment data are also extracted from NOMIS, and they are only available on a quarterly basis.

[15] A different approach would be to note that, if vacancies filled at Jobcentres described total matches, then M_n^V/U_n would be the average exit rate out of unemployment in month n, and hence U_n/M_n^V would be the average expected duration of unemployment. Computing this statistic implies an average duration of unemployment around 14.8 months. In contrast, the actual average duration of unemployment for this period is around 6.5 months. This ratio, [6.5]:[14.8] equals 0.44, and so suggests that Jobcentre vacancies account for 44% of the total in the United Kingdom.

(surprisingly) exactly the same as the one we find based on our constructed measure of hires.

One final issue concerns the nonstationarity of a few variables. Several of the series used are not stationary, as shown by the ADF statistics reported in the last two columns of Table 1, using 4 and 12 lags, respectively, in the ADF test. The matching structure defined above describes (very) short-run variations in matching rates due to short-run variations in labor market conditions. It cannot be used to explain long-run matching trends due to, say, changes in the matching technology, government policies, the composition of the workforce or regional migration. To focus on short-run variations on observed matching rates, our first (and simplest) approach is to include year dummies. However, it may be argued that the use of time dummies generates discontinuous breaks at arbitrary discontinuity points and at a frequency that may be potentially too low. An alternative that partly addresses these issues would consist in using a polynomial trend term instead of year dummies. The estimates we obtained including a cubic trend were very similar to those that include year dummies and are thus not reported.

Our second approach consists in estimating unemployment outflow equations on detrended data, as already done in the matching literature by Yashiv (2000). We thus filter all time series using a Hodrick and Prescott (1997) filter with a smoothing parameter equal to 14,400, and to preserve series means we add to the detrended series their sample averages. Although filtered series for all flow variables are almost indistinguishable from the raw series (there is virtually no trend in flow variables), it is interesting to compare the raw and filtered series for the stocks, as done in Figure 3. As the data show, the filtered series get rid of the downward trend in unemployment in the second half of the sample period and the increasing trend in vacancies, but very clearly reproduce the short-run fluctuations.[16]

Most of the discussion that follows focuses on the nonfiltered data with year dummies. At the end we discuss the results using filtered data instead and show that the insights and conclusions are qualitatively very similar.

5. RESULTS

5.1. *Random Matching.* Recall that given functional form $\lambda_n = \lambda^M(\bar{U}_n, \bar{V}_n; \theta)$ and parameters θ, then period n data and Equations (1)–(5) yield at-risk measures $\bar{U}_n$, $\bar{V}_n$ and period n expected matches

$$M_n(\theta) = \lambda^M(\bar{U}_n, \bar{V}_n; \theta)\bar{U}_n.$$

Assuming realized matches $M_n = M_n(\theta) + \varepsilon_n$ where ε_n are normally and independently distributed with mean zero and finite variance, the maximum likelihood estimator for θ minimizes $\sum_n [M_n - M_n(\theta)]^2$.

[16] We also detrended the data by log first differencing, and the results were very similar to those obtained under HP filtering and thus they are not reported here.

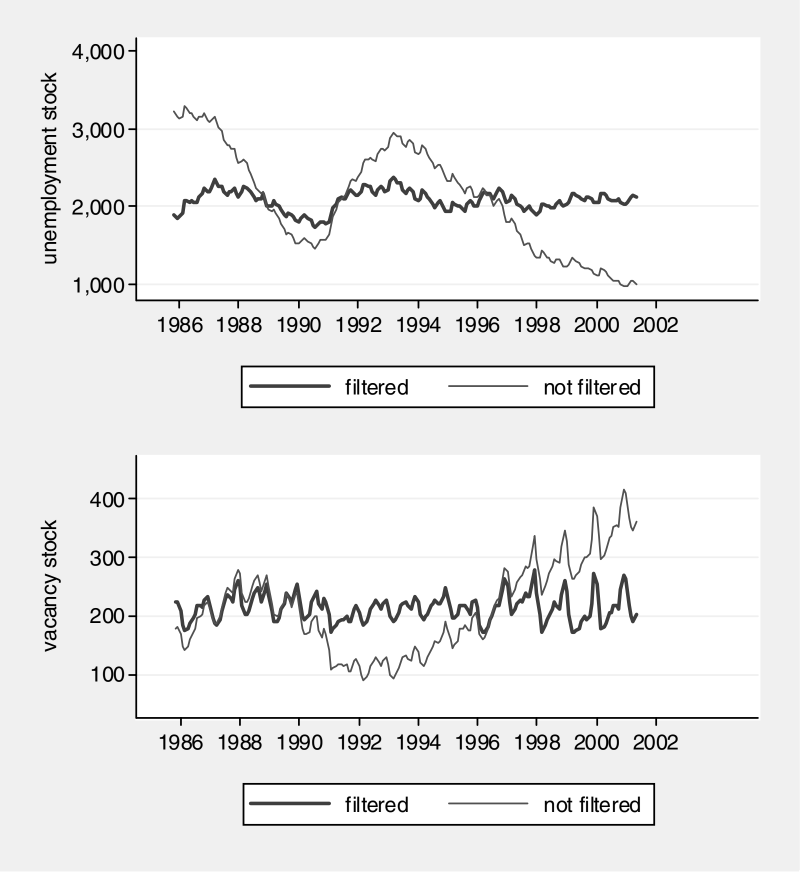

FIGURE 3

UNEMPLOYMENT AND VACANCY STOCKS; RAW AND HODRICK–PRESCOTT FILTERED SERIES

We identify the MLE using the following iterative procedure. Given initial parameter values θ_0 and for each period $n = 1, \ldots, 188$, step 1 solves numerically Equations (1)–(5) for "at-risk" values that we might denote $\bar{U}_n^0$, $\bar{V}_n^0$. With $\theta = \theta_0$, this implies expected matches

$$M_n^0(\theta) = \lambda^M(\bar{U}_n^0, \bar{V}_n^0; \theta)\bar{U}_n^0.$$

Given measures $\bar{U}_n^0$, $\bar{V}_n^0$, Step 2 identifies an updated value for θ by minimizing $\sum_n [M_n - M_n^0(\theta)]^2$. This updated value of θ is then used in step 1 to obtain new "at-risk" measures, and the procedure is iterated to convergence. It is straightforward to show this iterative procedure, if it converges, identifies the MLE.[17]

[17] To account for heteroskedasticity, the covariance matrix of $\hat{\theta}$ can be estimated as $\hat{V}(\hat{\theta}) = (\hat{X}'\hat{X})^{-1}(\hat{X}'\hat{\Omega}\hat{X})(\hat{X}'\hat{X})^{-1}$, where $\hat{X}$ is the matrix of partial derivatives of the regression function $M_n(\theta)$

TABLE 2

ESTIMATION RESULTS UNDER RANDOM MATCHING—RAW DATA

		1	2	3	4	5
$\ln \lambda_n$	Constant	−1.308	−1.163	−1.344	−1.308	−0.755
		(4.053)	(0.149)	(1.841)	(1.500)	(0.054)
	$\ln \bar{V}_n$	0.489	0.476	−0.148	–	–
		(0.148)	(0.058)	(0.064)		
	$\ln v_n$	–	–	0.746	0.613	0.600
				(0.056)	(0.053)	(0.020)
	$\ln \bar{U}_n$	−0.475	−0.476[a]	−0.563	−0.573	−0.600[a]
		(0.165)		(0.072)	(0.062)	
R^2		0.850	0.850	0.950	0.945	0.944
CRS[b]		0.003	–	0.068	0.308	–
Monthly dummies = 0[c]		119.5	144.4	164.9	173.0	189.2
Yearly dummies = 0[d]		22.7	26.7	107.0	104.1	104.6
ADF[e]		−3.680	−3.702	−3.364	−3.207	−3.143
Sample averages						
λ_n		0.170	0.166	0.193	0.211	0.210
$1/\lambda_n$		6.546	6.657	5.477	5.087	5.144
μ_n		0.703	0.703	1.054	1.127	1.095
δ_n		0.315	0.317	0.366	0.376	0.368
$1/(\mu_n + \delta_n)$		1.035	1.036	1.055	0.741	0.897

NOTES. Monthly data not seasonally adjusted. Dependent variable: unemployment outflow recorded at Jobcentres in Great Britain. Sample period: 1985:09–2001:04. All specifications include month and year dummies. Estimation method: maximum likelihood. Heteroskedastic-consistent standard errors (White, 1980) are reported in parentheses.

λ_n denotes the unemployment hazard rate. $1/\lambda_n$ denotes the predicted unemployment duration. μ_n denotes the vacancy hazard rate of being filled. δ_n denotes the vacancy hazard rate of withdrawal. $1/(\mu_n + \delta_n)$ denotes the predicted vacancy duration. No. Observations: 188. Source: NOMIS.

[a]Coefficient constrained to equal the value reported.

[b]Wald test, distributed as $\chi^2(1)$, of the hypothesis that the sum of the coefficients on $\ln \bar{V}_n$, $\ln v_n$ and $\ln \bar{U}_n$ is zero. Critical value at 5% significance level: $\chi^2(1) = 3.841$.

[c]Wald test, distributed as $\chi^2(11)$, of the hypothesis that monthly dummies are jointly zero. Critical value at 5% significance level: $\chi^2(11) = 19.675$.

[d]Wald test, distributed as $\chi^2(17)$, of the hypothesis that yearly dumies are jointly zero. Critical value at 5% singificance level: $\chi^2(17) = 27.587$.

[e]ADF statistics for the presence of a unit root in the estimated residuals. Critical value at 5% significance level: −2.88.

Given the identifying restrictions for random matching, Table 2 describes the ML estimates using various functional forms for $\lambda_n = \lambda^M(.)$. All specifications include year dummies, and as the data are not seasonally adjusted, monthly dummies, which turn out to be jointly significant in all specifications.[18]

with respect to right-hand side variables, and $\hat{\Omega}$ is a diagonal matrix with the nth diagonal element equal to the nth least squared residual (White, 1980), all evaluated at $\hat{\theta}$. To cater for small sample bias of such variance estimator, the nth diagonal element of $\hat{\Omega}$ can be further divided by $1 - \hat{X}_n(\hat{X}'\hat{X})^{-1}\hat{X}_n'$ (see Davidson and MacKinnon, 1993, p. 553). This is the approach that we adopt. Very similar results are obtained however using the uncorrected covariance estimator, $\hat{V}(\hat{\theta}) = \hat{\sigma}^2(\hat{X}'\hat{X})^{-1}$, where $\hat{\sigma}^2$ is the sum of squared residuals divided by the number of observations.

[18] The exact specification used for predicted matches is $M_n(\theta) = \lambda_n \bar{U}_n +$ dummies.

Column 1 assumes the standard Cobb–Douglas specification

$$\lambda_n = \exp[\alpha_0 + \alpha_1 \ln \bar{V}_n + \alpha_2 \ln \bar{U}_n]. \tag{15}$$

The coefficients on (time-aggregated) vacancies and unemployment have the expected sign and are significantly different from zero. Estimated matching elasticities around 0.5 are in line with previous matching function estimates (see Petrongolo and Pissarides, 2001), and constant returns to scale in the matching function are not rejected, given a virtually zero Wald test statistics on the restriction $\alpha_1 = -\alpha_2$. The extremely low value of this test statistic, however, together with a nonsignificant constant term in λ_n makes one doubt that the elasticities on $\bar{V}_n$ and $\bar{U}_n$ are separately identified. In Column 2, we impose constant returns to scale: The constant term is now precisely determined, and the goodness of fit remains unchanged. To compare outcomes with the sample means, note that if a worker matches at constant rate λ, the expected completed unemployment spell is $1/\lambda$. Outside of steady state, the average value of $1/\lambda_n$ over the entire sample would seem a useful proxy for the average unemployment spell. In Columns 1 and 2, the sample average of $1/\lambda_n$ is about six and a half months, which is very close to the actual unemployment duration during the sample period (6.657 months). Vacancies can either be filled at rate 0.7 by claimant unemployed or withdrawn at rate 0.3, implying 70% of vacancies advertised at Jobcentres are successfully filled. The predicted average completed vacancy spell (computed as the average value of $1/(\mu_n + \delta_n^V)$) is about one month, which slightly overstates by about 9 days the actual duration of filled vacancies.

Comparing these results with those obtained using OLS (see Equation (14)) finds that the estimated vacancy coefficient is much larger (0.49 instead of 0.25), is highly significant, and the fit is much improved, with an R^2 of 0.85 instead of 0.75. Ignoring temporal aggregation as in (14) implies thus a significant downward bias in the vacancy coefficient. The reason is that the initial vacancy stock V_n is a poor proxy for the total number of vacancies at risk over the month—it ignores the new vacancies that enter the market within the month. For example, Equation (2) with $\mu = 0.7$, $\delta^V = 0.3$ implies the appropriate vacancy at-risk measure is $\bar{V}_n = 0.73V_n + 0.37v_n$. As the average monthly inflow to stock ratio, v_n/V_n, is large (equal to 1.08), V_n is then a poor proxy for $\bar{V}_n$. Further as the unemployment outflow is highly correlated with the vacancy inflow during the month, correcting for the temporal aggregation bias results in a much better fit and a higher estimated vacancy coefficient.

Column 3 is a test of an overidentifying restriction—that random matching implies the matching rate of individual workers does not depend directly on the inflow of new vacancies (once one controls for time aggregation). Column 3 thus asks whether including the flow of new vacancies as an added explanatory variable for λ_n improves the fit. In fact the fit is not only much improved, the vacancy stock coefficient becomes wrong signed. Column 4 drops the vacancy stock term and the fit is essentially unchanged. In both Columns 3 and 4 constant returns in the matching function are not rejected, and this restriction is again imposed in Column 5.

Lagos (2000) provides a useful perspective for this result. He shows that even with stock flow matching at the microlevel, aggregation over locations may yield, in a *steady state,* aggregate matching behavior that appears consistent with a standard random matching function. Column 2 establishes that a standard random matching function fits the aggregate data well and implies constant returns. But the identifying equations are not consistent with the "out-of-steady-state" matching dynamics. The next set of results establish that the stock-flow identifying equations explain better the observed time series variation in unemployment outflow.

5.2. *Stock-Flow Matching.* The regression equation is now given by

$$M_n = \lambda_n \bar{U}_n + p_n u_n + \varepsilon_n,$$

where $\bar{U}_n$ is defined in (7), λ_n is defined in (11), and p_n is defined in (12).

The estimating procedure is again iterative as explained above, but we are now solving (5) and (7)–(12) for $\bar{U}_n$, $\bar{V}_n$, λ_n, μ_n, p_n, q_n, δ_n^V for each n, given θ_0, and θ is again obtained using ML techniques.

The results for stock-flow matching are reported in Table 3 under alternative specifications for $\lambda_n = \lambda^{\mathrm{SF}}(.)$ and $p_n = p^{\mathrm{SF}}(.)$. Recall that in contrast to random matching, stock-flow matching implies that λ_n depends on the vacancy inflow and not on the stock of vacancies. The pure job queueing hypothesis in addition predicts $p_n = 0$.

Column 1 adopts the functional form

$$\lambda_n = \exp\left(\alpha_0 + \alpha_1 \ln \bar{V}_n + \alpha_2 \ln v_n + \alpha_3 \ln \bar{U}_n\right) \tag{16}$$

whereas p_n is estimated as a constant parameter, and constrained to be nonnegative, i.e., $p_n = \exp(\beta_0)$. Moreover, as in no specification was the hypothesis of constant returns in the matching function rejected, we simply report estimates that impose constant returns, i.e., $\alpha_1 + \alpha_2 + \alpha_3 = 0$.

Consistent with stock-flow matching, Column 1 in Table 3 finds that λ_n is driven by the inflow of new vacancies and that the vacancy stock effect is not significantly different from zero (and wrong-signed). Column 2 drops the vacancy stock from the specification of λ_n and re-estimates. The results show that the exit rates of the longer term unemployed, λ_n, are driven by the inflow of new vacancies with an estimated elasticity around 0.75. Further, the pure job queueing hypothesis is rejected—the mean matching probability of the newly unemployed, p_n, is 0.44, and is significantly different from zero, with an estimated standard error of 0.061.[19]

Columns 3–5 consider a more general specification for

$$p_n = \exp\left(\beta_0 + \beta_1 \ln \bar{V}_n + \beta_2 \ln v_n + \beta_3 \ln u_n\right) \tag{17}$$

while leaving the specification of λ_n as in Column 2, which is consistent with the identifying assumptions. In Columns 3–5, we still get a positive p_n, but no

[19] Using the delta method: s.e. $(p_n) = \exp(\beta_0) * \text{s.e.}\ (\beta_0) = 0.061$.

TABLE 3

ESTIMATION RESULTS UNDER STOCK-FLOW MATCHING—RAW DATA

		1	2	3	4	5
$\ln \lambda_n$	Constant	−1.275	−1.258	−1.210	−1.246	−1.289
		(0.150)	(0.138)	(0.219)	(0.147)	(0.442)
	$\ln \bar{V}_n$	−0.124	–	–	–	–
		(0.090)				
	$\ln v_n$	0.857	0.753	0.789	0.778	0.727
		(0.112)	(0.066)	(0.094)	(0.074)	(0.126)
	$\ln \bar{U}_n$	−0.733[a]	−0.753[a]	−0.789[a]	−0.778[a]	−0.727[a]
$\ln p_n$	Constant	−0.871	−0.822	−0.849	−0.841	−0.823
		(0.194)	(0.155)	(0.209)	(0.189)	(0.319)
	$\ln \bar{V}_n$	–	–	−0.101	−0.121	
				(0.117)	(0.209)	
	$\ln v_n$	–	–	−0.040	–	0.068
				(0.305)		(0.414)
	$\ln u_n$			0.141[a]	0.121[a]	−0.068[a]
R^2		0.961	0.960	0.961	0.961	0.960
Monthly dummies = 0[b]		259.0	311.5	279.5	280.2	305.7
Year dummies = 0[c]		126.8	115.8	123.1	122.2	115.0
ADF[d]		−4.566	−5.488	−4.444	−4.703	−3.562
Sample averages						
λ_n		0.104	0.103	0.103	0.101	0.103
p_n		0.419	0.440	0.404	0.411	0.452
$(1-p_n)/\lambda_n$		6.176	6.107	6.479	6.506	5.969
μ_n		0.331	0.345	0.328	0.335	0.350
q_n		0.416	0.407	0.405	0.398	0.409
δ_n		0.333	0.331	0.329	0.329	0.325
$\frac{\mu_n(1-q_n)}{(\mu_n+q\delta_n^V)(\mu_n+\delta_n^V)}$		0.656	0.656	0.667	0.655	0.572

NOTES. Monthly data not seasonally adjusted. Dependent variable: unemployment outflow recorded at Jobcentres in Great Britain. Sample period: 1985:09–2001:04. All specifications include month and year dummies. Estimation method: maximum likelihood. Heteroskedastic-consistent standard errors (White, 1980) are reported in parentheses.

λ_n denotes the unemployment hazard rate. p_n denotes the initial matching rate of the unemployed. $(1-p_n)/\lambda_n$ denotes the predicted unemployment duration. μ_n denotes the vacancy hazard rate of being filled. δ_n denotes the vacancy hazard rate of withdrawal. q_n denotes the initial matching rate of vacancies. $\frac{\mu_n(1-q_n)}{(\mu_n+q_n\delta_n)(\mu_n+\delta_n)}$ denotes the predicted vacancy duration. No. Observations: 188.

SOURCE: NOMIS.

[a]Coefficient constrained to equal the value reported.

[b]Wald test, distributed as $\chi^2(11)$, of the hypothesis that monthly dummies are jointly zero. Critical value at 5% significance level: $\chi^2(11) = 19.675$.

[c]Wald test, distributed as $\chi^2(17)$, of the hypothesis that yearly dumies are jointly zero. Critical value at 5% singificance level: $\chi^2(17) = 27.587$.

[d]ADF statistics for the presence of a unit root in the estimated residuals. Critical value at 5% significance level: −2.88.

variables seem to explain it well. Column 3 is an overidentifying test for stock-flow matching: that the inflow of new vacancies does not explain p_n. This test is accepted—the vacancy inflow coefficient is insignificant (and wrong-signed). Column 4 is the "stock-flow" specification: that p_n depends on the vacancy stock but not the inflow. But the vacancy stock coefficient is also insignificant (and wrong signed). The results for p_n are therefore a little disappointing, but we note throughout that the estimates for λ_n are robust to these variations.

All these specifications provide an identical fit, with an R^2 of 0.96, which is higher than in all specifications in Table 1, even those that (inconsistently) include the vacancy inflow. Table 3 yields the following predicted sample averages for unemployed workers:

(i) $p_n \approx 0.4$; i.e., about 40% of workers entering the unemployment pool are on the short side of their markets and quickly become re-employed;
(ii) 60% of workers entering the unemployment pool are on the long side of their markets and so face an extended spell of unemployment. The sample average $1/\lambda_n$ suggests an average unemployment spell of around 10 months for these workers; and
(iii) in a steady state, the expected average completed unemployment spell is $(1 - p)/\lambda$. The sample average of $(1 - p_n)/\lambda_n \approx 6.1$ months, which is reasonably consistent with the actual average completed unemployment spell (6.5 months).

The predicted sample averages for vacancies are

(i) $q_n \approx 0.4$; i.e., about 40% of new vacancies are on the short side of their markets and are quickly filled (indeed 30% of new vacancies are filled on

TABLE 4

ESTIMATION RESULTS UNDER RANDOM MATCHING—HP FILTERED DATA

		1	2	3	4	5
$\ln \lambda_n$	Constant	−1.371	−0.839	−1.349	−1.320	−0.754
		(6.761)	(0.233)	(4.830)	(4.265)	(0.090)
	$\ln \bar{V}_n$	0.797	0.734	−0.203	–	–
		(0.185)	(0.167)	(0.125)		
	$\ln v_n$	–	–	0.714	0.685	0.681
				(0.072)	(0.066)	(0.066)
	$\ln \bar{U}_n$	−0.753	−0.734[a]	−0.487	−0.645	−0.681[a]
		(0.431)		(0.309)	(0.307)	
R^2		0.628	0.625	0.799	0.794	0.791
CRS[b]		0.019	–	0.007	0.043	–
Monthly dummies = 0[c]		124.5	153.7	119.2	123.7	143.3
Yearly dummies = 0[d]		12.9	12.0	26.5	24.4	30.9
ADF[e]		−3.374	−3.146	−3.252	−3.255	−3.215
Sample averages						
λ_n		0.154	0.152	0.175	0.177	0.175
$1/\lambda_n$		6.561	6.667	5.827	5.795	5.844
μ_n		0.646	0.634	0.773	0.776	0.738
δ_n		0.321	0.320	0.340	0.340	0.339
$1/(\mu_n + \delta_n)$		1.036	1.051	0.921	0.852	0.921

NOTE. See notes to Table 2.

the first day of being posted). Note this is not inconsistent with $p_n = 0.4$ for workers;

(ii) 60% of new vacancies are on the long side and must wait for suitable new workers to come onto the market. As $\mu_n \approx \delta_n^V$ a vacancy on the long side is as likely to be withdrawn as filled. A filled long-side vacancy has average spell $1/(\mu_n + \delta_n^V) = 6$ weeks; and

(iii) in a steady state, the average completed spell of a filled vacancy is $\frac{\mu(1-q)}{(\mu+q\delta^V)(\mu+\delta^V)}$. The sample average of $\frac{\mu_n(1-q_n)}{(\mu_n+q_n\delta_n)(\mu_n+\delta_n)}$ is 20 days, which is 2 days short of the duration of filled vacancies observed in the data.

5.3. *The Results Using HP Filtered Data.* We quickly discuss the results obtained when the data are first passed through an HP filter. Overall the results are qualitatively identical to those obtained on the raw data. For the random matching case, represented in Table 4, Columns 1 and 2 establish the random matching function provides a reasonably good fit of the data and constant returns are not rejected. But including vacancy inflow in Column 3 much improves the fit and the vacancy stock coefficient again becomes wrong-signed (though not significantly so).

TABLE 5

ESTIMATION RESULTS UNDER STOCK-FLOW MATCHING—HP FILTERED DATA

		1	2	3	4	5
$\ln \lambda_n$	Constant	−1.409	−1.371	−1.234	−1.298	−1.325
		(0.481)	(0.246)	(0.449)	(0.206)	(0.703)
	$\ln \bar{V}_n$	−0.089	–	–	–	–
		(0.358)				
	$\ln v_n$	1.218	1.170	0.906	0.883	0.969
		(0.296)	(0.280)	(0.220)	(0.223)	(0.142)
	$\ln \bar{U}_n$	−1.129[a]	−1.170[a]	−0.906[a]	−0.883[a]	−1.003[a]
$\ln p_n$	Constant	−0.467	−0.453	−0.672	−0.670	−0.562
		(0.108)	(0.125)	(0.285)	(0.228)	(0.290)
	$\ln \bar{V}_n$	–	–	−0.188	−0.226	
				(0.357)	(0.327)	
	$\ln v_n$	–	–	−0.177	–	−0.155
				(0.442)		(0.530)
	$\ln u_n$			0.365[a]	0.226[a]	0.155[a]
R^2		0.886	0.885	0.886	0.885	0.885
Monthly dummies = 0[b]		313.2	350.2	288.1	287.6	343.3
Yearly dummies = 0[c]		44.9	44.4	45.4	44.3	45.2
ADF[d]		−4.669	−6.756	−4.748	−4.976	−4.137
Sample averages						
λ_n		0.047	0.046	0.077	0.075	0.065
p_n		0.627	0.636	0.432	0.458	0.529
$(1 - p_n)/\lambda_n$		8.510	8.449	7.601	7.491	7.639
μ_n		0.387	0.392	0.271	0.288	0.334
q_n		0.201	0.197	0.335	0.325	0.278
δ_n		0.313	0.314	0.315	0.316	0.324
$\frac{\mu_n(1-q_n)}{(\mu_n+q_n\delta_n)(\mu_n+\delta_n)}$		1.001	1.001	0.660	0.611	0.473

NOTE. See notes to Table 3.

Table 5 estimates stock flow matching, restricting again to constant-returns specifications. As in Table 3, the estimates of λ_n are robust across all specifications and imply that the stock of longer-term unemployed workers match with the inflow of new vacancies. The elasticity of λ_n with respect to v_n is again higher than that obtained on the raw data, and tends to be not significantly different from one; i.e., the matching rate of the long-term unemployed would seem proportional to the vacancy inflow. The wide variation in estimated sample averages across these specifications, however, suggests these results are not as well identified as the above results using the raw (nonfiltered) data and time dummies.

6. CONCLUSION

Lagos (2000) considers an equilibrium trading framework where matching at the microlevel implies stock-flow matching, but aggregation over locations and a restriction to *steady state* implies that aggregate matching appears consistent with a standard random matching function. Our results are consistent with that view. The random matching function fits the aggregate data reasonably well and constant returns to matching are not rejected. Outside of steady state, however, the matching function approach, as typically used to explain equilibrium unemployment dynamics (e.g., Mortensen and Pissarides, 1994; Shimer, 2005),is inconsistent with the observed turnover dynamics. Specifically, unemployment outflows are too highly correlated with new vacancy inflows.

Obviously these results do not imply search frictions are unimportant in labor markets. An alternative search approach might instead assume a "good jobs"/"bad jobs" scenario, where vacancies for "good jobs" match quickly whereas vacancies for "bad jobs" do not. Unfortunately, such a model cannot be identified on our data. But even so, it is not clear in a search environment that "good" vacancies should necessarily match quickly. For example, all unemployed workers might desire the highly paid CEO vacancy, but the post still takes a long time to fill as shareholders search for the ideal candidate. Equilibrium matching with supermodular match payoffs implies high ability workers will tend to match with high quality vacancies, but being higher ability does not imply faster matching (e.g., Shimer and Smith, 2000; Burdett and Coles, 1999; Gautier and Teulings, 2004, with search frictions). Stock-flow matching instead distinguishes between "short-side" and "long-side" vacancies, where "short-side" vacancies are quickly filled by unemployed workers on the "long side." Such behavior is consistent with the fact that 30% of new vacancies are filled on the first day of being posted, a statistic that is hard to explain from a purely search perspective. Furthermore a "good jobs"/"bad jobs" scenario would suggest that estimated p_n should also be highly correlated with vacancy inflow. Tables 3 and 5 clearly demonstrate this is not the case.

Perhaps the major contribution of this article is that the matching behavior of the longer term unemployed is robustly identified—they wait for suitable new vacancies to come onto the market. The stock-flow view of matching is also useful in reconciling the McDonald's problem: that McDonald's invariably has vacancies and everyone knows where McDonald's is, so how can unemployment be frictional? It cannot be argued that McDonald's jobs are "bad" jobs, otherwise no one would work there.[20] Stock-flow matching instead implies that most unemployed workers are better qualified to do different work and prefer to wait for something more suitable to come onto the market. McDonald's instead hires from the inflow of workers into unemployment, hiring those who have a comparative advantage in working for them. The monopsonist then chooses wages to maximize expected profit given the entrant inflow into unemployment. This approach also provides insights into other market scenarios. For example, the house seller who does not find an immediate buyer might set a lower price to generate a quick sell. But he/she might instead maintain a high asking price and wait for a suitably interested new buyer to come onto the market.

Our results have important policy implications. Search effort in the standard search framework is a productive investment. But with job queueing by workers on the long side of the market, then job chase effort k_{it} is a pure rent seeking investment. Specifically greater aggregate job chase effort K_t does not affect the aggregate outcome: The vacancy on the short side of the market is always filled. Individual job chase effort only determines the probability with which each given individual gets the next available job. There is a large optimal unemployment insurance (UI) literature that argues that, with search frictions and unobserved search effort, unemployment benefit payments should be reduced with unemployment duration to encourage greater search effort (e.g., Shavell and Weiss, 1979). But with job chasing as described above, it is welfare reducing to distort UI payments by duration to induce greater "job chase effort." The relevant policy issue then is not whether unemployment benefits should be stopped after, say 6 months, 1 year, etc. Rather what is the optimal level of unemployment benefits, where unemployment benefits affect worker reservation values and so potentially lead to inefficiently low (or high) job acceptance rates (see Marimon and Zilibotti, 1999; Mortensen and Pissarides, 1999; Shimer and Werning, 2007, for related arguments within the standard search framework)?

Another related policy issue is whether unemployed job seekers should retrain or be encouraged to migrate to employment hot spots. However, it may be that capital is more mobile than labor (given that residential housing is immobile).[21] In that case, it may be more efficient to subsidize capital to move into unemployment black spots and the unemployed should then wait for new capital to arrive in town.

REFERENCES

Abraham, K. G., "Help Wanted Advertising, Job Vacancies and Unemployment," *Brookings Papers on Economic Activity* 1 (1987), 207–43.

[20] Although a popular term, the notion of a "good" or "bad" job is not particularly helpful. From the worker's perspective, a job with given characteristics is either well or badly compensated.

[21] Using information from the British Social Attitudes Surveys 1983–94, Oswald (1997) reports that "among currently unemployed house-owners who thought it would take them at least three months to find work, only 24% ... would be willing to move areas." In contrast, private renters are more geographically mobile and enjoy faster re-employment rates (e.g., Wadsworth, 1998).

BERMAN, E., "Help Wanted, Job Needed: Estimates of a Matching Function from Employment Service Data," *Journal of Labor Economics* 15 (1997), 251–91.

BLANCHARD, O. J., AND P. A. DIAMOND, "The Beveridge Curve," *Brookings Papers on Economic Activity* 1 (1989), 1–76.

BURDETT, K., AND J. VAN OURS, "Long Term Partnership Formation: Marriage and Employment," *Economic Journal*, 109 (1999), 307–34.

——, AND B. CUNNINGHAM, "Towards a Theory of Vacancies," *Journal of Labor Economics* 16 (1998), 445–78.

——, M. G. COLES, AND J. VAN OURS, "Temporal Aggregation Bias in Stock-Flow Models," CEPR Discussion Paper No. 967, 1994.

COLES, M.G., "Turnover Externalities with Marketplace Trading," *International Economic Review* 40 (1999), 851–67.

——, AND A. MUTHOO, "Strategic Bargaining and Competitive Bidding in a Dynamic Market Equilibrium," *Review of Economic Studies* 65 (1998), 235–60.

——, AND B. PETRONGOLO, "A Test Between Unemployment Theories Using Matching Data," CEPR Discussion Paper 3241, 2002.

——, AND E. SMITH, "Marketplaces and Matching," *International Economic Review* 39 (1998), 239–55.

DAVIDSON, R., AND J. MACKINNON, *Estimation and Inference in Econometrics* (Oxford: Oxford University Press, 1993)

EBRAHIMY, E., AND R. SHIMER, "Stock-Flow Matching," Mimeo, University of Chicago, 2006.

GAUTIER, P., AND C. TEULINGS, "The Right Man for the Job," *Review of Economic Studies* 71 (2004), 553–80.

GREGG, P., AND B. PETRONGOLO, "Non-random Matching and the Performance of the Labor Market," *European Economic Review* 49 (2005), 1087–2011.

——, AND J. WADSWORTH, "How Effective Are State Employment Agencies? Jobcentre Use and Job Matching in Britain," *Oxford Bulletin of Economics and Statistics* 58 (1996), 443–68.

HODRICK, R. J., AND E. C. PRESCOTT, "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit, and Banking* 29 (1997), 1–16.

JONES, S. R. G., AND W. C. RIDDELL, "Measurement of Unemployment," *Econometrica* 67 (1999), 147–62.

LAGOS, R., "An Alternative Approach to Search Frictions," *Journal of Political Economy* 108 (2000), 851–73.

MACHIN, S., AND V. CHRISTIAN, "A New Survey of Job Vacancies: The First Experimental Results," *Labor Market Trends* 110 (2002), 535–48.

——, AND A. MANNING, "The Causes and Consequences of Long Term Unemployment in Europe" in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics* (Amsterdam: North-Holland, 1999), 3085–3139.

MARIMON, R., AND F. ZILIBOTTI, "Unemployment vs. Mismatch of Talents: Reconsidering Unemployment Benefits," *Economic Journal* 109 (1999), 462–73.

MORTENSEN, D. T., AND C. PISSARIDES, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies* 61 (1994), 397–415.

——, AND ——, "Unemployment Responses to 'Skill-Biased' Technology Shocks: The Role of Labor Market Policy," *Economic Journal* 109 (1999), 242–65.

OSWALD, A., "The Missing Piece of the Unemployment Puzzle," Inaugural Lecture, University of Warwick, 1997.

PETRONGOLO, B., AND C. A. PISSARIDES, "Looking into the Black-Box: A Survey of the Matching Function," *Journal of Economic Literature* 39 (2001), 390–431.

PISSARIDES, C. A., "Unemployment and Vacancies in Britain," *Economic Policy* 3 (1986), 676–90.

——, *Equilibrium Unemployment Theory*, 2nd edition (Cambridge, MA: MIT Press, 2000).

SHAPIRO, C., AND J. E. STIGLITZ, "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review* 74 (1984), 433–44.

SHAVELL, S., AND L. WEISS, "The Optimal Payment of Unemployment Insurance Benefits over Time," *Journal of Political Economy* 87 (1979), 1347–62.

SHIMER, R., "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review* 95 (2005), 25–49.

——, AND L. SMITH, "Assortative Matching and Search," *Econometrica* 68 (2000), 343–69.

——, AND I. WERNING, "Reservation Wages and Unemployment Insurance," *Quarterly Journal of Economics* 122 (2007), 1145–85.

TAYLOR, C., "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 (1995), 837–55.

WADSWORTH, J., "Eyes Down for a Full House: Employment Polarisation and the Housing Market," *Scottish Journal of Political Economy* 45 (1998), 376–92.

WHITE, H., "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica* 48 (1980), 817–38.

YASHIV, E., "The Determinants of Equilibirum Unemployment," *American Economic Review* 90 (2000), 1297–1322.