
ELSEVIER
Volume 54S
September 2007
ISSN 0304-3932
Editor
ROBERT G. KING
Associate Editors
DAVID K. BACKUS
MARIANNE BAXTER
MARK J. BILS
JANICE EBERLY
MARTIN EICHENBAUM
SERGIO REBELO
STEPHEN WILLIAMSON
Advisory Board
ANDREW ABEL
ROBERT J. BARRO
THOMAS F. COOLEY
STANLEY FISCHER
MARVIN GOODFRIEND
LARS P. HANSEN
DAVID HUMPHREY
ROBERT E. LUCAS
BENNETT T. McCALLUM
ALLAN H. MELTZER
MAURICE OBSTFELD
ANNA J. SCHWARTZ
G.WILLIAM SCHWERT
ALAN C. STOCKMAN
MARK W. WATSON
CRC Editors
BENNETT T. McCALLUM
CHARLES I. PLOSSER
CRC Advisory Board
ANDREW ABEL
THOMAS F. COOLEY
JANICE EBERLY
MARVIN GOODFRIEND
SERGIO REBELO
ALAN C. STOCKMAN
STANLEY E. ZIN
Founding Editors
KARL BRUNNER
CHARLES I. PLOSSER
Journal of MONETARY ECONOMICS
Supplement issue:
October 20–21 2006 Research Conference on
'Microeconomic Adjustment and
Macroeconomic Dynamics'
Sponsored by the Swiss National Bank
(http://www.snb.ch) and
Study Center Gerzensee (www.szgerzensee.ch)
CONTENTS
F. Gourio and A.K. Kashyap
Investment spikes: New facts and a general equilibrium exploration 1
T. Sveen and L. Weinke
Lumpy investment, sticky prices, and the monetary transmission mechanism 23
L. Veldkamp and J. Wolfers
Aggregate shocks or aggregate information? Costly information and business cycle comovement 37
Published by ELSEVIER
in collaboration with the
WILLIAM E. SIMON GRADUATE SCHOOL OF BUSINESS ADMINISTRATION, UNIVERSITY OF ROCHESTER, and the
DEPARTMENT OF ECONOMICS, BOSTON UNIVERSITY

Available online at www.sciencedirect.com
ScienceDirect
Journal of Monetary Economics 54 (2007) 56–78
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Search frictions: Matching aggregate and establishment observations ☆

Russell Cooper[a,*], John Haltiwanger[b], Jonathan L. Willis[c]

[a]*Department of Economics, University of Texas at Austin and the NBER, Austin, TX 78712, USA*
[b]*Department of Economics, University of Maryland and the NBER, USA*
[c]*Research Department, Federal Reserve Bank of Kansas City, USA*

Received 15 September 2006; received in revised form 6 June 2007; accepted 7 June 2007
Available online 20 June 2007

Abstract

We estimate a search model to match hours, employment, vacancies and unemployment at the micro- and macrolevels. We establish a set of facts concerning the variability of unemployment and vacancies in the aggregate and the distribution of net employment growth and the comovement of hours and employment growth at the establishment level. A search model with non-convex costs of posting vacancies, establishment-level profitability shocks and a contracting framework to set hours provides a structure to understand these observations. The estimated model is able to capture both the aggregate and establishment-level facts.

JEL classifications: E24; J23; J64

Keywords: Labor search; Adjustment costs; Employment

☆We are grateful to an anonymous referee, Robert King, Shigeru Fujita, Murat Tasci and Eran Yashiv for helpful comments. Thanks also to Stephen Nickell, Julian Messina, participants at the 2006 Swiss National Bank/JME Conference and seminar participants at the Ohio State University, Washington University of St. Louis and the Federal Reserve Banks of Chicago, Dallas, Kansas City and New York for comments and suggestions. The authors thank the NSF for financial support. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

*Corresponding author. Tel.: +1 512 475 8704; fax: +1 512 471 3510.
E-mail addresses: cooper@eco.utexas.edu (R. Cooper), haltiwan@econ.umd.edu (J. Haltiwanger), jonathan.willis@kc.frb.org (J.L. Willis).

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2007.06.002

1. Motivation

Aggregate search models are routinely evaluated relative to macroeconomic flows without reference to observations of vacancy, employment and hours variations at the establishment level. The goal of this paper is to propose and estimate a model of labor adjustment at the microeconomic level that is consistent with observations at both the aggregate and the establishment levels.

A leading study of the aggregate implications of a search model is Shimer (2005). He argues that the standard search model, based on Mortensen and Pissarides (1994), fails to match certain key features of aggregate data on worker flows. There are equally important, though less frequently cited, facts coming from observations at the establishment level. At the microeconomic level, employment adjustment is sporadic with periods of inactivity in employment adjustment followed by relatively large adjustments in the number of workers.[1]

Our model extends the search model in a couple of important directions in order to match establishment-level observations. First, it includes a theory of a producer with multiple jobs. Second, we allow for fixed costs of posting vacancies at the establishment level. Third, we estimate a process for profitability shocks at the establishment level to match the observed data. These ingredients of the model allow us to match observed inactivity in employment flows at the establishment level.

Our search model allows for both search costs and firing costs. Both types of adjustment costs are quite successful in matching movements at the establishment and aggregate level. In fact, with the moments we have chosen, it is not possible to conclude that the model with vacancy-posting costs is superior to one with firing costs.

The estimated model does well matching aggregate facts. It does not suffer from the magnification puzzle highlighted by Shimer (2005). This is partly due to the role of the idiosyncratic shocks at the establishment level, which are not smoothed by aggregation due to non-convexities in the model. Further, even though the model matches moments on procyclical average wages, aggregate shocks do induce producers to create vacancies. Finally, average labor productivity is endogenous in our model, which reflects both aggregate and idiosyncratic shocks along with labor market frictions.

2. Facts

We present empirical evidence on key moments from both aggregate and establishment-level data to form the basis of our empirical analysis.[2] The evidence presented in this section is mostly drawn from the recent literature which we supplement from various databases from the U.S. statistical agencies. This evidence provides a rich empirical characterization of U.S. labor markets, which we use to estimate the structural parameters of our search model.

[1] See the discussion and references in Cooper et al. (2004) on job flows. Recent evidence on worker flows draws upon Davis et al. (2006b).

[2] Cooper et al. (2007) provide a considerably more detailed discussion of these facts.

Table 1
Unemployment and vacancies

	U	V	V/U	p
U.S. Data (Quarterly, 1951Q1–2003Q4)				
Standard deviation	0.19	0.20	0.38	0.02
Autocorrelation	0.94	0.95	0.95	0.89
Cross-correlations				
U		−0.89	−0.97	−0.42
V			0.97	0.37
V/U				0.40
Standard search model				
Standard deviation	0.01	0.02	0.03	0.02
Autocorrelation	0.85	0.74	0.81	0.81
Cross-correlations				
U		−0.87	−0.94	−0.94
V			0.99	0.99
V/U				0.99

Notes: For the data moments, the level of unemployment, U, is from the CPS, the level of vacancies, V, is from the Conference Board and average labor productivity, ALP, is real output per person from the BLS. These variables are seasonally adjusted and are log deviations from an HP trend.

2.1. Unemployment, vacancy and productivity dynamics

Table 1 summarizes the main findings on unemployment, vacancies and labor productivity, as in Shimer (2005).[3] The first part of the table reports U.S. data on unemployment, U, vacancies, V and average labor productivity, ALP.[4] The second part of the table, from Tasci (2006), reports moments from a "Standard Search Model."

Three features of the data in Table 1 deserve emphasis. First, the standard deviations of both unemployment and vacancies are about 10 times the standard deviation of the average labor productivity. Second, the data exhibit the Beveridge curve: the correlation between unemployment and vacancies is strongly negative. Finally, both unemployment and vacancies are highly serially correlated.

Comparing the two panels of Table 1, we see that the standard search model, as parameterized in Shimer (2005), is unable to create the volatility in unemployment and vacancies observed in the data. Interestingly, the model is able to capture the negative correlation between vacancies and unemployment, the Beveridge curve.

In what follows, we use key moments on the relationship between vacancies and unemployment consistent with the findings in Table 1. In our analysis we focus on *monthly* vacancy dynamics using the Job Openings and Labor Turnover Survey (JOLTS) data as opposed to the vacancy data from the Help Wanted Index (HWI) used by Shimer (2005).[5]

[3]This table was produced by Tasci (2006).

[4]These observations are quarterly, seasonally adjusted and detrended using a HP filter. Here unemployment is a level not a rate. See the discussion in Shimer (2005) for more data details.

[5]The volatility of vacancies and unemployment is smaller from JOLTS. This may reflect the sample period as well as the data used—many have noted the vacancy volatility from the Conference Board HWI series may be implausibly high—see, e.g., Davis et al. (2006a).

Table 2
Hours and employment adjustment: basic facts from the LRD

Moment	Plant	Aggregate
$\sigma(\Delta h)/\sigma(\Delta e)$	0.96	0.55
$corr(\Delta h, \Delta e)$	−0.29	0.55
$corr(\Delta h_{-1}, \Delta e)$	0.18	0.52

Notes: The data come from the Census Longitudinal Research Database (LRD). Δh is hours growth and Δe is employment growth. $\sigma(x)$ is the standard deviation of x. Seasonal and aggregate effects have been removed from establishment-level moments. Seasonal effects have been removed from aggregate-level moments.

2.2. Hours and employment dynamics

Table 2 presents quarterly evidence about aggregate and establishment-level hours and employment dynamics for production workers from the Longitudinal Research Database (LRD) for the period 1972–1980.[6] This information allows us to compute average hours per worker on a quarterly basis.

We use moments based on the growth (log first difference) of hours per worker and employment at the establishment level in Table 2 to focus on the patterns of adjustment of the labor input rather than the size distribution of the establishments in our sample. As year and seasonal effects have been removed, these moments characterize the aspects of the cross-sectional distribution of employment and hours growth.[7]

Three facts emerge from the moments reported in the first column of Table 2. First, the standard deviations of hours growth and employment growth are about the *same*. Second, hours growth and employment growth are *negatively* correlated. Finally, hours growth in one period is positively correlated with employment growth in the next period.

The second column of Table 2 is based on variations in the aggregate series created from the establishment-level data used in the first column. For the aggregate data, the correlation of hours and employment growth is positive and the standard deviation of employment growth *is almost twice* that of hours growth. This comparison with the microfacts demonstrates the value of using establishment-level moments to understand the interaction between hours and employment adjustment.

2.3. Employment growth, hires and separations

A new establishment-level survey that measures vacancies, hires, separations and employment growth is the JOLTS, which samples about 16,000 establishments per month. Respondents report hires and separations during the month, employment in the pay period covering the 12th of the month, and job openings at month's end. They also report quits, layoffs and discharges, and other separations (e.g., retirements). Recent analysis of the establishment-level data from JOLTS is reported in Davis et al. (2006a, b). Drawing from Figs. 6 and 7 of Davis et al. (2006b), Table 3 summarizes information about the patterns of

[6] The data set used is described in Cooper et al. (2007).

[7] If we compute these moments on a quarterly basis for each of the quarters in our sample, the variation over time is relatively small. Thus we interpret these as cross-sectional moments.

Table 3
Monthly net employment growth rate distribution

Net emp. growth	Share of emp.	Hires	Sep.	Net
< -0.10	0.040	0.025	0.291	−0.266
−0.10–−0.025	0.083	0.023	0.075	−0.052
−0.025–0.025	0.745	0.015	0.015	0.000
0.025–0.10	0.092	0.079	0.028	0.051
>0.10	0.040	0.296	0.041	0.266

Notes: The data come from JOLTS. Faberman (2005) provides a detailed discussion of the data set. See the detailed discussion in Davis et al. (2006b) regarding the measurement methods to insure the timing of net growth and hires and separations are consistent. The column labeled "net" is the average employment growth within each of the bins. These moments are size-weighted by employment share.

worker and job flows at the establishment level.[8] Net employment growth at the establishment level is characterized by five bins, listed in the first column. The second column shows the share of employment growth in each bin. The remaining columns decompose the employment growth into hires and separations. The column labeled "net" is the average employment growth within each of the bins. These moments are size-weighted by employment share.

A couple of facts stand out from the first three columns of the table and additional examination of the data. There is a significant amount of relatively small net employment adjustment: about 74% of the size-weighted observations entail net employment adjustment between −2.5% and 2.5% in an average month. These small adjustments are complemented by significant bursts of job creation and destruction: many establishments (8%) either contract or expand employment by more than 10% in a month.

There is also substantial inaction in employment adjustment. Approximately 80% of monthly establishment-level observations entail zero adjustment.[9] On a size-weighted basis, the inaction rate is 32%. Regarding vacancy posting, 45% of the size-weighted observations entail zero vacancies.[10]

A particularly challenging aspect of matching the patterns of net employment growth is the presence of inaction along with a substantial fraction of observations with relatively small adjustment. With zero adjustment in 32% of the size-weighted observations, the remainder of the 74.5% of the observations with employment growth between −2.5% and 2.5% are relatively small adjustments. The inaction is consistent with a model in which non-convexities are an important element of adjustment costs. But this model has difficulty explaining the small adjustments.[11]

[8]We thank these authors for the summary of the data points underlying these figures. Faberman (2005) provides a detailed discussion of the data set. See the detailed discussion in Davis et al. (2006b) regarding the measurement methods to insure that the timing of net growth and hires and separations are consistent.

[9]This statistic is from Davis et al. (2006b).

[10]This statistic is from Davis et al. (2006a).

[11]We interpret some of the very small adjustments as a form of noise reflecting factors outside of our model (and outside the scope of most of the search and matching literature). For example, in JOLTS an establishment is not supposed to report vacancy postings (job openings) that arise from workers returning from a temporary leave of absence or a hire of a worker onto the payroll who was previously a contract or temporary employee. In the

2.4. The cyclicality of real wages

It is well known that empirically quantifying the cyclicality of real wages (see, e.g., Abraham and Haltiwanger, 1995) is a challenge. Estimates of wage cyclicality are sensitive to sample selection, alternative measures of real wage data, and cyclical composition bias (Bils, 1987; Barsky et al., 1994). These measurement and conceptual issues have not been resolved, but it is apparent from the recent search and matching literature that accounting for the cyclicality of real wages is essential. Much of the debate between Shimer (2005), Hall (2005) and Hagedorn and Manovskii (2006) regarding the ability of standard search and matching models to account for the volatility and dynamics of unemployment and vacancies centers on what fraction of the driving processes is absorbed by variations in wages or the labor input.

The more recent literature (see, e.g., Hagedorn and Manovskii, 2006; Rudanko, 2006) focuses on a measure of real wages from the Bureau of Labor Statistics' (BLS) productivity statistics program. They construct output per person and output per hour in a manner consistent with data produced and released by the BLS. The BLS constructs a measure of hourly compensation based upon the quarterly income from the National Income and Product Accounts (NIPA) attributable to labor. In these papers, the elasticity of real wages (measured as the HP-detrended real hourly compensation series) with respect to labor productivity (also measured as the HP-detrended series) is about 0.45. Moreover, Hagedorn and Manovskii (2006) report that in using Panel Study of Income Dynamics (PSID) data and controlling for composition bias, they obtain an elasticity of real wages at the microlevel that is only slightly higher, about 0.47. To explore sensitivity of their results to this elasticity, Hagedorn and Manovskii (2006) use the 95% confidence interval of their real wage elasticities.[12]

3. Model

There are two types of agents in the model: producers and workers. Producers operate production sites which use labor as an input.[13] The labor input is total hours and thus combines the number of employees and hours worked per employee. There are both aggregate and producer-specific shocks which create revenue from the labor input.

Workers and producers are brought together through a search process. A worker who is matched with a producer has hours and compensation specified through a state-contingent contract. The worker may lose this job in a subsequent period, thus returning to a state of unemployment. Reflecting the search friction, workers without a job are assumed to find a new job with some probability each period. This probability is exogenous to the worker but is determined in equilibrium.

(*footnote continued*)
analysis that follows we attempt to match the fraction of employment at establishments with very small adjustments but do not distinguish between zero adjustment and very small adjustments.

[12] As discussed in Cooper et al. (2007) this elasticity is on the high side given the use of this particular real wage series.

[13] In this discussion, producers operate a production site and not a firm. This is consistent with our establishment-level observations and assumes that firms with multiple establishments operate them independently, at least with respect to employment decisions.

Producers have a set of workers with whom they have a contract at a point in time. In the short-run, the producer responds to variations in a profitability shock, which reflects both productivity and demand, through changes in hours worked per employee. The contract determines the response of hours and compensation to the shock.

Producers also can create vacancies and hence change the number of employees. The process of creating and filling vacancies entails adjustment costs. We allow for both fixed and variable costs of posting vacancies. The presence of these fixed costs distinguishes our model from the existing search literature and defines the boundaries of a producer. Empirically, these fixed costs are important for matching the observed inaction in the adjustment of the number of workers.

3.1. Workers

In general, workers are in one of two states, employed or unemployed. If unemployed, the workers enjoy leisure time and/or the fruits of home production, $b(a)$, which is allowed to depend on the level of aggregate productivity, captured by a. With a positive probability, unemployed workers become employed in the subsequent period. Formally, the value of unemployment for a worker is given by the following:

$$W^{\mathrm{u}}(a) = Z(b(a)) + \beta \mathrm{E}_{a'|a}[\phi(U, V)W^{\mathrm{e}}(a') + (1 - \phi(U, V))W^{\mathrm{u}}(a')],$$

where $Z(\cdot)$ is utility, β is the discount rate, and $\phi(\cdot)$ is the job finding rate that depends on the level of unemployment, U, and aggregate vacancies, V. Here there is an expectation operator associated with the future value of employment, $W^{\mathrm{e}}(a')$, and unemployment, $W^{\mathrm{u}}(a')$.[14]

Employed workers have a contract for the current period which governs their state-contingent compensation and hours worked. Workers do not save in equilibrium, and thus compensation, ω, and consumption, c, are identical. We specify utility of consumption and hours worked, h, as $Z(\omega - g(h))$.[15] Here $Z(\cdot)$ is strictly increasing and strictly concave and $g(\cdot)$ is strictly increasing and strictly convex.[16]

The value of employment is given by

$$W^{\mathrm{e}}(a) = \mathrm{E}[Z(\omega - g(h))] + \beta \mathrm{E}_{a'|a}[(1 - \delta)W^{\mathrm{e}}(a') + \delta W^{\mathrm{u}}(a')], \tag{1}$$

where δ is the rate of separation (quits plus fires). The first term is the expected utility given a contract, described below, with a producer. In the following period there is a probability, given by δ, of job loss leading to unemployment.

3.2. Producers

Producers have access to a technology which creates output from labor input. The revenue function is given by $a\varepsilon(eh)^{\alpha}$, where a is the aggregate (profitability) shock, ε is the

[14]The state vector for $W^{\mathrm{u}}(a)$ and $W^{\mathrm{e}}(a)$ highlights the dependence of these values on the aggregate state, a. Other state variables have been suppressed in the notation. Since, as discussed below, workers are indifferent between employment and unemployment, a is the only state variable which matters to them.

[15]The model can accommodate more general utility functions but this specification is particularly tractable.

[16]While workers do face uncertainty in (net) consumption, $\omega - g(h)$, workers all have the same level of net utility *ex post*. Thus workers have no incentive to trade state-contingent consumption. Further, there is no store of value in the model and thus compensation and consumption are the same.

producer-specific shock and total labor input is the product of the number of workers, e, and hours per worker, h. We allow for curvature in the revenue function, parameterized by α, which may capture diminishing returns to scale due to excluded fixed factors of production.

We assume two stages in the producer's problem. First, given the aggregate state, the producer contracts with its workers. Second, *ex post*, the producer-specific shock is realized, and state-contingent hours are determined given the contract.

3.2.1. Setting a contract

A contract is $\Upsilon = (\omega(S), h(S))$ for all S, where $S = (a, \varepsilon, e, \theta)$ is the establishment's state. The state vector includes $\theta \equiv V/U$, which measures the tightness of labor markets. As explained in more detail below, this aggregate variable summarizes the state of the labor market. Therefore, producers use this state variable to predict the ease of hiring workers.

In the state-contingent contract, $\omega(S)$ is compensation and $h(S)$ is hours worked. The contract allows compensation and hours to be fully state contingent. In terms of timing, the contract is determined given (a, e) but prior to the determination of ε. All workers with a given producer get the same contract since they are identical and have the same outside option of unemployment.

For the contracting process, assume producers make a "take it" or "leave it" offer to workers.[17] This implies that employed workers get no surplus. Therefore, the value of employment in equilibrium is independent of the producer with whom the worker has a job.

The producer selects the contract to maximize profits, π,

$$\pi(a, \varepsilon_{-1}, e) = \max_{\Upsilon} \mathrm{E}_{\varepsilon|\varepsilon_{-1}}[a\varepsilon(eh(S))^{\alpha} - e\omega(S)], \tag{2}$$

where the expectation is over the idiosyncratic component of profitability. The constraint is that the expected utility from the contract not be less than the outside option of unemployment, $W^{\mathrm{e}}(a) \geqslant W^{\mathrm{u}}(a)$ where $W^{\mathrm{e}}(a)$ is given in (1). In equilibrium, these values depend on aggregate productivity through $b(a)$ but do not depend on labor market tightness since $W^{\mathrm{e}}(a) = W^{\mathrm{u}}(a)$ for all a.

As the worker's participation constraint binds, $W^{\mathrm{e}}(a) = W^{\mathrm{u}}(a)$ for all a, $\mathrm{E}_{\varepsilon} Z(\omega(S) - g(h(S))) = Z(b(a))$ for all S. Given the risk aversion of the workers, optimal risk sharing implies that marginal utility is independent of the realized value of ε. Thus compensation and hours satisfy the following the condition for all S:

$$Z(\omega(S) - g(h(S))) = Z(b(a)). \tag{3}$$

Interestingly, producer heterogeneity is present in compensation levels, and hours reflect producer-specific state variables and shocks. This implies that there is a non-degenerate cross-sectional distribution of (ω, h) but a degenerate cross-sectional distribution of utility levels given a.

Variations in the aggregate state, acting through $b(a)$, influence the terms of the contract. In this way, the model produces movements in the aggregate wage without the complexity

[17] This simplification reduces the state space of the problem. As discussed in Hagedorn and Manovskii (2006), the allocation of bargaining weight has important implications for the behavior of search models.

of allowing the workers to have bargaining power.[18] This added feature is important for matching observations on aggregate wage movements and for tempering the response of vacancies and unemployment to aggregate shocks.

3.2.2. Determining hours

Once ε is realized, hours are determined by the contract. With (3) holding for all S, workers are fully compensated for hours variations. Given (a, ε, e), the producer chooses a level of hours subject to (3):

$$\pi(a, \varepsilon, e) = \max_h \{a\varepsilon(eh)^\alpha - eg(h) - eb(a)\}.$$

The hours choice satisfies

$$\alpha a\varepsilon(eh)^{\alpha-1} = g'(h). \tag{4}$$

This first-order condition generates a state-dependent policy function, $h(a, \varepsilon, e)$. Holding (a, ε) fixed, it is clear that h falls as e increases.

3.2.3. Determining the level of employment

The level of employment is determined by the vacancy-posting decision of the producer.[19] The recruiting decision is made knowing $s \equiv (a, \varepsilon_{-1}, e_{-1}, \theta)$ where e_{-1} is the inherited stock of workers and ε_{-1} is the shock last period that is used to predict the current one. The state vector s is similar to S, except for the timing of e and ε.

$Q(s)$ is the value of the establishment in state s and is given by

$$Q(s) = \max\{Q^v(s), Q^n(s), Q^f(s)\}. \tag{5}$$

In this optimization problem, $Q^v(s)$, $Q^n(s)$ and $Q^f(s)$ relate to the vacancy posting, no-adjustment and firing options. The value of posting vacancies is given by

$$Q^v(s) = \max_v \mathrm{E}_{e,\varepsilon}[\pi(a, \varepsilon, e)] - F_v - C_v(v) + \beta E[Q(s')].$$

There are two types of costs of posting vacancies in the model: a fixed cost component, F_v, and a variable cost component, C_v. A familiar interpretation of this type of specification is based upon recruiting in Economics. The fixed cost appears in the form of reading numerous files, flying a committee to interview and so forth. The variable cost is related to the number of interviews and fly-outs. In terms of matching the moments, these two costs are relevant for capturing inaction, through F_v, and partial adjustment, through C_v.

When the producer posts v vacancies, the evolution of employment is

$$e = e_{-1}(1 - q) + H(U, V)v, \tag{6}$$

where q is the quit rate and $H(\cdot)$ is the rate at which vacancies are filled. In this formulation, the vacancy-filling rate depends on the level of unemployment, U, and the number of vacancies, V. This is where the aggregate state of the economy influences the magnitude of the adjustment cost for employment. For some specifications, labor market

[18] We are grateful to the referee and Borghan Narajabad for suggesting this addition to the model. Acemoglu and Hawkins (2006) analyze a model where firms have multiple workers with bargaining. As they note, this approach has many challenges and their model does not have the rich features of hours per worker and adjustment costs that are the focus of our paper.

[19] Here vacancies must be reposted each period. See Fujita and Ramey (2005) for a model, where vacancies are a state variable.

tightness makes employment adjustment more costly, and thus more of the variation in labor input occurs in hours worked.

The value of firing workers is given by

$$Q^f(s) = \max_f \mathrm{E}_\varepsilon[\pi(a, \varepsilon, e_{-1}(1-q) - f)] - F_f - C_f(f) + \beta \mathrm{E}[Q(s')].$$

Here the level of employment reflects quits and fires. There are fixed, F_f, and variable costs, $C_f(f)$, of firing workers.

The value of inaction is given by

$$Q^n(s) = \mathrm{E}_\varepsilon[\pi(a, \varepsilon, e_{-1}(1-q))] + \beta \mathrm{E}[Q(s')].$$

Here inaction means no vacancy posting and no firing so that employment at the establishment level falls due to quits.

We assume that any profits realized by producers are consumed by entrepreneurs who own the production process. These agents are risk neutral and thus are the natural suppliers of insurance to workers. Producers discount at the same rate, β, as do workers.

For the exogenous state variables, we assume that aggregate profitability and producer-specific profitability follow autoregressive processes:

$$\ln a_t = \rho_a \ln a_{t-1} + v_{a,t}, \quad v_a \sim \mathrm{N}(0, \sigma_a),$$
$$\ln \varepsilon_t = \rho_\varepsilon \ln \varepsilon_{t-1} + v_{\varepsilon,t}, \quad v_\varepsilon \sim \mathrm{N}(0, \sigma_\varepsilon).$$

3.3. Equilibrium

An equilibrium for this economy requires optimization by producers and workers and consistency conditions. One component of equilibrium is an optimal labor contract which solves (2) subject to the participation constraint of the workers. A second element is a state-contingent hours schedule which solves (4). A third component is a decision rule for employment adjustment which solves (5). Fourth, there is a decision rule for workers entailing acceptance or rejection of the contract.

With regards to the consistency conditions, the vacancy-filling rate, which appears in the employment transition constraint (6) in the optimization problem of the producers, depends on θ, an aggregate variable determined in equilibrium. This function, which is taken as given in the optimization problem of the producer, must be consistent with the relationship generated by the model and the data. As described below, this consistency is enforced in our estimation.

Finally, unemployment follows $U' = (1-U)\delta(U,V) + (1-\phi(U,V))U$ where $\delta(U,V)$ is the separation rate (quits plus layoffs) and $\phi(U,V)$ is the job finding rate.

4. Estimation

The key parameters in our study are those determining the costs of posting vacancies and firing as well as the driving process for the shocks at the establishment level. These parameters are estimated through a simulated method of moments procedure. The estimation entails finding the vector of structural parameters, Λ, to minimize the (weighted) distance between moments from the data, Γ^{d}, and moments produced from a simulation of the model given a vector of parameters, $\Gamma^{s}(\Lambda)$.[20] Our estimate of Λ minimizes $£(\Lambda)$ where

$$£(\Lambda) \equiv (\Gamma^{d} - \Gamma^{s}(\Lambda))\Xi(\Gamma^{d} - \Gamma^{s}(\Lambda))' \tag{7}$$

and Ξ is a weighting matrix.[21]

This minimization problem is solved by simulation to create a mapping from Λ to the moments. Given vector Λ, we solve the producer's dynamic optimization problem using value function iteration.[22] From this and the solution to (4), we generate policy functions at the producer level for employment, vacancies and hours. Using these policy functions, we create a simulated data set at the producer level. To produce $\Gamma^{s}(\Lambda)$, we compute the microeconomic moments directly from the simulated data and aggregate this data to obtain flows of vacancies and unemployment.

The simulated data set consists of 8000 establishments simulated over 360 months. The results are robust to increasing the number of establishments and time periods.

4.1. Functional forms

To characterize and estimate an equilibrium for this model, we rely on some particular functional forms. The parameters of these functions are estimated.

In equilibrium, the wage, $\omega(S)$, satisfies $Z(\omega(S) - g(h(S))) = Z(b(a))$ for all S. We parameterize the disutility of work, $g(h)$, so that $\omega(S) = b(a) + \omega_0 h(S)^{\zeta}$ is the compensation required to guarantee the utility level $Z(b(a))$ in all states, S. Here ζ is important for determining the utility cost of variations in hours. Generally, if ζ is low, then variations in hours are inexpensive and much of the adjustment of the labor input is on the intensive margin rather than through variations in the number of workers.

We allow the workers' value of leisure to depend on the aggregate shock, a. Thus we assume $b(a) = b_0 a^{b_1}$. The parameter b_1 thus governs the sensitivity of the worker's outside option to aggregate profitability.

For the hiring and firing of workers, we assume that the variable cost of posting vacancies is given by $C(v) = c_{0,v} v^{c_{1,v}}$. The variable cost of firing workers is similarly parameterized by $C(f) = c_{0,f} f^{c_{1,f}}$.

Following the literature, we specify the matching function with constant returns to scale:

$$m = \mu U^{\gamma} V^{1-\gamma} = \mu V \theta^{-\gamma}, \tag{8}$$

where $\theta \equiv V/U$ measures the tightness of labor markets. From (8) we obtain two additional functions: the vacancy-filling rate for producers, $H = m/V$, and the job finding rate for workers, $\phi = m/U$. Using the specification of the match rate in (8),

$$H = \mu\theta^{-\gamma}. \tag{9}$$

[20]The discussion below provides detail on the exact elements of Λ.

[21]In the discussion which follows, Ξ is an identity matrix that produces consistent estimates of Λ. Since the data come from different sources, computing a variance–covariance matrix directly from the data to create Ξ is not feasible.

[22]The variables in the state space, $s \equiv (a, \varepsilon_{-1}, e_{-1}, \theta)$, are placed on a discretized grid with 7, 21, 200, and 7 points, respectively. For the exogenous state variables the points were spread equally in terms of the cumulative distribution function of the variables. We follow Tauchen (1986) to construct the transition matrices.

Using (8) again,

$$\phi = \mu\theta^{1-\gamma}. \tag{10}$$

The solution of the producers optimization problem requires the vacancy-filling rate in the transition equation for employment, (6). Using (9), the vacancy-filling rate depends on the current state of labor market tightness, θ. Thus, the solution of the producer's optimization problem requires knowledge of the evolution of θ in the equilibrium of the model economy. Given the heterogeneity of producers, forecasting labor market tightness requires knowledge of the cross-sectional distribution of workers and producer-specific profitability as well as the aggregate variables. This is a computationally complex forecasting problem.

4.1.1. Estimated functions

This section reports our estimates of the parameters for the vacancy-filling rate function. This estimation is done directly from the data and is thus outside of the solution of the producer's dynamic optimization problem.

We construct a monthly vacancy series using JOLTS. The monthly unemployment data are from the Current Population Survey (CPS). Relative to Shimer (2005), our data are higher frequency and we use the JOLTS vacancy series. For monthly labor productivity, we construct a series using the Industrial Production Index from the FRB and total hours using employment and hours data from the BLS. Labor market tightness, θ, is the log of the vacancy–unemployment ratio. All series are converted to logarithms and then HP filtered.

The primary relationship estimated directly from these data is the matching function, where we regress the logarithm of the match rate on the logarithm of labor market tightness. In keeping with a large part of the literature we impose constant returns to scale. With this restriction we estimate $\gamma = 0.36$, which is considerably lower than the estimate of 0.72 reported in Shimer (2005) and closer to the estimate of 0.24 reported in Hall (2005). The difference in estimates may reflect the use of the JOLTS data rather than the Conference Board vacancy numbers and the use of different sample periods. The logarithm of the constant is estimated at 0.0072. Hence μ from (8) is estimated at 1.0072.[23]

To check on these estimates, we measure the monthly job finding rate of workers at 0.61. Using (10), the estimate of $\gamma = 0.36$ and the mean value of $\theta = 0.46$ imply $\mu = 1.0009$. These results are very close to the point estimate from the regression.

To simplify the analysis, we assume a form of bounded rationality or limited information by producers: they forecast θ using an AR(1) process. Thus θ appears in the state vector of the producer. We estimate an AR(1) representation of θ from our data. From that regression, the AR(1) coefficient on θ is 0.93 with an R^2 of 0.96. The simple autoregressive structure does an outstanding job of capturing the dynamics of labor market tightness in our sample.

4.1.2. Calibrated parameters

We calibrate a subset of our parameters, summarized in Table 4. As the model is monthly, we set $\beta = 0.9966$, which represents a 4% annualized rate of return. The

[23] The standard error on the estimate of γ and μ equals 0.02.

Table 4
Calibrated parameters

Relationship	Parameter
Discount rate (β)	0.9966
Curvature of profit function (α)	0.65
Elasticity of disutility of hours (ζ)	2.90
Serial correlation of aggregate shocks (ρ_a)	0.95
Standard deviation of innovation to aggregate shocks (σ_a)	0.0016
Curvature of matching function (γ)	0.36

parameters, b_0 and ω_0 are set to match average establishment size and average hours.[24] The value of α is set based on the value estimated in Cooper et al. (2004). The calibrated value of ζ follows Caballero and Engel (1993). This wage elasticity governs the cost of adjusting hours and is relevant for matching the relative standard deviation of hours growth to employment growth.

The common component of the profitability shock, a, is modeled as a log normal AR(1) process. We set the serial correlation, ρ_a, and standard deviation of the innovation to the process, σ_a, to be roughly consistent at a monthly frequency with quarterly estimates in the literature using aggregate and manufacturing data.[25]

4.1.3. Moments

The parameter vector we estimate includes: $\Lambda = (F_v, c_{0,v}, F_f, c_{0,f}, b_1, \rho_\varepsilon, \sigma_\varepsilon)$. The first two parameters represent the cost of posting vacancies and the second two are firing costs. The parameters $(\rho_\varepsilon, \sigma_\varepsilon)$ characterize the log normal AR(1) process for the establishment-specific profitability shocks, where σ_ε is the standard deviation of the innovation to the process. The final parameter, b_1, captures the sensitivity of the worker's value of leisure to variations in aggregate profitability.

We separate the moments into five categories. To enforce *equilibrium*, the estimated model must mimic the regression results for AR(1) representation of labor market tightness, denoted as ρ_θ. The inclusion of the AR(1) coefficient for labor market tightness implies that the beliefs of the producers in the model, which conform with the empirically observed serial correlation, match the data. In this way, we also ensure that we have an equilibrium: the beliefs of the producers are mutually consistent. Second, we focus on *unemployment and vacancies* where the key moments are $\sigma(U)$, $\sigma(V)$ and $corr(U, V)$. Third, we look at *hours and employment* where the key moments are $corr(\Delta e, \Delta h)$ and $\sigma(\Delta h)/\sigma(\Delta e)$.[26] Fourth, we focus on *worker and job flows*. The key moments are from the distribution of Δe reported in Table 3. Finally, we consider *aggregate wages and*

[24]For our model, the value of b_0 influences the size of producers, not the difference in utility between employed and unemployed agents. We set $b_0 = 1.12$ and $\omega_0 = 0.000013$ so that the average establishment size is 60 and hours equal 40 on average. The average size of 60 is consistent with the data set, described in Cooper et al. (2004), used to estimate α.

[25]We choose $\rho_a = 0.95$ because it lies between the typical quarterly aggregate *TFP* estimate of 0.95, which is 0.98 at a monthly frequency, and the Cooper and Willis (2002) estimate of a quarterly profitability shock for the manufacturing sector, which becomes 0.90 at the monthly frequency. We choose $\sigma_a = 0.0016$ to be in a similar range as found in the literature.

[26]In matching these moments we take time aggregation into account.

productivity, where the key moment, denoted by $\omega|ALP$, is the coefficient from a regression of (log) average hourly wages on (log) average labor productivity. This moment has no structural interpretation.

We choose these moments largely because they characterize basic aspects of worker and job flows at both the microeconomic and aggregate levels. This is in accord with the point of our analysis: to investigate a search model capable of jointly explaining both microeconomic and macroeconomic facts.

4.2. Results

We estimate our model under four cases: two with vacancy-posting costs and two with firing costs. While the model focuses on the significance of search and thus vacancy-posting costs, we also study the implications of firing costs. This approach provides a perspective on the role of search costs and also on the more general topic of the source of labor adjustment costs in general. Each of the two vacancy-posting cost specifications includes a fixed cost of posting vacancies in conjunction with either a linear ($c_{1,v} = 1$) or quadratic ($c_{1,v} = 2$) cost of posting vacancies. Each of the two firing cost cases includes a fixed cost of firing in conjunction with either a linear ($c_{1,f} = 1$) or quadratic ($c_{1,f} = 2$) firing cost.

The results are reported in the following two tables. The parameter estimates are reported in Table 5. The moments of the models relative to data are summarized in Table 6 and discussed in Section 5.1.

Table 5 shows the parameter estimates for the four specifications, the adjustment costs as a percentage of gross profits (AC) and the fit of the model from (7). For both vacancy-posting and firing costs, the fixed and quadratic specifications do not do as well as the fixed and linear adjustment cost cases. The specification with linear and fixed costs of posting vacancies, hereafter VC, fits the moments essentially as well as the specification with linear and fixed firing costs, hereafter FC. We have not been able to improve the fit using a specification with both vacancy posting and firing costs. Note that this does not mean the model with vacancy-posting costs is observationally equivalent to the model with firing

Table 5
Estimation results: parameters

Specification	F_v	$c_{0,v}$	F_f	$c_{0,f}$	b_1	ρ_ε	σ_ε	AC	£(Λ)
Vacancy-posting costs									
Fixed, quadratic	0.056	0.001	0	0	0.369	0.338	0.186	0.550	0.059
Fixed, linear (VC)	0.135	0.007	0	0	0.501	0.331	0.227	0.775	0.024
Firing costs									
Fixed, quadratic	0	0	0.220	0.858	0.575	0.894	0.115	0	0.048
Fixed, linear (FC)	0	0	0.112	0.024	0.487	0.459	0.196	0.235	0.024

Notes: The table reports estimation results from four different specifications of the model: fixed and quadratic vacancy-posting costs, fixed and linear vacancy-posting costs, fixed and quadratic firing costs, and fixed and linear firing costs. $\{F_v, c_{0,v}\}$ are vacancy-posting costs; $\{F_f, c_{0,f}\}$ are firing costs; b_1 is the elasticity of the worker's outside option with respect to aggregate profitability; $\{\rho_\varepsilon, \sigma_\varepsilon\}$ parameterize the idiosyncratic profitability shock; AC are average adjustment costs paid as a percentage of gross profits; and £(Λ) is the minimization statistic from the estimation.

Table 6
Estimation results: moments

Model	Wage and ALP	AR(1) of θ	Unemployment and vacancies			LRD		JOLTS: Δe				
	$\omega\|ALP$	ρ_θ	$\sigma(U)$	$\sigma(V)$	$corr(U,V)$	$\sigma(\Delta h)/\sigma(\Delta e)$	$corr(\Delta h,\Delta e)$	$\Delta e \leqslant -0.10$	$-0.10 < \Delta e \leqslant -0.025$	$-0.025 < \Delta e \leqslant 0.025$	$0.025 < \Delta e \leqslant 0.10$	$0.10 < \Delta e$
Data	0.45	0.93	0.09	0.12	−0.95	0.96	−0.30	0.04	0.08	0.75	0.09	0.04
Vacancy-posting costs												
Fixed, quadratic	0.41	0.89	0.13	0.08	−0.82	0.92	−0.43	0.06	0.11	0.72	0.00	0.11
Fixed, linear (*VC*)	0.44	0.90	0.13	0.10	−0.90	0.99	−0.36	0.05	0.07	0.79	0.00	0.08
Firing costs												
Fixed, quadratic	0.39	0.94	0.12	0.07	−0.91	0.88	−0.26	0.00	0.00	0.88	0.04	0.07
Fixed, linear (*FC*)	0.45	0.94	0.12	0.08	−0.93	0.97	−0.39	0.01	0.00	0.78	0.10	0.10

Notes: $\omega|ALP$ is the coefficient from a regression of HP-filtered quarterly (log) real hourly compensation on HP-filtered log quarterly average output per hour using the same series as Hagedorn and Manovskii (2006). Unemployment, U, comes from the CPS and vacancies, V, are directly calculated from JOLTS at a monthly frequency. ρ_θ is the serial correlation of labor market tightness ($\theta = (V/U)$). Δh is establishment-level hours growth and Δe is establishment-level employment growth from the Census Longitudinal Research Database (LRD). The last five columns represent the size-weighted distribution of employment growth across establishments from JOLTS. The last four rows of the table report simulated data moments from the four cases estimated in Table 5.

costs. Looking at the distribution of net employment growth, the specifications clearly have very different implications for these moments.

For *VC*, the estimated costs of adjustment (paid) are around 0.77% of gross profits. The costs are much lower in *FC* since workers quit at an exogenous rate and thus producers can avoid this cost. The idiosyncratic profitability shocks are serially correlated and much more variable than aggregate shocks. The value of b_1 in both specifications is around 0.5, which indicates a strong response of the value of leisure to aggregate profitability.

5. Evaluation of results

This section summarizes our findings. We first discuss how well the estimated model matches key moments. We then discuss other implications of the estimated model.

5.1. Explaining the moments

Table 6 summarizes the moment implications for the cases. Both the *VC* and *FC* specifications do well matching moments in many dimensions. Both models, largely through b_1, are able to match the regression coefficient of wages on average labor productivity, $\omega|ALP$. With regard to aggregate facts, the models match the variability of unemployment and vacancies and produce a Beveridge curve. At the establishment level, the estimated models match the relative volatility of hours and employment growth as well as the negative correlation between hours and workers.

Further, both models reproduce the serial correlation in θ from the data. Thus in both cases, the beliefs of producers about the evolution of labor market tightness and thus the vacancy-filling rate are consistent with the data and the outcome of the model economy.

Differences between the *VC* and *FC* specifications emerge in the distribution of employment growth at the producer level. The model with vacancy-posting costs predicts too little hiring and an excessive level of net employment growth smaller than 10% since there are no firing costs. Conversely, the model with firing costs produces fewer burst of firing but some intermediate rates of positive net employment growth. Neither model alone is fully capable of capturing the entire employment growth distribution.

5.2. Inspecting the mechanism

This is a rather rich model, and the mapping from parameters to moments is not immediately clear. To help build further intuition about the models mechanics, we present two figures related to key moments. For the simulations underlying these figures, we used the *VC* specification.

Fig. 1 shows simulation results at the establishment level. We use this figure to show how an establishment responds to variations in profitability in the presence of fixed costs of posting vacancies.

Three points are illustrated in Fig. 1. First, hours and profitability are highly positively correlated over much of the sample. When the producer is subject to a large increase in profitability, such as in period 28, hours respond immediately. This variability of hours comes both from the timing of the model and the estimated adjustment costs so that small variations in employment are not profitable. Thus the producer adjusts to shocks by varying hours.

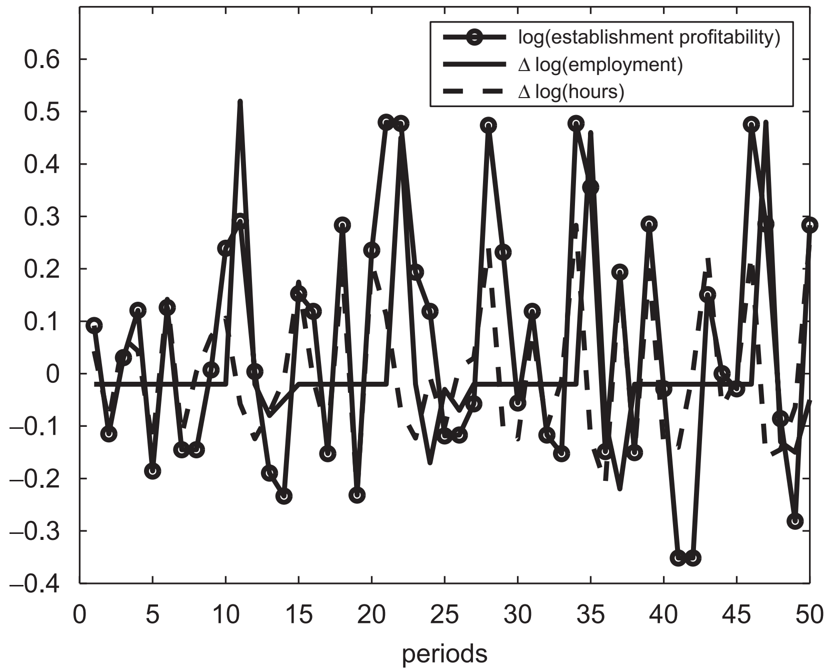

Fig. 1. Employment and hours: establishment-level growth rates.

Second, employment does not always respond to variations in profitability. The large shock in period 28 does not lead to employment adjustment. Yet, employment does respond in period 35 of the simulation. During the interim, the size of the workforce fell due to quits and thus the producer found it profitable to adjust employment until period 35.

Third, the model produces a negative correlation at the establishment level between hours and employment growth in large part due to the timing assumption on employment adjustment. For example, the producer makes its vacancy decision in period 10 *before* observing the large, positive shock to profitability. The producer's only contemporaneous response is through an increase in average hours worked of nearly 10%. In the subsequent period, the firm posts enough vacancies to increase employment by 50%, adjusting to the recent profitability shock as well as hiring to replace the quits that occurred during the preceding periods of inaction. With the large increase in the employment in period 11, average hours worked are cut, which results in a negative correlation between hours and employment growth.

Fig. 2 shows simulation results for unemployment, vacancies and the aggregate component of the profitability shock. These variables are obtained by the aggregation of the establishment-level data from the same panel simulation used for Fig. 1. There are two cases shown here: $b_1 = 0.501$, its estimated value, and b_1 near 0.[27] The comparison illustrates the role of b_1 for our analysis.

The Beveridge curve is apparent in these simulations. When there is a positive aggregate shock, such as around period 6, there is an immediate response in the creation of vacancies. Unemployment falls as vacancies are filled. The strength of this response depends partly on the cost of creating vacancies and on the rate in which vacancies are filled.

[27]The actual value of b_1 was positive but very close to zero for computational reasons.

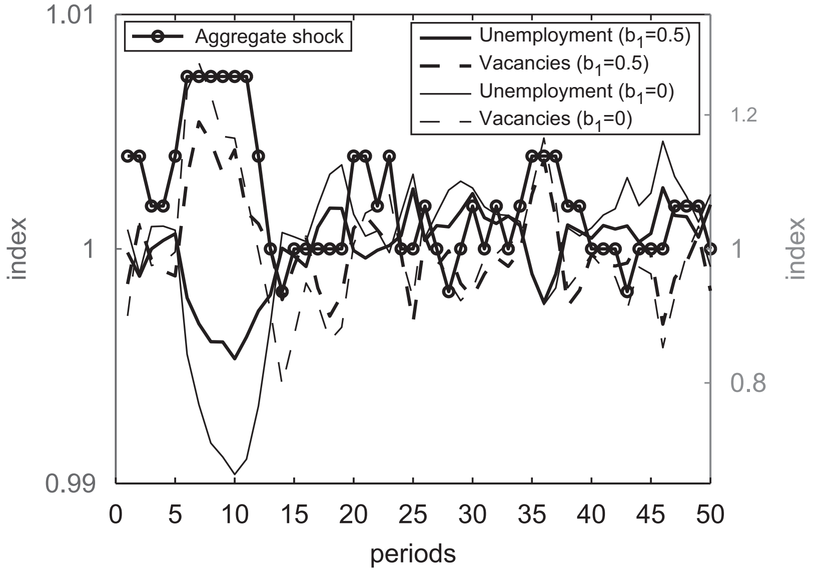

Fig. 2. Unemployment and vacancies (aggregate).

The main difference from making the value of leisure (home production) sensitive to the aggregate state, i.e., $b_1 = 0.50$, is in the volatility of the unemployment and vacancy series. From Fig. 2, it is clear that both unemployment and vacancies respond much less to the aggregate profitability shock when $b_1 = 0.50$ compared to b_1 near zero.

5.3. *Total factor productivity vs. average labor productivity*

One of the more interesting points of difference between the empirical literature on aggregate search models and the macroeconomic literature based on the stochastic growth model is the measurement of productivity. The search literature largely looks at average labor productivity (*ALP*), while the macroeconomics literature focuses on total factor productivity *TFP* as an exogenous shock.

In this paper, we are closer to the tradition of the stochastic growth model. We have treated profitability as exogenous where the profitability shock summarizes both technological and demand factors influencing the revenues of a producer. In our competitive model, revenues are the same as output so the profitability shock is most naturally interpreted as a variation in technology.[28]

If labor is freely mobile (no adjustment costs and no rigidities due to timing assumptions) between production sites, then the cross-sectional distribution of *ALP* is degenerate. Further, fluctuations in the average product of labor reflect the dependence of worker opportunities on aggregate profitability, b_1 in the model.

Generally though, *TFP* and *ALP* are not the same due to: (i) frictions in the adjustment process and (ii) idiosyncratic shocks. To see these influences, think about two extreme economies. In one, suppose that labor flows freely across producers, and in the second,

[28]Allowing monopolistic competition provides an interpretation of α as including a markup and the shocks to revenue including a relative demand disturbance. As the model studied here does not include product differentiation, α and the shocks should be more narrowly interpreted.

Table 7
Unemployment, vacancies and average labor productivity

Moment	JOLTS data	Estimated model
$\frac{\sigma(U)}{\sigma(a)}$	–	28.04
$\frac{\sigma(V)}{\sigma(a)}$	–	21.49
$\frac{\sigma(U)}{\sigma(ALP)}$	7.64	43.36
$\frac{\sigma(V)}{\sigma(ALP)}$	10.34	33.21
$corr(\theta, ALP)$	−0.70	0.44
$corr(a, ALP)$	–	0.60

Notes: Unemployment, U, comes from the CPS and vacancies, V, are directly calculated from JOLTS at a monthly frequency. θ represents labor market tightness ($\theta = (V/U)$). Monthly average labor productivity, ALP, is computed using industrial production as a measure of output and total hours for production workers as the measure of labor input. a represents the aggregate profitability shock in the model. $\sigma(x)$ is the standard deviation of x. The variables are in logs and are HP filtered.

suppose there are frictions in labor flows. Suppose the distribution of idiosyncratic shocks is the same in the two economies and is fixed over time.

As aggregate *TFP* varies, *ALP* varies in both of these economies. For fixed *TFP* in the second economy, *ALP* increases as labor flows from less productive to more productive producers. Thus, variations in *ALP* generally reflect both *TFP* and frictions.

ALP, however, is much easier to measure and thus plays a prominent role in the empirical literature. One of the advantages of our simulation environment is that we can use our model to generate a measure of *ALP*, for producer i in period t as $a_t\varepsilon_{i,t}(e_{i,t}h_{i,t})^{\alpha-1}$. The value of α used for our analysis is 0.65. In the context of the model, $\alpha<1$ reflects the presence of fixed factors of production such as managerial ability, structures and predetermined components of the equipment stock.

With these differences between *TFP* and *ALP* in mind, we return to a discussion of additional moments. The puzzle posed by Shimer (2005) concerns the standard deviation of unemployment and vacancies relative to *ALP*, as shown in Table 1. Our estimation, in contrast, has not focused on this moment *per se* but rather we target the absolute standard deviations on unemployment and vacancies.

Given our estimated model, we simulate average labor productivity and relate it to unemployment and vacancies. Our results are summarized in Table 7 along with relevant moments from the JOLTS data. For the simulated data, we compute *ALP* at the establishment level. The aggregate measure is total output divided by aggregate hours. For the actual data, *ALP* is computed using industrial production as a measure of output and total hours, *eh*, for production workers as the measure of labor input. The standard deviations of unemployment and vacancies are similar in magnitude to those reported in Shimer (2005) despite differences in data (JOLTS vs. Conference Board) and frequency (monthly vs. quarterly average).

From this table there is clearly a substantial difference between a and *ALP*: the correlation is 0.60. As discussed earlier, this reflects the interaction of frictions in labor flows and the idiosyncratic shocks in the estimated model.

There are two measures of the standard deviations of U and V relative to productivity. Looking at a, both unemployment and vacancies are substantially more volatile than productivity. This is also true when we use ALP as a productivity measure. The magnification puzzle isolated by Shimer (2005) is not present in this model. Indeed, in comparing the amplification in the model vs. the data when we use ALP, we have too much amplification.

The volatility of unemployment and vacancies relative to the variability in a stems from a few features of the labor market in our model. First, producers have all of the bargaining power. Second, the outside option of workers depends on a through the parameter b_1. Recall that b_1 is estimated in the model and is useful in matching the relationship between compensation and productivity, the $\omega|ALP$ moment.

To see the linkages, an increase in a leads to an increase in the profitability of hiring, which creates incentives on both the extensive and the intensive margins: a larger fraction of producers to post vacancies and post more vacancies. But, as $b_1>0$, the outside option of the workers is more valuable and thus average compensation rises with a. This effect dampens incentives for vacancy creation.

Also note that the time-series volatility of ALP is much less than the volatility of a. This is consistent with the intuition above that, in effect, with the estimated value of b_1 the economy is operating on a relatively flat labor supply curve so that average productivity fluctuates less than the aggregate shock.

Hagedorn and Manovskii (2006) argue that the elasticity of wages to average labor productivity is one for Shimer (2005). To study the importance of this elasticity for our results, we estimated a version of our model with $\omega|ALP = 1.0$ instead of $\omega|ALP = 0.45$. This change in the moments increased the estimate of b_1 from 0.5 to 0.87. For this case, $\sigma(U)/\sigma(p) = 26.80$ and $\sigma(V)/\sigma(p) = 24.7$. Relative to Table 8, there is less but still abundant amplification. It makes sense there is less amplification since wages absorb more of the productivity fluctuations in this case. However, even with the same elasticity of aggregate wages to productivity as assumed in Shimer (2005), we do not have an amplification puzzle. As emphasized, there are a substantial number of differences in our specification relative to Shimer (2005) including allowing for idiosyncratic profit shocks with endogenous job destruction, hours per worker variation, fixed and linear vacancy posting and firing costs, multiworker firms with decreasing returns, and endogenous labor productivity. Further work is required to isolate how each of these differences matter individually or interact but we note that Mortensen and Nagypal (2006) have some of the same components (e.g., time varying job destruction and vacancy-posting costs) and they find a mitigation of the amplification puzzle.

The behavior of ALP and labor market tightness poses another challenge to the analysis. As shown in Table 7, there is a negative correlation between labor market tightness and average labor productivity in the sample period data. But in the estimated model, as reported in the table, $corr(\theta, p)>0$. In Shimer (2005), the relationship between productivity and labor market tightness is not stable across sub-samples of the data set. For his overall sample, as in our model, the correlation is positive, but for the sub-samples covered by JOLTS this correlation is also negative.

This difference between model and data may reflect the fact that average labor productivity is endogenous. Thus reverse causality may be an important consideration. Understanding the instability between ALP and labor market tightness, as well as the factors influencing the interaction, remains an area for further work.

Table 8
Effects of reducing σ_ε

σ_ε	$corr(U, V)$	$\Delta e \leqslant -0.10$	$-0.10 < \Delta e \leqslant -0.025$	$< -0.025\Delta e \leqslant 0.025$
0.227	−0.90	0.052	0.071	0.794
0.170	−0.84	0.018	0.059	0.849
0.113	−0.68	0.000	0.035	0.892
0.057	−0.54	0	0.001	0.925

Notes: σ_ε is the standard deviation of the innovation for the idiosyncratic profitability shock. U is aggregate unemployment, and V is aggregate vacancies. The last three columns represent the size-weighted distribution of employment growth across simulated establishments from a model with fixed and linear vacancy-posting costs.

5.4. Idiosyncratic shocks

An important part of the specification of the model is the presence of producer-specific profitability shocks. The standard deviation of the innovation to the idiosyncratic component of profitability is estimated at 0.227. Table 8 shows how some of the key moments respond to a reduction in σ_ε from this estimated value to 0.057. A key lesson from these results is that the moments of the model economy do depend on the distribution of establishment shocks: these shocks are not smoothed out by aggregation.

As is evident from the table, the correlation between unemployment and vacancies becomes less negative for lower values of σ_ε. This is indicative of the fact that parameters describing microeconomic objects can impact aggregate variables in this framework. Further, the distribution of employment adjustment is more condensed as σ_ε falls. This makes sense since there are fewer large draws of the idiosyncratic shock and so, given adjustment costs, less variability in employment growth. The latter interacts with the aggregate dynamics since with a lower variance of idiosyncratic shocks there are fewer firms simultaneously in the creation and destruction ranges to respond to aggregate shocks. With less simultaneous creation and destruction, an aggregate shock has less of a simultaneous (but opposite) effect on unemployment and vacancies.

Overall, the results of Table 8 help to illustrate the importance of simultaneously matching the micro and the macromoments. To match the micromoments, both adjustment costs and a high variance of idiosyncratic shocks must be present. Moreover, without the high variance of idiosyncratic shocks but with the adjustment costs, we cannot match the macromoments such as the Beveridge curve.

6. Conclusions

The goal of this paper has been to study the implications of a search model for observed movements in hours, employment, vacancies and unemployment, drawing upon establishment and aggregate data. The microevidence guides and disciplines the models built to match aggregate observations, and the models based on the establishment level ought to be challenged to match aggregates.

Our framework is an extension of the standard search model with three features. First, in order to match establishment-level observations, we introduce non-convexities into the process of posting vacancies along with convex adjustment costs. Second, we introduce hours variations into the search model through an *ex ante* labor contracting structure.

Both of these features require us to create a non-trivial model of the producer. Finally, as in Mortensen and Pissarides (1994), establishment-specific shocks play a prominent role in the analysis and in the moments generated by the model.

The returns on these novel modeling features accrue from insights into the costs of vacancies and firing and the ability of the model to resolve some of the puzzling aspects of the data. In that regard, we find the following fixed and linear adjustment costs are necessary to match observations. In the estimation, the model with vacancy-posting cost and the one with firing costs do equally as well matching key moments. The estimated model does not suffer from the amplification problem highlighted in Shimer (2005), even with a high elasticity of wages to productivity. Based on the non-convexities, the model does match some of the inaction and bursts reported for vacancies and employment adjustment at the establishment level.

There are many extensions to consider, intended to further shrink the gap between the model and facts about labor markets.[29] One important issue is the role of quits in the model. Extending the model to include stochastic and endogenous quits may be important for matching observations on the frequency of zero net employment growth, about 30% of observations, and the frequency of zero vacancies, about 45% of observations. With deterministic quits and non-convex adjustment costs, it is not possible to explain inaction in both vacancies and net employment growth. A second issue is the role of capital in the model. The interaction of capital adjustment and labor adjustment costs remains of considerable interest.

References

Abraham, K., Haltiwanger, J., 1995. Real wages and the business cycle. Journal of Economic Literature 33 (3), 1215–1264.

Acemoglu, D., Hawkins, W., 2006. Equilibrium unemployment in a generalized search model. ⟨http://troi.cc.rochester.edu/~whawkin2/papers/acehaw.pdf⟩.

Barsky, R., Solon, G., Parker, J., 1994. Measuring the cyclicality of real wages: how important is the composition bias? Quarterly Journal of Economics 109, 1–25.

Bils, M., 1987. Real wages over the business cycle: evidence from panel data. Journal of Political Economy 85, 666–689.

Caballero, R., Engel, E., 1993. Microeconomic adjustment hazards and aggregate dynamics. Quarterly Journal of Economics 108, 313–358.

Cooper, R., Willis, J., 2002. The cost of labor adjustment: inferences from the gap. Federal Reserve Bank of Kansas City, Working Paper #02-11.

Cooper, R., Haltiwanger, J., Willis, J., 2004. Dynamics of labor demand: evidence from plant-level observations and aggregate implications. NBER Working Paper #10297.

Cooper, R., Haltiwanger, J., Willis, J., 2007. Implications of search frictions: matching aggregate and establishment-level observations. NBER Working Paper #13115.

Davis, S., Faberman, R., Haltiwanger, J., 2006a. The establishment-level behavior of vacancies and hiring. Mimeo.

Davis, S., Faberman, R., Haltiwanger, J., 2006b. The flow approach to labor markets: new data sources and micro-macro links. Journal of Economic Perspectives 20 (3), 3–26.

Faberman, R.J., 2005. Analyzing the JOLTS hires and separations data. In: Proceedings of the 2005 American Statistical Association.

Fujita, S., Ramey, G., 2005. The dynamic beveridge curve. Federal Reserve Bank of Philadelphia, Working Paper #05-22. ⟨http://www.phil.frb.org/files/wps/2005/WP05-22.pdf⟩.

[29]More extensive discussion of extensions and implications is provided in Cooper et al. (2007).

Hagedorn, M., Manovskii, I., 2006. The cyclical behavior of equilibrium unemployment and vacancies revisited. Mimeo, University of Frankfurt and University of Pennsylvania. ⟨http://www.econ.upenn.edu/~manovski/papers/BCUV.pdf⟩.

Hall, R., 2005. Employment fluctuations with equilibrium wage stickiness. American Economic Review 95, 53–69.

Mortensen, D., Nagypal, E., 2006. More on unemployment and vacancy fluctuations. Northwestern University. Mimeo.

Mortensen, D.T., Pissarides, C.A., 1994. Job creation and job destruction in the theory of unemployment. Review of Economic Studies 61, 397–415.

Rudanko, L., 2006. Idiosyncratic and aggregate risk in a frictional labor market. Mimeo.

Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. American Economic Review 95, 25–49.

Tasci, M., 2006. On the job search and labor reallocation. ⟨http://www.eco.utexas.edu/graduate/Tasci/ojsandlmr.pdf⟩.

Tauchen, G., 1986. Finite state Markov-chain approximation to univariate and vector autoregressions. Economics Letters 20, 177–181.