The Cyclical Behavior of Equilibrium Unemployment and Vacancies: Evidence and Theory*

Robert Shimer
Department of Economics
University of Chicago
shimer@uchicago.edu

March 31, 2004

Abstract

This paper argues that the textbook search and matching model cannot generate the observed business-cycle-frequency fluctuations in unemployment and job vacancies in response to shocks of a plausible magnitude. In the U.S., the standard deviation of the vacancy-unemployment ratio is 18 times as large as the standard deviation of average labor productivity, while the search model predicts that the two variables should have nearly the same volatility. A shock that changes average labor productivity primarily alters the present value of wages, generating only a small movement along a downward sloping Beveridge curve (unemployment-vacancy locus). A shock to the separation rate generates a counterfactually positive correlation between unemployment and vacancies. In both cases, the shock is only slightly amplified and the model exhibits virtually no propagation.

*A previous version of this paper was entitled 'Equilibrium Unemployment Fluctuations'. I thank Daron Acemoglu, Robert Barro, Olivier Blanchard, V. V. Chari, Joao Gomes, Robert Hall, Dale Mortensen, Christopher Pissarides, two anonymous referees, the editor Richard Rogerson, and numerous seminar participants for comments that are incorporated throughout the current draft. This material is based upon work supported by the National Science Foundation under grants SES-0079345 and SES-0351352. I am grateful to the Alfred P. Sloan Foundation for financial support, to the Federal Reserve Bank of Minneapolis for its hospitality while I revised a previous version of this paper, and to Mihai Manea and especially Sebastian Ludmer for excellent research assistance.

1 Introduction

In recent years, the Mortensen-Pissarides search and matching model has become the standard theory of equilibrium unemployment (Mortensen and Pissarides 1994, Pissarides 2000). The model is attractive for a number of reasons: it offers an appealing description of how the labor market functions; it is analytically tractable; it has rich and generally intuitive comparative statics; and it can easily be adapted to study a number of labor market policy issues, such as unemployment insurance, firing restrictions, and mandatory advanced notification of layoffs. Given these successes, one might expect that there would be strong evidence that the model is consistent with key business cycle facts. On the contrary, I argue in this paper that the model cannot explain the cyclical behavior of two of its central elements, unemployment and vacancies. In U.S. data, both are highly variable and strongly negatively correlated. Equivalently, the model cannot explain the strong procyclicality of the probability that an unemployed worker finds a job, the hiring rate.

I focus on two sources of shocks, changes in labor productivity and changes in the separation rate. In a one sector model, a change in labor productivity is most easily interpreted as a technology or supply shock. But in a multi-sector model, a preference or demand shock changes the relative price of goods, which induces a change in real labor productivity as well. Thus these shocks represent a broad set of possible impulses.

An increase in labor productivity relative to the value of non-market activity and to the cost of advertising a job vacancy makes unemployment relatively expensive and vacancies relatively cheap.[1] The market substitutes towards vacancies, and the increased hiring pulls down the unemployment rate, moving the economy along a downward sloping Beveridge curve (vacancy-unemployment locus). But the increase in hiring shortens unemployment duration, raising workers' threat point in wage bargaining, and therefore raising the present value of wages in new jobs. Higher wages absorb most of the productivity increase, eliminating the incentive for vacancy creation. As a result, fluctuations in labor productivity have little impact on the unemployment, vacancy, and hiring rates.

An increase in the separation rate does not affect the relative value of unemployment and vacancies, and so leaves the vacancy-unemployment (v-u) ratio essentially unchanged. Since the increase in separations reduces employment duration, the unemployment rate increases, and so therefore must vacancies. As a result, fluctuations in the separation rate induce a counterfactually positive correlation between unemploy-

[1] The interpretation in this paragraph and its sequel builds on discussions with Robert Hall.

ment and vacancies.

Section 2 presents the relevant business cycle facts: unemployment u is strongly countercyclical, vacancies v are equally strongly procyclical, and the correlation between the two variables is -0.88 at business cycle frequencies. As a result, the v-u ratio is procyclical and volatile, with a standard deviation around its trend equal to 0.38 log points. To provide further evidence in support of this finding, I examine the rate at which unemployed workers find jobs, the hiring rate. If the process of pairing workers with jobs is well-described by an increasing, constant returns to scale matching function $m(u, v)$, as in Pissarides (1985), the hiring rate $m(u, v)/u$ should be an increasing function of the v-u ratio. I use unemployment duration data to measure the hiring rate directly, providing evidence that this is the case. The standard deviation of fluctuations in the hiring rate around trend is 0.18 log points and the correlation with the v-u ratio is 0.88. Finally I look at the two proposed impulses. The separation rate is only weakly countercyclical and moderately volatile, with a standard deviation about trend equal to 0.11 log points. Average labor productivity is weakly procyclical and even more stable, with a standard deviation about trend of 0.02 log points.

In Section 3, I extend the Pissarides (1985) search and matching model to allow for aggregate fluctuations. I introduce two types of shocks: labor productivity shocks raise output in all matches but do not affect the rate at which employed workers lose their job; and separation shocks raise the rate at which employed workers become unemployed but do not affect the productivity in surviving matches. In equilibrium, there is only real economic decision, firms' calculation of whether to open a new vacancy. The equilibrium vacancy rate depends on the unemployment rate, on labor market tightness, and on the expected present value of wages in new employment relationships. Wages, in turn, are determined by Nash bargaining, at least in new matches. In principle, the wage in old matches may re-bargained in the face of aggregate shocks or may fixed by a long-term employment contract. Section 3.1 describes the basic model, while Section 3.2 derives a forward-looking equation for the v-u ratio in terms of model parameters.

Section 3.3 performs simple comparative statics in some special cases. For example, if there are no aggregate shocks, the v-u ratio is an implicit function of current and future labor market conditions and it is possible to make some simple computations analytically. I show that the elasticity of the v-u ratio with respect to the difference between labor productivity and the value of non-market activity or 'leisure' is barely in excess of 1 for reasonable parameter values. To reconcile this with the data, one must assume that the value of leisure is nearly equal to labor productivity, so market work provides little utility. The separation rate has an even smaller impact on the v-u ratio, with an elasticity of -0.09 according to the comparative statics. Moreover, while shocks to labor productivity at least induce a negative correlation between unemployment and vacancies, separation shocks cause both variables to increase, which tends to generate a positive correlation between the two variables. Similar results obtain in some other special cases.

Section 3.4 calibrates the model to match the data along as many dimensions as possible and Section 3.5 presents the results. The exercise confirms the quantitative predictions of the comparative statics. If the economy is hit only by productivity shocks, it moves along a downward sloping Beveridge curve, but empirically plausible movements in labor productivity result in tiny fluctuations in the v-u ratio. Moreover, labor productivity is perfectly correlated with the v-u ratio, indicating that the model has almost no internal propagation mechanism. If the economy is hit only by separation shocks, the v-u ratio is stable in the face of large unemployment fluctuations, so vacancies are countercyclical. Equivalently, the model-generated Beveridge curve is upward-sloping.

Section 3.6 explores the extent to which the Nash bargaining solution is responsible for these results. First I examine the behavior of wages in the face of labor productivity and separation shocks. An increase in labor productivity encourages firms to create vacancies. The resulting increase in the hiring rate puts upward pressure on wages, soaking up virtually of the shock. A decrease in the separation rate also induces firms to create more vacancies, again putting upward pressure on wages and minimizing the impact on the v-u ratio and hiring rate. On the other hand, I examine a version of the model in which workers' bargaining power is stochastic. Small fluctuations in bargaining power induce induce realistic movements in the v-u ratio while inducing only a moderately countercyclical real wage, with standard deviation of 0.01 log points around trend.

Section 4 provides another angle from which to view the model's basic shortcoming. I consider a centralized economy in which a social planner decides how many vacancies to create in order to maximize the present value of market and non-market income net of vacancy creation costs. The decentralized and centralized economies behave identically if the matching function is Cobb-Douglas in unemployment and vacancies, a generalization of Hosios (1990). But if unemployment and vacancies are more substitutable, fluctuations are amplified in the centralized economy, essentially because the shadow wage is less procyclical. Empirically it is difficult to measure the substitutability of unemployment and vacancies in the matching function, and so hard to tell whether observed fluctuations are optimal.

Section 5 reconciles this paper with a number of existing studies that claim standard search and matching models are consistent with the business cycle behavior of labor markets. Finally, the paper concludes in Section 6 by suggesting some modifications to the model that might deliver rigid wages and thereby do a better job of matching the empirical evidence on vacancies and unemployment.

It is worth emphasizing one important feature of the Pissarides (1985) model which I use throughout this paper: workers are risk-neutral and supply labor inelastically. In the absence of search frictions, employment would be constant even in the face of productivity shocks. This distinguishes the present model from those based on intertemporal labor supply decisions (Lucas and Rapping 1969). Thus I am asking about the extent to which the combination of search frictions and aggregate shocks can generate plausible fluctuations in unemployment and vacancies. Whether a model with an elastic labor supply can provide a satisfactory explanation of the observed fluctuations in these two variables remains an open question.[2]

2 U.S. Labor Market Facts

This section discusses the time series behavior of unemployment u, vacancies v, hiring rates h, separation rates s, and labor productivity p in the United States. Table 1 summarizes the detrended data.

2.1 Unemployment

The unemployment rate is the most commonly used cyclical indicator of job search activity. In an average month from 1951 to 2003, 5.67 percent of the U.S. labor force was out of work, available for work, and actively seeking work. This time series exhibits considerable temporal variation, falling as low as 2.6 percent in 1953 and 3.4 percent in 1968 and 1969, but reaching 10.8 percent in 1982 and 1983 (Figure 1). Some of these fluctuations are almost certainly due to demographic and other factors unrelated to business cycles. To highlight business-cycle-frequency fluctuations, I take the log deviation of the level of unemployment from an extremely low frequency trend, a Hodrick-Prescott (HP) filter with smoothing parameter 10^5 using quarterly data. The log ratio of unemployment to its trend has a standard deviation of 0.19, so unemploy-

[2] There are reasons to be concerned. If unemployed workers do not particularly want jobs during recessions, they will tend to drop out of the labor force. Veracierto (2002) emphasizes this possibility.

ment is often as much as 38 percent above or below trend. Detrended unemployment also exhibits considerable persistence, with quarterly autocorrelation 0.94.

There is some question as to whether unemployment or the employment-population ratio is a better indicator of job search activity. Advocates of the latter view, for example Cole and Rogerson (1999), argue that the number of workers moving directly into employment from out-of-the-labor force is as large as the number who move from unemployment to employment (Blanchard and Diamond 1990). On the other hand, there is ample evidence that unemployment and nonparticipation are distinct economic conditions. Juhn, Murphy, and Topel (1991) show that almost all of the cyclical volatility in prime-aged male nonemployment is accounted for by unemployment. Flinn and Heckman (1983) show that unemployed workers are significantly more likely to find a job than nonparticipants, although Jones and Riddell (1999) argue that other variables also help to predict the likelihood of finding a job. In any case, since labor market participation is procyclical, the employment-population ratio is a more cyclical measure of job search activity, worsening the problems highlighted in this paper.

It is also conceivable that when unemployment rises, the amount of job search activity per unemployed worker declines so much that aggregate search activity actually falls. There is both direct and indirect evidence against this hypothesis. As direct evidence, one would expect that a reduction in search intensity could be observed as a decline in the number of job search methods used or a switch towards towards less time-intensive methods. An examination of Current Population Survey (CPS) data indicates no cyclical variation in the number or type of job search methods utilized. Indirect evidence comes from estimates of matching functions, which universally find that an increase in unemployment is associated with an increase in the number of matches (Petrongolo and Pissarides 2001). If job search activity declined sharply when unemployment increased, the matching function would be measured as decreasing in unemployment. I conclude that aggregate job search activity is positively correlated with unemployment.

2.2 Vacancies

The flip side of unemployment is job vacancies. The Job Openings and Labor Turnover Survey (JOLTS) provides an ideal empirical definition: "A job opening requires that 1) a specific position exists, 2) work could start within 30 days, and 3) the employer is actively recruiting from outside of the establishment to fill the position. Included are full-time, part-time, permanent, temporary, and short-term openings. Active recruiting means that the establishment is engaged in current efforts to fill the opening, such as advertising in newspapers or on the Internet, posting help-wanted signs, accepting applications, or using similar methods."[3] Unfortunately, JOLTS only began in December 2000 and comparable data had never previously been collected in the U.S.. Although there are too few observations to look systematically at this time series, its behavior has been instructive. In the first month of the survey, the non-farm sector maintained a seasonally adjusted 4.66 million job openings. This number fell rapidly during 2001, and averaged just 2.91 million in 2002 and 2003. This decline in job openings during a period with high unemployment, depicted in Figure 2, suggests that job vacancies are procyclical.

To obtain a longer time series, I use a standard proxy for vacancies, the Conference Board help-wanted advertising index, measured as the number of help-wanted advertisements in 51 major newspapers.[4] A potential shortcoming is that help-wanted advertising is subject to low frequency fluctuations that are only tangentially related to the labor market: the Internet may have reduced firms' reliance on newspapers as a source of job advertising; newspaper consolidation may have increased advertising in surviving newspapers; and Equal Employment Opportunity laws may have encouraged firms to advertise job openings more extensively. Fortunately, a low frequency trend should remove the effect of these and other secular shifts. Figure 3 shows the help wanted advertising index and its trend. Notably, the decline in the de-trended help-wanted index closely tracks the decline in job openings measured directly from JOLTS during the period when the latter time series is available (Figure 2).

Figure 4 shows a scatter plot of the relationship between the cyclical component of unemployment and vacancies, the 'Beveridge curve'. The correlation of the percentage deviation of unemployment and vacancies from trend is -0.89 between 1951 and 2003.[5] Moreover, the standard deviation of the cyclical variation in unemployment and vacancies is almost identical, between 0.19 and 0.20, so the product of unemployment and vacancies is nearly acyclical. The v-u ratio is therefore extremely procyclical, with a standard deviation of 0.38 around its trend.

[3] This definition comes from the Bureau of Labor Statistics news release, July 30, 2002, available at `http://www.bls.gov/jlt/jlt_nr1.pdf`.

[4] Abraham (1987) discusses this measure in detail. From 1972 to 1981, Minnesota collected state-wide job vacancy data. Abraham (1987) compares this with Minnesota's help-wanted advertising index and shows that the two series track each other very closely through two business cycles and ten seasonal cycles.

[5] Abraham and Katz (1986) and Blanchard and Diamond (1989) discuss the U.S. Beveridge curve. Abraham and Katz (1986) argue that the negative correlation between unemployment and vacancies is inconsistent with Lilien's (1982) sectoral shifts hypothesis, and instead indicates that business cycles are driven by aggregate fluctuations. Blanchard and Diamond (1989) conclude that at business cycle frequencies, shocks generally drive the unemployment and vacancy rates in the opposite direction.

2.3 Hiring Rate

An implication of the procyclicality of the v-u ratio is that the hazard rate for an unemployed worker to find a job, his hiring rate, should be lower during a recession. Assume that the number of newly hired workers is given by an increasing and constant returns to scale matching function $m(u,v)$, depending on the number of unemployed workers u and the number of vacancies v. Then the probability that any individual unemployed worker finds a job, the average transition rate from unemployment to employment, is $h \equiv \frac{m(u,v)}{u} = m(1,\theta)$, where $\theta \equiv v/u$ is the vacancy-unemployment ratio. The hiring rate h should therefore move together with the v-u ratio.

Gross worker flow data can be used to measure the hiring rate directly, and indeed both the unemployment to employment and nonparticipation to employment transition rates are strongly procyclical (Blanchard and Diamond 1990, Bleakley, Ferris, and Fuhrer 1999, Abraham and Shimer 2001). There are two drawbacks to this approach. First, the requisite public use data set is only available since 1976, and so using this data would require throwing away half of the available data. Second, measurement and classification error lead a substantial overestimate of gross worker flows (Abowd and Zellner 1985, Poterba and Summers 1986), the magnitude of which cannot easily be computed. Instead, I infer the hiring rate from the dynamic behavior of the unemployment level and average unemployment duration. Let d_t denote mean unemployment duration measured in months. Then assuming all unemployed workers find a job with probability h_t in month t and no unemployed worker exits the labor force,

$$d_{t+1} = \frac{(1+d_t)(1-h_t)u_t + \big(u_{t+1} - (1-h_t)u_t\big)}{u_{t+1}}.$$

The numerator is the number of unemployed workers in period t who fail to find a job times the mean unemployment duration of those workers, $1 + d_t$, plus the number of newly unemployed workers in period $t+1$, each of whom has an unemployment duration of 1 month. This is divided through by the number of unemployed workers in month $t+1$ to get mean unemployment duration in that month. Equivalently,

$$h_t = 1 - \frac{(d_{t+1} - 1)u_{t+1}}{d_t u_t}. \tag{1}$$

In steady state, $u_t = u_{t+1}$ and $d_t = d_{t+1}$, so the right hand side reduces to the inverse of unemployment duration, a familiar relationship. If unemployment were constant between months t and $t+1$, the hiring rate could be expressed as a function of unemployment duration alone, $1 - \frac{d_{t+1}-1}{d_t}$. The correlation between this time series and h_t

constructed using equation (1) is 0.97, so time variation in the number of unemployed workers accounts for almost none of the time variation in the hiring rate. More generally, I use mean unemployment duration and the number of unemployed workers, both constructed by the BLS from the CPS, to compute h_t from 1951 to 2003.[6] Figure 5 shows the results. The monthly hazard rate averaged 0.343 from 1951 to 2003. After detrending with the usual low-frequency HP filter, the correlation between h_t and θ_t at quarterly frequencies is 0.88, although h_t is about half as variable as θ_t. Given that both measures are crudely yet independently constructed, this correlation is remarkable and strongly suggests that a matching function is a useful way to approach U.S. data.

One can in fact use the detrended data to estimate a matching function $m(u, v)$. Data limitations force me to impose two restrictions on the estimated function. First, because unemployment and vacancies are strongly negatively correlated, it is difficult to tell empirically whether $m(u, v)$ exhibits constant, increasing, or decreasing returns to scale. But in their literature survey, Petrongolo and Pissarides (2001) conclude that most estimates of the matching function cannot reject the null hypothesis of constant returns; I therefore estimate $h_t = h(\theta_t)$, consistent with a constant returns to scale matching function. Figure 6 shows the raw data for the hiring rate h_t and the v-u ratio θ_t, a nearly linear relationship when both variables are expressed as log deviations from trend. Second, I impose that the matching function is Cobb-Douglas, so $\log h_t = \log \mu + (1 - \alpha) \log \theta_t$ for some unknown parameters α and μ. Again, the data are uninformative as to whether this is a reasonable restriction.[7] I estimate the matching function using data on the log deviation from trend for the hiring rate and the v-u ratio. Depending on exactly how I control for autocorrelation in the residuals, I estimate values of α between 0.60 and 0.65. For example, if the residual is an AR(1), which appears to be a reasonable approximation, I get $\alpha = 0.618$ with a standard error of 0.025. Allowing for higher order autocorrelation raises α to 0.65.

2.4 Separation Rate

I can also deduce the behavior of the separation rate from data on employment, unemployment, and unemployment duration. The number of unemployed workers next

[6]The BLS measures unemployment duration in weeks. I convert this to a monthly measure by multiplying each data point by 12/52.

[7]Consider the CES matching function $\log h_t = \log \mu + \frac{1}{\rho} \log \left(\alpha + (1 - \alpha)\theta_t^{\rho}\right)$. Cobb-Douglas corresponds to limiting case of $\rho = 0$. When I estimate the CES function using non-linear least squares and correct for first order autocorrelation, I get a point estimate of $\rho = 0.59$ with a standard error of 0.84.

month satisfies

$$u_{t+1} = (1 - h_t)u_t + s_t e_t,$$

the sum of the number of unemployed workers who are not hired and the number of employed workers e_t who separate in the current month. Solving this for s_t and eliminating h_t using (1) gives

$$s_t = \frac{(d_t - d_{t+1} + 1)u_{t+1}}{d_t e_t}. \tag{2}$$

The monthly separation rate thus constructed is shown in Figure 7. It averages 0.0198 from 1951 to 2003, so jobs last on average for about three years. Although fluctuations in the log deviation of the separation rate from trend are fairly large, with a standard deviation of 0.11, the correlation with other labor market variables, including unemployment, is relatively weak.

The strong procyclicality of the hiring rate and relatively weak procyclicality of the separation rate might appear to contradict Blanchard and Diamond's (1990) conclusion that "the amplitude in fluctuations in the flow out of employment is larger than that of the flow into employment." This is easily reconciled. Blanchard and Diamond look at the *number* of people entering or exiting employment in a given month, $h_t u_t$ or $s_t e_t$, while I focus on the *probability* that an individual switches employment states, h_t and s_t. Although the probability of entering employment h_t declines sharply in recessions, this is almost exactly offset by the increase in unemployment u_t, so that the number of people exiting unemployment is essentially acyclic. Viewed through the lens of an increasing matching function $m(u, v)$, this is consistent with the independent evidence that vacancies are strongly procyclical.

2.5 Labor Productivity

The final important empirical observation is the weak procyclicality of labor productivity, measured as real output per person in the non-farm business sector. The BLS constructs this quarterly time series as part of its Major Sector Productivity and Costs program. The output measure is based on the National Income and Product Accounts, while employment is constructed from the BLS establishment survey, the Current Employment Statistics. This series offers two advantages compared with total factor productivity: it is available quarterly since 1948; and it better corresponds to the concept of labor productivity in the subsequent models, which do not include capital.

Figure 8 shows the behavior of labor productivity and Figure 9 compares the cyclical components of the v-u ratio and labor productivity. There is a positive correlation between the two time series and some evidence that labor productivity leads the v-u ratio by about one year, with a maximum correlation of 0.57.[8] But the most important fact is that labor productivity is stable, never deviating by more than six percent from trend. In contrast, the v-u ratio has twice risen to .5 log points about its trend level and six times fallen by .5 log points below trend.

It is possible that the measured cyclicality of labor productivity is reduced by a composition bias, since less productive workers are more likely to lose their jobs in recessions. I offer two responses to this concern. First, there is a composition bias that points in the opposite direction: labor productivity is higher in more cyclical sectors of the economy, e.g. durable goods manufacturing. And second, a large literature on real wage cyclicality has reached a mixed conclusion about the importance of composition biases (Abraham and Haltiwanger 1995). Solon, Barsky, and Parker (1994) provide perhaps the strongest evidence that labor force composition is important for wage cyclicality, but even they argue that accounting for this might double the measured variability of real wages. This paper argues that the search and matching model cannot account for the cyclical behavior of vacancies and unemployment unless labor productivity is at least ten times as volatile as the data suggests, so composition bias is at best an incomplete explanation.

3 Search and Matching Model

I now examine whether search theory can reconcile the strong procyclicality of the vacancy-unemployment ratio and the hiring rate with the weak procyclicality of labor productivity and countercyclicality of the separation rate.. The model I consider is essentially a stochastic version of the Pissarides matching model with exogenous separations (Pissarides 1985, Pissarides 2000).

3.1 Model

I start by describing the exogenous variables that drive fluctuations. Labor productivity p and the separation rate s follow a first order Markov process in continuous time. A shock hits the economy according to a Poisson process with arrival rate σ, at which

[8]From 1951 to 1985, the contemporaneous correlation between detrended labor productivity and the v-u ratio was 0.57 and the peak correlation was 0.74. From 1986 to 2003, however, the contemporaneous and peak correlations are negative, -0.35 and -0.43, respectively. This has been particularly noticeable during the last three years of data. An exploration of the cause of this change goes beyond the scope of this paper.

point a new pair (p', s') is drawn from a state dependent distribution. Let $\mathbb{E}_{p,s} X_{p',s'}$ denote the expected value of an arbitrary variable X following the next aggregate shock, conditional on the current state (p, s). I assume that this conditional expectation is finite, which is ensured if the state space is compact. At every point in time, the current values of productivity and the separation rate are common knowledge.

Next I turn to the economic agents in the economy, a measure 1 of risk-neutral, infinitely-lived workers and a continuum of risk-neutral, infinitely-lived firms. All agents discount future payoffs at rate $r > 0$.

Workers can either be unemployed or employed. An unemployed worker gets flow utility z from non-market activity ('leisure') and searches for a job. An employed worker earns an endogenous wage but may not search. I discuss wage determination shortly.

Firms have a constant returns to scale production technology that uses only labor, with labor productivity at time t given by the stochastic realization $p(t)$. In order to hire a worker, a firm must maintain an open vacancy at flow cost c. Free entry drives the expected present value of an open vacancy to zero. A worker and a firm separate according to a Poisson process with arrival rate governed by the stochastic separation rate $s(t)$, leaving the worker unemployed and the firm with nothing.

Let $u(t)$ denote the endogenous unemployment rate,[9] $v(t)$ denote the endogenous measure of vacancies in the economy, and $\theta(t) \equiv v(t)/u(t)$ denote the v-u ratio at time t. The flow of matches is given by a constant returns to scale function $m(u(t), v(t))$, increasing in both arguments. This implies that an unemployed worker finds a job according to a Poisson process with time-varying arrival rate $h(\theta(t)) \equiv m(1, \theta(t))$ and that a vacancy is filled according to a Poisson process with time varying arrival rate $q(\theta(t)) \equiv m(\theta(t)^{-1}, 1) = \frac{h(\theta(t))}{\theta(t)}$.

I assume that in every state of the world, labor productivity $p(t)$ exceeds the value of leisure z, so there are bilateral gains from matching. There is no single compelling theory of wage determination in such an environment, and so I follow the literature and assume that when a worker and firm first meet, the expected gains from trade are divided according to the Nash bargaining solution. The worker can threaten to become unemployed and the firm can threaten to end the job. The present value of surplus beyond these threats is divided between the worker and firm, with the worker keeping a fraction $\beta \in (0, 1)$ of the surplus, her "bargaining power". I make almost no assumptions about what happens to wages after this initial agreement, except that

[9] With the population of workers constant and normalized to one, the unemployment rate and unemployment level are identical in this model. I therefore use these terms interchangeably.

the worker and firm manage to exploit all the joint gains from trade. For example, the wage may be re-bargained whenever the economy is hit with a shock. Alternatively, it may be fixed at its initial value until such time as the firm would prefer to fire the worker or the worker would prefer to quit, whereupon the pair resets the wage so as to avoid an unnecessary and inefficient separation.

3.2 Characterization of Equilibrium

I look for an equilibrium in which the v-u ratio depends only on the current value of p and s, $\theta_{p,s}$.[10] Given the state-contingent v-u ratio, the unemployment rate evolves according to a standard backward looking differential equation,

$$\dot{u}(t) = s(t)(1 - u(t)) - h(\theta_{p(t),s(t)})u(t), \tag{3}$$

where $(p(t), s(t))$ is the aggregate state at time t. A flow $s(t)$ of the $1 - u(t)$ employed workers become unemployed, while a flow $h(\theta)$ of the $u(t)$ unemployed workers find a job. An initial condition pins down the unemployment rate and the aggregate state at some date t_0.

I characterize the v-u ratio using a recursive equation for the joint value to a worker and firm of being matched in excess of breaking up as a function of the current aggregate state, $V_{p,s}$.

$$rV_{p,s} = p - \big(z + h(\theta_{p,s})\beta V_{p,s}\big) - sV_{p,s} + \lambda\big(\mathbb{E}_{p,s}V_{p',s'} - V_{p,s}\big). \tag{4}$$

Appendix A derives this equation from more primitive conditions. The first two terms represent the current flow surplus from matching. If the pair is matched, they produce p units of output. If they were to break up the match, free entry implies the firm would be left with nothing, while the worker would become unemployed, getting leisure z and a probability $h(\theta_{p,s})$ of contacting a firm, in which event the worker would keep a fraction β of the match value $V_{p,s}$. Next, there is a flow probability s that the worker and firm separate, destroying the match value. Finally, an aggregate shock arrives at rate λ, resulting in an expected change in match value $\mathbb{E}_{p,s}V_{p',s'} - V_{p,s}$.

Another critical equation for the match value comes from firms' free entry condition. The flow cost of a vacancy c equals the flow probability that the vacancy contacts a worker times the resulting capital gain, which by Nash bargaining is equal to a fraction

[10] It is straightforward to show in a deterministic version of this model that there is no other equilibrium, e.g. one in which θ depends on the unemployment rate. See Pissarides (1985).

$1-\beta$ of the match value $V_{p,s}$:

$$c = q(\theta_{p,s})(1-\beta)V_{p,s}. \tag{5}$$

Eliminating current and future values of $V_{p,s}$ from (4) using (5) gives

$$\frac{r+s+\lambda}{q(\theta_{p,s})} + \beta\theta_{p,s} = (1-\beta)\frac{p-z}{c} + \lambda\mathbb{E}_{p,s}\frac{1}{q(\theta_{p',s'})}, \tag{6}$$

which implicitly defines the v-u ratio as a function of the current state (p,s).[11] This equation can easily be solved numerically, even with a large state vector. This simple representation of the equilibrium of a stochastic version of the Pissarides (1985) model appears to be new to the literature.

3.3 Comparative Statics

In some special cases, equation (6) can be solved analytically to get a sense of the quantitative results implied by this analysis. First, suppose there are no aggregate shocks, $\lambda = 0$.[12] Then the state-contingent v-u ratio satisfies

$$\frac{r+s}{q(\theta_{p,s})} + \beta\theta_{p,s} = (1-\beta)\frac{p-z}{c}$$

The elasticity of the vacancy-unemployment ratio with respect to 'net labor productivity' $p-z$

$$\frac{r+s+\beta h(\theta_{p,s})}{(r+s)(1-\eta(\theta_{p,s}))+\beta h(\theta_{p,s})},$$

where $\eta(\theta) \in [0,1]$ is the elasticity of $h(\theta)$. This is large only if workers' bargaining power β is small and the elasticity η is close to zero. But with reasonable parameter values, it is close to 1. For example, think of a time period as equal to one month, so the average hiring probability is approximately 0.34 (Section 2.3), the elasticity $\eta(\theta)$ is approximately $1-\alpha = 0.38$ (Section 2.3 again), the average separation probability is approximately 0.02 (Section 2.4), and the interest rate is about 0.004. Then if workers' bargaining power β is equal to $\alpha = 1-\eta(\theta)$, the so-called Hosios (1990) condition for

[11] A similar equation obtains in the presence of aggregate variation in the value of leisure z, the cost of a vacancy c, or workers' bargaining power β.

[12] In Shimer (2003), I show that if there are no aggregate shocks and neither workers nor firms discount future payoffs, the model has similar comparative statics under much weaker assumptions. For example, the matching function can exhibit increasing or decreasing returns to scale and there can be an arbitrary idiosyncratic process for productivity.

efficiency,[13] the elasticity is 1.04. Lower values of β yield a slightly higher elasticity, say 1.19 when $\beta = 0.1$, but only at $\beta = 0$ does the elasticity of the v-u ratio with respect to $p - z$ rise appreciably, to 1.61. It would take extreme parameter values for this elasticity to exceed 2. This implies that unless the value of leisure is close to labor productivity, the v-u ratio is likely to be unresponsive to changes in the labor productivity.

I can similarly compute the elasticity of the v-u ratio with respect to the separation rate,

$$\frac{-s}{(1-\eta(\theta))(r+s)+\beta h(\theta)}.$$

Substituting the same numbers into this expression gives -0.088. Doubling the separation rate would have have a scarcely-discernible impact on the v-u ratio.

Finally, one can examine the independent behavior of vacancies and unemployment. In steady state, equation (3) holds with $\dot{u} = 0$. If the matching function is Cobb-Douglas, $m(u, v) = \mu u^{\alpha} v^{1-\alpha}$, this implies

$$v_{p,s} = \left(\frac{s(1-u_{p,s})}{\mu u_{p,s}^{\alpha}}\right)^{\frac{1}{1-\alpha}}.$$

An increase in labor productivity raises the v-u ratio which lowers the unemployment rate and hence raises the vacancy rate. Vacancies and unemployment should move in opposite directions in response to such shocks. But an increase in the separation rate scarcely affects the v-u ratio. Instead, it raises both the unemployment and vacancy rates, an effect that is likely to produce a counterfactually positive correlation between unemployment and vacancies.

I can perform similar analytic exercises by making other simplifying assumptions. For example, suppose that each vacancy contacts an unemployed worker at a constant Poisson rate μ, independent of the unemployment rate, so $q(\theta) = \mu$. Given the risk-neutrality assumptions, this is equivalent to assuming that firms must pay a fixed cost $\frac{c}{\mu}$ in order to hire a worker. Then even with aggregate shocks, equation (6) yields a static equation for the v-u ratio:

$$\frac{r+s}{\mu} + \beta\theta_{p,s} = (1-\beta)\frac{p-z}{c}.$$

[13] Section 4 shows that the Hosios condition carries over to the stochastic model.

In this case, the elasticity of the v-u ratio with respect to net labor productivity is

$$\frac{r+s+\beta\mu\theta}{\beta\mu\theta}.$$

The elasticity of the v-u ratio with respect to the separation rate is $\frac{-s}{\beta\mu\theta}$. Since $h(\theta) = \mu\theta$, one can again pin down all the parameter values except workers' bargaining power β. Using the same parameter values as above, including $\beta = 0.62$, I obtain elasticities of 1.11 and 0.095, almost unchanged from the case with no shocks. More generally, unless β is nearly equal to zero, both elasticities are very small.

At the opposite extreme, suppose that each unemployed worker contacts a vacancy at a constant Poisson rate μ, independent of the vacancy rate, so $h(\theta) = \mu$ and $q(\theta) = \mu/\theta$. Also assume that the separation rate s is constant and average labor productivity p is a Martingale, $\mathbb{E}_p p' = p$. With this matching function, equation (6) is linear in current and future values of the v-u ratio:

$$\left(\frac{r+s+\lambda}{\mu}+\beta\right)\theta_p = (1-\beta)\frac{p-z}{c}+\frac{\lambda}{\mu}\mathbb{E}_p\theta_{p'}.$$

It is straightforward to verify that the v-u ratio is linear in productivity, and therefore $\mathbb{E}_p\theta_{p'} = \theta_p$. This in turn implies that the elasticity of the v-u ratio with respect to net labor productivity is always equal to 1, regardless of workers' bargaining power. I conclude that with a wide range of parameterizations, the v-u ratio θ should be approximately proportional to net labor productivity $p - z$. I turn next to a serious quantitative evaluation of this proposition.

3.4 Calibration

In this section, I parameterize the model to match the time series behavior of the U.S. unemployment rate. The most important question is the choice of the Markov process for labor productivity and separations. Appendix C develops a discrete state space model which builds on a simple Poisson process corresponding to the theoretical analysis in Section 3.2. I define an underlying variable y that lies on a finite ordered set of points. When a Poisson shock hits, y either moves up or down by one point. The probability of moving up is decreasing in the current value of y, which ensures that y is mean reverting. The stochastic variables are then expressed as functions of y.

Although I also use the discrete state space model in my simulations as well, it is almost exactly correct and significantly easier to think about the behavior of the extrinsic shocks by discussing a related continuous state space model. I express the state variables as functions of an Ornstein-Uhlenbeck process (See Taylor and Karlin 1998, Section 8.5). Let y satisfy

$$dy = -\gamma y dt + \sigma db,$$

where b is a standard Brownian motion. Here $\gamma > 0$ is a measure of persistence, with higher values indicating faster mean reversion, and $\sigma > 0$ is the instantaneous standard deviation. This process has some convenient properties: y is conditionally and unconditionally normal; it is mean reverting, with expected value converging asymptotically to zero; and asymptotically its variance converges unconditionally to $\frac{\sigma^2}{2\gamma}$.

I consider two different cases. In the first, the separation rate is constant and productivity satisfies $p = z + e^y(p^* - z)$, where y is an Ornstein-Uhlenbeck process with parameters γ and σ, and $p^* > z$ is a measure of long-run average productivity. Since $e^y > 0$, this ensures $p > z$. In the second case, productivity is constant and separations satisfy $s = e^y s^*$, where again y follows an Ornstein-Uhlenbeck process and now $s^* > 0$ is a measure of long-run average separation rate. In both cases, the stochastic process is reduced to three parameters, γ, σ, and either p^* or s^*.

I now proceed to explain the choice of the other parameters, starting with the case of stochastic productivity. I follow the literature and assume that the matching function is Cobb-Douglas,

$$h(\theta) = \theta q(\theta) = \mu\theta^\alpha.$$

This reduces the calibration to ten parameters: the productivity parameter p^*, the value of leisure z, workers' bargaining power β, the discount rate r, the separation rate s, the two matching function parameters α and μ, the vacancy cost c, and the mean reversion and standard deviation of the stochastic process, γ and σ.

Without loss of generality, I normalize the productivity parameter to $p^* = 1$. I choose the standard deviation and persistence of the productivity process to match the empirical behavior of labor productivity. This requires setting $\sigma = 0.017$ and $\gamma = 0.004$. An increase in the volatility of productivity σ has a nearly proportional effect on the volatility of other variables, while the persistence of the stochastic process γ scarcely affects the reported results. For example, suppose I reduce γ to 0.001, so productivity is nearly a random walk. After HP filtering the model generated data, the persistence and magnitude of the impulse is virtually unchanged compared with the baseline parameterization. This is because it is difficult to distinguish small values of γ in a finite data set. But reassuringly, the detrended behavior of unemployment and vacancies is also scarcely affected by increasing the persistence of labor productivity.

I set the value of leisure to $z = 0.4$. Since mean labor income in the model is 0.993, this lies at the upper end of the range of income replacement rates in the United States if interpreted entirely as an unemployment benefit.

I normalize a time period to be one quarter, and therefore set the discount rate to $r = 0.012$, equivalent to an annual discount factor of 0.953. The analysis in Section 2.4 suggests a quarterly separation rate of $s = 0.06$, so jobs last for about three years. This is somewhat lower than Abowd and Zellner's (1985) finding that 3.42 percent of employed workers exit employment during a typical month between 1972 and 1982, after correcting for classification and measurement error. It is also somewhat lower that measured turnover rates in the JOLTS, although some separations in that survey reflect job-to-job transitions, a possibility that is absent from this model.

Using the matching function estimates from Section 2.3, I set the elasticity parameter to $\alpha = 0.62$. This lies well within the range of estimates that Petrongolo and Pissarides (2001) report. I also set workers' bargaining power β to the same value, 0.62. Although the results reported here are insensitive to the value of that parameter, I show in Section 4 that if $\alpha = \beta$, the 'Hosios Rule', the decentralized equilibrium maximizes a well-posed social planner's problem.

I use the final two parameters, the matching function constant μ and the vacancy cost c, to pin down the average hiring rate and the average v-u ratio. As reported in Section 2.3, a worker finds a job with a 0.343 probability per month, so the flow arrival rate of job offers $\mu\theta^{1-\alpha}$ should average approximately 1 on a quarterly basis. I do not have a long time series with the level of the v-u ratio but fortunately the model offers one more normalization. Equation (6) implies that doubling c and multiplying μ by a factor 2^α divides the v-u ratio θ in half, doubles the rate at which firms contact workers $q(\theta)$ but does not affect the rate at which workers find jobs. In other words, the average v-u ratio is intrinsically meaningless in the model. I choose to target a mean v-u ratio of 1, which requires setting $\mu = 1.03$ and $c = 0.334$.

Finally, I work on a discrete grid with $2n + 1 = 2001$ points, which closely approximate Gaussian innovations. This implies that Poisson arrival rate of shocks is $\lambda = n\gamma = 4$ times per quarter. Again, Appendix C discusses the relationship between the continuous and discrete state space models.

In the case of shocks to the separation rate, I change only the stochastic process so as to match the empirical results discussed in Section 2.4. Productivity is constant and equal to 1, while the mean separation rate is $s^* = 0.06$. I set $\sigma = 0.109$ and $\gamma = 0.450$, a much less persistent stochastic process. This leaves the average v-u ratio and average hiring rate virtually unchanged. With this parameterization, the economy is hit by approximately $\lambda = n\gamma = 450$ shocks per quarter. Table 2 summarizes the parameter choices in the two simulations.

I use equation (6) to find the state-contingent v-u ratio $\theta_{p,s}$ and then simulate the model. That is, starting with an initial unemployment rate and aggregate state at time 0, I use a pseudo-random number generator to calculate the arrival time of the first Poisson shock. I compute the unemployment rate when that shock arrives, generate a new aggregate state using the discrete-state-space mean-reverting stochastic process described in Appendix C, and repeat. At the end of each period (quarter), I record the aggregate state and the unemployment rate.

I throw away the first 1000 'quarters' of data. I then use the model to generate 212 data points, corresponding to quarterly data from 1951 to 2003, and detrend the model-generated data using an HP filter with the usual smoothing parameter 10^5.[14] I repeat this 10,000 times, giving me good estimates of both the mean of the model-generated data and the standard deviation across model-generated observations. The standard deviation provides me with a good sense of how precisely the model predicts the value of a particular variable.

3.5 Results

Table 3 reports the results from simulations of the model with labor productivity shocks. Along some dimensions, notably the co-movement of unemployment and vacancies, the model performs remarkably well. The empirical correlation between these two variables is -0.89, the noted Beveridge curve, while the model generates an average correlation of -0.88, an economic and statistically insignificant difference. The model also generates the correct autocorrelation for unemployment, although the behavior of vacancies is somewhat off-target. In the data, vacancies are as persistent and volatile as unemployment, while in the model the autocorrelation of vacancies is significantly lower than that of unemployment, while the standard deviation of vacancies is twice as large as the standard deviation of unemployment fluctuations around trend. It is likely that anything that makes vacancies a state variable, such as planning lags or an adjustment cost in vacancy creation, would increase its persistence and reduce its volatility, bringing the model more in line with the data. Fujita (2003) develops a model that adds these realistic features.

But the real problem with the model lies in the volatility of the vacancies and

[14] The quantitative behavior of the model is insensitive to whether I detrend the data, although this affects the choice of impulses. In an earlier version of this paper, I worked directly with the model generated data and reached very similar conclusions (Shimer 2003).

unemployment, or more succinctly, in the volatility of the v-u ratio θ and the hiring rate h. In a reasonably calibrated model, the v-u ratio is only ten percent as volatile as in U.S. data. This is exactly the result predicted from the deterministic comparative statics in Section 3.2. A 1 percent increase in labor productivity p from its average value of 1.01 raises net labor productivity $p - z$ by about 1.66 percent. Using the deterministic model, I argued before that the elasticity of the v-u ratio with respect to net labor productivity is about 1.04 with this choice of parameters, giving a total elasticity of θ with respect to p of approximately $1.66 \times 1.04 = 1.73$ percent. In fact, the standard deviation of $\log \theta$ around trend is about 1.71 times as large as the standard deviation of $\log p$, a quantitatively insignificant difference. Similarly, the hiring rate is eight percent as volatile in the model as in the data.

Not only is there little amplification, but there is also no propagation of the labor productivity shock in the model. The contemporaneous correlation between labor productivity, the v-u ratio, and the hiring rate is 1.00. In the data, the contemporaneous correlation between the first two variables is 0.38 and the v-u ratio lags labor productivity by about one year, with an even lower correlation between labor productivity and the hiring rate.

Table 4 reports the results from the model with shocks to the separation rate. These introduce an almost perfectly positive correlation between unemployment and vacancies, an event that has essentially never been observed at business cycle frequencies (see Figure 3). As a result, separation shocks produce almost no variability in the v-u ratio or the hiring rate. Again, this is consistent with the back-of-the-envelope calculations performed in Section 3.2, where I argued that the elasticity of the v-u ratio with respect to the separation rate should be approximately -0.088. According to the model, the ratio of the standard deviations is about 0.063 and the two variables are strongly negatively correlated.

One might be concerned that the disjoint analysis of labor productivity and separation shocks masks some important interaction between the two impulses. Modelling an endogenous increase in the separation rate due to low labor productivity, as in Mortensen and Pissarides (1994), goes beyond the scope of this paper. Instead, I introduce perfectly negatively correlated labor productivity and separation shocks into the basic model. More precisely, I allow both labor productivity and the separation rate to be functions of the same latent variable y; since both functions are nonlinear, the correlation is slightly larger than -1.

I start with the parameterization of the model with only labor productivity shocks and introduce volatility in the separation rate. Table 5 shows the results from a calibration with equal standard deviations in the log deviation from trend of the separation rate and labor productivity. The behavior of vacancies in the model is now far from the data, with an autocorrelation of 0.38 (compared to 0.94 empirically) and a correlation with unemployment of -0.37 (-0.89). The difference between model and data is highly significant both economically and statistically; across 10,000 sample paths, each 212 quarters long, the correlation between unemployment and vacancies never fell below -0.58. Moreover, although cyclical fluctuations in the separation rate boosts the volatility of unemployment considerably, it has a negligible effect on the cyclical volatility of the v-u ratio and hiring rate, which remain at ten percent of their empirical values. Smaller fluctuations in the separation rate naturally have a smaller effect, while realistically large fluctuations in the separation rate induce a strong positive correlation between unemployment and vacancies (0.96) even in the presence of correlated productivity shocks.

To summarize, the stochastic version of the Pissarides (1985) model confirms that separation shocks induce a positive correlation between unemployment and vacancies. It also confirms that, while labor productivity shocks are qualitatively consistent with a downward sloping Beveridge curve, the search model does not substantially amplify the extrinsic shocks and so labor productivity shocks induce only very small movements along the curve.

3.6 Wages

Until this point, I have assumed that the surplus in new matches is divided according to a generalized Nash bargaining solution, but have made no assumption about the division of surplus in old matches. Although this is sufficient for determining the response of unemployment and vacancies to exogenous shocks, it does not pin down the timing of wage payments. In this section, I introduce an additional assumption, that the surplus in all matches, new or old, is always divided according to the Nash bargaining solution, and so wages are renegotiated following each aggregate shock. This stronger restriction pins down the wage as a function of the aggregate state, $w_{p,s}$. This facilitates a more detailed discussion of wages, which serves two purposes. First, I have argued that flexibility of the present value of wage payments is critical for the many of the results emphasized in this paper. Modelling wages further illuminates that point. And second, it enables me to relate this paper to a literature that examines whether search models can generate rigid wages. Appendix B proves that a continually renegotiated wage solves

$$w_{p,s} = (1-\beta)z + \beta(p + c\theta_{p,s}). \tag{7}$$

This generalizes equation (1.20) in Pissarides (2000) to a stochastic environment.

Consider first the effect of a separation shock on the wage. An increase in the separation rate s induces a slight decline in the v-u ratio (see Table 4), which in turn, by equation (7), reduces wages slightly. Although the direct effect of the shock lowers firms' profits by shortening the duration of matches, the resulting decline in wages partially offsets this. In net, the drop in the v-u ratio is small, and since the increase in separation raises the unemployment rate by shortening unemployment duration, the vacancy rate increases as well. In addition, equation (7) indicates that a separation shock has little effect on the wage since it has little effect on the v-u ratio.

Second, consider a productivity shock. For simplicity, assume that the value of leisure is zero, so an increase in p causes approximately a proportional increase in the v-u ratio (see Table 3). Equation (7) implies that the wage must also increase by the same proportion. This soaks up most of the productivity shock, giving little incentive for firms to create new vacancies. Hence there is a modest increase in vacancies and decrease in unemployment in response to a large productivity shock. The response of wages is nearly proportional to the productivity shock, at least when the value of leisure is small.

To fully understand the importance of wages for the the v-u ratio, it is useful to consider a version of the model in which labor productivity and the separation rate are constant at $p = 1$ and $s = 0.06$, but workers' bargaining power β changes stochastically. An increase in β reduces the profit from creating vacancies, which puts downward pressure on the v-u ratio and upward pressure on the unemployment ratio. It is difficult to know exactly how much variability in β is reasonable, but one can ask how much wage variability is required in order to generate the observed volatility in the v-u ratio. I assume β is a function of the latent mean-reverting variable y, $\beta = \Phi\big(y + \Phi^{-1}(\alpha)\big)$, where Φ is the cumulative standard normal distribution. If y were constant at zero, this implies $\beta = \alpha$, but more generally β is simply bounded between 0 and 1. I set the standard deviation of y to $\sigma = 0.123$ and the mean reversion parameter to $\gamma = 0.04$. Although this implies very modest fluctuations in wages—the standard deviation of detrended log wages, computed as in equation (7), is just 0.01—the calibrated model generates the observed volatility in the v-u ratio, with persistence similar to that in the model with labor productivity shocks. Table 6 shows the complete results. I conclude that minor modifications to the behavior of wages may significantly improve the performance of the model.

If wages are bargained in new matches but then not continually renegotiated, this analysis is inapplicable. Nevertheless, the frequency of wage negotiation does not affect the expected present value of wage payments in new matches, but only changes the timing of wage payments. An increase in productivity or decrease in separation raises the present value of wage payments in new jobs and therefore has little effect on the v-u ratio. An increase workers' bargaining power in new employment relationship induces a large reduction in vacancies and in the v-u ratio.

4 Optimal Vacancy-Unemployment Fluctuations

Another way to highlight the role played by the Nash bargaining assumption is to examine a centralized economy in which it is possible to sidestep the wage-setting issue entirely.[15] Consider a hypothetical social planner who chooses a state-contingent v-u ratio in order to maximize the present discounted value of output net of vacancy creation costs. The planner's problem is represented recursively as

$$rW(p,s,u) = \max_{\theta}\Big(zu + p(1-u) - cu\theta + W_u(p,s,u)\big(s(1-u) - uh(\theta)\big) \\ + \lambda \mathbb{E}_{p,s}\big(W(p',s',u) - W(p,s,u)\big)\Big).$$

Instantaneous output is equal to z times the unemployment rate u plus p times the employment rate minus c times the number of vacancies $v \equiv u\theta$. The value changes gradually as the unemployment rate adjusts, with $\dot{u} = s(1-u) - uh(\theta)$, and suddenly when an aggregate shock changes the state from (p,s) to (p',s') at rate λ.

It is straightforward to verify that in the solution to this problem, the Bellman value W is linear in the unemployment rate, $W_u(p,s,u) = \frac{-c}{h'(\theta_{p,s})}$, and the v-u ratio satisfies

$$\frac{r+s+\lambda}{h'(\theta_{p,s})} - \theta_{p,s}\left(1 - \frac{h(\theta_{p,s})}{\theta_{p,s}h'(\theta_{p,s})}\right) = \frac{p-z}{c} + \lambda \mathbb{E}_{p,s}\left(\frac{1}{h'(\theta_{p',s'})}\right).$$

This implicitly defines the optimal $\theta_{p,s}$, independent of the unemployment rate.

[15] A number of papers examine a 'competitive' search economy, in which firms can commit to wages before hiring workers and can increase their hiring rate by promising higher wages (Peters 1991, Montgomery 1991, Moen 1997, Shimer 1996, Burdett, Shi, and Wright 2001). It is by now well-known that a competitive search equilibrium maximizes output, essentially by creating a market for job applications. This discussion of output maximizing search behavior therefore also pertains to these models.

With a Cobb-Douglas matching function $m(u,v) = \mu u^{\alpha} v^{1-\alpha}$, this reduces to

$$\frac{r+s+\lambda}{q(\theta_{p,s})} + \alpha\theta_{p,s} = (1-\alpha)\frac{p-z}{c} + \lambda\mathbb{E}_{p,s}\left(\frac{1}{q(\theta_{p',s'})}\right),$$

a special case of equation (6), with workers' bargaining power β equal to the elasticity α. This generalizes the Hosios (1990) condition for efficiency of the decentralized equilibrium to an economy with stochastic productivity and separation rates. Since the numerical example in Section 3.5 assumed a Cobb-Douglas matching function with $\alpha = \beta$, the equilibrium allocation described in that section solves the social planner's problem. Conversely, if those parameter values describe the U.S. economy, the observed degree of wage rigidity is inconsistent with output maximization.

With other matching functions, the link between the equilibrium with wage bargaining and the solution to the planner's problem is broken. At one extreme, if unemployment and vacancies are perfect substitutes, i.e. $h(\theta) = \alpha_u + \alpha_v\theta$, then the output-maximizing v-u ratio is infinite whenever $\alpha_v(p-z) > c(r+s+\alpha_u)$ and is zero if the inequality is reversed. With near-perfect substitutability, the output-maximizing v-u ratio is very sensitive to current productivity. This implies small productivity shocks generate large movements in the unemployment rate. On the other hand, if unemployment and vacancies are perfect complements, $h(\theta) = \min\langle\alpha_u, \alpha_v\theta\rangle$, the v-u ratio never strays from the efficient ratio $\frac{\alpha_u}{\alpha_v}$. With imperfect complements, the impact of productivity shocks on the v-u ratio is muffled but not eliminated.

The economics behind these theoretical findings is simple. An increase in labor productivity relative to the value of non-market activity and the cost of advertising a vacancy induces a switch away from the expensive activity, unemployment, and towards the relatively cheap activity, vacancies. The magnitude of the switch depends on how substitutable unemployment and vacancies are in the job search process. If they are strong complements, substitution is nearly impossible and the v-u ratio barely changes. If they are strong substitutes, substitution is nearly costless, and the v-u ratio is highly procyclical.

In the decentralized economy, the extent of substitution between unemployment and vacancies is not governed by technology (the matching function) but rather by the bargaining solution. The Nash bargaining solution effectively corresponds to a moderate degree of substitutability, the Cobb-Douglas case. If wages were more rigid, an increase in productivity would induce more vacancy creation and less unemployment, analogous to a centralized environment with a high elasticity of substitution in the matching function.

The substitutability of unemployment and vacancies is an empirical issue. Blanchard and Diamond (1989) use nonlinear least squares to estimate a Constant Elasticity of Substitution (CES) matching function on U.S. data. Their point estimate for the elasticity of substitution is 0.74, i.e. slightly less substitutable than the Cobb-Douglas case, although they cannot reject the Cobb-Douglas elasticity of 1. As footnote 7 describes, my data suggests a significantly higher elasticity of substitution, 2.5, although the standard errors are so large that I cannot reject either the possibility of a Cobb-Douglas matching function or of perfect substitutability. Whether the observed movements in unemployment and vacancies are optimal therefore remains an open question.

5 Related Literature

There is a large literature that explores whether the search model is consistent with the cyclical behavior of labor markets. Some papers look at the implications of the model for the behavior of various stocks and flows, including the unemployment and vacancy rates, but do not examine the implicit magnitude of the exogenous impulses. Others assume that business cycles are driven by fluctuations in the separation rate s. These papers either impose exogenously or derive within the model a counterfactually constant v-u ratio θ. A third group of papers has tried but failed to reconcile the procyclicality of the v-u ratio with extrinsic shocks of a plausible magnitude.

Papers by Abraham and Katz (1986), Blanchard and Diamond (1989), and Cole and Rogerson (1999) fit into the first category, matching the behavior of labor market stocks and flows by sidestepping the magnitude of impulses. For example, Abraham and Katz (1986) argue that the downward sloping Beveridge curve is inconsistent with models in which unemployment is driven by fluctuations in the separation rate, notably Lilien's (1982) sectoral shifts model. That leads them to advocate an alternative in which unemployment fluctuations are driven by aggregate disturbances, e.g. productivity shocks. Unfortunately, they fail to examine the magnitude of shocks needed to deliver the observed shifts along the Beveridge curve. Blanchard and Diamond (1989) also focus on the negative correlation between unemployment and vacancies, but they do not model the supply of jobs and hence do not explain why there are so few vacancies during recessions. Instead, they assume the total stock of jobs follows an exogenous stochastic process. This paper pushes the cyclicality of the v-u ratio to the front of the picture. Likewise, Cole and Rogerson (1999) argue that the Mortensen and Pissarides (1994) model can match a variety of business cycle facts, but they do so in a reduced form model that treats fluctuations in the job finding rate, and hence implicitly in the v-u ratio, as exogenous.

The second group of papers, including work by Pries (2004), Ramey and Watson (1997), Den Haan, Ramey, and Watson (2000), and Gomes, Greenwood, and Rebelo (2001), assume that employment fluctuations are largely due to time-variation in the separation rate, minimizing the role played by the observed cyclicality of the v-u ratio. These papers typically deliver rigid wages from a search model, consistent with the findings in Section 3.6. Building on the ideas in Hall (1995), Pries (2004) shows that a brief adverse shock that destroys some old employment relationships can generate a long transition period of high unemployment as the displaced workers move through a number of short-term jobs before eventually finding their way back into long-term relationships. During this transition process, the v-u ratio remains constant, since aggregate economic conditions have returned to normal. Equivalently, the economy moves along an upward sloping Beveridge curve during the transition period, in contradiction to the evidence. Ramey and Watson (1997) argue that two-sided asymmetric information generates rigid wages in a search model. But in their model, shocks to the separation rate are the only source of fluctuations in unemployment. The job finding rate $h(\theta)$ is exogenous and constant, which is equivalent to assuming that vacancies are proportional to unemployment. This is probably an important part of the explanation for why their model produces rigid wages. Den Haan, Ramey, and Watson (2000) show that fluctuations in the separation rate amplify productivity shocks in a model similar to the one examined here; however, they do not discuss the cyclical behavior of the v-u ratio. It is unlikely that they are successful in matching the empirical volatility of this key variable. Similarly, Gomes, Greenwood, and Rebelo (2001) sidestep the v-u issue by looking at a model in which the job finding rate is exogenous and constant, i.e. vacancies are proportional to unemployment. Again, this helps keep wages relatively rigid in their model.

Mortensen and Pissarides (1994) is probably the best known paper in this literature. In their three state 'illustrative simulation', the authors introduce, without comment, enormous productivity or leisure shocks into their model. Average labor productivity minus the value of leisure $p - z$ is approximately three times as high in the good state as in the bad state.[16] This paper confirms that in response to such large shocks, the v-u ratio should also be about three times as large in the good state as in the bad state, but argues that there is no evidence for these large shocks in the data. Even if one

[16]This calculation would be easy in the absence of heterogeneity, i.e. if their parameter σ were equal to zero. Then $\bar{p} - z$ would take on three possible values: 0.022, 0.075, and 0.128, for a six-fold difference in $\bar{p} - z$ between the high and low states.

accepts the magnitude of the implied impulses, Mortensen and Pissarides (1994) still only delivers a correlation of -0.26 between unemployment and vacancies, far lower than the empirical value of -0.90. This is probably because of the tension between productivity shocks, which put the economy on a downward-sloping Beveridge curve, and endogenous movements in the separation rate, which have the opposite effect. Merz (1995) and Andolfatto (1996) both put the standard search model into a real business cycle framework with intertemporal substitution of leisure, capital accumulation, and other extensions. Neither paper can match the negative correlation between unemployment and vacancies and both papers generate real wages that are too flexible in response to productivity shocks. Thus these papers encounter the problem I highlight in this paper, although they do not emphasize this shortcoming of the search model. Finally, in contemporaneous work, Hall (2002) highlights the same issues that I emphasize here. Hall (2004) proposes a solution: real wages are determined largely by a backward looking social norm, and hence scarcely change over the business cycle.

6 Conclusion

I have argued in this paper that a search and matching model in which wages are determined by Nash bargaining cannot generate substantial movements along a downward sloping Beveridge curve in response to shocks of a plausible magnitude. A labor productivity shock primarily results in higher wages, with little effect on the v-u ratio. A separation shock generates an increase in both unemployment and vacancies. It is important to stress that this is not an attack on the search approach to labor markets, but rather a critique of the commonly-used Nash bargaining assumption for wage determination. An alternative wage determination mechanism that generates more rigid wages in new jobs (measured in present value terms) will amplify the effect of productivity shocks on the v-u ratio, helping to reconcile the evidence and theory. Countercyclical movements in workers' bargaining power provide one such mechanism, at least in a reduced form sense.

If the matching function is Cobb-Douglas, the observed behavior of the v-u ratio is not socially optimal, but it is optimal if the elasticity of substitution between unemployment and vacancies in the matching function is large. The estimates of a CES matching function are imprecise, so it is unclear whether observed wages are 'too rigid'.

One way to generate more rigid wages in a theoretical model is to introduce considerations whereby wages affect the worker turnover rate. For example, in the Burdett and Mortensen (1998) model of on-the-job search, firms have an incentive to offer high wages in order to attract workers away from competitors and to reduce employees' quit rate. The distribution of productivity affects an individual firm's wage offer and vacancy creation decision in complex ways, breaking the link between average labor productivity and the v-u ratio. In particular, a shift in the productivity distribution that leaves average labor productivity unchanged may have appreciably affect average wages and hence the equilibrium v-u ratio.

Another possibility is to drop some of the informational assumptions in the standard search model.[17] Suppose that when a worker and firm meet, they draw an idiosyncratic match-specific productivity level from some distribution F. Workers and firms know about aggregate variables, including the unemployment rate and the distribution F, but only the firm knows the realized productivity level. Bargaining proceeds as follows: with probability $\beta \in (0, 1)$, a worker makes a take-it-or-leave-it wage demand, and otherwise the firm makes a take-it-or-leave-it offer. Obviously the firm extracts all the rents from the employment relationship when it makes an offer. But if the uninformed worker makes the offer, she faces a tradeoff between demanding a higher wage and reducing her risk of unemployment, so the wage depends on the hazard rate of the distribution F. This again breaks the link between average labor productivity and the equilibrium v-u ratio. Exploring whether either of these models, or some related model, delivers substantial fluctuations in the v-u ratio in response to plausible impulses remains a topic for future research.

[17] Ramey and Watson (1997) develop a search model with two-sided asymmetric information. Because they assume workers' job finding rate is exogenous and acyclic, their results are not directly applicable to this analysis, although their methodology may prove useful.

Appendix

A Derivation of the Equation for Surplus (4)

For notational simplicity alone, assume the wage payment depends only on the aggregate state, $w_{p,s}$, not on the history of the match. I return to this issue at the end of this section. Define $U_{p,s}$, $E_{p,s}$, and $J_{p,s}$ to be the state-contingent present value of an unemployed worker, employed worker, and filled job, respectively. They are linked recursively by:

$$rU_{p,s} = z + h(\theta_{p,s})\big(E_{p,s} - U_{p,s}\big) + \lambda\big(\mathbb{E}_{p,s}U_{p',s'} - U_{p,s}\big) \tag{8}$$

$$rE_{p,s} = w_{p,s} - s\big(E_{p,s} - U_{p,s}\big) + \lambda\big(\mathbb{E}_{p,s}E_{p',s'} - E_{p,s}\big) \tag{9}$$

$$rJ_{p,s} = p - w_{p,s} - sJ_{p,s} + \lambda\big(\mathbb{E}_{p,s}J_{p',s'} - J_{p,s}\big) \tag{10}$$

Equation (8) states that the flow value of an unemployed worker is equal to her value of leisure z plus the probability she finds a job $h(\theta_{p,s})$ times the resulting capital gain $E - U$ plus the probability of an aggregate shock times that capital gain. Equation (9) expresses a similar idea for an employed worker, who receives a wage payment $w_{p,s}$ but loses her job at rate s. Equation (10) provides an analogous recursive formulation for the value of a filled job. Note that a firm is left with nothing when a filled job ends.

Sum equations (9) and (10) and then subtract equation (8), defining $V_{p,s} \equiv J_{p,s} + E_{p,s} - U_{p,s}$:

$$rV_{p,s} = p - z - h(\theta_{p,s})\big(E_{p,s} - U_{p,s}\big) - sV_{p,s} + \lambda\big(\mathbb{E}_{p,s}V_{p',s'} - V_{p,s}\big). \tag{11}$$

In addition, the Nash bargaining solution implies that the wage is set so as to maximize the Nash product $\big(E_{p,s} - U_{p,s}\big)^{\beta} J_{p,s}^{1-\beta}$, which gives

$$\frac{E_{p,s} - U_{p,s}}{\beta} = V_{p,s} = \frac{J_{p,s}}{1-\beta}. \tag{12}$$

Substituting for $E - U$ in equation (11) yields equation (4).

If I allow wages to depend in an arbitrary manner on the history of the match, this would affect the Bellman values E and J; however, the wage, and therefore the history-dependence, would drop out when summing the Bellman equations for E and J, assuming matches end only when the surplus is negative. In other words, the match surplus V is unchanged and in particular is not history dependent, regardless of the frequency of wage renegotiation.

B Derivation of the Wage Equation

Assume that wages are continually renegotiated, so the wage only depends on the current aggregate state (p, s). Eliminate current and future values of J from equation (10) using equation (12)

$$w_{p,s} = p - (r + s + \lambda)(1 - \beta)V_{p,s} + \lambda \mathbb{E}_{p,s}(1 - \beta)V_{p',s'}.$$

Similarly, eliminate current and future values of V using (5):

$$w_{p,s} = p - \frac{(r + s + \lambda)c}{q(\theta_{p,s})} + \lambda \mathbb{E}_{p,s}\frac{c}{q(\theta_{p',s'})}$$

Finally, replace the last two terms using equation (6) to get equation (7).

C The Stochastic Process

The text describes a continuous state space approximation to the discrete state space model used in both the theory and simulations. Here I describe the discrete state space model and show that it asymptotes to an Ornstein-Uhlenbeck process.

Consider a random variable y that is hit with shocks according to a Poisson process with arrival rate λ. The initial value of y lies on a discrete grid,

$$y \in Y \equiv \{-n\Delta, -(n-1)\Delta, \dots, 0, \dots, (n-1)\Delta, n\Delta\},$$

where $\Delta > 0$ is the step size and $2n + 1 \geq 3$ is the number of grid points. When a shock hits, the new value y' either moves up or down by one grid point:

$$y' = \begin{cases} y + \Delta \\ y - \Delta \end{cases} \text{with probability} \begin{cases} \frac{1}{2}\left(1 - \frac{y}{n\Delta}\right) \\ \frac{1}{2}\left(1 + \frac{y}{n\Delta}\right) \end{cases}$$

Note that although the step size is constant, the probability that $y' = y + \Delta$ is smaller when y is larger, falling from 1 at $y = -n\Delta$ to zero at $y = n\Delta$.

It is trivial to confirm that $y' \in Y$, so the state space is discrete. To proceed further, define $\gamma \equiv \lambda/n$ and $\sigma \equiv \sqrt{\lambda}\Delta$. For any fixed $y(t)$, I examine the behavior of $y(t + h)$ over an arbitrarily short time period h. For sufficiently short h, the probability that two Poisson shocks arrive is negligible, and so $y(t+h)$ is equal to $y(t)$ with probability $1-h\lambda$, has increased by Δ with probability $\frac{h\lambda}{2}\left(1-\frac{y}{n\Delta}\right)$, and has decreased by Δ with probability $\frac{h\lambda}{2}\left(1+\frac{y}{n\Delta}\right)$. Adding this together shows

$$\mathbb{E}\left(y(t+h)-y(t)|y(t)\right)=-\frac{h\lambda}{n}y(t)=-h\gamma y(t).$$

Next, the conditional variance of $y(t+h)-y(t)$ can be decomposed into

$$\operatorname{Var}\left(y(t+h)-y(t)|\,y(t)\right)=\mathbb{E}\left((y(t+h)-y(t))^2|y(t)\right)-\left(\mathbb{E}\big(y(t+h)-y(t)|y(t)\big)\right)^2.$$

The first term evaluates to $h\lambda\Delta^2$ over a sufficiently short time interval h, since it is equal to Δ^2 if a shock, positive or negative, arrives and zero otherwise. The second term is $(h\gamma y(t))^2$, and so is negligible over a short time interval h. Thus

$$\operatorname{Var}\left(y(t+h)-y(t)|\,y(t)\right)=h\lambda\Delta^2=h\sigma^2.$$

Putting this together, we can represent the stochastic process for y as

$$dy=-\gamma y dt+\sigma dx,$$

where for $t>0$, the expected value of $x(t)$ given $x(0)$ is $x(0)$ and the conditional variance is t. This is similar to a Brownian motion, except that the innovations in x are not Gaussian, since y is constrained to lie on a discrete grid.

Now suppose one changes the three parameters of the stochastic process, the step size, arrival rate of shocks, and number of steps, from (Δ,λ,n) to $\left(\Delta\sqrt{\varepsilon},\frac{\lambda}{\varepsilon},\frac{n}{\varepsilon}\right)$ for any $\varepsilon>0$. It is easy to verify that this does not change either the autocorrelation parameter $\gamma=\lambda/n$ or the instantaneous variance $\sigma=\sqrt{\lambda}\Delta$. But as $\varepsilon\to 0$, the distribution of the innovation process x converges to a normal by the Central Limit Theorem. Equivalently, y converges to an Ornstein-Uhlenbeck process.[18] This observation is also useful for computation. It is possible to find a solution on a coarse grid and then to refine the grid by decreasing ε without substantially changing the results.

[18]Notably, for large n it is extraordinarily unlikely that the state variable reaches its limiting values of $\pm n\Delta$. The unconditional distribution of the state variable is approximately normal with mean zero and standard deviation $\sigma/\sqrt{2\gamma}=\Delta\sqrt{n/2}$. The limiting values of the state variables therefore lie $\sqrt{2n}$ standard deviations above and below the mean. If $n=1000$, as is the case in the simulations, one should expect to observe such values approximately once in 10^{436} periods.

References

Abowd, John, and Arnold Zellner (1985): "Estimating Gross Labor-Force Flows," *Journal of Business and Economic Statistics*, 3(3), 254–283.

Abraham, Katharine (1987): "Help-Wanted Advertising, Job Vacancies, and Unemployment," *Brookings Papers on Economic Activity*, 1, 207–243.

Abraham, Katharine, and John Haltiwanger (1995): "Real Wages and the Business Cycle," *Journal of Economic Literature*, 33, 1215–1264.

Abraham, Katharine, and Lawrence Katz (1986): "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?," *Journal of Political Economy*, 94(3), 507–522.

Abraham, Katharine, and Robert Shimer (2001): "Changes in Unemployment Duration and Labor-Force Attachment," in *The Roaring Nineties: Can Full Employment be Sustained?*, ed. by Alan Krueger, and Robert Solow, chap. 8, pp. 367–420. Russell Sage Foundation, New York.

Andolfatto, David (1996): "Business Cycles and Labor-Market Search," *American Economic Review*, 86(1), 112–132.

Blanchard, Olivier, and Peter Diamond (1989): "The Beveridge Curve," *Brookings Papers on Economic Activity*, 1, 1–60.

——— (1990): "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity*, 2, 85–143.

Bleakley, Hoyt, Ann Ferris, and Jeffrey Fuhrer (1999): "New Data on Worker Flows During Business Cycles," *New England Economic Review*, pp. 49–76.

Burdett, Kenneth, and Dale Mortensen (1998): "Equilibrium Wage Differentials and Employer Size," *International Economic Review*, 39, 257–274.

Burdett, Kenneth, Shouyong Shi, and Randall Wright (2001): "Pricing and Matching with Frictions," *Journal of Political Economy*, 109(5), 1060–1085.

Cole, Harold, and Richard Rogerson (1999): "Can the Mortensen-Pissarides Matching Model Match the Business-Cycle Facts?," *International Economic Review*, 40(4), 933–959.

Coles, Melvyn (2001): "Equilibrium Wage Dispersion, Firm Size, and Growth," *Review of Economic Dynamics*, 4(1), 159–187.

Den Haan, Wouter, Garey Ramey, and Joel Watson (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90(3), 482–498.

Flinn, Christopher, and James Heckman (1983): "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States," *Journal of Labor Economics*, 1, 28–42.

Fujita, Shigeru (2003): "The Beveridge Curve, Job Creation and the Propagation of Shocks," University of California at San Diego Mimeo.

Gomes, Joao, Jeremy Greenwood, and Sergio Rebelo (2001): "Equilibrium Unemployment," *Journal of Monetary Economics*, 48, 109–152.

Hall, Robert (1995): "Lost Jobs," *Brookings Papers on Economic Activity*, 1, 221–273.

——— (2002): "Modern Theory of Unemployment Fluctuations: Empirics and Policy Applications," Mimeo prepared for *American Economic Review, Papers and Proceedings*.

——— (2004): "Employment Fluctuations with Equilibrium Wage Stickiness," Mimeo.

Hosios, Arthur (1990): "On the Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies*, 57(2), 279–298.

Jones, Stephen, and Craig Riddell (1999): "The Measurement of Unemployment: An Empirical Approach," *Econometrica*, 67, 147–161.

Juhn, Chinhui, Kevin Murphy, and Robert Topel (1991): "Why Has the Natural Rate of Unemployment Increased over Time?," *Brookings Papers on Economic Activity*, 2, 75–142.

Lilien, David (1982): "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy*, 90(4), 777–793.

Lucas, Robert, and Leonard Rapping (1969): "Real Wages, Employment, and Inflation," *Journal of Political Economy*, 57, 721–754.

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

Moen, Espen (1997): "Competitive Search Equilibrium," *Journal of Political Economy*, 105(2), 385–411.

Montgomery, James (1991): "Equilibrium Wage Dispersion and Interindustry Wage Differentials," *Quarterly Journal of Economics*, 106(1), 163–179.

Mortensen, Dale, and Christopher Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, 397–415.

Peters, Michael (1991): "Ex Ante Price Offers in Matching Games: Non-Steady States," *Econometrica*, 59(5), 1425–1454.

Petrongolo, Barbara, and Christopher Pissarides (2001): "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, 39(2), 390–431.

Pissarides, Christopher (1985): "Short-Run Equlibrium Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review*, 75, 676–690.

——— (2000): *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, second edn.

Poterba, James, and Lawrence Summers (1986): "Reporting Errors and Labor Market Dynamics," *Econometrica*, 54(6), 1319–1338.

Pries, Michael (2004): "Persistence of Employment Fluctuations: A Model of Recurring Job Loss," *Review of Economic Studies*, 1(71), 193–215.

Ramey, Gary, and Joel Watson (1997): "Contractual Fragility, Job Destruction, and Business Cycles," *Quarterly Journal of Economics*, 112(3), 873–911.

Shimer, Robert (1996): "Contracts in a Frictional Labor Market," Mimeo.

——— (2003): "The Cylical Behavior of Equilibrium Unemployment and Vacancies: Evidence and Theory," NBER Working Paper 9536.

Solon, Gary, Robert Barsky, and Jonathan Parker (1994): "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?," *Quarterly Journal of Economics*, 109, 1–25.

Taylor, Howard, and Samuel Karlin (1998): *An Introduction to Stochastic Modeling*. Academic Press, third edn.

Veracierto, Marcelo (2002): "On the Cyclical Behavior of Employment, Unemployment and Labor Force Participation," Federal Reserve Bank of Chicago working paper 2002-12.

Summary Statistics, quarterly U.S. data, 1951 to 2003

		u	v	v/u	h	s	p
Standard Deviation		0.192	0.204	0.384	0.175	0.109	0.021
Quarterly Autocorrelation		0.937	0.939	0.941	0.871	0.589	0.881
	u	1	−0.878	−0.967	−0.914	0.328	−0.414
	v	—	1	0.971	0.792	−0.449	0.346
Correlation Matrix	v/u	—	—	1	0.878	−0.403	0.391
	h	—	—	—	1	−0.031	0.286
	s	—	—	—	—	1	−0.474
	p	—	—	—	—	—	1

Table 1: Seasonally adjusted unemployment u is constructed by the BLS from the Current Population Survey (CPS). The seasonally adjusted help-wanted advertising index v is constructed by the Conference Board. The hiring rate h and separation rate s are constructed from seasonally adjusted employment, unemployment, and mean unemployment duration, all computed by the BLS from the CPS, as explained in equations (1) and (2). u, v, h, and s are quarterly averages of seasonally adjusted monthly series. Average labor productivity p is seasonally adjusted real average output per person in the non-farm business sector, constructed by the Bureau of Labor Statistics (BLS) from the National Income and Product Accounts and the Current Employment Statistics. All variables are expressed as log-deviations from an HP trend with smoothing parameter 10^5.

	Source of Shocks	
Parameter	Productivity	Separation
productivity p	stochastic	1
separation rate s	0.06	stochastic
discount rate r	0.012	0.012
value of leisure z	0.4	0.4
matching function $q(\theta)$	$1.03\theta^{-0.62}$	$1.03\theta^{-0.62}$
bargaining power β	0.62	0.62
cost of vacancy c	0.334	0.334
standard deviation σ	0.017	0.109
autoregressive parameter γ	0.004	0.450
grid size $2n+1$	2001	2001

Table 2: Parameter values in simulations of the dynamic stochastic model. The text provides details on the stochastic process for productivity and for the separation rate.

Labor Productivity Shocks

		u	v	v/u	h	p
Standard Deviation		0.012 (0.002)	0.025 (0.003)	0.036 (0.005)	0.014 (0.002)	0.021 (0.003)
Quarterly Autocorrelation		0.949 (0.015)	0.801 (0.052)	0.878 (0.035)	0.878 (0.035)	0.878 (0.035)
	u	1	−0.879 (0.031)	−0.944 (0.016)	−0.944 (0.016)	−0.943 (0.016)
	v	—	1	0.987 (0.003)	0.987 (0.003)	0.986 (0.003)
Correlation Matrix	v/u	—	—	1	1.000 (0.000)	0.999 (0.001)
	h	—	—	—	1	0.999 (0.001)
	p	—	—	—	—	1

Table 3: Results from simulating the dynamic model with stochastic labor productivity. All variables are reported as log deviations from an HP trend with smoothing parameter 10^5. Bootstrapped standard errors—the standard deviation across 10,000 model simulations—are reported in parentheses. The text provides details on the stochastic process for productivity.

Separation Rate Shocks

		u	v	v/u	h	s
Standard Deviation		0.085 (0.008)	0.079 (0.007)	0.007 (0.001)	0.003 (0.000)	0.108 (0.008)
Quarterly Autocorrelation		0.815 (0.029)	0.813 (0.029)	0.580 (0.057)	0.580 (0.057)	0.581 (0.056)
Correlation Matrix	u	1	0.999 (0.000)	−0.813 (0.025)	−0.814 (0.025)	0.815 (0.025)
	v	—	1	−0.783 (0.029)	−0.783 (0.029)	0.785 (0.029)
	v/u	—	—	1	1.000 (0.000)	−0.998 (0.000)
	h	—	—	—	1	−0.998 (0.001)
	s	—	—	—	—	1

Table 4: Results from simulating the dynamic model with a stochastic separation rate. All variables are reported as log deviations from an HP trend with smoothing parameter 10^5. Bootstrapped standard errors—the standard deviation across 10,000 model simulations—are reported in parentheses. The text provides details on the stochastic process for the separation rate.

Labor Productivity and Separation Rate Shocks

		u	v	v/u	h	s	p
Standard Deviation		0.032 (0.005)	0.013 (0.001)	0.038 (0.006)	0.015 (0.002)	0.021 (0.003)	0.021 (0.003)
Quarterly Autocorrelation		0.949 (0.015)	0.382 (0.074)	0.879 (0.035)	0.878 (0.035)	0.878 (0.035)	0.878 (0.035)
Correlation Matrix	u	1	−0.361 (0.060)	−0.945 (0.016)	−0.944 (0.016)	0.945 (0.016)	−0.943 (0.016)
	v	—	1	0.644 (0.034)	0.644 (0.034)	−0.644 (0.034)	0.642 (0.035)
	v/u	—	—	1	0.999 (0.001)	−1.000 (0.000)	0.998 (0.003)
	h	—	—	—	1	−0.999 (0.001)	0.999 (0.001)
	s	—	—	—	—	1	−0.998 (0.002)
	p	—	—	—	—	—	1

Table 5: Results from simulating the dynamic model with stochastic but perfectly correlated labor productivity and separation rate. All variables are reported as log deviations from an HP trend with smoothing parameter 10^5. Bootstrapped standard errors—the standard deviation across 10,000 model simulations—are reported in parentheses. The text provides details on the stochastic process.

Bargaining Power Shocks

		u	v	v/u	h	w
Standard Deviation		0.128 (0.021)	0.267 (0.038)	0.381 (0.058)	0.145 (0.022)	0.010 (0.004)
Quarterly Autocorrelation		0.943 (0.017)	0.786 (0.056)	0.865 (0.038)	0.865 (0.038)	0.851 (0.043)
Correlation Matrix	u	1	-0.850 (0.035)	-0.929 (0.019)	-0.926 (0.020)	0.870 (0.046)
	v	—	1	0.985 (0.003)	0.985 (0.002)	-0.899 (0.046)
	v/u	—	—	1	1.000 (0.001)	-0.920 (0.044)
	h	—	—	—	1	-0.920 (0.044)
	w	—	—	—	—	1

Table 6: Results from simulating the dynamic model with a stochastic bargaining power. All variables are reported as log deviations from an HP trend with smoothing parameter 10^5. Bootstrapped standard errors—the standard deviation across 10,000 model simulations—are reported in parentheses. The text provides details on the stochastic process for the workers' bargaining power.

Quarterly U.S. Unemployment, in millions, and Trend, 1951–2003

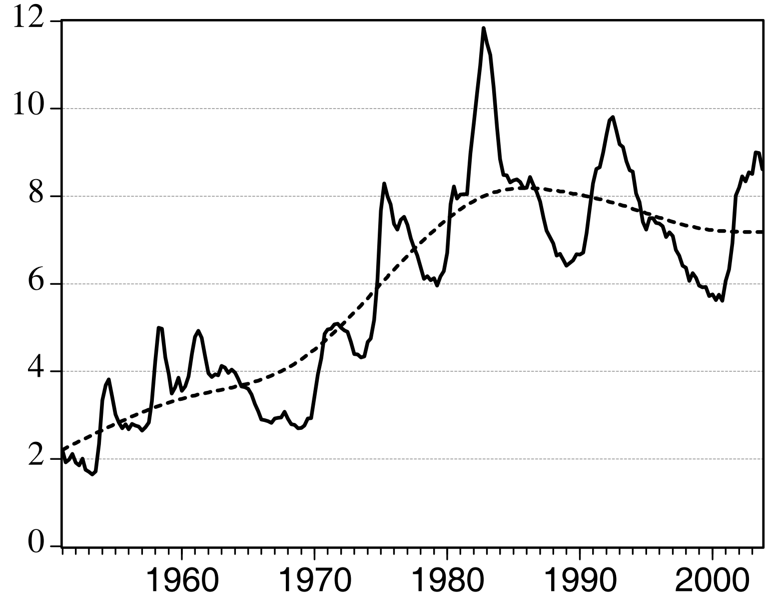

Figure 1: Unemployment is a quarterly average of the seasonally adjusted monthly series constructed by the BLS from the CPS, survey home page `http://www.bls.gov/cps/`. The trend is an HP filter of the quarterly data with smoothing parameter 10^5.

Two Measures of U.S. Job Vacancies, 2000Q4–2003Q4

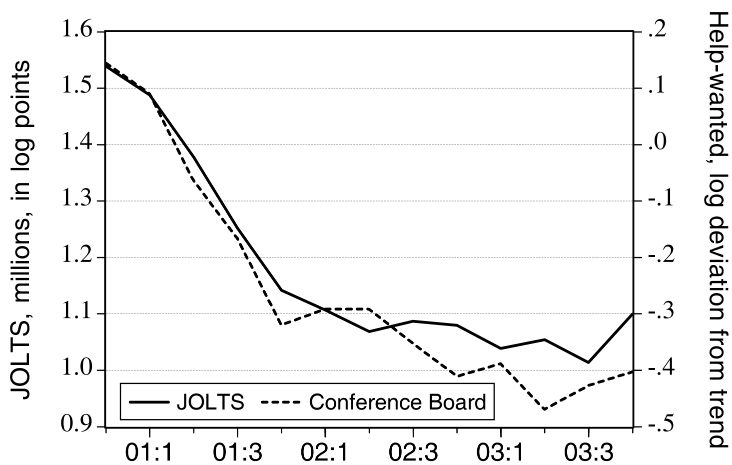

Figure 2: The solid line shows the logarithm of the number of job openings in millions, measured by the BLS from the JOLTS, survey homepage `http://www.bls.gov/jlt`, seasonally adjusted using using the Census X12 algorithm, and averaged over a quarter. The dashed line shows the log deviation from trend of the quarterly average, seasonally adjusted Conference Board help-wanted advertising index.

Quarterly U.S. Help Wanted Advertising Index and Trend, 1951–2003

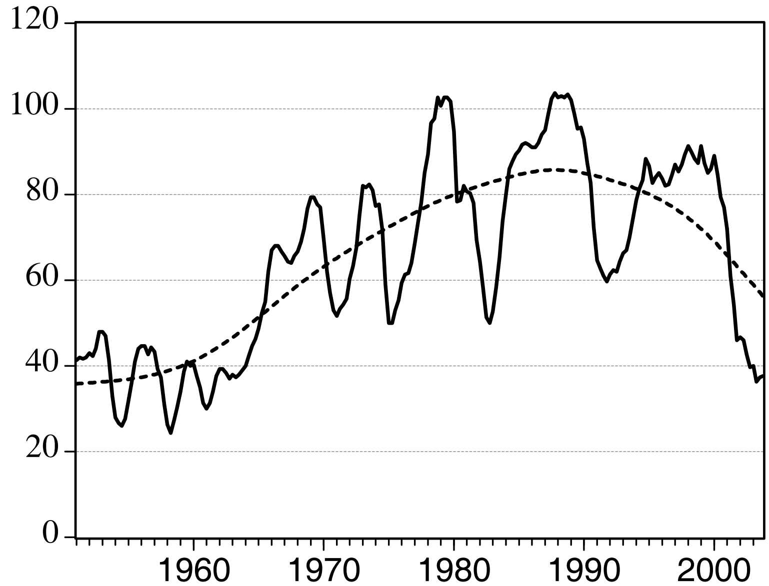

Figure 3: The help-wanted advertising index is a quarterly average of the seasonally adjusted monthly series constructed by the Conference Board with normalization 1987 = 100. The data were downloaded from the Federal Reserve Bank of St. Louis database at `http://research.stlouisfed.org/fred2/data/HELPWANT.txt`. The trend is an HP filter of the quarterly data with smoothing parameter 10^5.

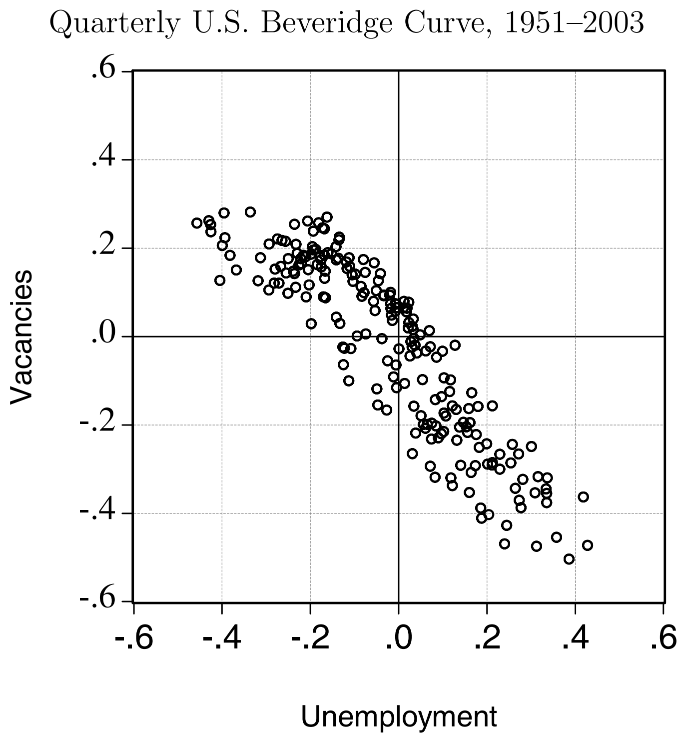

Figure 4: Unemployment is constructed by the BLS from the CPS. The help-wanted advertising index is constructed by the Conference Board. Both are quarterly averages of seasonally adjusted monthly series and are expressed as log deviations from an HP filter with smoothing parameter 10^5.

Monthly Hiring Probability for Unemployed Workers, 1951–2003

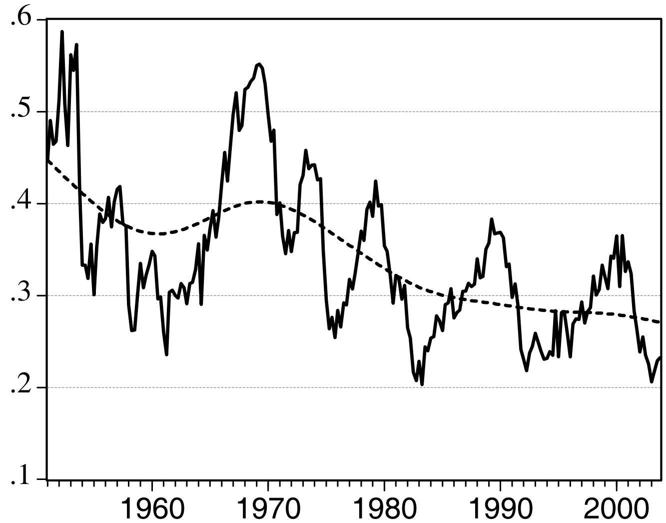

Figure 5: The hiring rate is computed using equation (1), with unemployment and unemployment duration data constructed and seasonally adjusted by the BLS from the CPS, survey home page `http://www.bls.gov/cps/`. It is expressed as a quarterly average of monthly data. The trend is an HP filter of the quarterly data with smoothing parameter 10^5.

Quarterly U.S. Matching Function, 1951–2003

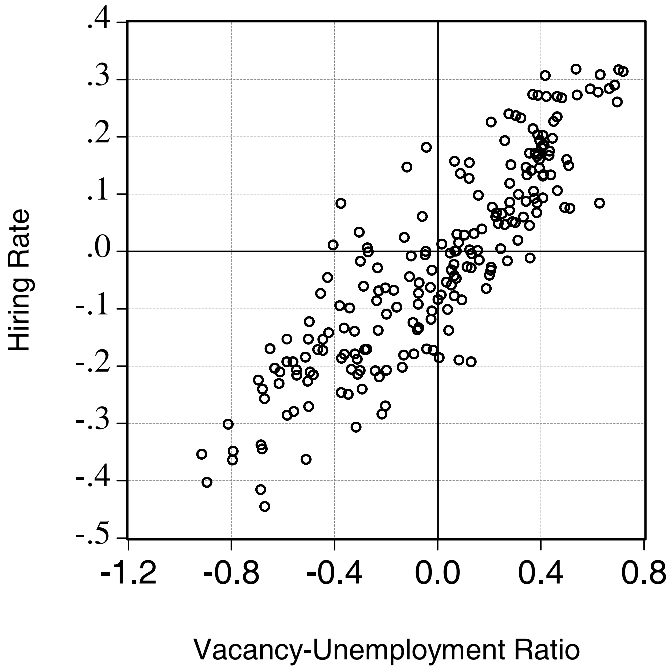

Figure 6: The v-u ratio is constructed by the BLS from the CPS and by the Conference Board. The hiring rate is constructed using equation (1) and BLS data from the CPS. Both are quarterly averages of seasonally adjusted monthly series and are expressed as log deviations from an HP filter with smoothing parameter 10^5.

Monthly Separation Probability for Employed Workers, 1951–2003

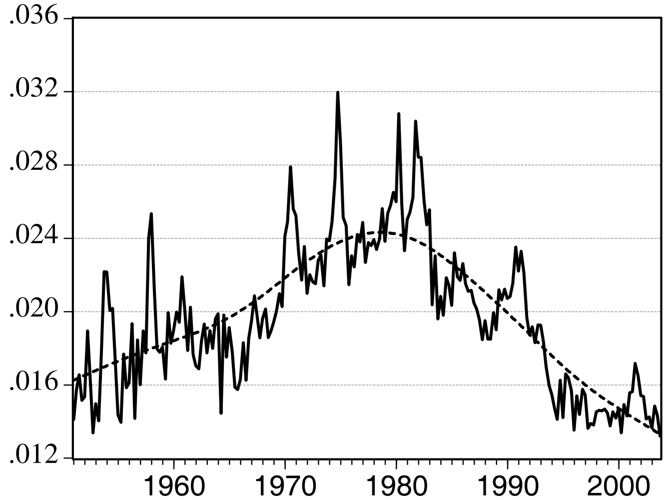

Figure 7: The separation rate is computed using equation (2), with employment, unemployment, and unemployment duration data constructed and seasonally adjusted by the BLS from the CPS, survey home page `http://www.bls.gov/cps/`. It is expressed as a quarterly average of monthly data. The trend is an HP filter of the quarterly data with smoothing parameter 10^5.

Quarterly U.S. Average Labor Productivity and Trend, 1951–2003

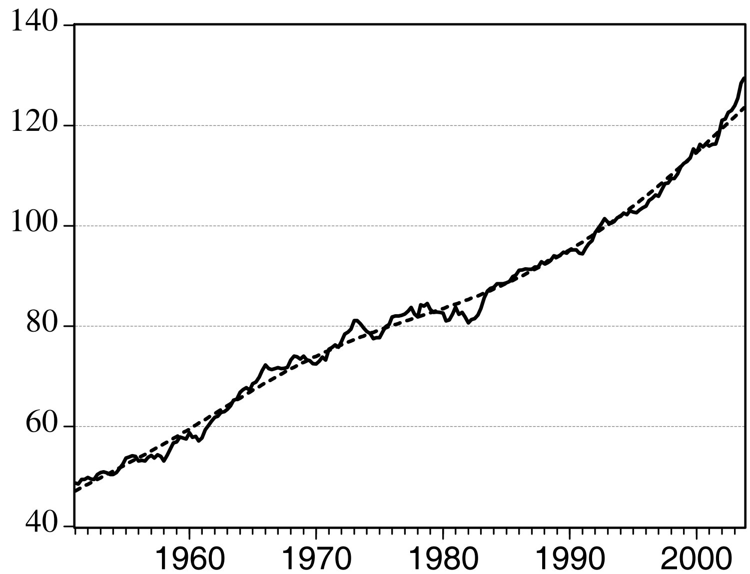

Figure 8: Real output per person in the non-farm business sector, constructed by the BLS Major Sector Productivity and Costs program, survey home page `http://www.bls.gov/lpc/`, 1992 = 100. The trend is an HP filter of the quarterly data with smoothing parameter 10^5.

Quarterly U.S. Vacancy-Unemployment Ratio and Average Labor Productivity, 1951–2003

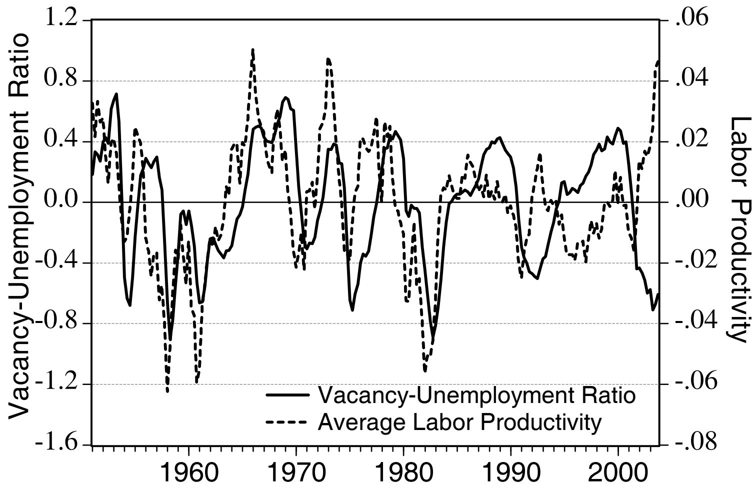

Figure 9: Unemployment is constructed by the BLS from the CPS. The help-wanted advertising index is constructed by the Conference Board. Both are quarterly averages of seasonally adjusted monthly series. Labor productivity is real average output per worker in the non-farm business sector, constructed by the BLS Major Sector Productivity and Costs program. The v-u ratio and labor productivity are expressed as log deviations from an HP filter with smoothing parameter 10^5.