Economic Quarterly—Volume 95, Number 2—Spring 2009—Pages 101–120

Estimating a Search and Matching Model of the Aggregate Labor Market

Thomas A. Lubik

The search and matching model has become the workhorse for labor market issues in macroeconomics. It is a conceptually attractive framework as it provides a rationale for the existence of equilibrium unemployment, such that workers who would be willing to work for the prevailing wage cannot find a job. By focusing on the search and matching aspect, that is, workers searching for jobs, firms searching for workers, and both sides being matched with each other, the model also provides a description of employment flows in an economy. Moreover, the search and matching model is tractable enough to be integrated into standard macroeconomic models as an alternative to the perfectly competitive Walrasian labor market model.

However, the search and matching framework has been criticized, most notably by Shimer (2005), for being unable to match key labor market statistics, chiefly the volatility of unemployment and job vacancies. This observation has generated a large amount of research intended to remedy this "puzzle." Most of this literature is largely theoretical and based on calibration. Only recently have there been efforts to more formally study the quantitative implications of the entire search and matching framework. This article is among the first attempts to take a search and matching model to the data in a full-information setting.[1]

In this article I contribute to these efforts by estimating a small search and matching model using Bayesian methods. My focus is mainly on the actual

The author thanks Andreas Hornstein, Marianna Kudlyak, Devin Reilly, Pierre Sarte, and Michael Krause for many helpful comments and discussions. The views expressed in this article do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: thomas.lubik@rich.frb.org.

[1] Other recent contributions are Trigari (2004); Christoffel, Küster, and Linzert (2006); Gertler, Sala, and Trigari (2008); and Krause, López-Salido, and Lubik (2008).

parameter estimates and the implied sources of business cycle fluctuations. Calibrating the search and matching model tends to be problematic since some of the model parameters, such as the bargaining share or the value of staying unemployed, are difficult to pin down. Hence, much of the arguments about the empirical usefulness of the search and matching model center around alternative calibrations. This paper provides some perspective on this issue by adopting a full-information approach in estimating the model. Parameters are selected so as to be consistent with the co-movement patterns in the full data set as seen through the prism of the theoretical model.

My main finding is that the structural parameters of the model are generally tightly estimated and robust across various empirical specifications that include different sets of observables and shock processes. Parameters associated with the matching process tend to be more stable than those associated with the search process. However, I also find that the estimates are consistent with an emerging consensus on the search and matching model (e.g., Hornstein, Krusell, and Violante 2005 and Hagedorn and Manovskii 2008) that emphasizes a low bargaining power but a high outside option for a worker. On a more cautionary note, I show that the most important determinant of labor market dynamics are exogenous movements in the match efficiency, which essentially acts as a residual in an adjustment equation for unemployment. This finding casts doubt on the viability of the search and matching framework to provide a theory for labor market *dynamics*.

In a larger sense, this article also deals with the issue of identification in structural general equilibrium models. I use the term "identification" loosely in that I ask whether the theoretical model contains enough restrictions to back out parameters from the data. In that respect, the search and matching framework performs reasonably well. But identification also has a dimension that may be more relevant for the theoretical modeler. I show that specific parameters, such as the worker benefit or search costs, can vary widely across specifications, and thus are likely not identified in either an econometric or economic sense. I also argue that they capture the stable behavior of an underlying structure. They therefore adapt to a change in the environment and might be better described as reduced-form coefficients.

The article proceeds as follows. In the next section, I lay out a simple search and matching model, followed by a discussion in Section 2 of the empirical strategy and the data used. In Section 3, I present the benchmark estimation results, discuss the estimated dynamics, and investigate the sources of business cycle fluctuations. Section 4 contains various robustness checks that change the set of observables and the exogenous shocks. Section 5 concludes.

1. A SIMPLE SEARCH AND MATCHING MODEL

I develop a simple search and matching model in which the labor market is subject to frictions. Workers and firms cannot meet instantaneously but must go through a time-consuming search process. The costs of finding a partner give rise to rents that firms and workers share between each other. Thus, wages are the outcome of a bargaining process and are not determined competitively. The labor market set-up is embedded in a simple general equilibrium framework with optimizing consumers and firms that serves as a data-generating process for aggregate time series. The model is otherwise standard and has been extensively studied in the literature.[2] I first describe the optimization problems of households and firms, followed by a discussion of the labor market and wage determination.

The Household

Time is discrete and the time period is one quarter. The economy is populated by a continuum of households. Each household consists of a continuum of workers of measure one. Individual households send out their members to the labor market, where they search for jobs when unemployed, and supply labor services when employed. During unemployment the afflicted household member receives government benefits, while all employed workers earn the wage rate. Total income is shared equally among all members. I hereby follow the literature and abstract from heterogeneity in asset holdings and consumption of individual workers and households (see Merz 1995).[3] In what follows I drop any household- and worker-specific indices.

The intertemporal utility of a representative household is

$$E_t \sum_{j=t}^{\infty} \beta^{j-t} \left[\frac{C_j^{1-\sigma} - 1}{1-\sigma} - \chi_j n_j \right], \tag{1}$$

where C is aggregate consumption and $n \in [0, 1]$ is the fraction of employed household members, which is determined in the matching market for labor services and is not subject to the household's control. $\beta \in (0, 1)$ is the discount factor and $\sigma \geq 0$ is the coefficient of relative risk aversion. χ_t is an exogenous stochastic process, which I refer to as a labor shock. Note that in the benchmark version I assume that households value leisure, which is subject to stochastic shifts. As we will see later on, this affects wage determination and the interpretation of the parameter estimates.

The representative household's budget constraint is

$$C_t + T_t = w_t n_t + (1 - n_t) b + \Pi_t. \tag{2}$$

[2] Pissarides (2000) gives an excellent overview of the search and matching framework.

[3] Trigari (2006) gives a concise description of the assumptions required for this construct.

The household receives unemployment benefits, b, from the government, which are financed by a lump-sum tax, T. Π_t are profits that the household receives as the owner of the firms. w is the wage paid to each employed worker. The sole problem of the household is to determine the consumption path of its members. There is no explicit labor supply choice since the employment status of the workers is the outcome of the matching process. Since the household's program does not involve any intertemporal decision, the first-order condition is simply

$$C_t^{-\sigma} = \lambda_t, \tag{3}$$

where λ_t is the Lagrange multiplier on the budget constraint.

The Labor Market

The household supplies labor services to firms in a frictional labor market. Search frictions are captured by a matching function $m(u_t, v_t) = \mu_t u_t^{\xi} v_t^{1-\xi}$, which describes the outcome of the search process. Unemployed job seekers, u_t, and vacancies, v_t, are matched at rate $m(u_t, v_t)$, where $0 < \xi < 1$ is the match elasticity of the unemployed, and the stochastic process, μ_t, affects the efficiency of the matching process. The aggregate probability of filling a vacancy (taken parametrically by the firms) is $q(\theta_t) = m(v_t, u_t)/v_t$, where $\theta_t = v_t/u_t$ is labor market tightness. I assume that it takes one period for new matches to become productive and that both old and new matches are destroyed at a constant rate. The evolution of employment, defined as $n_t = 1 - u_t$, is then given by

$$n_t = (1 - \rho)\left[n_{t-1} + v_{t-1} q(\theta_{t-1})\right], \tag{4}$$

where $0 < \rho < 1$ is the constant separation rate that measures inflows into unemployment.

The Firms

The firm sector is populated by monopolistically competitive firms that produce differentiated products. This assumption deviates from the standard search and matching framework, which lets atomistic firms operate in a perfectly competitive environment. I introduce this modification to be able to analyze the effects of mark-up variations on labor market dynamics, as suggested by Rotemberg (2008). The firms' output is demanded by households with a preference for variety that results in downward-sloping demand curves. Thus, a typical firm faces a demand function:

$$y_t = \left(\frac{p_t}{P_t}\right)^{-1-\varepsilon_t} Y_t, \tag{5}$$

where y_t is firm production (and its demand), Y_t is aggregate output, p_t is the price set by the firm, and P_t is the aggregate price index. The stochastic process, ε_t, is the time-varying demand elasticity. I assume that all firms behave symmetrically and suppress firm-specific indices. The firm's production function is

$$y_t = A_t n_t^{\alpha}. \tag{6}$$

A_t is an aggregate technology process and $0 < \alpha \leq 1$ introduces curvature into production. This implicitly assumes that capital is fixed and firm-specific.

The firm chooses its desired number of workers, n_t, the number of vacancies, v_t, to be posted, and its optimal price, p_t, by maximizing the intertemporal profit function

$$E_t \sum_{j=t}^{\infty} \beta^{j-t} \lambda_j \left[p_j \left(\frac{p_j}{P_j} \right)^{-(1+\varepsilon_j)} Y_j - w_j n_j - \frac{\kappa}{\psi} v_j^{\psi} \right], \tag{7}$$

subject to the employment accumulation equation (4) and the demand function (5). Profits are evaluated at the household's discount factor in terms of marginal utility, λ_t. Following Rotemberg (2008), I assume that vacancy posting is subject to cost, $\frac{\kappa}{\psi} v_t^{\psi}$, where $\kappa > 0$ and $\psi > 0$. For $0 < \psi < 1$, posting costs exhibit decreasing returns while costs are increasing for $\psi > 1$. The standard case in the literature with fixed vacancy costs is given by $\psi = 1$.

The first-order conditions are

$$\tau_t = \alpha \frac{y_t}{n_t} \frac{\varepsilon_t}{1 + \varepsilon_t} - w_t + (1 - \rho) E_t \beta_{t+1} \tau_{t+1} \text{ and} \tag{8}$$

$$\kappa v_t^{\psi - 1} = (1 - \rho) q(\theta_t) E_t \beta_{t+1} \tau_{t+1}, \tag{9}$$

where $\beta_{t+1} = \beta \lambda_{t+1} / \lambda_t$ is a stochastic discount factor and τ_t is the Lagrange multiplier associated with the employment constraint. It represents the current-period marginal value of a job. This is given by a worker's marginal productivity, net of wage payments, and the expected value of the worker in the next period if the job survives separation.

Since hiring is costly, firms spread employment adjustment over time. Firms that hire workers today reap benefits in the future since lower hiring costs can be expended otherwise. This is captured by the second condition, which links the expected benefit of a vacancy in terms of the marginal value of a worker to its cost, given by the left-hand side. Note that this is adjusted by the job creation or hiring rate, $q(\theta_t) = m_t \left(\frac{v_t}{u_t} \right)^{-\xi}$. Firms are more willing to post vacancies, the higher the probability is that they can find a worker. Moreover, vacancy posting also depends positively on the worker's expected marginal value, τ_{t+1}, (and thus productivity and wages) and on the elasticity of posting costs.

Combining these two equations results in the job creation condition typically found in the literature:

$$\frac{\kappa v_t^{\psi-1}}{q(\theta_t)} = (1-\rho)E_t\beta_{t+1}\left[\alpha\frac{y_{t+1}}{n_{t+1}}\frac{\varepsilon_{t+1}}{1+\varepsilon_{t+1}} - w_{t+1} + \frac{\kappa v_{t+1}^{\psi-1}}{q(\theta_{t+1})}\right]. \qquad (10)$$

The left-hand side captures effective marginal hiring costs, which a firm trades off against the surplus over wage payments it can appropriate and against the benefit of not having to hire someone next period.

Wage Determination

Wages are determined as the outcome of a bilateral bargaining process between workers and firms. Since the workforce is homogeneous without any differences in skill, each worker is marginal when bargaining with the firm. Both parties choose wage rates to maximize the joint surplus generated from their employment relationship: Surpluses accruing to the matched parties are split to maximize the weighted average of the individual surpluses. It is common in the literature to assume a bargaining function, S, of the following type:

$$S_t \equiv \left(\frac{1}{\lambda_t}\frac{\partial \mathcal{W}_t(n_t)}{\partial n_t}\right)^{\eta}\left(\frac{\partial \mathcal{J}_t(n_t)}{\partial n_t}\right)^{1-\eta}, \qquad (11)$$

where $\eta \in [0, 1]$ is the workers' weight, $\frac{\partial \mathcal{W}_t(n_t)}{\partial n_t}$ is the marginal value of a worker to the household's welfare, and $\frac{\partial \mathcal{J}_t(n_t)}{\partial n_t}$ is the marginal value of the worker to the firm.[4]

The latter term is given by the firm's first-order condition with respect to employment, Eq. (8), where we define $\tau_t = \frac{\partial \mathcal{J}_t(n_t)}{\partial n_t}$. The marginal utility value, $\frac{\partial \mathcal{W}_t(n_t)}{\partial n_t}$, can be found by comparing the options available to the worker in terms of a recursive representation. If the worker is employed, he contributes to household value by earning a wage, w_t. However, he suffers disutility from working, χ_t (which is simply the exogenous preference shifter), and forfeits the outside option payments, b. This is weighted against next period's expected utility. The marginal value of a worker is thus given by

$$\frac{\partial \mathcal{W}_t(n_t)}{\partial n_t} = \lambda_t w_t - \lambda_t b - \chi_t + \beta E_t\frac{\partial \mathcal{W}_{t+1}(n_{t+1})}{\partial n_{t+1}}\frac{\partial n_{t+1}}{\partial n_t}. \qquad (12)$$

Using the employment equation (4), I can then substitute for $\frac{\partial n_{t+1}}{\partial n_t} = (1-\rho)\left[1-\theta_t q(\theta_t)\right]$. Furthermore, note that real payments are valued at the marginal utility, λ_t.

[4] Detailed derivations of the bargaining solutions and the utility values can be found in Trigari (2006) and Krause and Lubik (2007).

Taking derivatives of (11) with respect to the bargaining variable, w_t, results in the standard optimality condition for wages:

$$(1-\eta)\frac{1}{\lambda_t}\frac{\partial \mathcal{W}_t(n_t)}{\partial n_t} = \eta \frac{\partial \mathcal{J}_t(n_t)}{\partial n_t}. \tag{13}$$

Substituting the marginal utility values results, after lengthy algebra, in an expression for the bargained wage:

$$w_t = \eta \left[\alpha \frac{y_t}{n_t} \frac{\varepsilon_t}{1+\varepsilon_t} + \kappa v_t^{\psi-1} \theta_t \right] + (1-\eta)\left[b + \chi_t c_t^{\sigma} \right]. \tag{14}$$

As is typical in models with surplus sharing, the wage is a weighted average of the payments accruing to workers and firms, with each party appropriating a fraction of the other's surplus. The bargained wage also includes mutual compensation for costs incurred, namely hiring costs and the utility cost of working. The bargaining weight determines how close the wage is to either the marginal product or to the outside option of the worker, the latter of which has two components, unemployment benefits and the consumption utility of leisure.

Closing the Model

I assume that government benefits, b, to the unemployed are financed by lump-sum taxes, T, and that the government runs a balanced budget, $T_t = (1-n_t)b$. The social resource constraint is, therefore,

$$C_t + \frac{\kappa}{\psi} v_t^{\psi} = Y_t. \tag{15}$$

In the aggregate, employment evolves according to the law of motion:

$$n_t = (1-\rho)\left[n_{t-1} + \mu_{t-1} u_{t-1}^{\xi} v_{t-1}^{1-\xi} \right]. \tag{16}$$

The model description is completed by specifying the properties of the shocks, namely the technology shock, A_t, the labor shock, χ_t, the demand shock, ε_t, and the matching shock, μ_t. I assume that the logarithms of these shocks follow independent AR(1) processes with coefficients ρ_i, $i \in (A, \chi, \varepsilon, \mu)$ and innovations $\epsilon_i \sim N\left(0, \sigma_i^2\right)$.

2. EMPIRICAL APPROACH

Most papers in the labor market search and matching literature that take a quantitative perspective rely on calibration methods and concentrate on the model's ability to replicate a few key statistics. One issue with such an approach is that information on some model parameters is difficult to come by. The bargaining parameter, η, and the worker's outside option, b, are prime examples. Much of the debate on the viability of search and matching as a description of the labor market centers around the exact values of these parameters (Shimer 2005; Hagedorn and Manovskii 2008). In this article, I therefore take an encompassing, but somewhat agnostic, perspective on the model's empirical implication. I treat the model as a data-generating process for a large set of aggregate variables. My focus is on the actual parameter estimates, the implied adjustment dynamics, and the contribution of various driving forces to labor market movements.

Methodology

I estimate the model using Bayesian methods. First, I log-linearize the nonlinear model around a deterministic steady state and write the linearized equilibrium conditions in a state-space form. The resulting linear rational expectations model can then easily be solved by methods such as in Sims (2002). The model thus describes a data-generating process for a set of aggregate variables. Define a vector of model variables, X_t, and a data vector of observable variables, Z_t. The state-space representation of the model can then be written as

$$X_t = \Gamma X_{t-1} + \Psi \epsilon_t \text{ and} \tag{17}$$

$$Z_t = \Phi X_t, \tag{18}$$

where Γ and Ψ are coefficient matrices, the elements of which are typically nonlinear functions of the structural parameters, and Φ is a selection matrix that maps the model variables to the observables. ϵ_t collects the innovations of the shocks.

In applications, there are typically more variables than observables. The empirical likelihood function can therefore not be computed in the standard manner since the algorithm has to account for the evolution of the model variables not in the data set. This can easily be done using the Kalman filter, which implicitly constructs time series for the unobserved variables. A second concern for the modeler is to ensure that there is enough independent variation in the model to be able to explain the data. In order to avoid this potential stochastic singularity, there have to be at least as many sources of uncertainty in the empirical model as there are observables. This imposes a choice upon the researcher that can affect the estimation results in a nontrivial manner.

In the benchmark specification, I treat the model as a data-generating process for four aggregate variables: unemployment, vacancies, wages, and output. A potential pitfall is that unemployment and vacancies are highly negatively correlated in the data and may therefore not contain enough independent variation to be helpful in identifying parameters. Moreover, the employment equation (4) implies that these two variables co-move perfectly. With both unemployment and vacancy data used in the estimation, this relationship would most likely be violated. Hence, I need to introduce an additional source of variation to break this link, which I do by making the match efficiency parameter an exogenous process. I choose not to include consumption since my focus is on the labor market aspects of the model; nor do I use data on, for instance, the hiring rate, $q(\theta_t)$, since the model implies that it is a log-linear function of u_t and v_t.[5]

The use of four series of observables requires the inclusion of at least four independent sources of variation. Researchers not only have to rely on standard shocks such as technology or variations in market power (i.e., shocks to the demand elasticity, ε), but they often have to introduce disturbances that may be considered nonstandard.[6] This can take the form of converting fixed parameters into exogenous stochastic processes, such as the shock to the match efficiency, μ, used above. Shocks can also capture "wedges" between marginal rates of substitution (Hall 1997), such as the one between the real wage and the marginal product of labor, that the model would otherwise not be able to explain. The labor shock, χ, is an example of this approach.

In order to implement the Bayesian estimation procedure, I employ the Kalman filter to evaluate the likelihood function of the observable variables, which I then combine with the prior distribution of the model parameters to obtain the posterior distribution. The posterior distribution is evaluated numerically by employing the random-walk Metropolis-Hastings algorithm. Further details on the computational procedure are discussed in Lubik and Schorfheide (2005) and An and Schorfheide (2007).

Data

For the estimation, I use observations on four data series: unemployment, vacancies, wages, and output. I extract quarterly data from the Haver Analytics database. The data set covers a sample from 1964:1–2008:4. The starting date of the sample is determined by the availability of the earnings series. Unemployment is measured by the unemployment rate of over-16-year-olds. The series for vacancies is the index of help-wanted ads in the 50 major metropolitan areas. I capture real wages by dividing average weekly earnings in private nonfarm employment by the GDP deflator in chained 2,000$. The output series is real GDP in chained 2,000$. I convert the output series to per-capita terms by scaling with the labor force. All series are passed through

[5] I analyze the implications of changing the set of observables in a series of robustness exercises, where I also address the tight link between unemployment and vacancies.

[6] An alternative is to use shocks to the measurement equation in the state-space representation of the model. While this is certainly a valid procedure, these measurement errors lack clear economic interpretation. In particular, structural shocks are part of the primitive of the theoretical model and agents respond to them. Measurement errors, however, are only relevant for the econometrician and do not factor into the agents' decision problem.

Table 1 Prior Distributions

Definition	Parameter	Density	Mean	Std. Dev.
Discount factor	β	Fixed	0.99	—
Labor elasticity	α	Fixed	0.67	—
Demand elasticity	ε	Fixed	10.00	—
Relative risk aversion	σ	Gamma	1.00	0.10
Match elasticity	ξ	Beta	0.70	0.15
Match efficiency	μ	Gamma	0.60	0.15
Separation rate	ρ	Beta	0.10	0.02
Bargaining power of the worker	η	Uniform	0.50	0.25
Unemployment benefit	b	Beta	0.40	0.20
Elasticity of vacancy creation	ψ	Gamma	1.00	0.50
Scaling factor on vacancy creation	κ	Gamma	0.05	0.01
AR-coefficients of shocks	ρ_i	Beta	0.90	0.05
Standard deviation of shocks	σ_i	Inverse Gamma	0.01	1.00

the Hodrick-Prescott filter with smoothing parameter 1,600 and are demeaned prior to estimation.

Prior

I choose priors for the Bayesian estimation based on the typical values used in calibration studies. I assign share parameters a Beta distribution with support on the unit interval, and I use Gamma distributions for real-valued parameters. I roughly distinguish between two groups of parameters—those associated with production and preferences, and labor market parameters. I choose tight priors for the former, but fairly uninformative priors for most of the latter because the literature lacks independent evidence or disagreement. The priors are reported in Table 1.

I set the discount factor, β, at a value of 0.99. The labor input elasticity, α, is kept fixed at 0.67, the average labor share in the U.S. economy, while the demand elasticity, ε, is set to a mean value of 10, which implies a steady-state mark-up of 10 percent, a customary value in the literature.[7] I choose a reasonably tight prior for the intertemporal substitution elasticity, σ, centered on one. The priors of the matching function parameters are chosen to be consistent with the observed job-finding rate of 0.7 per quarter (Shimer 2005). This leads to a prior mean of 0.7 for the match elasticity, ξ, and of 0.6 for the match efficiency, μ. I allow for a reasonably wide coverage interval as these values are not uncontroversial in calibration exercises. Similarly, I set

[7] Estimating the model by allowing for variation in the fixed parameters shows virtually no differences in the estimates. Using marginal data densities as measures of goodness of fit, I find that the preferred specification is for an unrestricted α. The differences in posterior odds are tiny, however, and it is well known that they are sensitive to minor specification changes.

the mean exogenous separation rate at $\rho = 0.1$ with a standard deviation of 0.02.

I choose to be agnostic about the bargaining parameter, η. Calibration studies have used a wide range of values, most of which center around 0.5. Since I am interested in how much information on η is in the data, which matters for determining the volatility of wages and labor market tightness, I choose a uniform prior over the unit interval. Similarly, the value of the outside option of the worker is crucial to the debate on whether the search and matching model is consistent with labor market fluctuations (Hagedorn and Manovskii 2008). Consequently, I set b at a mean of 0.4 with a very wide coverage region.

The prior mean for the vacancy posting elasticity, ψ, is 1 with a large standard deviation. Linear posting cost is the standard assumption in the literature, but I allow here for both concave and convex recruiting costs as in Rotemberg (2008). The scale parameter in the vacancy cost function is tightly set to $\kappa = 0.05$. Finally, we specify the exogenous stochastic processes in the model as AR(1) processes with a prior mean on the autoregressive parameters of 0.90 and the innovations as having inverse-gamma distributions with typical standard deviations. Moreover, I normalize the means of the productivity process, A_t, and of the labor shock, χ_t, at 1, while the means of the other shock processes are structural parameters to be estimated.

3. BENCHMARK RESULTS

Parameter Estimates

I report posterior means and 90 percent coverage intervals in Table 2. Three parameter estimates stand out. First, the posterior estimate of η is almost zero with a 90 percent coverage region that is concentrated and shifted away considerably from the prior. This implies that firms can lay claim to virtually their entire surplus (and are therefore quite willing to create vacancies) while workers are just paid the small outside benefit, b, and compensation for the disutility of working (see Eq. [14]). Moreover, the disutility of working has an additional cyclical component via the labor shocks. In order to balance this so that wages do not become excessively volatile and thus stymie vacancy creation, the estimation algorithm adjusts the contribution of the marginal product downward, which reduces the bargaining parameter even further.

Second, the posterior estimate of the benefit parameter $b = 0.18$ is moved away considerably from the prior without much overlap with the prior coverage regions. The posterior is also much more concentrated, which indicates that the data are informative. Thus, this estimate seems to indicate that the model resolves the Shimer puzzle in favor of smooth wages to stimulate vacancy posting, and not through a high outside option of the worker. Recall that Hagedorn and Manovskii (2008) suggest values of b as high as 0.9, to which the

Table 2 Posterior Estimates: Benchmark Model

		Prior Mean	Posterior Mean	90 Percent Interval
Relative risk aversion	σ	1.00	0.72	[0.62, 0.79]
Match elasticity	ξ	0.70	0.74	[0.68, 0.82]
Scaling factor matching function	m	0.60	0.81	[0.58, 0.99]
Separation rate	ρ	0.10	0.12	[0.09, 0.15]
Bargaining power	η	0.50	0.03	[0.00, 0.07]
Benefit	b	0.40	0.18	[0.12, 0.22]
Vacancy cost elasticity	ψ	1.00	2.53	[1.92, 3.54]
Vacancy creation cost	κ	0.05	0.05	[0.03, 0.06]

posterior distribution assigns zero probability. This reasoning is misleading, however, as some parameters may be specific to the environment they live in. The benefit parameter, b, is a case in point. In the model it is introduced as payment a worker receives when unemployed. What matters for wage determination, however, is the overall outside option of the worker, which in my model is $b + \chi_t c_t^{\sigma}$. That is, it includes the endogenous disutility of working. This becomes an issue of how to interpret the large variations in this parameter that are reported in both the calibration and the estimation literature. For instance, Trigari (2004) reports a value of $b = 0.03$ in an estimated model that includes a utility value of leisure over both an extensive and intensive labor margin, while Gertler, Sala, and Trigari (2008) find $b = 0.98$ in a framework without these elements.

The discussion thus indicates that the generic parameter, b, is not structural per se, but rather a reduced-form coefficient that captures part of the outside option of the worker relevant for explaining wage dynamics. Its value varies with the other components of the outside option. To get a sense of the magnitude of the latter, I compute $b + \chi c^{\sigma}$ at the posterior mean and find 0.74 with a 90 percent coverage region of [0.56, 0.88]. In the end, this does give support to the argument in Hagedorn and Manovskii (2008) that a high outside option of the worker is needed to match vacancy and unemployment dynamics via smooth wages. The caveat for calibration studies is that values for b cannot be taken off the shelf but should be chosen to match, for instance, wage dynamics.

The third surprising estimate is the vacancy posting elasticity, ψ, with a posterior mean of 2.53, which is also considerably shifted away from the prior. This makes vacancy creation more costly to the firm since marginal postings costs are increasing in the level of vacancies, and therefore labor market tightness. This estimate is substantially different from what is typically assumed in the calibration literature. In most papers, vacancy creation costs are linear, i.e., $\psi = 1$. Rotemberg (2008) even assumes values as low as $\psi = 0.2$. A likely explanation for this high value is that it balances potentially

Table 3 Measures of Fit: Benchmark Model

	Data	**Model**
Overall fit		
MDD	736.20	667.50
Second moments		
$\sigma(y)$	1.61	1.67
$\sigma(u)/\sigma(y)$	7.53	6.49
$\sigma(v)/\sigma(y)$	9.13	7.81
$\sigma(\theta)/\sigma(y)$	14.56	4.36
$\sigma(w)/\sigma(y)$	0.65	0.48
$\rho(u, v)$	−0.89	−0.36

"excessive" vacancy creation that is driven by a low η and by the exogenous shocks.

Estimates for the other labor market parameters are much less dramatic and show substantial overlap with the priors. The posterior means of the matching function parameters are in line with other values in the literature, although the match elasticity, ξ, of 0.74 is at the high end of the range typically considered. However, there is significant probability mass on the more typical values. The estimate of the level parameter, κ, in the vacancy cost function simply replicates the prior, and would therefore not be identified in a purely econometric sense. The estimate of the intertemporal substitution elasticity, σ, as 0.72 is not implausible, and it is reasonably tight and different from the prior. The autoregressive coefficients of the shocks (not reported) are largely clustered around 0.8, which suggests that the model does generate enough of an internal propagation mechanism to capture the still substantial persistence in the filtered data.

I also assess the overall fit of the model, and report some statistics in Table 3. I first compare the structural model to a VAR(2) estimated on the same four data series. There is typically no expectation that a small-scale model such as this can match the overall fit of an unrestricted VAR. This is confirmed by a comparison of the marginal data densities (MDD).[8] While the fit of the structural model is clearly worse than the VAR, and would therefore be rejected in a Bayesian posterior odds as the preferred model, it appears to be at least in the ballpark. Perhaps a more interesting measure is how well the estimated model matches unconditional second moments in the data. I compute various statistics from simulation of the estimated model with parameters set at their posterior means. The model is reasonably successful in matching these statistics. The volatility of HP-detrended output is captured quite well, which is

[8] The MDD is the value of the posterior distribution with the estimated parameters integrated out. It is akin to the value of the maximized likelihood function in a frequentist framework.

not surprising since the technology process, A_t, is identified as the residual in the production function and therefore adapts to the properties of output. The relative standard deviations of unemployment and vacancies are also close to the data, although the volatility of tightness is still considerably off. Finally, wages are less volatile than in the data, which contributes to the relative success of capturing vacancy dynamics. The estimated model is less successful in capturing the high negative correlation between unemployment and vacancies in the data, the so-called Beveridge curve. These findings should not be overinterpreted, however, since the empirical model is designed to capture the data well simply by virtue of the exogenous shocks. An example of this is the presence of the matching shock, which can act as a residual in the employment equation. Consequently, this relative goodness of fit does not invalidate the argument in Shimer (2005), which is based on a single second moment, the volatility of tightness, and a single shock to labor productivity.

I can draw a few conclusions at this point. First, the structural labor market model captures the data reasonably well, in particular the high volatilities of unemployment and vacancies and the relative smoothness of wages. The parameters for the matching process are tightly estimated and close to those found in the calibration and nonstructural estimation literature. There is more discrepancy in the parameters that affect wage bargaining. The bargaining power of the worker is found to be almost zero, while the outside option of the worker is fairly high. The estimates thus confirm the reasoning in Hornstein, Krusell, and Violante (2005), but they also suggest that specific parameters should not be interpreted as strictly structural. Furthermore, the posterior estimates raise questions about the extent to which the performance of the model is due to the inherent dynamics of the search and matching model or whether it is largely explained by the exogenous shocks. I delve further into this issue in the next section.

Variance Decompositions

I now compute variance decompositions in order to investigate the most important driving forces of the business cycle as seen through the model. The results are reported in Table 4. The table shows that in the estimated model unemployment and vacancies are exclusively driven by demand and matching shocks. In the case of unemployment, the matching shock essentially takes the role of a residual in the employment equation (4), which confirms the impression formed above in the comparison of simulated and data moments. This illustrates the model's lack of an endogenous propagation mechanism, as emphasized by Shimer (2005), and the overall fit of the employment equation. Similarly, the demand shock mainly operates through the job creation condition (10) as it affects the expected value of a job.

Table 4 Variance Decompositions: Benchmark Model

	Technology	Labor	Demand	Matching
U	0.00	0.00	0.08	0.92
	[0.00, 0.00]	[0.00, 0.00]	[0.01, 0.14]	[0.76, 0.99]
V	0.00	0.06	0.55	0.38
	[0.00, 0.00]	[0.00, 0.14]	[0.41, 0.67]	[0.25, 0.51]
W	0.32	0.10	0.43	0.15
	[0.15, 0.45]	[0.04, 0.17]	[0.24, 0.50]	[0.05, 0.26]
Y	0.71	0.04	0.04	0.21
	[0.55, 0.87]	[0.01, 0.08]	[0.01, 0.08]	[0.06, 0.32]

Employment and vacancy dynamics thus appear to be largely independent from the rest of the model. An interesting implication of this finding is that search and matching models that do not include either shock offer an incomplete characterization of business cycle dynamics in the sense that their contribution would be attributed to other disturbances. An altogether more critical view would be that the search and matching framework does not present a theory for unemployment dynamics since they are explained exclusively by the residual in the definitional equation (4). In other words, unemployment in the data can be described by a persistent AR(1) process, which is introduced by the matching shock. The intrinsic persistence component, i.e., lagged employment and via the endogenous components of the matching function, on the other hand, does not seem to matter as it likely imposes restrictions that are violated in the data.

The picture for the other variables is more balanced: 70 percent of output variations are explained by the technology shock and 21 percent by the matching shock because of its influence on employment dynamics. Demand and technology shocks explain most of the wage dynamics, with the matching shock coming in a distant third. It is perhaps surprising that the labor shock does not matter more as it directly affects wages through the outside option of the worker. Moreover, it appears directly only in the wage equation (14) and thus could be thought of as a residual, similar to the matching shock. The variance decomposition would, however, support the idea that the wage equation is reasonably well specified and that the need for a residual shock, designed to capture the unexplained components of wage dynamics, is small.

4. ROBUSTNESS CHECKS

I now perform three robustness checks to assess the stability of parameter estimates across specifications and to analyze the dependence of estimates and variance decompositions on the specific choice of observables and shocks. The first robustness check uses the same set of observables as the benchmark, but introduces an AR(1) preference shock to the discount factor, $\beta^t \zeta_t$, instead of the labor shock, χ_t. This changes the model specification in two places: The discount factor in the job creation condition (10) now has an additional time-varying component, and the time preference shock essentially replaces the leisure preference shock in the wage equation (14). Since this specification and the benchmark use the same set of observables, I can directly compare the marginal data densities. The time preference shock specification would be preferred with an MDD of 673.4. However, there are only small differences (not reported) in the posterior means and the 90 percent coverage regions of the two specifications overlap considerably. As in the case of the labor shock, the preference shock plays only a minor role in explaining business cycle dynamics. It does, however, reduce the importance of the demand shock, ε_t, in driving vacancies and wages. Its contributions are now, respectively, 0.42 and 0.29. This indicates that it may be difficult to disentangle the effects of a shock to the mark-up (which I labeled a "demand" shock) from those of movements in the intertemporal utility function.

In the second robustness check, I remove one series from the set of observables. By excluding unemployment I can leave out the shock to match efficiency, μ_t. This allows me to assess to what extent the model is able to replicate vacancy and unemployment dynamics without relying on movements in the residual. The prior specification is as before. Selected results are reported in Table 5. The estimates are, in many respects, strikingly different. The bargaining parameter, η, is still very close to zero, while the benefit parameter, b, is close to the prior mean, but also more concentrated. The total value of the implied outside option is now 0.92 and thus matches the calibrated value in Hagedorn and Manovskii (2008). The apparent reason is that in the benchmark model, the matching shock played a crucial role in explaining unemployment and vacancy dynamics. Without it, the estimation algorithm has to compensate, and it does so in the direction suggested by these authors: a low value of η and a high value of b. This impression is also supported by the decline in the vacancy cost elasticity. The table also reports selected variance decompositions for unemployment, vacancies, and the wage. The term in brackets below the entry denotes the largest contributor to the variation in the respective variable. The contribution of the matching shock to vacancy dynamics in the baseline version now gets captured by technology, which explains 39 percent, but the demand shock still explains 51 percent. Movements in wages are now largely captured by the technology shock, while the demand shock remains important with a contribution of 32 percent.

I also experiment with removing the output series from the set of observables. I then estimate the model for technology, matching, and labor shocks. The removal of the demand shock has the most pronounced effect on the variance decomposition as the previous contribution of movements in the

Table 5 Posterior Estimates: Robustness Checks

Specification		**Estimates**			**Variance Decompositions**	
	η	b	ψ	U	V	W
Benchmark	0.03	0.18	2.53	0.92	0.55	0.43
	[0.00, 0.07]	[0.12, 0.22]	[1.92, 3.54]	(Match.)	(Demand)	(Demand)
Data: V_t, W_t, Y_t	0.02	0.39	1.67	—	0.51	0.48
	[0.00, 0.04]	[0.32, 0.46]	[1.49, 1.88]	—	(Demand)	(Tech.)
Data: U_t, V_t, W_t	0.07	0.23	1.22	0.94	0.61	0.71
	[0.01, 0.18]	[0.10, 0.34]	[0.99, 1.65]	(Match.)	(Tech.)	(Tech.)
Data: U_t, V_t	0.10	0.21	1.45	0.95	0.89	—
	[0.01, 0.25]	[0.08, 0.40]	[0.99, 2.01]	(Match.)	(Tech.)	—

mark-up gets transferred to the technology process. It now explains, respectively, 61 percent and 71 percent of the variations in vacancies and wages. The removal of the output series has no marked effect on the parameter estimates compared to the benchmark. Obviously, including output helps pin down the technology process but is otherwise not crucial for pinning down the structural parameters.

The third robustness check only uses data on unemployment and vacancies, the exogenous shocks being technology and matching. The predictions from the estimated model are fairly clear-cut. Unemployment dynamics are driven by the matching shock, while vacancy dynamics are driven by the technology shock. The parameter estimates are consistent with the results from the previous specifications. However, the coverage regions are noticeably wider and closer to the prior distributions, which reflect the reduction in information when fewer data series are used.

I can now summarize the findings from the robustness exercise as follows. The parameter estimates of the search and matching model are fairly consistent across specifications. In particular, the parameters associated with the matching process, i.e., the match elasticity, ξ, the match efficiency, μ, and the separation rate, ρ, do not show much variation and are close to the values reported in other empirical studies. The other parameters exhibit more variation, in particular the benefit parameter, b. Its estimated value is heavily influenced by both the empirical specification of the model as well as the theoretical structure, and should therefore be properly considered a reduced-form coefficient rather than a structural parameter. Furthermore, the different estimates of the vacancy cost elasticity, ψ, suggest that a model with linear creation cost is misspecified.

Overall, the model matches the data and the second moments reasonably well. Much of this success is, however, due to the incidence of specific shocks. Unemployment dynamics, for instance, are captured almost exclusively by movements in the match efficiency, which acts as a residual in the equation defining how unemployment evolves. This calls into question whether the restrictions imposed by the theoretical search and matching model hold in the data and whether the model provides a reasonable theory of labor market dynamics. The estimates also show that shocks that are not typically considered in the calibration literature, such as the matching or the demand shock, are important in capturing model dynamics, while others, such as preference shocks, play only a subordinate role.

5. CONCLUSION

I estimate a typical search and matching model of the labor market on aggregate data using Bayesian methods. The structural estimation of the full model allows me to assess the viability of the model as a plausible description of labor market dynamics, taking into account all moments of the data and not just selected covariates. The findings in this article are broadly consistent with the literature and would support continued use of the search and matching framework to analyze aggregate labor market issues. However, the article also shows that the relative success of this exercise relies on atypical shock processes that may not have economic justification, such as variations in the match efficiency. An alternative interpretation would be that the shock proxies for a missing component in the employment. A prime candidate would be endogenous variations in the separation rate. The article has also attempted to make inroads into the issue of identification in structural general equilibrium models, mainly by means of extensive robustness checks with respect to alternative data and shocks. Research into this issue is still in its infancy since simple measures of identification in nonlinear models of this kind are not easy to come by.

REFERENCES

An, Sungbae, and Frank Schorfheide. 2007. "Bayesian Analysis of DSGE Models." *Econometric Reviews* 26: 113–72.

Christoffel, Kai Philipp, Keith Küster, and Tobias Linzert. 2006. "Identifying the Role of Labor Markets for Monetary Policy in an Estimated DSGE Model." Discussion Paper Series 1: Economic Studies 2006,17, Deutsche Bundesbank, Research Centre.

Gertler, Mark, Luca Sala, and Antonella Trigari. 2008. "An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining." *Journal of Money, Credit, and Banking* 40 (December): 1,713–64.

Hagedorn, Marcus, and Iourii Manovskii. 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited." *American Economic Review* 98 (September): 1,692–706.

Hall, Robert E. 1997. "Macroeconomic Fluctuations and the Allocation of Time." *Journal of Labor Economics* 15 (January): S223–50.

Hornstein, Andreas, Per Krusell, and Giovanni L. Violante. 2005. "Unemployment and Vacancy Fluctuations in the Matching Model: Inspecting the Mechanism." Federal Reserve Bank of Richmond *Economic Quarterly* 91 (Summer): 19–51.

Krause, Michael U., and Thomas A. Lubik. 2007. "The (Ir)relevance of Real Wage Rigidity in the New Keynesian Model with Search Frictions." *Journal of Monetary Economics* 54 (April): 706–27.

Krause, Michael U., David López-Salido, and Thomas A. Lubik. 2008. "Inflation Dynamics with Search Frictions: A Structural Econometric Analysis." *Journal of Monetary Economics* 55 (July): 892–916.

Lubik, Thomas A., and Frank Schorfheide. 2005. "A Bayesian Look at New Open Economy Macroeconomics." *NBER Macroeconomics Annual*, 313–66.

Merz, Monika. 1995. "Search in the Labor Market and the Real Business Cycle." *Journal of Monetary Economics* 36 (November): 269–300.

Pissarides, Christopher. 2000. *Equilibrium Unemployment Theory*. Cambridge, Mass.: MIT Press.

Rotemberg, Julio J. 2008. "Cyclical Wages in a Search-and-Bargaining Model with Large Firms." In *NBER International Seminar on Macroeconomics 2006*. Chicago: University of Chicago Press, 65–114.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (March): 25–49.

Sims, Christopher A. 2002. "Solving Linear Rational Expectations Models." *Computational Economics* 20 (October): 1–20.

Trigari, Antonella. 2004. "Equilibrium Unemployment, Job Flows, and Inflation Dynamics." European Central Bank Working Paper 304.

Trigari, Antonella. 2006. "The Role of Search Frictions and Bargaining for Inflation Dynamics." Manuscript.