

The Dynamic Beveridge Curve

Shigeru Fujita and Garey Ramey*

August 14, 2005

This version: November 29, 2005

Abstract

In aggregate U.S. data, exogenous shocks to labor productivity induce highly persistent and hump-shaped responses to both the vacancy-unemployment ratio and employment. We show that the standard version of the Mortensen-Pissarides matching model fails to replicate this dynamic pattern due to the rapid responses of new job openings. We extend the model by introducing a sunk cost for creating job positions, motivated by the well-known fact that worker turnover exceeds job turnover. In the matching model with sunk costs, new job openings react sluggishly to shocks, leading to highly realistic dynamics.

JEL codes: E32, J63, J64

Keywords: Unemployment, Vacancies, Labor Adjustment, Matching

*Fujita: Federal Reserve Bank of Philadelphia; e-mail: shigeru.fujita@phil.frb.org. Ramey: University of California, San Diego; e-mail: gramey@ucsd.edu. For helpful comments, we are grateful to Valerie Ramey, Wouter den Haan, and seminar participants at Federal Reserve Bank of Philadelphia, Bank of Japan, Kansai University, Michigan State University, NBER-EFRSW meeting, University of Tokyo, UCSD, the Wharton School, Yokohama National University. Michael Conlow provided research assistance. The views expressed here are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1 Introduction

The Mortensen-Pissarides job matching model has become the standard framework for analyzing macroeconomic labor adjustment. Recent research has focused on the ability of this model to account for the high volatility of unemployment using empirically reasonable shocks to labor productivity.¹ A largely overlooked question concerns whether the model can account for the *dynamic* characteristics of labor market adjustment over the business cycle.² For example, aggregate employment is widely viewed as a lagging indicator of the business cycle by policy makers and business analysts. Business cycles are also linked to a characteristic pattern of unemployment and vacancy movements. As shown in Figure 1, this pattern involves long counterclockwise loops in the space of vacancies and unemployment, with the vacancy-unemployment ratio rising in booms and falling in recessions.³

It is reasonable to imagine that the Mortensen-Pissarides model captures these features of aggregate labor market adjustment, since it incorporates a sluggish labor reallocation process that generates qualitatively plausible movements in unemployment and vacancies.⁴ This paper shows, however, that the model in its standard form cannot account for the cyclical pattern of unemployment and vacancies when calibrated to U.S. data, nor can it capture the associated dynamics of employment adjustment. Modifying the model by introducing a sunk cost for the creation of new job positions, however, dramatically improves the performance of the model in matching the empirical

¹See Costain and Reiter (2005), De Bock (2005), Hagedorn and Manovskii (2005), Krause and Lubik (2004), Hall (2005), Mortensen and Nagypal (2005), Rudanko (2005), Shimer (2005) and Silva and Toledo (2005).

²Papers that evaluate dynamic aspects of the job matching model include Bastos (2004), Cole and Rogerson (1999), Fujita (2003), Fujita (2004), Menzio (2004), Mortensen (1994) and Yashiv (2005).

³The negative empirical relationship between unemployment and vacancies is known as the Beveridge curve in recognition of the contribution of W.H. Beveridge. For a background discussion on the Beveridge curve, see Abraham (1987), Blanchard and Diamond (1989), Jackman, Layard, and Pissarides (1989) and the references therein. Bleakley and Fuhrer (1997) provide an updated empirical evaluation of the Beveridge curve. In this paper we focus on the dynamic characteristics of vacancy-unemployment relationship at business cycle frequencies.

⁴Pissarides (2000) (pp. 26-33) provides a qualitative analysis of the the dynamic properties of the model.

dynamics.

We begin our analysis by estimating a three-variable VAR at quarterly frequency for the U.S. consisting of aggregate labor productivity, the vacancy-unemployment ratio and employment. In the context of the matching model, the vacancy-unemployment ratio reflects the intensity of search activities in the matching market, and thus it influences employment through match formation. Procyclical movements in the vacancy-unemployment ratio capture the adjustment pattern depicted in Figure 1. This specification allows us to trace the dynamic effects of labor productivity shocks on matching intensity, as measured by the vacancy-unemployment ratio, and employment. Note that in the context of the matching model, our identified cyclical productivity shock can be broadly interpreted as resulting from impulses to demand or technology.⁵

The estimates reveal that the productivity driving process identified in our VAR has a high positive correlation with matching intensity and employment at lags of 0-2 quarters and 1-3 quarters, respectively. Correspondingly, the impulse responses for the two variables exhibit hump shapes, with matching intensity and employment peaking after four and five quarters, respectively. In the period of the shock, matching intensity jumps by about a third of the way toward its peak response, while employment does not jump.

We next consider the ability of the Mortensen-Pissarides model to account for these patterns. A discrete-time version of the model is calibrated to U.S. data, and simulated data are obtained using a productivity driving process that approximates the estimated process. The resulting dynamic correlations are sharply peaked at zero lag, in stark contrast to the empirical correlations. The impulse responses derived from the simulated data exhibit excessively large jumps in the period of the shock, and the subsequent responses fail to capture the empirical hump shapes. Thus, the standard matching model fails to capture the dynamic characteristics of the data.

The poor performance of the standard model can be traced to the rapid adjustment of new job openings following a shock. Since firms in the model incur no costs for creating job positions, they can freely enter and exit the vacancy pool in response to changes in productivity, causing vacancies to closely track the driving process.

⁵Shimer (2005) offers a similar interpretation concerning shocks to labor productivity.

Evidence suggests, however, that job positions behave differently from vacancies. Davis et al. (1996), for example, emphasize that jobs themselves are created and destroyed at a rate substantially lower than the rate at which firms open up and fill vacancies.⁶ This makes sense if refilling existing job openings is more costly than creating new ones, pointing to the presence of sunk costs for creation of job positions. Such costs could tend to slow the adjustment of positions in response to shocks, leading to sluggishness in vacancy adjustment.

To assess the importance of this mechanism for the dynamic performance of the matching model, we extend the model in a very simple way by requiring potential entrant firms to pay a cost when they create new job positions. The cost of the marginal position is assumed to be increasing in the number of positions created a period, reflecting diseconomies at the firm or aggregate level. Positions are durable in the sense that they remain active, whether filled or unfilled, until they are eliminated by obsolescence. In particular, when a worker quits, the firm can post a vacancy to refill the position without incurring the creation cost. This provides a meaningful distinction between job flows and worker flows.

Simulated data from the calibrated sunk cost model reproduce the dynamic correlations for matching intensity and employment with great precision. Impulse responses exhibit realistic hump shapes, with the correct timing of peaks. The initial jump in matching intensity closely matches the empirical estimate, while the counterfactual jump in employment is significantly reduced relative to the standard model.

The superior performance of the sunk cost model is explained by the fact that, due to increasing marginal creation costs, the response of new openings following a shock is spread out over time. Moreover, durability of job positions means that changes in new openings are not rapidly reversed. Thus, the vacancy pool continues to respond to the shock even as productivity returns to the steady state. In this way, sunk costs for creating job positions, working through the adjustment of new openings, give rise to realistic labor market adjustment.

⁶Drawing on several sources of information, we calculate below that nearly 19 percent of employed workers separate from their jobs over the ensuing quarter, while the rate of job destruction itself is only about 8 percent per quarter. Thus, the “churning” of workers across filled job positions amounts to over 10 percent of employment each quarter.

Related literature. In their pioneering work, Blanchard and Diamond (1989, 1990) study the dynamic adjustment of vacancies and unemployment using VAR models that incorporate the labor force as a third variable. The impulse responses estimated by these authors capture the hump-shaped responses of vacancies and unemployment, consistent with our empirical analysis using the single matching intensity variable. Fujita (2004) studies vacancy persistence using an identification approach that builds on the method of Blanchard and Diamond (1989, 1990), and he shows that the Mortensen-Pissarides model cannot generate realistic persistence. In the present paper, we utilize a different identification procedure, based on direct use of labor productivity data, and we also consider employment dynamics. Further, we show that the deficiencies of the model can be corrected by introducing a sunk job creation cost.

Fujita (2003), Hornstein, Krusell, and Violante (2004) and Yashiv (2005) consider job matching models in which firms bear costs for adjusting the stock of job positions. In Fujita (2003), sluggish adjustment results from an initial job planning stage, while Yashiv (2005) introduces sluggishness by means of a standard convex adjustment cost. Neither of these papers considers costs of creating job positions, nor do they develop the distinction between job flows and worker flows. Hornstein, Krusell, and Violante (2004) focus on steady-state economic growth rather than cyclical adjustment.

The remainder of the paper proceeds as follows. Section 2 presents the empirical findings from the estimated VAR, and the extended version of the matching model is developed in Section 3. Section 4 describes the empirical implementation of the model, including its calibration, and results are presented in Section 5. Section 6 concludes.

2 Cyclical shocks and adjustment

In this section we assess the cyclical characteristics of unemployment, vacancies and employment for the U.S. between 1951 and 2004. The first step is to determine an appropriate measure of cyclical shocks. Shimer (2005) has argued that average labor productivity can be used as a cyclical indicator for purposes of evaluating the matching model. This is compelling in that productivity incorporates a broad range of factors that influence the returns to employment. Productivity can itself be influenced by the labor market, however, so we must account for potential endogeneity in evaluating the

effect of productivity on labor market variables.

We next specify an empirical relationship between productivity and the labor market variables. A large body of empirical work has demonstrated that hiring flows are well approximated by a constant returns-to-scale function of unemployment and job vacancies.⁷ This implies that hiring may be regarded as a function of the vacancy-unemployment ratio. Thus, it makes sense to summarize the relationship between unemployment and vacancies in terms of the vacancy-unemployment ratio; we refer to the latter as *matching intensity*. This leads to the following reduced-form VAR:

$$A(L) \begin{bmatrix} \ln z_t \\ \ln vu_t \\ \ln emp_t \end{bmatrix} = \begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^{vu} \\ \varepsilon_t^e \end{bmatrix}, \quad (1)$$

where $\ln z_t$, $\ln vu_t$ and $\ln emp_t$ denote the logs of labor productivity, the vacancy-unemployment ratio and the employment-population ratio, respectively; ε_t^z , ε_t^{vu} and ε_t^e are the reduced-form residuals of the three equations; and $A(L)$ is a lag polynomial matrix, with $A(0)$ being the identity matrix. Each variable is detrended by regressing on Chebyshev polynomial time trends. The sample period is 1951:Q1 to 2004:Q3.⁸

To identify the exogenous component of productivity, we use the recursive ordering $\ln z_t$, $\ln vu_t$ and $\ln emp_t$, so that innovations to labor productivity are treated as exogenous with respect to matching intensity and employment in the current quarter.⁹ The

⁷See Petrongolo and Pissarides (2001) for a detailed survey of research on the estimation of matching functions.

⁸Labor productivity is measured as real GDP divided by civilian employment, 16 years and over. For the vacancy-unemployment ratio, we use the index of newspaper help-wanted advertising divided by the number of unemployed, 16 years and over. The employment-population ratio consists of civilian employment, 16 years and over, divided by civilian noninstitutional population. Data were obtained from the FRED II database maintained by the Federal Reserve Bank of St. Louis. The sample period is restricted by the availability of the help-wanted index. The model was estimated for all combinations of lag lengths (up to four) and polynomial trends (up to fourth order). Based on the AIC, we choose lag lengths of three quarters and third-order time trends in each equation.

⁹Our identification of cyclical shocks can be motivated by the reasonable restriction that innovations to the vacancy-unemployment ratio should not have a positive contemporaneous effect on productivity. Among the six possible orderings of the three variables, this restriction eliminates the three orderings in which $\ln vu_t$ is ordered before $\ln p_t$. The results for the remaining three orderings are nearly

exogenous productivity series, denoted by $\ln \hat{z}_t$, is determined by

$$\hat{A}_{11}(L) \ln \hat{z}_t = \hat{\varepsilon}_t^z, \quad (2)$$

where $\hat{A}_{11}(L)$ is the estimated value of the polynomial in the first row and first column of $A(L)$ and $\hat{\varepsilon}_t^z$ indicates the fitted residuals from (1). Note that $\ln vu_t$ and $\ln emp_t$ do not enter in (2). In other words, $\ln \hat{z}_t$ is a generated series from the first equation in the VAR system (1), with the restriction that the coefficients on the lagged values of $\ln vu_t$ and $\ln emp_t$ are equated to zero. This allows us to extract the exogenous cyclical component of $\ln z_t$ by suppressing the feedbacks to productivity from matching intensity and employment that are associated with the identified productivity shock.

We trace out the dynamic effects of the identified productivity shock on matching intensity and employment by regressing $\ln vu_t$ and $\ln emp_t$ on the identified productivity shock $\hat{\varepsilon}_t^z$ along with the lagged values of $\ln vu_t$, $\ln emp_t$ and $\ln \hat{z}_t$:

$$B(L) \begin{bmatrix} \ln vu_t \\ \ln emp_t \end{bmatrix} = C(L) \ln \hat{z}_{t-1} + D \hat{\varepsilon}_t^z + \begin{bmatrix} \eta_t^{vu} \\ \eta_t^e \end{bmatrix}, \quad (3)$$

where $B(0)$ is an identity matrix; η_t^{vu} and η_t^e are the innovations to $\ln vu_t$ and $\ln emp_t$ in the above system; and $B(L)$, $C(L)$, and D indicate a polynomial matrix, a polynomial vector and a real vector, respectively.¹⁰

The conditional correlations of $\ln vu_t$ and $\ln emp_t$ with $\ln \hat{z}_t$ at various leads and lags provide a summary measure of the dynamic relationship between the identified productivity process and the other two variables. These correlations can be computed by assuming that the productivity shock is the only source of the variations in the three variables, i.e., the innovations in (3), η_t^{vu} and η_t^e , are suppressed. Figure 2 reports the correlations graphically. Observe that matching intensity and employment are highly correlated with lagged values of exogenous productivity, with peak correlations at lags of 0-2 quarters and 1-3 quarters, respectively.

Figure 3 draws the impulse response functions to a one-standard-deviation productivity shock based on the system (3). As seen in the top panel, $\ln \hat{z}_t$ jumps by about

indistinguishable from one another. For concreteness, we focus on one particular choice from among these three.

¹⁰In this estimation we use lag lengths of three quarters in order to maintain consistency with the previous specification.

0.7 percent as a result of the shock, then returns monotonically to its steady state after oscillating slightly for two quarters.¹¹ The dynamic pattern of unemployment and vacancy adjustment is captured in the middle panel as a hump-shaped response of the variable $\ln vu_t$. Vacancies initially jump relative to unemployment by about 5 percent. The vacancy-unemployment ratio continues to rise rapidly for four quarters, with a peak at roughly 12 percent above the steady state, after which it falls fairly rapidly for eight quarters or so. The variable $\ln emp_t$, in the bottom panel, does not jump in the period of the shock, but otherwise its response closely mimics the response of $\ln vu_t$, with a one-quarter lag and a peak of about 0.35 percent above the steady state. This indicates that the adjustment of employment is closely tied to the dynamics of unemployment and vacancies.

We now link these results back to adjustments of unemployment and vacancies themselves. The matching intensity variable $\ln vu_t$ can be decoupled into separate vacancy and unemployment components by estimating a bivariate VAR in which the fitted exogenous productivity process and its innovations appear as independent variables:

$$X(L) \begin{bmatrix} \ln vac_t \\ \ln unemp_t \end{bmatrix} = F(L) \ln \hat{z}_{t-1} + G \hat{\varepsilon}_t^z + \begin{bmatrix} \eta_t^v \\ \eta_t^u \end{bmatrix},$$

where $\ln vac_t$ and $\ln unemp_t$ represent the logs of vacancies and unemployment, respectively; η_t^v and η_t^u are innovations to vacancies and unemployment, respectively; and $X(L)$, $F(L)$ and G indicate a polynomial matrix, a polynomial vector and a real vector, respectively.¹²

Figure 4 depicts the responses of vacancies and unemployment to a one-standard-deviation shock to productivity. Observe that vacancies and unemployment respond in opposite directions, but by the same magnitude, in the period of the shock, with vacancies rising by 2 percent and unemployment falling by 2 percent. The variables move as rough mirror images of each other over the next 12 quarters, with unemployment decreasing at a slightly faster rate in the first year. The vacancy response peaks

¹¹Note that the estimated impulse response for $\ln \hat{z}_t$ is very close to the one generated by the technology process $\ln z_t = .95 \ln z_{t-1} + \varepsilon_t$, with $\sigma_z = .007$, that is standard in RBC analysis.

¹²Unemployment and vacancies are detrended by regressing on third-order Chebyshev polynomial time trends prior to the above estimation. The lag length is again set to three quarters as in the previous estimations.

at 5.5 percent in the third period following the shock, and the unemployment response troughs at -6.5 percent in the fourth quarter. It follows that just under half of the overall adjustment of matching intensity over the first 12 quarters is accounted for by changes in vacancies. By the 13th quarter, vacancies have substantially returned to their steady state, while unemployment maintains an extended gap of about one percentage point below its steady state.

3 Matching Model with Costly Job Positions

In this section we present a modification of the standard matching model in which firms incur a sunk cost to create new job positions. Once a position is created, it continues to exist, either filled or unfilled, until it is eliminated by obsolescence. Firms post vacancies to fill newly-created positions. Furthermore, when a worker leaves a filled position for reasons other than obsolescence, the firm may post a vacancy to refill the position.

3.1 Model description

The model consists of a unit mass of workers and an infinite mass of firms. In any given period, workers may be either matched with a firm or in the worker matching pool searching for a match. Firms are either matched with workers, in the vacancy matching pool searching for a match, or in a pool of potential entrant firms that are not actively searching. Matched worker-firm pairs engage in production, while workers and firms in the matching pools seek to form new matched pairs. Potential entrant firms may choose to create new job openings in any period.

Matching and separation. While in the worker matching pool, a worker receives a flow benefit of b per period, which may be interpreted as utility from leisure. Each firm in the vacancy pool pays a posting cost of c per period. The net number of new matches created in period t is given by the matching function $m(u_t, v_t)$, where u_t and v_t indicate the sizes of the period t worker and vacancy matching pools, respectively. The u_t pool includes workers who separate from their employers at the start of the period;

thus, these workers can be matched with new employers during the period. This allows us to account for the large volume of monthly employer-to-employer flows in our calibration. Similarly, the v_t pool includes job positions that experience separation at the start of the period for reasons other than obsolescence.¹³ The function m satisfies the customary properties.¹⁴

At the start of period t , each worker-firm match undergoes two distinct separation hazards. First, separation occurs with probability λ^d as a consequence of *obsolescence*. In this case, the worker enters the period t matching pool and the firm enters the pool of potential entrants. Second, separation may occur for reasons that do not destroy the job position; we which we refer to this generically as a *quit*. For continuing matches, quits occur with probability λ^q , while they occur with the higher probability λ^n for matches that were newly-formed in the preceding period; this reflects the higher observed separation hazards among low-tenure employment relationships. Following a quit, the worker enters the period t worker matching pool and the firm enters the period t vacancy matching pool.

If the worker-firm match survives the separation hazards, then the agents negotiate a contract that divides the period t surplus according to the Nash bargaining solution, where π gives the worker's bargaining weight and the threat point is a quit. Given that the worker and firm agree to continue, the match incurs a flow capital cost of κ and engages in production in period t . The output of the match is given by the productivity level z_t , which evolves according to a Markov process.

Creation of job positions. Potential entrant firms create new job positions in period t after observing z_t . The cost of creating a position depends on the total number of positions created in the period. Let $K(e_t)$ give the marginal creation cost, where e_t indicates the number of new job openings in period t , and $K(0) = 0$ is assumed. We distinguish between two cases. First, if $K(e_t) = 0$ for all e_t , then the model reduces

¹³The latter job positions may reenter the vacancy pool if the value of a vacancy is nonnegative, which will be true in the equilibria that we consider.

¹⁴That is, m is twice continuously differentiable, strictly increasing and strictly concave in each argument, and homogeneous of degree one. Further, $m(0, v_t) = m(u_t, 0) = 0$ and $m(u_t, v_t) \leq \min\{u_t, v_t\}$ for all u_t and v_t .

to a discrete-time version of the standard matching model, described in Chapter 1 of Pissarides (2000), for example; we call this the *standard matching model*. Second, if $K(e_t)$ is strictly increasing, then creating a new job position entails a positive sunk cost in addition to the cost of vacancy posting; this is referred to as the *sunk cost model*.

In the sunk cost model, the positive creation cost represents costs of adjusting physical and organizational capital that are associated with new job positions. These costs rise with the number of new positions created, due to diseconomies at the level of individual firms or in the aggregate. This generates an incentive to smooth out the creation of job positions. Alternatively, job openings may be tied to new business opportunities that must be identified through a search process. In order to locate a greater number of opportunities, more intensive search is required, leading to an increase in the creation cost for marginal job positions.¹⁵

3.2 Equilibrium

Let $\theta_t = v_t/u_t$ denote the size of the vacancy matching pool relative to the worker matching pool. For a worker who begins period t in the worker matching pool, the probability of ending period t in a job match is

$$\frac{m(u_t, v_t)}{u_t} = m(1, \theta_t) = p(\theta_t),$$

where we have made use of the linear homogeneity of the function m . Let S_t indicate the value of surplus for a match that survives the separation hazards in period t . A worker in the worker matching pool receives the benefit b along with a proportion of the surplus from any match made in period t that survives into period $t + 1$. Thus, the expected present value of current and future receipts for an unmatched worker is given by

$$U_t = b + \beta E_t \left[p(\theta_t)(1 - \lambda^d)(1 - \lambda^n)\pi S_{t+1} + U_{t+1} \right], \quad (4)$$

where β indicates the discount factor.

¹⁵Caballero and Hammour (1994, 1996) consider adjustment costs as a motivation for increasing marginal job creation costs. See Fujita and Ramey (2004) for a model in which search for new investment opportunities underlies sluggish adjustment.

Similarly, for a firm that begins period t in the vacancy matching pool, the probability of ending period t in a job match is

$$\frac{m(u_t, v_t)}{v_t} = m\left(\frac{1}{\theta_t}, 1\right) = q(\theta_t),$$

and the expected present value of the firm's current and future receipts is

$$V_t = -c + \beta E_t \left[q(\theta_t)(1 - \lambda^d)(1 - \lambda^n)(1 - \pi)S_{t+1} + (1 - \lambda^d)V_{t+1} \right]. \quad (5)$$

Note that the firm receives the outside option value V_{t+1} only if separation occurs as a consequence of a quit, whereas the worker receives U_{t+1} after either obsolescence or a quit. This follows from the fact that the job ceases to exist following obsolescence.

A job match that survives the separation hazards in period t obtains the following expected present value:

$$N_t = z_t - \kappa + \beta E_t \left[(1 - \lambda^d)(1 - \lambda^q)S_{t+1} + U_{t+1} + (1 - \lambda^d)V_{t+1} \right]. \quad (6)$$

Equilibrium surplus is thus defined by

$$S_t = N_t - U_t - V_t. \quad (7)$$

Standard model. In the standard model, entry into the vacancy pool entails zero sunk cost. Thus, competitive pressure from potential entrant firms will drive the equilibrium value of a vacancy to zero, i.e., $V_t = 0$ must hold for all t . Using (5), this gives rise to the following free entry condition:

$$\frac{c}{q(\theta_t)} = \beta E_t \left[(1 - \lambda^d)(1 - \lambda^n)(1 - \pi)S_{t+1} \right]. \quad (8)$$

Equations (4) and (6)-(8) determine equilibrium paths of U_t , N_t , S_t and θ_t for a given process z_t . In particular, vacancies adjust freely in each period in order to maintain the relationship $v_t/u_t = \theta_t$.

The equilibrium solution determines the following law of motion for u_t :

$$u_t = (1 - p(\theta_{t-1}))u_{t-1} + [\lambda^d + (1 - \lambda^d)\lambda^q](1 - u_{t-1}) + [\lambda^d + (1 - \lambda^d)\lambda^n]p(\theta_{t-1})u_{t-1}. \quad (9)$$

The worker matching pool u_t consists of three terms. First, workers who were in u_{t-1} and did not match with a firm in period $t - 1$ remain in u_t . Second, workers who were

not in u_{t-1} will enter u_t if they separate at the start of period t , which occurs with probability $[\lambda^d + (1 - \lambda^d)\lambda^q]$. Finally, workers who were in u_{t-1} and were matched in period $t - 1$ will enter u_t if their new match separates at the start of period t ; the probability of this is $[\lambda^d + (1 - \lambda^d)\lambda^n]$. Note that our timing assumptions allow a worker to flow between employers within a period, based on separating from one match at the start of the period and forming another during the period.¹⁶

Sunk cost model. In the sunk cost model, entry into the vacancy matching pool is constrained by the cost of creating new job positions. The free entry condition (8) is replaced by

$$V_t = K(e_t). \quad (10)$$

The implied law of motion for v_t is

$$v_t = (1 - q(\theta_{t-1}))(1 - \lambda^d)v_{t-1} + (1 - \lambda^d)\lambda^q(1 - u_{t-1}) + (1 - \lambda^d)\lambda^n q(\theta_{t-1})v_{t-1} + e_t. \quad (11)$$

The vacancy matching pool v_t consists of four terms. The first three terms correspond to those of (9), modified by the fact that obsolete job positions do not return to v_t . The fourth term is new openings e_t . Equilibrium paths of U_t , V_t , N_t , S_t , u_t , v_t and e_t for the sunk cost model are determined by equations (4)-(7), (9)-(11), and the identity $v_t/u_t = \theta_t$, given the z_t process and the predetermined variables u_{t-1} and v_{t-1} . Importantly, the vacancy matching pool becomes a state variable in the sunk cost model, in contrast to the standard model, where vacancies are a jump variable.

¹⁶An alternative specification would allow for on-the-job search by a subset of the employed workers, giving rise to direct flows from one employer to another without an intervening unemployment spell. We are currently investigating this specification, and our findings will be included in the next version of the paper.

4 Empirical Implementation

4.1 Functional forms

We adopt the following functional forms. For the matching function we use the form introduced by den Haan et al. (2000):

$$m(u_t, v_t) = \frac{u_t v_t}{(u_t^l + v_t^l)^{1/l}}. \quad (12)$$

Since this function always lies below unity, it does not require truncation, in contrast to the standard Cobb-Douglas specification.

The stochastic process for productivity assumes the standard first-order autoregressive form:

$$\ln z_t = (1 - \rho) \ln z + \rho \ln z_{t-1} + \varepsilon_t,$$

where ε_t is normally distributed with mean zero and standard deviation σ_ε . Note that the realization of ε_t is observed by all workers and firms in the economy at the start of each period t .

Finally, the creation cost function $K(e_t)$ is specified as

$$K(e_t) = \bar{K} e_t, \quad (13)$$

where $\bar{K} = 0$ in the standard model, and $\bar{K} > 0$ in the sunk cost model.

4.2 Calibration

Our calibration of the model utilizes monthly data on flows and stocks of workers and job vacancies, combined with quarterly job flow and aggregate income data. Moreover, the empirical evaluation makes use of quarterly output data. Thus, measured monthly flow rates will be time aggregated to obtain appropriate quarterly rates, and our calibrated model will be analyzed at quarterly frequency.

Separation rates. The monthly average separation rate for job matches can be calculated as the average of monthly outflows from job matches divided by total employment. Monthly outflows consist of worker transitions to new job matches, unemployment, and out of the labor force. Fallick and Fleischman (2004) have recently provided

measures of aggregate U.S. stocks and flows of workers, derived from the Current Population Survey (CPS) for 1994 and 1996-2003, that include these three transitions. The numbers reported by Fallick and Fleischman imply outflows from unemployment that are inconsistent with steady state, so we adjust them slightly to achieve consistent steady-state flows and stocks. Adjusted flows and stocks are reported in Table 1. The units represent percentages of noninstitutional civilian population aged 16 and over. Observe that 4.15 new matches are created each month, while total employment is 63.15; thus, the implied steady-state separation rate is $4.15/63.15 = .0657$ per month.

This figure may be corroborated using Anderson and Meyer (1994), who consider a large panel of workers derived from U.S. state unemployment insurance records for 1978-1984. In their sample, 17.23 percent of job matches experience a permanent separation each quarter, which translates into an average monthly separation rate of .0611.

Anderson and Meyer (1994) also find that 43.42 percent of job matches observed during a quarter have duration of less than one year, which cannot hold unless these newer matches experience separation at rates in excess of .0611. In fact, separation rates for new hires decline sharply with job tenure; see Farber (1999), for example. Moreover, as noted by Pries (2004), failure to account for higher separation rates among new matches may cause the dynamics of labor adjustment to be significantly misstated.

To incorporate this factor in a simple way, the separation rate is assumed to be higher in the first month of a match.¹⁷ We choose separation rates in the first and subsequent months so that, in the steady state, 43.42 percent of matches have duration of less than one year when the flow of new matches is 4.15 per month. This yields monthly separation rates of .3663 and .0446 in the first and subsequent months, respectively, which translate into the following quarterly separation rates in the first

¹⁷This calibration procedure matches the distribution of employment durations, but not the duration-specific survival probabilities, since separations are crowded into the first month. Fitting the survival probabilities would require match-specific separation rates, and the number of categories of employed workers would correspondingly expand. In this paper we have opted instead for the simplest specification that allows new matches to separate at a higher rate.

quarter and subsequent quarters:

$$\lambda^d + (1 - \lambda^d)\lambda^n = .4212, \quad (14)$$

$$\lambda^d + (1 - \lambda^d)\lambda^q = .1279. \quad (15)$$

Job finding rate. As reported in Table 1, 36.6 percent of new matches at the end of a month are formed with workers who were employed at the start of the month, and an equal percentage are formed with workers who were out of the labor force at the start of the month. Since all job vacancies are included in the vacancy matching pool, the relevant worker matching pool must take account of all workers who would accept these jobs, including those who start the month employed or out of the labor force.

First consider the workers who are recorded as out of the labor force, but are still available for work, at the start of each month. Table 1 shows that the flow of workers from out of the labor force into employment or unemployment is 2.45 per month, so the steady state number of these available workers must be at least 2.45. In addition, some number of “discouraged” workers remains out of the labor force at the end of the month. Evidence from the 1994 CPS, discussed in Castillo (1998, p. 36), indicates that this group amounts to anywhere between 22.5 and 77.5 percent of the officially unemployed population. For simplicity, we assume that the number of available workers is equal to the number of unemployed workers, which is 3.4 in the steady state. This implies that the discouraged worker population amounts to 27.9 percent of official unemployment, lying at the low end of Castillo’s range.

Next consider the workers who start the month employed. Assume that all workers who separate from employment within the month are employed, unemployed or available for work at the end of the month, i.e., movement between the available and unavailable components of the out-of-labor-force population is ignored. Of the 4.15 workers who separate during the month, 1.6 are employed at the end of the month, yielding a monthly job finding rate of $1.6/4.15 = .3855$ in the steady state. Correspondingly, monthly job finding rates for unemployed and available workers are $.95/3.4 = .2794$ and $1.6/3.4 = .4706$, respectively.

To calculate appropriate quarterly job finding rates, we must account for separations and rematchings that occur at monthly frequency within the quarter. The above

monthly separation and job finding rates may be used to construct a monthly transition matrix for worker who move between the states of employed at a new employer, employed at the same employer, unemployed, and out of labor force but available. Based on this matrix, we calculate that a worker who separates from a match at the start of a quarter will end the quarter in one of the two employed states with probability .6413. The corresponding figures for workers who are unemployed and available at the start of the quarter are .5834 and .6579, respectively. Averaging over the three groups gives an overall job finding rate of $p(\theta) = .6285$ per quarter.

Vacancy filling rate. To determine the vacancy filling rate, we use evidence from the Job Opening and Labor Turnover Survey (JOLTS). JOLTS measures the vacancy rate as follows:

$$vac_rate = \frac{v^m}{1 - u^m + v^m},$$

where $v^m = (1 - q(\theta))v$ gives the number of vacancies, and $u^m = (1 - p(\theta))u$ gives the number of workers, remaining in the matching pool at the end of a quarter. According to the JOLTS, the average vacancy rate is 2.5 percent over the period of December 2000 through May 2005. Based on the weakness of the U.S. labor market between 2001 and 2003, it seems reasonable to view this as an underestimate, so we specify a slightly higher value $vac_rate = .03$. Further, we can use the previously-calculated separation and job finding rates, together with the law of motion (9), to calculate steady state values of u , m and u^m . After solving the above equation for v^m , we obtain $q(\theta) = \frac{m}{v} = \frac{m}{v^m + m} = .8541$.

Obsolescence rate. Because of the free entry condition $V_t = 0$, the equilibrium conditions for the standard model are influenced only by the combined separation rates (14) and (15). Thus, we do not need to separately measure λ^d for purposes of evaluating the standard model. Equivalently, job creation and destruction are not distinguishable from worker turnover in the standard model.

In the sunk cost model, in contrast, $V_t > 0$ will hold in equilibrium, meaning that firms have an incentive to refill job positions following quits. This introduces a distinction between job flows and worker flows, and makes it necessary to determine the job obsolescence rate. To accomplish this, we use job destruction data collected by

the Business Employment Dynamics (BED) survey. According to Faberman (2004), the quarterly job destruction rate in the private sector averaged around 8 percent over the period 1990 through 2003. We equate this measured job destruction rate to the corresponding magnitude in the model:

$$\lambda^d + (1 - \lambda^d)(1 - q(\theta)) \left[\frac{m}{1 - u + m} \lambda^n + \frac{1 - u}{1 - u + m} \lambda^q \right] = .08. \quad (16)$$

The first term on the left-hand side measures job obsolescence, while the second term reflects jobs that become vacant as a result of quits, and remain vacant at the end of the period. Since job destruction in the BED is computed from employment changes over a quarter, both sources of employment change must be included. Using the previously determined values of $q(\theta)$, u , and m , equations (14), (15), and (16) may be solved for the values $\lambda^d = .063$, $\lambda^n = .382$ and $\lambda^q = .069$.¹⁸

Other parameters. Having determined the steady-state values of u , v and m , the efficiency parameter of the matching function l is obtained using the matching function (12), yielding $l = 2.413$. We adopt the standard values $\beta = .99$ and $\pi = .50$ for the discount factor and worker bargaining weight. Productivity is normalized by setting $z = 1$, and the parameters ρ and σ_ε are selected to match the dynamics generated by our Section 2 estimates.

For the standard model, we make use of the value function equations (4)-(7) in the steady state, together with the free entry condition $V = 0$, to specify the remaining parameters. The flow posting cost c is determined by (8). The flow capital cost κ is chosen so that the per-period wage payment to the worker matches the measured labor share of .65:

$$[\lambda^d + (1 - \lambda^d)\lambda^q]\pi S + (1 - \beta)U = .65. \quad (17)$$

The unemployment benefit b is selected so that $b = .65(z - \kappa)$, i.e., the unemployment benefit is 65 percent of net match output. This lies between the values of 40 percent and 94.3 percent suggested by Shimer (2005) and Hagedorn and Manovskii (2005),

¹⁸In view of the free entry condition $V_t = 0$, the equilibrium conditions for the standard model are influenced only by the combined separation rates $\lambda^d + (1 - \lambda^d)\lambda^n$ and $\lambda^d + (1 - \lambda^d)\lambda^q$. Thus, we do not need to separately measure λ^d for purposes of evaluating the standard model.

respectively, and close to the value of 75 percent proposed by Costain and Reiter (2005). This procedure yields the values $c = .193$, $\kappa = .293$ and $b = .460$.

For the sunk cost model, we substitute the numerical values for λ^d , λ^q , u , v and $q(\theta)$ into (11) to calculate the steady state value for new openings $e = .058$. Furthermore, for a given value of c , the parameters κ and b are chosen to satisfy (17) and $b = .65(z - \kappa)$, while \bar{K} is chosen to satisfy the free entry condition $V = \bar{K}e$. Although there is little direct evidence as to the appropriate level of c , the responsiveness of e_t to changes in z_t is highly sensitive to c . Thus we may use second moment properties of the model to pin down this parameter. From the VAR model (1), the estimated standard deviation of $\ln vu$, conditional on $\ln \hat{z}_t$, is .299. We choose c so that the standard deviation of the log vacancy-unemployment ratio in the simulated data, based on end-of-period stocks $v_t^m = (1 - q(\theta_t))v_t$ and $u_t^m = (1 - p(\theta_t))u_t$, lies as close as possible to the empirical value. This yields the values $\kappa = .288$, $b = .463$, $\bar{K} = .682$ and $c = .183$. The implied standard deviation of the vacancy-unemployment ratio is .303. Table 2 summarizes the choices of parameter values for both models.

5 Results

The standard and sunk cost models are solved using the nonlinear global projection method called the collocation parameterized expectation algorithm (PEA); see Christiano and Fisher (2000) for a general discussion of this method. Details of the procedure are given in the Appendix. The second moment statistics of the model economies are based on 500 simulated samples. Each sample consists of 400 periods, where the first 200 observations are ignored to randomize initial conditions, and the last 200 observations are used to compute the statistics.

Recall that measured unemployment in the simulated data is based on the end-of-period stock u_t^m , rather than the worker matching pool u_t , since the latter includes workers who begin and end the period in job matches. Similarly, measured vacancies are based on v_t^m rather than v_t . The measured log vacancy-unemployment ratio, corresponding to the empirical variable $\ln vu_t$, is thus given by

$$\ln(v_t^m/u_t^m) = \ln \frac{(1 - q(\theta_t))v_t}{(1 - p(\theta_t))u_t}.$$

We measure log employment as $\ln n_t^m$, where $n_t^m = 1 - u_t^m$; this corresponds to $\ln emp_t$.

5.1 Empirical evaluation

Dynamic correlations. Figure 5 presents the dynamic correlations of productivity with matching intensity and employment, calculated from simulated data, for the standard and sunk cost models, along with the empirical correlations originally reported in Figure 2. The productivity-matching intensity correlations generated by the standard model, shown in the upper panel, exhibit an unrealistically sharp peak at zero lag, and fail to capture the flatness of the empirical correlations between lags of zero and two quarters. Further, as seen in the lower panel, the model makes the counterfactual prediction of almost perfect contemporaneous correlation between productivity and employment, and the correlations at lags of 2 and 3 quarters are too low. These results clearly show that the standard model does not generate realistic dynamics.

The sunk cost model, in contrast, yields correlation profiles that essentially duplicate the patterns seen in the data: the simulated correlations are flat over the correct ranges of lags, with the correct phasing. Notably, the sunk cost model does not produce the sharp peaks at zero lag that are characteristic of the standard model. Quantitatively, the correlations of matching intensity with current and lagged productivity are somewhat too high in the sunk cost model, as may be seen in the upper panel of Figure 5. The model produces a remarkably close match with the productivity-employment correlations, however, as the lower panel shows.

Impulse responses. The deficiencies of the standard model can be further illustrated using impulse responses. Figure 6 overlays the model-based impulse responses for a one standard deviation shock to productivity with the VAR-based impulse responses reported in Figure 3. The top panel shows the dynamics of our calibrated productivity driving process in comparison with the estimated process. The middle panel indicates the response of matching intensity. On impact, the simulated variable $\ln(v_t^m/u_t^m)$ jumps upward by over twice as much as the estimated response. Following this, it returns monotonically to the steady state. Thus, the standard model significantly overstates the effect of the shock on impact, and entirely misses the subsequent

hump-shaped response pattern. Reflecting this behavior, the simulated variable $\ln n_t^m$ jumps upward by a large amount in the period of impact, and its subsequent upward movements are slight and short-lived. While the presence of matching frictions introduces some persistence into the employment response, the amount of added persistence is quantitatively tiny in the standard model.

Impulse responses for the sunk cost model, together with the empirical impulse responses, are graphed in Figure 7. As seen in the middle panel, the jump in $\ln(v_t^m/u_t^m)$ matches the empirical value closely, and the subsequent response also displays a realistic hump-shaped pattern: matching intensity rises for four quarters, then returns to the steady state. The third panel shows that the response of $\ln n_t^m$ continues to exhibit the counterfactual upward jump in the period of impact, but the size of the jump is about half that of the standard model.¹⁹ Moreover, the subsequent dynamics closely match the empirical hump shape. Overall, the sunk cost model does a much better job capturing the dynamic characteristics of the employment adjustment process.

5.2 Sources of propagation

Vacancy adjustment and new openings. The dynamic properties of the standard and sunk cost models are linked to the behavior of vacancies. This is illustrated in Figure 8, which considers the vacancy adjustments associated with the calibrated standard and sunk cost models. The upper panel depicts the impulse response of the measured vacancy stock v_t^m for a one-standard-deviation productivity shock, and the bottom panel shows the corresponding net changes in the vacancy stock. In this case, we express the variables as *level* deviations from their steady-state values, since we wish to decompose the net changes into separate gross outflows and inflows.

In the standard model, the vacancy stock spikes upward in the period of the shock, and then decreases monotonically in tandem with productivity. This means the initial large inflow of vacancies is immediately reversed by outflows, as the bottom panel

¹⁹Under our timing assumptions, the productivity shock induces a contemporaneous change in entry into the vacancy pool, and the job matches generated by these added vacancies are counted as higher employment at the end of the period. The measured jumps in both matching intensity and employment would be reduced by a measurement convention that incorporated some averaging of beginning- and end-of-period levels.

shows. Since there is no sunk cost for creating job positions, potential entrant firms can freely enter the vacancy pool to compete away rents following a shock. Consequently, the vacancy stock adjusts directly with productivity.

In the sunk cost model, the initial upward jump is much smaller, and vacancies continue to rise for several periods following the shock. The bottom panel of Figure 8 shows how the increases in vacancies gradually diminish, in contrast to the abrupt adjustment exhibited by the standard model. This smooth adjustment of the changes in vacancies underlies the persistent responses of the vacancy stock, matching intensity and employment exhibited by the sunk cost model.

Smooth adjustment arises because the cost of creating marginal job positions becomes greater as more positions are created. Entrant firms therefore spread out their new openings, causing vacancies to continue rising even as productivity declines toward the steady state. Moreover, the equilibrium value of a job position, whether filled or unfilled, is positive in the presence of sunk costs. Job positions thus become durable once they are created, meaning that entrant firms will not choose to leave the vacancy matching pool after they have entered it, either initially or following a quit. This contributes to sluggish vacancy adjustment. Similar reasoning applies with respect to negative productivity shocks: a lower volume of new openings reduces marginal creation costs, causing entrant firms to spread out their responses; and durability means that reductions in new openings have a more persistent effect on the vacancy stock.

The quantitative importance of new openings can be seen by decomposing the net changes in the measured vacancy stock into separate gross outflows and inflows, using (11):

$$\Delta v_t^m = \underbrace{-\lambda^d v_{t-1}}_{\text{obsolescence}} - \underbrace{\left[m_t - (\lambda^d + (1 - \lambda^d)\lambda^n) m_{t-1} \right]}_{\text{hires}} + \underbrace{(1 - \lambda^d)\lambda^q (1 - u_{t-1})}_{\text{repostings}} + \underbrace{e_t}_{\text{new openings}} .$$

Observe that the net changes comprise gross outflows due to obsolescence and hires, together with gross inflows due to repostings following quits and new openings. Figure 9 plots the impulse responses for these gross flows, which make up the impulse response for the sunk cost model shown in the bottom panel of Figure 8. The graph

clearly shows that vacancy adjustment is almost entirely driven by new openings and hires. Furthermore, the inflows from new openings lead the outflows from hires, as a consequence of the matching process. In the four quarters following the shock, new openings inflows exceed hiring outflows, accounting for the vacancy increases observed in Figure 8.

Empirical evidence. Although new openings play a crucial role in shaping vacancy stock behavior in the sunk cost model, the available vacancy data do not permit a direct empirical assessment of this role. JOLTS, however, does provide information about vacancy stocks, quits, layoffs and hires for 2001:Q1 to 2005:Q1, and this allows us to obtain an indirect measure of new openings via the following stock-flow relationship:

$$vac_t = (1 - \lambda^d)vac_{t-1} + quits_t - hires_t + e_t. \quad (18)$$

Our calibration of the model suggests that $\lambda^d = .063$ provides a reasonable estimate of the vacancy withdrawal rate. We can combine this figure with the JOLTS data to impute an estimate of e_t from (18).

Consider the first quarter of 2001, which we view as the most typical within the limited JOLTS sample. For this quarter, the ratio of cumulative hires to end-of-quarter vacancy stock is about 3.5. Our imputed inflow of new vacancies amounts to 1.5 times the end-of-quarter stock. Thus, new openings amount to nearly half of total hires within the quarter. Figure 10 plots indices of imputed new openings, vacancies, hires and quits, treating 2001:Q1 as the base period. The four series fluctuate by comparable amounts over the sample period, and, in particular, new openings exhibit significant variability. Moreover, new openings move strongly upward in 2003:Q2, leading the upward movement of vacancies by four quarters. The upward movements of hires and quits lag those of new openings and are less steep. Based on this limited evidence, it appears that new openings adjust sooner and by a greater magnitude in comparison with the other components.

5.3 Amplification of shocks

Recent literature has focused on the ability of the matching model to amplify productivity shocks in order to explain the volatility of the vacancy-unemployment ratio. This question may be addressed using our calibrations of the standard and sunk cost models. Table 3 presents the standard deviations of matching intensity, employment, vacancies and unemployment obtained from the detrended data, along with corresponding measurements from simulated data for the standard and sunk cost models. The standard deviation of the productivity process is 0.019 in all three cases.

Under the calibrated parameter values, the volatilities in the standard model exceed the empirical levels for all variables except unemployment. As for the sunk cost model, recall that our calibration procedure matches the standard deviation of $\ln(v_t^m/u_t^m)$ to that of $\ln vu_t$. The standard deviations of the four variables are all lower than those of the standard model, but employment and vacancies remain somewhat more volatile than in the data.

Shimer (2005), Hagedorn and Manovskii (2005) and Mortensen and Nagypal (2005) have stressed the importance of the parameter b in determining the amplification of productivity shocks. To assess the sensitivity of amplification to our choice of b , we reevaluate the standard model using Shimer's suggested value of $b = .40$, with the other parameters adjusted to maintain the calibration requirements.²⁰ Under the alternative parameter values, the volatility of matching intensity is reduced to 0.199 from 0.363. Broadly speaking, our results thus confirm that the matching model may generate insufficient volatility of the vacancy-unemployment ratio, but the amplification mechanism is not nearly so weak as suggested by Shimer (2005). This is evidently due to two factors. First, we allow for a fixed flow overhead cost κ that works like b to amplify productivity shocks; see Mortensen and Nagypal (2005). Second, including employer-to-employer flows in the worker matching pool reduces the stock of unmatched workers relative to the total volume of matching activity. This serves to raise the variability of the unmatched worker stock for a given level of matching intensity variability.²¹

²⁰In particular, κ and c have been changed to .245 and .352, respectively, while the other parameters remain the same.

²¹These two factors appear to be quite robust sources of amplification in the matching model. We are currently investigating their precise importance, particularly employer-to-employer flows, in

6 Conclusion

This paper has evaluated the cyclical dynamics of job matching intensity, as measured by the vacancy-unemployment ratio, and employment, where business cycles are driven by exogenous shocks to labor productivity. The two variables respond to shocks in similar ways, with employment responses lagging matching intensity responses by one quarter. Both variables display the “hump-shaped” dynamics that are commonly observed throughout the empirical business cycle literature.

We show that the standard matching model, as exemplified by Pissarides (2000, ch. 1), cannot account for the observed dynamic patterns. Because potential entrant firms are able to respond easily to shocks, the vacancy pool adjusts in tandem with productivity. As a consequence, most of the adjustment of matching intensity and employment occurs in the period of the shock, leading to counterfactual dynamic correlations and impulse responses. Introducing a sunk cost for creating job positions spreads out the response of new openings to shocks, leading to realistic dynamic behavior. Our results suggest that sunk creation costs may play a central role in explaining cyclical adjustment. Moreover, the modified matching model can generate highly realistic employment dynamics without resort to any kind of consumption-smoothing mechanism.

The paper has relied on a number of simplifying assumptions that we believe are worth evaluating in future research. First, diseconomies in new job creation, associated with increasing marginal creation costs, could be considered in greater detail. These may arise from explicit costs of adjustment at the establishment or firm level, limited availability of key capital inputs, or technical constraints associated with R&D activity. Aggregate adjustment may be influenced by entry and exit of establishments. These factors may introduce important additional sources of propagation, including the possibility of longer-run feedbacks from the labor market to productivity. Relatedly, the assumed equivalence of newly-created and preexisting job positions can be modified by incorporating a vintage structure, whereby new jobs enjoy higher productivity. This would permit the endogenous obsolescence of jobs and the turnover of workers to be generating amplification, and we will include our findings in the next version of the paper.

considered as separate flows within a common framework.²² Finally, we have ignored the effects of cyclical variation in the relative sizes of the pools of unemployed workers, workers out of the labor force but available for work, and workers out of the labor force and unavailable. Changes in the characteristics of these pools may, however, represent another important source of longer-run propagation effects.

7 Appendix

7.1 Solution Method

Standard Model. For the standard model, the aggregate state of the economy for period t is a set of variables $\{m_{t-1}, u_{t-1}, z_t\}$.²³ We set the aggregate state space to $[m - 0.1, m + 0.1] \times [u - 0.1, u + 0.1] \times [\exp(-4\sqrt{\sigma_z^2/(1 - \rho^2)}), \exp(4\sqrt{\sigma_z^2/(1 - \rho^2)})]$, where m and u are the steady-state values of the number of new matches and unemployment. We parameterize the right-hand side of equations (4), (6) and (8) as a tensor product of second-order Chebyshev polynomials of the three state variables. Note that each function has 27 ($= 3^3$) unknowns, and thus there is a total of 81 ($= 27 \times 3$) unknowns to be determined.

Consider starting at an arbitrary grid point in the state space. For an initial guess of the unknown parameters of the approximating functions, we use the unemployment law of motion (9) to obtain the next period unemployment u_t . Using the approximating function for the right-hand side of the free entry condition (8) and the initial guess of its unknown parameters, we can obtain v_t . The u_t and v_t allow us to determine the number of matches formed m_t . Once we obtain the next period values of the state variables,

²²Aghion and Howitt (1994), Caballero and Hammour (1994) and Mortensen and Pissarides (1998), for example, analyze endogenous obsolescence in models that combine embodied technological progress with search/matching frictions. None of those papers distinguish between worker and job turnover. In recent work, Hornstein, Krusell, and Violante (2004) adopt a specification similar to ours for purposes of analyzing the unemployment experiences of the U.S. and Europe. They focus on comparison of steady states, however, rather than cyclical adjustment.

²³Note that m_{t-1} is in the set of period- t state variables because new matches and preexisting matches are subject to different separation hazard rates. In the sunk cost model below, v_{t-1} is also a state variable, and therefore we do not need to treat m_{t-1} as a separate state variable.

we compute the conditional expectations appearing in the value functions from the distribution of the productivity shock ϵ_t . The conditional expectations associated with the productivity shock are computed via the Gauss-Hermite quadrature with 5 nodes.

The conditional expectations for each value function are evaluated at 27 grid points that are chosen by finding zeros of Chebyshev polynomials of each of the three state variables, and taking all possible combinations of the zeros. The new set of coefficients of the approximating functions are obtained by equating the right-hand side of equations (4), (6) and (8) to the values of the approximating functions at the 27 grid points. Since there are 27 coefficients in each approximating function, this uniquely pins down the new set of coefficients. This process is iterated until convergence of the 81 Chebyshev coefficients is achieved. The convergence criterion is set to 10^{-8} . Finally, we simulate a long time series (200,000 observations) using the obtained solution functions in order to check that the economy remains within the specified state space.

Sunk Cost Model. We solve the sunk cost model in a similar way. The period- t state variables in this economy consist of $\{u_{t-1}, v_{t-1}, z_t\}$. The state space is defined as $[v - 0.1, v + 0.1] \times [u - 0.1, u + 0.1] \times [\exp(-4\sqrt{\sigma_z^2/(1-\rho^2)}), \exp(4\sqrt{\sigma_z^2/(1-\rho^2)})]$, where v and u are the steady-state values of vacancies and unemployment. This time we parameterize the right-hand side of equations (4), (5) and (6), again as a tensor product of second-order Chebyshev polynomials of the state variables. Because there are three state variables and three functions to be parameterized, there is the same number of unknowns ($= 81$).

Using the initial guess of the set of unknown parameters for the parameterized equation (5), the entry conditions (10) and (13) pin down the entry level e_t at each grid point. Using the law of motion for vacancies (11), we can then obtain the next period value of vacancies v_t corresponding to the initial grid point. The next period unemployment u_t is also obtained from equation (9). Given the distribution of the productivity shock ϵ_t , we can compute the conditional expectation of the right-hand side of equations (4), (5) and (6). Each of the three conditional expectations is evaluated at 27 grid points, and the new set of coefficients is obtained by equating the right-hand side of equations (4), (5) and (6) to the values of the approximating functions at those grid points. The iteration process continues until convergence of the coefficients is

achieved.

References

- Katherine Abraham. Help-wanted advertising, job vacancies, and unemployment. *Brookings Papers on Economic Activity*, 82(1):207–43, June 1987.
- Philippe Aghion and Peter Howitt. Growth and unemployment. *Review of Economic Studies*, 61:477–494, 1994.
- Patricia M. Anderson and Bruce D. Meyer. The extent and consequences of job turnover. *Brookings Papers on Economic Activity; Microeconomics*, 1994:177–244, 1994.
- Fabiano Bastos. Persistent employment fluctuations and the structure of search models: A frequency-domain perspective. November 2004. Draft, University of Maryland.
- Olivier Blanchard and Peter Diamond. The Beveridge curve. *Brookings Papers on Economic Activity*, 1:1–60, 1989.
- Olivier Blanchard and Peter Diamond. The cyclical behavior of the gross flows of U.S. workers. *Brookings Papers on Economic Activity*, 2:85–155, 1990.
- Hoyt Bleakley and Jeffrey Fuhrer. Shifts in the Beveridge curve, job matching and labor market dynamics. *New England Economic Review*, pages 3–19, Sept./Oct. 1997.
- Ricardo Caballero and Mohamad Hammour. The cleansing effect of recessions. *American Economic Review*, 85(5):1350–1368, December 1994.
- Ricardo Caballero and Mohamad Hammour. On the timing and efficiency of creative destruction. *Quarterly Journal of Economics*, 1996.
- Monica D. Castillo. Persons outside the labor force who want a job. *Monthly Labor Review*, 121(7):34–42, 1998.
- Lawrence Christiano and Jonas Fisher. Algorithms for solving dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control*, 24(8):1179–1232, July 2000.

- Harold Cole and Richard Rogerson. Can the Mortensen-Pissarides matching model match the business-cycle facts? *International Economic Review*, 40(4):933–959, November 1999.
- James Costain and Michael Reiter. Business cycles, unemployment insurance, and the calibration of matching models. July 2005. Draft.
- Steven Davis, John Haltiwanger, and Scott Schuh. *Job Creation and Destruction*. MIT Press, 1996.
- Reinout De Bock. Embodied technical change and the persistence of vacancies. February 2005. Draft, Northwestern University.
- Wouter den Haan, Garey Ramey, and Joel Watson. Job destruction and propagation of shocks. *American Economic Review*, 90(3):482–98, June 2000.
- Jason Faberman. Gross job flows over the past two business cycles: Not all ‘recoveries’ are created equal. *BLS Working Paper Series*, (372), June 2004.
- Bruce Fallick and Charles A. Fleischman. Employer-to-employer flows in the U.S. labor market: The complete picture of gross worker flows. *Finance and Economics Discussion Series 2004-34*, Board of Governors of the Federal Reserve System, 2004.
- Henry Farber. Mobility and stability: Dynamics of job change in labor markets. *Handbook of Labor Economics*, 1999. Edited by Orley Ashenfelter and David Card.
- Shigeru Fujita. The Beveridge curve, job creation and the propagation of shocks. 2003. Draft, UC San Diego.
- Shigeru Fujita. Vacancy persistence. *Federal Reserve Bank of Philadelphia Working Paper*, 04(23), 2004.
- Shigeru Fujita and Garey Ramey. Costly implementation of technical progress and business cycles. 2004. Draft, UC San Diego.
- Marcus Hagedorn and Iourii Manovskii. The cyclical behavior of equilibrium unemployment and vacancies revisited. April 2005. Draft, University of Frankfurt and University of Pennsylvania.

- Robert E. Hall. Employment fluctuations with equilibrium wage stickiness. *American Economic Review*, 95(1):25–49, March 2005.
- Andreas Hornstein, Per Krusell, and Giovanni Violante. Vintage capital in frictional labor markets. May 2004. Draft.
- Richard Jackman, R Layard, and Christopher Pissarides. On vacancies. *Oxford Bulletin of Economics and Statistics*, 51(4):377–394, 1989.
- Michael Krause and Thomas Lubik. On-the-job search and the cyclical dynamics of the labor market. September 2004. Draft, Tilberg University and Johns Hopkins University.
- Guido Menzio. High frequency wage rigidity. November 2004. Draft, Northwestern University.
- Dale Mortensen. The cyclical behavior of job and worker flows. *Journal of Economic Dynamics and Control*, 18:1121–1142, 1994.
- Dale Mortensen and Eva Nagypal. More on unemployment and vacancy fluctuations. 2005. Draft, Northwestern University.
- Dale Mortensen and Christopher Pissarides. Technological progress, job creation, and job destruction. *Review of Economic Dynamics*, 1:733–753, 1998.
- Barbara Petrongolo and Christopher Pissarides. Looking into the black box: A survey of the matching function. *Journal of Economic Literature*, 39:390–431, June 2001.
- Christopher Pissarides. *Equilibrium Unemployment Theory*. MIT Press, second edition, 2000.
- Michael Pries. Persistence of employment fluctuations: A model of recurring job loss. *Review of Economic Studies*, 71(1):193–215, January 2004.
- Leena Rudanko. Labor market dynamics under long term wage contracting and incomplete markets. June 2005. Draft, University of Chicago.

Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49, March 2005.

José I. Silva and Manuel Toledo. Labor turnover costs and the behavior of vacancies and unemployment. June 2005. Draft, Universidad Autónoma de Barcelona and Central Bank of Venezuela.

Eran Yashiv. Forward-looking hiring behavior and the dynamics of the aggregate labor market. January 2005. Draft, Tel Aviv University.

Table 1: Labor market transition matrix

		End of Month				
		E^s	E^n	U	N	
Start of Month	$E^s \cup E^n$	59.0	1.60	0.85	1.70	63.15
	U		0.95	1.70	0.75	3.40
	N		1.60	0.85	31.00	33.45
		59.0	4.15	3.40	33.45	

Notes: Computed from Fallick and Fleischman, Table 2 (2001, p11). E^s : same employer, E^n : new employer, U : unemployed, N : out of labor force. Outflows from U in their table are too large to be consistent with steady state. To adjust for this, we subtracted 0.05 from $U \rightarrow E^n$ flows and $U \rightarrow N$ flows, and added 0.05 to $E^n \rightarrow U$ and $N \rightarrow U$ flows. Units are percentages of civilian noninstitutional population aged 16 and over.

Table 2: Parameter Values


symbol	description	calibrated value	
		standard	sunk cost
λ^d	rate of obsolescence	0.063	same
λ^n	separation rate for new matches	0.382	same
λ^q	separation rate for preexisting matches	0.069	same
l	efficiency parameter of the matching function	2.413	same
β	discount factor	0.99	same
π	bargaining weight of workers	0.5	same
z	steady-state value of labor productivity	1.0	same
ρ	autoregressive parameter of $\ln z_t$	0.93	same
σ_z	standard deviation of the productivity shock	0.007	same
c	flow vacancy posting cost	0.193	0.183
κ	Overhead cost	0.293	0.288
b	utility from leisure	0.460	0.463
\bar{K}	parameter of the creation cost function	n.a.	0.682

Table 3: Volatilities

data	$\ln v_t/u_t$	$\ln emp_t$	$\ln v_t$	$\ln u_t$
	0.299	0.0087	0.131	0.179
model	$\ln v_t^m/u_t^m$	$\ln n_t^m$	$\ln v_t^m$	$\ln u_t^m$
standard	0.363	0.0161	0.225	0.141
sunk cost	0.303	0.0134	0.181	0.119



Notes: Volatilities of the empirical data are conditional standard deviations.

Figure 1: Beveridge curve



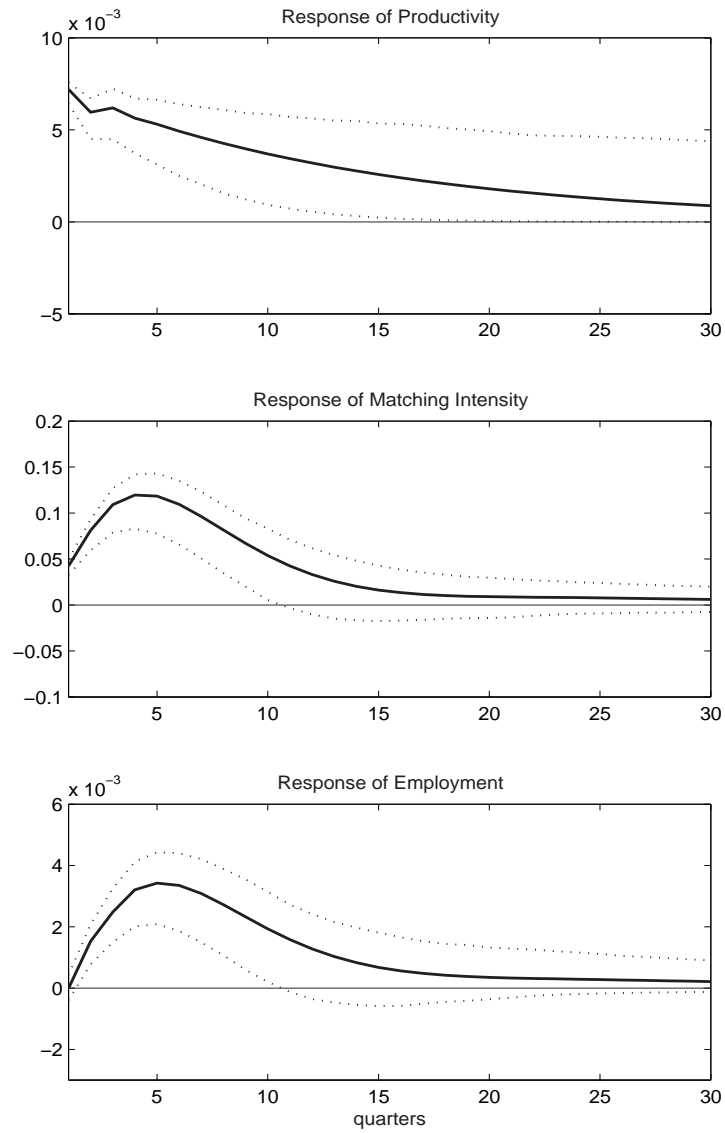
Notes: Quarterly averages of seasonally adjusted monthly data. The vacancy rate is the number of newspaper help-wanted ads divided by the sum of employment and help-wanted ads. Sample period: 1951Q1-2004Q3.

Figure 2: Empirical dynamic correlations



Notes: Plotted are conditional correlations.

Figure 3: VAR-based impulse responses to one-s.d. productivity shock



Notes: Dotted lines are 90% confidence bands computed via Monte-Carlo simulations with 1,000 replications.

Figure 4: Responses of vacancies and unemployment to one-s.d. productivity shock




Figure 5: Comparison of dynamic correlations


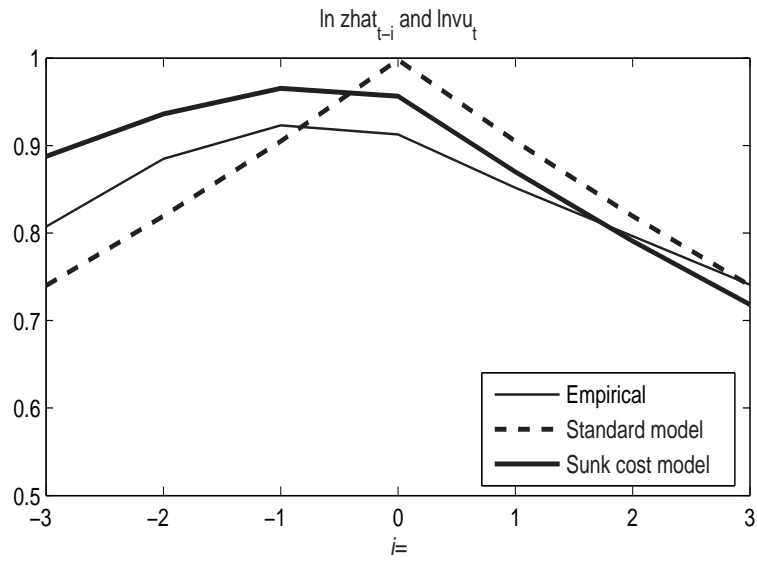


Figure 6: Comparison of the impulse responses: standard model




Figure 7: Comparison of the impulse responses: sunk cost model

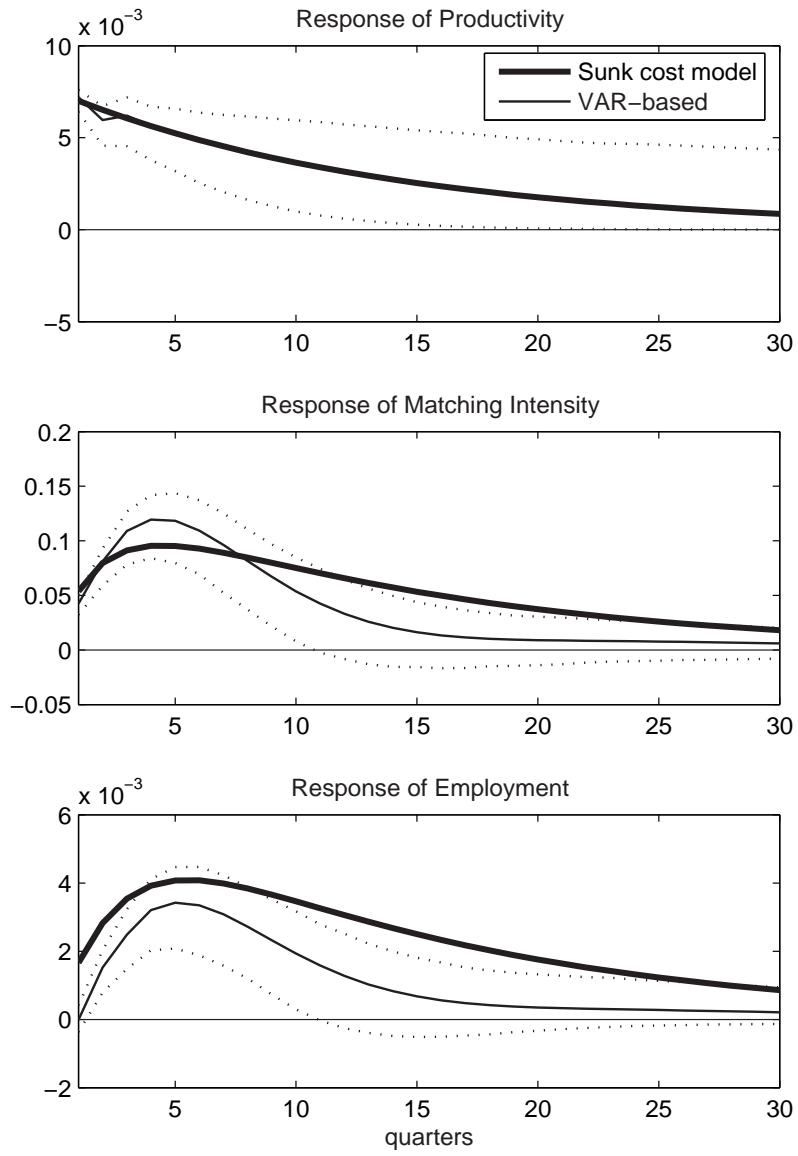
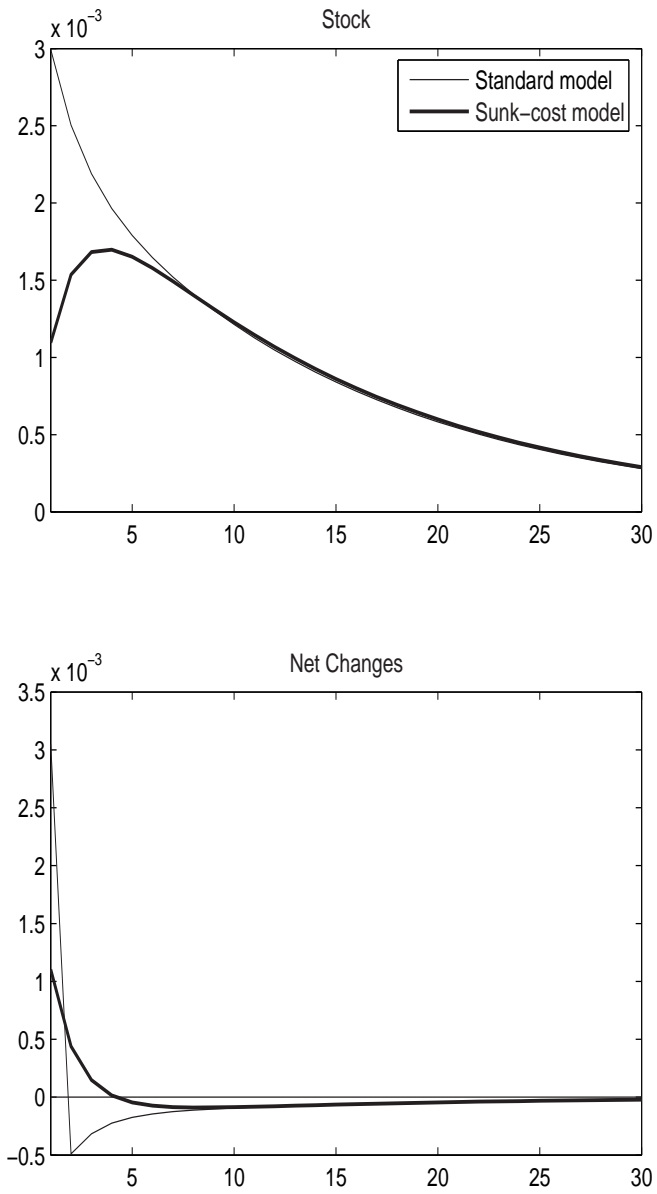



Figure 8: Vacancy responses



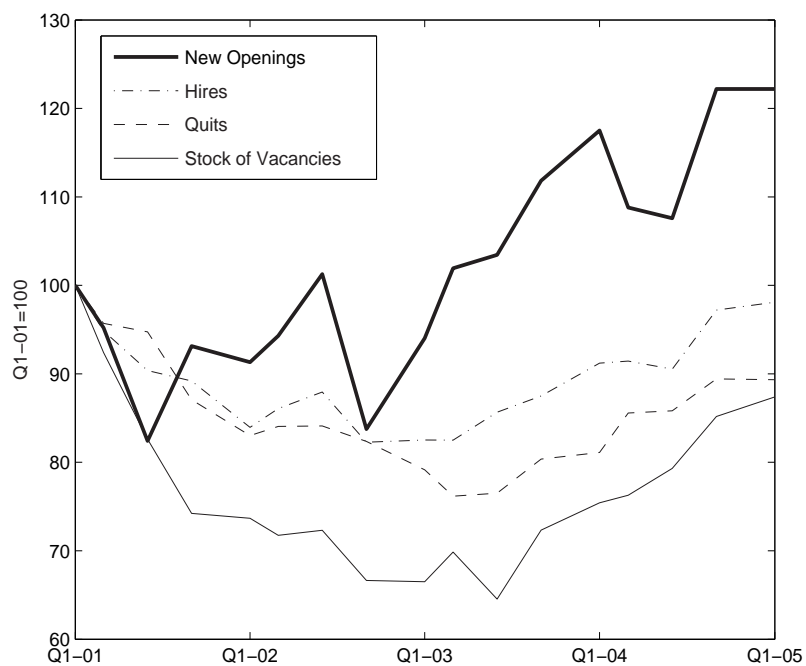
Notes: Level deviations from the steady-state values.

Figure 9: Gross flows of vacancies in the sunk cost model



Notes: Level deviations from the steady-state values. Outflows are plotted as negative values.

Figure 10: JOLTS data



Notes: New postings are imputed from the vacancy stock-flow relationship (18) using the JOLTS data on quits, hires and end-of-the-period stock of vacancies, and the calibrated value of the withdrawal rate λ^d . The above figure plots the indices that treat 2001:Q1 as the base period.