

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

Robert Shimer
Department of Economics
Princeton University
and NBER
shimer@princeton.edu

February 14, 2003

Abstract

This paper argues that a broad class of search models 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 vacancy-unemployment ratio is 20 times as volatile as average labor productivity, while under weak assumptions, search models predict that the vacancy-unemployment ratio and labor productivity have nearly the same variance.

I establish this claim both using analytical comparative statics in a very general deterministic search model and using simulations of a stochastic version of the model. I show that 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 job destruction 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. I reconcile these findings with an existing literature and argue that the source of the model's failure is lack of wage rigidity, a consequence of the assumption that wages are determined by Nash bargaining.

*This is a major revision of 'Equilibrium Unemployment Fluctuations'. I thank Daron Acemoglu, Olivier Blanchard, V. V. Chari, Joao Gomes, Robert Hall, Dale Mortensen, Christopher Pissarides, Richard Rogerson, and numerous other seminar participants for comments. This material is based upon work supported by the National Science Foundation under grant number 0079345. 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.

I focus on two sources of shocks, changes in labor productivity and changes in the job destruction 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 (Hall 2003). 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 job creation pulls down the unemployment rate, moving the economy along a downward sloping Beveridge curve (unemployment-vacancy locus). But the increase in job creation 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 or vacancy rates.

An increase in the job destruction rate does not affect the relative price of unemployment and vacancies, and so leaves the vacancy-unemployment ratio essentially unchanged. Since the increase in job destruction reduces employment duration, the unemployment rate increases, and so therefore must vacancies. As a result, fluctuations in the job destruction rate induce a counterfactually positive correlation between unemployment and vacancies.

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

I establish the model's inability to generate realistic fluctuations in unemployment and vacancies in several steps. Section 2 presents the relevant business cycle facts: unemployment is strongly countercyclical, vacancies are equally strongly procyclical, and the correlation between the two variables is -0.9 at business cycle frequencies. I also show that the magnitude of fluctuations in the vacancy-unemployment ratio is large, with the ratio frequently rising or falling by 100 percent or more within a few years. On the other hand, average labor productivity is much more stable, rarely moving by more than 5 percent during the business cycle. These facts have been noted before, and so I focus my attention on an attempt to convince the reader that they are not measurement artifacts, but rather represent a real phenomenon that should be explained.

In Section 3, I perform comparative statics in a search and matching model with rich microeconomic heterogeneity but no aggregate uncertainty and no discounting. I allow individual productivity to follow an arbitrary stochastic process and I do not impose constant returns to scale on the matching function. I make only two critical assumptions: wages are determined by Nash bargaining, at least when a worker and firm first meet; and a free entry condition determines the number of vacancies. The model has four critical parameters: average labor productivity \bar{p} , the value of non-market activity z , the cost of maintaining an open vacancy c , and workers' bargaining power in wage negotiations β . I show that for a fixed value of the bargaining parameter, the vacancy-unemployment ratio is nearly proportional to $\frac{\bar{p}-z}{c}$, where the constant of proportionality depends on β . In particular, unless the percentage difference between average labor productivity and the value of non-market activity is small, so it does not much matter whether individuals work in the market, moderate movements in average labor productivity, in the value of non-market activity, or in the cost of a vacancy cannot generate even a small fraction of the observed fluctuations along the Beveridge curve. Countercyclical movements in workers' bargaining power appear to be a more promising source of fluctuations, but are difficult to interpret in an axiomatic bargaining model.

Section 4 shows that the comparative static results carry over to a stochastic model with aggregate uncertainty and discounting. I allow for 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 job destruction shocks raise the rate at which employed workers become unemployed but do not affect the productivity in surviving matches. This section also maintains the assumption of Nash bargaining over wages in new matches, but continues to allow for the possibility that a worker and firm

implicitly or explicitly agree to a long-term contract within the match. I specialize the model in some dimensions, imposing a constant returns to scale matching function and assuming that productivity is the same in all matches at any point in time. In other words, Section 4 studies a stochastic version of the Pissarides (1985) model.²

I first derive a forward-looking equation for the vacancy-unemployment ratio in terms of model parameters. I then calibrate the model to match the data along as many dimensions as possible. The calibration confirms the quantitative predictions of the comparative statics in the deterministic model. If the economy is hit only by productivity shocks, it moves along a downward sloping Beveridge curve, but matching the amplitude of the observed fluctuations requires introducing labor productivity shocks that are at least an order of magnitude larger than those in the data. Moreover, labor productivity is perfectly correlated with the vacancy-unemployment ratio, indicating that the model has almost no internal propagation mechanism. If the economy is hit only by job destruction shocks, the vacancy-unemployment ratio is stable in the face of large unemployment fluctuations, so vacancies are countercyclical. Equivalently, the model-generated Beveridge curve is upward-sloping.

Section 5 argues that the model's basic shortcoming is the Nash bargaining assumption, which implies that the wage in new jobs varies substantially in response to aggregate labor market conditions. I show this by considering a related centralized economy. 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.

Section 6 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 7 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.

²To my knowledge, this is the first paper that quantitatively explores a stochastic search and matching model with many aggregate states; for example, Mortensen and Pissarides (1994) assumes that there are only three states.

2 U.S. Labor Market Facts

This section discusses the time series behavior of unemployment, vacancies, and labor productivity 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 2001, 5.7 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 ratio of the unemployment rate to an extremely low frequency trend, a Hodrick-Prescott (HP) filter with smoothing parameter 10^5 using quarterly data. The ratio of the unemployment rate to its trend has a standard deviation of 0.19, so the unemployment rate is often as much as 38 percent above or below trend. Detrended unemployment also exhibits considerable persistence, with quarterly autocorrelation 0.92.

There is some question as to whether the unemployment rate 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 the unemployment rate 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 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 the unemployment rate is associated with an increase in the number of matches (Petrongolo and Pissarides 2001). If job search activity declined sharply with the unemployment rate, the matching function would be measured as decreasing in the unemployment rate. I conclude that aggregate job search activity is positively correlated with the unemployment rate.

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 during recession that began in March 2001 is instructive. During the first year of the survey, firms had 3.94 million job openings during an average month. This declined by 19 percent, to 3.19 million, in the second year,⁴ with the decline most noticeable (27 percent) when comparing the relatively strong six month period from December 2000 to May 2001 with the relatively weak period from December 2001 to May 2002. This suggests 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.⁵ A potential shortcoming is that help-wanted

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

⁴The decline in the help-wanted advertising index during the comparable period was somewhat larger, 27 percent.

⁵Abraham (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.

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 2 shows the help wanted advertising index and its trend.

Figure 3 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 -0.90 between 1951 and 2001.⁶ Moreover, the magnitude of the cyclical variation in unemployment and vacancies is almost identical, between 0.18 and 0.19, so the product of unemployment and vacancies is nearly acyclical. In other words, when the cyclical component of the unemployment rate falls from 6 to 5 percentage points, the cyclical component of vacancies rises by approximately 17 percent as well. The vacancy-unemployment ratio is therefore extremely procyclical, with a standard deviation of 0.35 around its trend.

An implication of the procyclicality of the vacancy-unemployment ratio is that it should be harder to find a job 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 unemployment rate 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 $\lambda(\theta) \equiv \frac{m(u, v)}{u} = m(1, \theta)$, where $\theta \equiv v/u$ is the vacancy-unemployment ratio. Procyclicality of vacancies is then equivalent to procyclicality of the job finding rate $\lambda(\theta)$.

Gross worker flow data can be used to measure the job finding 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). Alternatively, the job finding rate can be inferred from the dynamic behavior of the unemployment rate and average unemployment duration. Let d_t denote mean unemployment duration measured in months. Then assuming all unemployed workers find a job with probability $\lambda(\theta_t)$ in

⁶Abraham 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.

month t ,

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

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,

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

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. Out of steady state, I infer the job finding rate $\lambda(\theta_t)$ from the Bureau of Labor Statistics (BLS) monthly measurements of mean unemployment duration and the unemployment rate, based on the Current Population Survey. The correlation between cyclical component of $\lambda(\theta_t)$ and u_t in quarterly data is -0.89 , and the coefficient of variation on the job finding rate is 0.17 , only slightly smaller than the coefficient of variation on the unemployment rate, 0.19 . The job finding rate is indeed strongly procyclical.

The procyclicality of the job finding 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, while I focus on the *probability* that an individual enters or exits employment given her current employment state. Although the probability of entering employment declines sharply in recessions, this is almost exactly offset by the increase in the unemployment rate, so that the number of people finding jobs is essentially acyclic. Again, with an increasing matching function $m(u, v)$, this is only possible if procyclicality of vacancies offsets countercyclicality of unemployment.

2.3 Labor Productivity

The third important empirical observation is the weak procyclicality of labor productivity, measured as real output per hour 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 aggregate hours are 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 4 shows the behavior of labor productivity and Figure 5 compares the cyclical components of the vacancy-unemployment ratio and labor productivity. There is a positive correlation between the two time series and some evidence that labor productivity leads the vacancy-unemployment ratio by about one year, with a maximum correlation of 0.57. But the most important fact is that labor productivity is stable, only once moving more than five percent away from trend. In contrast, the vacancy-unemployment ratio has twice risen to double its trend level and five times fallen by at least fifty percent below trend, most recently at the end of 2001.

It is possible that the measured cyclicity 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 cyclicity 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 cyclicity, 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 Deterministic Model

I now examine whether search theory can reconcile the strong procyclicality of the vacancy-unemployment ratio with the weak procyclicality of labor productivity. The model I consider is a generalization of the textbook Mortensen-Pissarides matching model (Pissarides 1985, Mortensen and Pissarides 1994, Pissarides 2000).

The economy consists of a measure 1 of risk-neutral, infinitely-lived workers and a continuum of risk-neutral, infinitely-lived firms. Time is continuous and I focus throughout this section on steady states. I assume workers and firms discount future payoffs at a common rate $r > 0$, but focus on limiting results as $r \rightarrow 0$. 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 a wage but may not search. I discuss wage determination shortly. Firms have a constant returns to scale production technology that uses only labor. 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.

Let u denote the unemployment rate and v denote the measure of vacancies in the economy. The flow of meetings is given by a function $m(u, v)$,⁷ increasing in both arguments. The meeting function may exhibit decreasing, constant, or increasing returns to scale. An arbitrary unemployed worker finds a job according to a Poisson process with arrival rate $m(u, v)/u$ and an arbitrary vacancy contacts a worker according to a Poisson process with arrival rate $m(u, v)/v$.

When a worker and firm first meet, they realize an idiosyncratic match productivity level $p \sim F(p)$ with support on a subset of $[0, \hat{p}]$. Thereafter, productivity is hit by a shock with Poisson arrival rate $\delta(p) > 0$, with new productivity $p' \sim G(p'|p)$. I assume that for all $p \in [0, \hat{p}]$, $G(0|p) = 0$ and $G(\hat{p}|p) = 1$, so productivity always remains within these bounds. These productivity shocks may represent an underlying stochastic process for productivity. For example, Pissarides (1985) assumes that the initial productivity of a match is always equal to \hat{p} , while subsequent shocks leave the match with productivity 0. Mortensen and Pissarides (1994) assume new matches have productivity \hat{p} , while following a shock, the new productivity level is drawn from a distribution $G(p')$, independent of current productivity p . Alternatively, the shocks may represent a learning process about match quality (Pries 2001). For example, the initial expected productivity of a match may be at an intermediate level $p \in (0, \hat{p})$. After the first shock hits, the worker and firm realize the actual productivity, either \hat{p} or 0 with probability π and $1 - \pi$, respectively, where $p = \pi\hat{p}$. If productivity is equal to 0, the match is immediately destroyed. If it is equal to \hat{p} , the match continues until another (low probability) shock destroys it. This process can easily be generalized to allow for more gradual learning (Jovanovic 1979, Moscarini 2002).

There are generally bilateral gains from matching. I 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 no assumption about what happens to wages

⁷Not all meetings result in matches in this model, and so I distinguish between the ‘meeting’ and ‘matching’ functions. In the stochastic model in Section 4, all meetings result in matches, and so I use the terms interchangeably.

after this initial agreement. For example, the wage may be fixed forever at its initial value or it may be rebargained whenever the match is hit with a shock. Matches break up whenever the surplus falls to zero.

3.1 Characterization of Equilibrium

I provide a partial characterization of the equilibrium of this economy using Bellman equations. Let U denote the expected present value of an unemployed worker and $S(p)$ denote the expected present value of surplus in a match with productivity p . Then

$$rU = z + \frac{m(u, v)}{u} \int_0^{\hat{p}} \max\langle \beta S(p), 0 \rangle dF(p). \quad (1)$$

The flow value of an unemployed worker comes from her leisure z plus the probability that she meets a firm times the expected capital gain, which is obtained by integrating her share β of the match surplus $S(p)$ over the density of productivity in new positive surplus matches. If match surplus is negative, $S(p) < 0$, the meeting does not result in a match. Similarly, free entry implies

$$c = \frac{m(u, v)}{v} \int_0^{\hat{p}} \max\langle (1 - \beta)S(p), 0 \rangle dF(p). \quad (2)$$

For the sake of completeness, I also express the match surplus $S(p)$ recursively as

$$rS(p) = p - rU + \delta(p) \int_0^{\hat{p}} \max\langle -S(p), S(p') - S(p) \rangle dG(p'|p);$$

however, this equation is not used in the analysis below. The current surplus is equal to the amount produced p in excess of the worker's flow value if unemployed rU . The match is shocked at rate $\delta(p)$, in which case the new productivity level p' is drawn from the conditional distribution $G(p'|p)$. The match then ends if $S(p') < 0$, resulting in a capital gain $-S(p)$, or it continues if $S(p') > 0$, in which case the capital gain is $S(p') - S(p)$.

3.2 Productivity, Vacancies, and Unemployment

There is no analytic solution to the model at this level of generality. This section instead develops a key relationship between average labor productivity and the vacancy-unemployment ratio in a limiting case, $r \rightarrow 0$. Since in practice discounting accounts both for the rate of time preference and for the impermanence of any match, and

since empirically the average match dissolution rate is significantly higher than the discount rate, numerical simulations of the model show that setting r equal to zero is quantitatively unimportant.

Eliminate $m(u, v) \int_0^{\hat{p}} \max(S(p), 0) dF(p)$ from (1) using (2):

$$rU = z + \frac{\beta cv}{(1 - \beta)u} \equiv \bar{y} \quad (3)$$

Equation (3) has been noted before in special cases of this model, e.g. equation (1.19) in Pissarides (2000), but its quantitative implications appear to have been ignored. Consider hypothetically an unemployed worker who is offered a payoff y forever, in return for staying out of this economy. If the worker accepts the offer, the present value of her income is y/r , and so the offer is accepted if $y/r > U$, or equivalently, by equation (3), if $y > \bar{y}$; it is rejected if $y < \bar{y}$; and the worker is indifferent if $y = \bar{y}$. Note that this choice is unaffected by the discount rate, since \bar{y} is independent of r , and still holds in the special case without discounting, $r = 0$. But in this special case, the worker only cares about her average flow income.⁸ This implies that if $r = 0$, a worker's average flow income is \bar{y} . Since in an economy without discounting, free entry drives a firm's average flow income to zero,⁹ \bar{y} also represents aggregate flow income.

Another expression for aggregate flow income is

$$\bar{y} = uz + (1 - u)\bar{p} - vc, \quad (4)$$

where \bar{p} is average labor productivity. There are u unemployed workers, each of whom earns leisure z , $1 - u$ employed workers with average productivity \bar{p} , and v vacancies, each imposing a flow cost c . Eliminating \bar{y} between equations (3) and (4) gives:

$$\frac{\bar{p} - z}{c} = \frac{((1 - \beta)u + \beta)v}{(1 - \beta)(1 - u)u}. \quad (5)$$

To my knowledge, this reduction of a broad class of equilibrium search models to a single equation is novel.

Although it is not *a priori* obvious that comparative statics shed light on the dynamic behavior of a model, I put aside this concern temporarily and interpret equation (5) as if it says something about business cycles; numerical simulations of the

⁸Stronger criterion, e.g. the overtaking criterion, may be useful in this case, but all imply that if one income stream has a higher average flow payoff than another, the worker will prefer the former to the latter.

⁹Firms lose money when creating vacancies and earn profits later. If firms discount future profits, their average flow income must therefore be positive. But if $r = 0$, the gains and losses cancel out not only intertemporally but also in the cross-section when the economy is in steady state.

stochastic model in Section 4 will validate this presumption. I compute the right hand side of equation (5) quarter-by-quarter using historical U.S. data on unemployment and vacancies. The solid lines in Figure 6 plot the results for four different constant values of β . In each case, the right hand side is strongly procyclical and highly volatile. In fact, for values of β in excess of 0.1, it is roughly proportional to the vacancy-unemployment ratio, indicated by the dashed line in each figure,

$$\frac{v}{u} \approx \kappa \frac{\bar{p} - z}{c}, \quad (6)$$

where κ is a constant of proportionality. This implies that it takes large fluctuations in the difference between average labor productivity and the value of leisure, $\bar{p} - z$, or in the cost of a vacancy, c , in order to generate the observed movements in the vacancy-unemployment ratio. For example, if average labor productivity is typically double the value of leisure, an enormous five percent increase in average labor productivity raises $\bar{p} - z$ by ten percent and hence the vacancy-unemployment ratio moves by a similar amount. But since the coefficient of variation on the vacancy-unemployment ratio is 35 percent in U.S. data, the consequences of this shock for unemployment and vacancies would be scarcely discernible.

This problem goes away if average labor productivity is only a few percent larger than the value of leisure, but such a solution relies on an unappealing assumption that market and non-market activities are almost equally productive. Another possibility is that the value of leisure or the cost of a vacancy is volatile and countercyclical. Absent the possibility of directly measuring these quantities, this too seems an unsatisfactory resolution. I conclude from the comparative statics that fluctuations in labor productivity, the value of leisure, or the vacancy cost are not promising impulses for generating substantial variation in the vacancy-unemployment ratio in this model.

Alternatively, if labor productivity, the value of leisure, and the cost of a vacancy are constant, changes in workers' bargaining power β may induce fluctuations in vacancies and unemployment. Solving equation (5) for β gives

$$\beta = 1 - \frac{v}{(1-u)\left(v + u\frac{\bar{p}-z}{c}\right)}.$$

I again use actual data on unemployment and vacancies. This time, I assume that $\frac{\bar{p}-z}{c} = 1282$, a constant level so that workers' bargaining power is, on average from 1951 to 2001, equal to 0.5. Figure 7 shows the requisite bargaining power shocks, which are strongly negatively correlated with the vacancy-unemployment ratio. Superficially, this might seem to be an unlikely source of business cycle fluctuations; if

bargaining power is actually correlated with the business cycle, most people would guess that it is procyclical. Still, a serious evaluation of the possibility that bargaining power fluctuations are relevant at business cycle frequencies requires a deeper explanation of wage determination within the employment relationship. Section 5 argues that a reasonable modification to the search and matching model effectively induces countercyclical bargaining power in a reduced-form model.

To understand the separate behavior of unemployment and vacancies, and hence movements along the Beveridge curve, it is necessary to impose more structure on the model. Let $\bar{\delta}$ denote the average rate at which employed workers lose their job, ρ denote the fraction of meetings that result in matches, and assume that the meeting function exhibits constant returns to scale,¹⁰ $\frac{m(u,v)}{u} = m(1, \frac{v}{u}) \equiv \lambda(\frac{v}{u})$, with $\lambda(0) = 0$ and λ concave. Then in steady state,

$$\rho\lambda(\frac{v}{u})u = \bar{\delta}(1 - u). \quad (7)$$

Taking the approximation (6) as an exact relationship, we can solve for the unemployment and vacancy rates:

$$u = \frac{\bar{\delta}}{\bar{\delta} + \rho\lambda(x)} \text{ and } v = \frac{\bar{\delta}x}{\bar{\delta} + \rho\lambda(x)}, \quad (8)$$

where $x \equiv \kappa \frac{\bar{p}-z}{c}$. It is easy to confirm that the unemployment rate is decreasing in x and the vacancy rate is increasing in x for fixed values of $\bar{\delta}$ and ρ and for a fixed meeting function λ . This implies that an increase in labor productivity induces the observed downward sloping Beveridge curve. Indeed, if $\lambda(\frac{v}{u}) = \mu\sqrt{\frac{v}{u}}$, the steady state equation (7) tells us that $uv = (\bar{\delta}(1 - u)/\rho\mu)^2$, and hence is nearly constant in the face of variations in the composite parameter x , consistent with the empirical evidence. On the other hand, an increase in the average job destruction rate $\bar{\delta}$ or decrease in the matching probability ρ raises both the unemployment and vacancy rates if x is unchanged. Since the U.S. has never experienced a business cycle with an upward sloping Beveridge curve, the model suggests it is unlikely that such shocks are important at these frequencies.

It is worth pausing to ask which features of the search and matching model are behind the results summarized in equation (5). The most significant assumption is that the match surplus in new jobs is divided proportionally between the worker and firm via Nash bargaining. Consider the effect of a productivity increase. Because the

¹⁰Petrongolo and Pissarides (2001) argue that the existing body of empirical evidence suggests matching functions exhibit constant returns to scale.

supply of jobs is perfectly elastic, firms respond to an increase in profits by creating more vacancies. This raises unemployed workers' meeting rate and hence the value of unemployment, which is workers' threat point when bargaining. Firms anticipate having to pay higher wages and are reluctant to create many new jobs. In the end, most of the productivity increase accrues to workers in the form of higher wages, leaving little left over to shift the economy along the Beveridge curve. That is, wage flexibility in new jobs is the critical reason why the model does not generate large movements in the vacancy-unemployment ratio in response to a productivity shock.

4 Stochastic Model

This section generalizes the model to allow for aggregate shocks and discounting but specializes it by removing most of the microeconomic heterogeneity. I assume that at any point in time, all jobs have a common productivity $p > z$ and end exogenously at rate δ , the job destruction rate.¹¹ I further simplify the analysis by imposing constant returns to scale on the meeting function $m(u, v)$. Workers meet firms and accept jobs at rate $\frac{m(u, v)}{u} \equiv \lambda(\frac{v}{u})$ and firms meet and hire workers at rate $\frac{m(u, v)}{v} \equiv q(\frac{v}{u})$, both functions of the vacancy-unemployment ratio $\theta \equiv \frac{v}{u}$. Note that all meetings result in matches and so I sometimes refer to m as a 'matching', rather than a 'meeting', function.

I introduce aggregate shocks by allowing average labor productivity and the job destruction rate to follow a first order Markov processes.¹² A shock hits the economy according to a Poisson process with arrival rate s , at which point a new pair (p', δ') is drawn from a state dependent distribution. Let $\mathbb{E}_{p, \delta} X_{p', \delta'}$ denote the expected value of an arbitrary variable X following the next aggregate shock, conditional on the current state (p, δ) . I assume that this conditional expectation is finite, which is ensured if the state space is compact. Moreover, I impose that $p \geq z$ in every history, which guarantees that meetings always result in matches. At every point in time, the current values of productivity and the job destruction rate are common knowledge. The model is otherwise unchanged.

From the comparative statics exercise, I expect that an increase in average labor productivity p will raise the vacancy rate and reduce the unemployment rate, but

¹¹In terms of the deterministic model, the distribution of new jobs F puts all its weight on a single productivity level p , while the distribution of old jobs $G(p'|p)$ puts all its weight on $0 < z$, so old jobs are 'endogenously' destroyed following the first shock.

¹²Recall that average labor productivity may change due to a supply shock in a one-sector model or a demand shock in a multi-sector model (Hall 2003).

quantitatively by only a small amount. An increase in the job destruction rate δ will scarcely affect the vacancy-unemployment ratio because average labor productivity is unchanged. Instead, both unemployment and vacancies will increase.

4.1 Characterization of Equilibrium

I look for an equilibrium in which the vacancy-unemployment ratio depends only on the current value of p and δ , $\theta_{p,\delta}$.¹³ I characterize the equilibrium using a recursive equation for match surplus, the joint value to a worker and firm of being matched in excess of breaking up, as a function of the current aggregate state, $S_{p,\delta}$.

$$rS_{p,\delta} = p - (z + \lambda(\theta_{p,\delta})\beta S_{p,\delta}) - \delta S_{p,\delta} + s(\mathbb{E}_{p,\delta}S_{p',\delta'} - S_{p,\delta}). \quad (9)$$

Appendix A derives this equation from more primitive conditions. The first two terms represent the current flow surplus. 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 $\lambda(\theta_{p,\delta})$ of contacting a vacancy, in which event the worker would keep a fraction β of the match surplus $S_{p,\delta}$. Next, there is a flow probability δ that the match ends exogenously, destroying the surplus. Finally, an aggregate shock arrives at rate s , resulting in an expected capital gain $\mathbb{E}_{p,\delta}S_{p',\delta'} - S_{p,\delta}$.

Another critical equation for the match surplus 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 $1 - \beta$ of the match surplus $S_{p,\delta}$:

$$c = q(\theta_{p,\delta})(1 - \beta)S_{p,\delta}. \quad (10)$$

Eliminating current and future values of $S_{p,\delta}$ from (9) using (10) gives

$$\frac{r + \delta + s}{q(\theta_{p,\delta})} + \beta\theta_{p,\delta} = (1 - \beta)\frac{p - z}{c} + s\mathbb{E}_{p,\delta}\frac{1}{q(\theta_{p',\delta'})}, \quad (11)$$

which implicitly defines the vacancy-unemployment ratio as a function of the current state. 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

¹³It 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).

(1985) model appears to be new to the literature.

An equation for the evolution of the unemployment rate closes the model:

$$\dot{u}(t) = \delta(t)(1 - u(t)) - \lambda(\theta_{p(t),\delta(t)})u(t), \quad (12)$$

where $(p(t), \delta(t))$ is the aggregate state at time t . An initial condition pins down the unemployment rate and the aggregate state at some date t_0 .

In some special cases, equation (11) can be solved analytically. First, 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 equation (11) yields a static equation for the vacancy-unemployment ratio:

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

Moreover, if the interest rate is equal to zero, $r = 0$, this equation reduces to

$$\frac{p - z}{c} = \frac{(\beta + (1 - \beta)u_{p,\delta}^*)\theta_{p,\delta}}{(1 - \beta)(1 - u_{p,\delta}^*)},$$

where $u_{p,\delta}^* \equiv \frac{\delta}{\delta + \mu\theta_{p,\delta}}$ is the state-contingent steady state unemployment rate from equation (12). With this matching function, the comparative statics in the deterministic model (equation 5) are virtually identical to the nonstationary behavior of the stochastic model. I do not use this restriction on the matching function in the numerical work below, however, because it implies too little volatility in vacancies.

At the opposite extreme, suppose that each unemployed worker contacts a vacancy at a constant Poisson rate μ , independent of the vacancy rate, so $\lambda(\theta) = \mu$ and $q(\theta) = \mu/\theta$. Also assume that the job destruction rate δ is constant and average labor productivity p is a Martingale, $\mathbb{E}_p p' = p$. With this matching function, equation (11) is linear in current and future values of the vacancy-unemployment ratio:

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

It is straightforward to verify that the vacancy-unemployment ratio is linear in productivity, and therefore $\mathbb{E}_p \theta_{p'} = \theta_p$. In the limiting case of $r = 0$, this equation can be expressed as

$$\frac{p - z}{c} = \frac{(\beta + (1 - \beta)u^*)\theta_p}{(1 - \beta)(1 - u^*)},$$

where $u^* \equiv \frac{\delta}{\delta + \mu}$ is the steady state unemployment rate from equation (12). The relationship between average labor productivity and the vacancy-unemployment ratio once again coincides with the findings from the comparative statics of the deterministic model. But a shortcoming of this parameterization is that productivity shocks do not affect the job finding rate μ or job losing rate δ , and hence do not affect the unemployment rate. My numerical work uses an intermediate assumption on the elasticity of the matching function.

4.2 Parameterization

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 productivity and job destruction. Appendix B develops a discrete state space model which builds on a simple Poisson process corresponding to the theoretical analysis in Section 4.1. 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

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

where z 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, job destruction is constant and productivity satisfies $p = (1 - e^y)z + e^y p^*$, 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 job destruction satisfies $\delta = e^y \delta^*$, where again y follows an Ornstein-Uhlenbeck process and now $\delta^* > 0$ is a measure of long-run average job destruction. In both cases, the stochastic

process is reduced to three parameters, γ , σ , and either p^* or δ^* .

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,

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

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 job destruction rate δ , 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 set the value of leisure to $z = 0.4$. Interpreted as an unemployment benefit, this lies at the upper end of the range of income replacement rates in the United States.¹⁴ The deterministic model suggests that much higher values of z (i.e. $z \rightarrow p^*$) will significantly increase the impact of a productivity shock. I set the bargaining parameter to $\beta = 1/2$, a standard assumption in a literature that lacks any evidence to the contrary. This has little effect on the cyclical behavior of the model economy.

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. I use two data sources to pin down the job destruction rate δ . First, Abowd and Zellner (1985) find that 3.42 percent of employed workers exit employment during a typical month between 1972 and 1982, after correcting for classification and measurement error. Second, the Job Openings and Labor Turnover Survey constructs an employer-based measure of labor turnover. From December 2000 to November 2002, 3.36 percent of employed workers left their current employer in a typical month. Although some of these undoubtedly moved to another employer,¹⁵ it is reassuring how similar this number is to Abowd and Zellner's (1985) estimates. In the model with productivity shocks, I fix the job destruction rate at $\delta = 0.1$ per quarter.

I set the elasticity of the matching function with respect to vacancies at $\alpha = 1/2$, intermediate to the two special cases discussed at the end of Section 4.1. This generates roughly equal but opposite fluctuations in unemployment and vacancies in response to productivity shocks. I use the next two parameters, the matching function constant μ and the vacancy cost c , to pin down the average unemployment rate and the average

¹⁴Mean labor income in the model is 1.02.

¹⁵In an average month, more than half of the turnover, 1.88 percent of employed workers, quit their current job, while only 1.23 percent were laid off or discharged. The remaining workers left for other reasons, e.g. retirement or maternity leave. Other evidence suggests that many of the quits do not move directly to another job.

vacancy rate. The data indicate an unemployment rate just below six percent on average. I do not have direct evidence on vacancy rates, but fortunately the model offers one more normalization. Equation (11) implies that doubling c and multiplying μ by a factor 2^α divides the vacancy-unemployment ratio θ in half, doubles the worker-finding rate $q(\theta)$, and has no effect on the wage or the job finding rate $\lambda(\theta)$. In other words, the average vacancy rate is intrinsically meaningless. I choose to target a mean vacancy-unemployment ratio of 1, which requires setting $\mu = 1.7$ and $c = 0.54$.

I choose the standard deviation and persistence of the productivity process to match two final facts, the standard deviation of unemployment and the correlation between unemployment and vacancies, the Beveridge curve relationship. The link between the standard deviation of productivity and the standard deviation of unemployment is clear. The persistence of productivity affects the correlation between unemployment and vacancies because if productivity is less persistent, vacancies are more volatile, reducing the correlation with unemployment. To match the data, I set $\sigma = 0.161$ and $\gamma = 0.08$.¹⁶ Finally, I work on a discrete grid with 2001 points, which closely approximate Gaussian innovations.¹⁷ This implies that Poisson arrival rate of shocks is $s = n\gamma = 80$ times per quarter. Again, Appendix B discusses the relationship between the continuous and discrete state space models.

In the case of shocks to the job destruction rate, I change only the standard deviation of the stochastic process, reducing it to $\sigma = 0.076$. Productivity is constant and equal to 1, while the mean job destruction rate $\delta^* = 0.1$. Simulations show that this leaves the mean, standard deviation, and first autocorrelation of the unemployment rate virtually unchanged. For reasons already discussed and emphasized further below, it is impossible to match the correlation between unemployment and vacancies in the economy with job destruction shocks. Table 2 summarizes the parameter choices in the two simulations.

4.3 Results

This section examines the joint behavior of unemployment and vacancies in response to productivity and job destruction shocks. The main conclusion is that the analytic comparative static results closely approximate the numerical dynamic stochastic results.

¹⁶The quarterly autocorrelation of productivity is approximately $1 - \gamma = 0.92$, lower than in standard real business cycle models (Cooley and Prescott 1995).

¹⁷The results are remarkably similar using a much coarser grid, including only three points. Only third and higher moments are substantially affected by the discreteness of the state space. Details are available upon request.

I use equation (11) to find the state-contingent vacancy-unemployment ratio $\theta_{p,\delta}$ 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 stochastic mean reverting process described in Appendix B, and repeat. At the end of each period (quarter), I record the aggregate state and the unemployment rate. After 100,000 periods, I calculate a set of summary statistics. Table 3 reports the average results from 100 such simulations, i.e. a total of 10 million ‘quarters’ of data encompassing approximately 800 million Poisson shocks. The top panel depicts the effect of productivity shocks, while the bottom panel shows the effect of job destruction shocks.

The comovement of unemployment and vacancies in response to productivity shocks, shown in the top panel of Table 3, is largely consistent with the empirical evidence summarized in Table 1; however, recall that I chose parameters to match the mean and coefficient of variation for unemployment and the correlation between unemployment and vacancies. The model predicts that vacancies are slightly more volatile than unemployment, while the data suggests they have nearly the same volatility. This can be mended by increasing the share of vacancies in the matching function, α . A more serious concern lies with the persistence of the two variables. The model predicts unemployment is very persistent, with a quarterly first order autocorrelation of 0.96, while vacancies are much less persistent, with first order autocorrelation 0.84. In the data, both have an intermediate level of persistence, with vacancies slightly more persistent than unemployment. It is likely that anything that makes vacancies a state variable, such as planning lags or an adjustment cost in job creation, would increase its persistence and reduce its volatility, bringing the model more in line with the data.

The real problem with the model lies in the magnitude of the underlying productivity shocks. According to the data, the coefficient of variation of average labor productivity is about 1.8 percent. The stochastic model requires productivity shocks that are 14 times as large in order to generate the observed fluctuations in unemployment. The source of this lack of amplification can be understood in terms of equation (5) from the deterministic model. A one percent increase in average labor productivity raises the gap between market and non-market activity, $p - z$, by about $\frac{1\%}{\bar{p}-z} = 1.54$ percent, where $\bar{p} = 1.05$ is the mean level of productivity. Equation (5) suggests the response of the vacancy-unemployment ratio should be of a similar order of magnitude; in fact, the coefficient of variation on the vacancy-unemployment ratio is 1.65 times as large as the coefficient of variation on labor productivity. 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 and the vacancy-unemployment ratio is 1.00. In the data, the contemporaneous correlation is 0.40 and the vacancy-unemployment ratio lags labor productivity by about one year.

Turning this around, during the past 50 years, productivity has never increased by more than 5 percent above trend. According to the model, such an extreme event should increase the vacancy-unemployment ratio to perhaps 8 percent above trend. It follows that if the model is correct, almost all of the fluctuations in the vacancy-unemployment ratio during the post-war period were not a consequence of productivity shocks.

The bottom panel of Table 3 shows that job destruction shocks induce an almost perfectly positive correlation between unemployment and vacancies, an event that has essentially never been observed at business cycle frequencies (see Figure 2). As a result, job destruction shocks produce almost no variability in the vacancy-unemployment ratio. This is a dynamic reflection of equation (8) in the deterministic model. Assume there is no change in the constant x , which summarizes the behavior of average labor productivity, the value of leisure, the cost of a vacancy, and workers' bargaining power. Then an increase in the job destruction rate raises both unemployment and vacancies.

To summarize, the stochastic model confirms that job destruction shocks induce a positive correlation between unemployment and vacancies. It also confirms that, while productivity shocks are qualitatively consistent with a downward sloping Beveridge curve, the search model does not substantially amplify the extrinsic shocks and so productivity shocks induce only very small movements along the curve. These results are consistent with the intuition developed in the deterministic model.

4.4 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, $w_{p,\delta}$. 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 C proves that a continually renegotiated wage solves

$$w_{p,\delta} = (1 - \beta)z + \beta(p + c\theta_{p,\delta}). \quad (13)$$

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

Consider first the effect of a job destruction shock on the wage. An increase in the job destruction rate δ induces a slight decline in the vacancy-unemployment ratio (see Table 3), which in turn, by equation (13), 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 vacancy-unemployment ratio is small, and since the increase in job destruction raises the unemployment rate by shortening unemployment duration, the vacancy rate increases as well. In addition, equation (13) indicates that a job destruction shock will have little effect on the wage since it has little effect on the vacancy-unemployment 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 vacancy-unemployment ratio. Equation (13) 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.

If wages are 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 (where wage negotiation is required). It only changes the timing of wage payments. An increase in productivity or decrease in job destruction raises the present value of wage payments in new jobs and therefore has little effect on the vacancy-unemployment ratio.

5 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.¹⁸ Consider a hypothetical social planner who chooses a state-contingent vacancy-unemployment ratio in order to maximize the present discounted value of output net of vacancy creation costs. In the stochastic model, the planner's problem is represented recursively as

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

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 $\dot{u}(t) = \delta(1 - u(t)) - u(t)\lambda(\theta)$, and suddenly when an aggregate shock changes the state from (p, δ) to (p', δ') at rate s .

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

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

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

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

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

a special case of equation (11), 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 job destruction rates. Since the numerical example in Section 4.3 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 unem-

¹⁸A 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.

ployment and vacancies are perfect substitutes, i.e. $\lambda(\theta) = \alpha_u + \alpha_v\theta$, then the output-maximizing vacancy-unemployment ratio is infinite whenever $\alpha_v(p - z) > c(r + \delta + \alpha_u)$ and is zero if the inequality is reversed. With near-perfect substitutability, the output-maximizing vacancy-unemployment 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, $\lambda(\theta) = \min\langle \alpha_u, \alpha_v\theta \rangle$, the vacancy-unemployment ratio never strays from the efficient ratio $\frac{\alpha_u}{\alpha_v}$. With imperfect complements, the impact of productivity shocks on the vacancy-unemployment ratio is muffled but not eliminated.

The economics behind these theoretical findings is quite 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 vacancy-unemployment ratio barely changes. If they are strong substitutes, substitution is nearly costless, and the vacancy-unemployment 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. To my knowledge, no one else has examined this elasticity of substitution. Given the poor quality of the data Blanchard and Diamond (1989) use, one might be concerned that this result could be overturned. But taking it as a fact, optimal movements along the Beveridge curve should be slightly less pronounced than the those predicted by the model with Nash bargaining and much less pronounced than those in the data.

6 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 job destruction rate δ . These papers either impose exogenously or derive within the model a counterfactually constant vacancy-unemployment ratio θ . A third group of papers has tried but failed to reconcile the procyclicality of the vacancy-unemployment 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 job destruction 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 cyclicity of the vacancy-unemployment 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 vacancy-unemployment ratio, as exogenous.

The second group of papers, including work by Pries (2001), 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 job destruction rate, minimizing the role played by the observed cyclicity of the vacancy-unemployment ratio. These papers typically deliver rigid wages from a search model, consistent with the findings in Section 4.4. Building on the ideas in Hall (1995), Pries (2001) 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 vacancy-unemployment 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 job destruction rate are the only source of fluctuations in unemployment. The job finding rate $\lambda(\theta)$ 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 job destruction rate amplify productivity shocks in a model similar to the one examined here; however, they do not discuss the cyclical behavior of the vacancy-unemployment 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 vacancy-unemployment 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.¹⁹ This paper confirms that in response to such large shocks, the vacancy-unemployment 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 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 job destruction rate, which have the opposite effect. Merz (1995) and Andolfatto (1996) both put the standard search model into a real business cycle framework with risk aversion, 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

¹⁹This 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.

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 (2003) proposes one solution: sectoral demand shocks can generate large movements in real sectoral productivity and hence in the sectoral vacancy-unemployment ratio.

7 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 productivity shock primarily results in higher wages, with little effect on the vacancy-unemployment ratio. A job destruction 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 vacancy-unemployment ratio, helping to reconcile the evidence and theory.

If the matching function is Cobb-Douglas, the observed behavior of the vacancy-unemployment 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 existing estimate of a CES matching function suggests that the Cobb-Douglas assumption is approximately right, and so observed wages are ‘too rigid’.

One way to generate more rigid wages in a theoretical model is to drop some of the informational assumptions in the standard search model.²⁰ For example, suppose workers know about aggregate variables, including the unemployment rate and the aggregate productivity distribution F and G , but they do not know how productive they are in a particular job. Also assume workers can make take-it-or-leave-it wage demands, which firms accept if the worker asks for less than her productivity. On the margin, a worker faces a tradeoff between demanding a higher wage and reducing her risk of unemployment. As a result, an optimal wage demand depends on the hazard rate of the productivity distribution. Vacancy creation, on the other hand, depends

²⁰Ramey 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.

on the expected value of productivity in excess of workers' optimal wage demand. In such a model, a shift in the productivity distribution may change firms' incentive to create jobs while having little effect on wages, or vice versa. In other words, asymmetric information can break the link between the vacancy-unemployment ratio and the wage. A model in which firms cannot verify workers' outside opportunities, e.g. their value of leisure or alternative wage offers, delivers similar predictions. At this point, it is unclear whether either model delivers wage rigidity, in the sense that large wage changes are infrequently observed, or if they simply weaken the correlation between wages and the vacancy-unemployment ratio.

Another possibility is to modify the standard search model so as to make wages at least partially backward-looking. 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 labor turnover. Burdett and Mortensen show that this results in steady state wage dispersion even if all workers and jobs are identical. To my knowledge, no one has analyzed the out-of-steady state behavior of this model.²¹ Intuitively, wage offers are backward looking, because the cost of luring a worker away from her current employer depends on the existing wage distribution, and forward looking, because the likelihood that a worker quits depends on the wage offers she receives in the future. Both effects help keep wages low in expansions and high in recessions, although it will take further research to see whether this mechanism is quantitatively significant.

²¹But see Coles (2001) for work on a related model.

Appendix

A Derivation of the Equation for Surplus (9)

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

$$rU_{p,\delta} = z + \lambda(\theta_{p,\delta})(E_{p,\delta} - U_{p,\delta}) + s(\mathbb{E}_{p,\delta}U_{p',\delta'} - U_{p,\delta}) \quad (14)$$

$$rE_{p,\delta} = w_{p,\delta} - \delta(E_{p,\delta} - U_{p,\delta}) + s(\mathbb{E}_{p,\delta}E_{p',\delta'} - E_{p,\delta}) \quad (15)$$

$$rJ_{p,\delta} = p - w_{p,\delta} - \delta J_{p,\delta} + s(\mathbb{E}_{p,\delta}J_{p',\delta'} - J_{p,\delta}) \quad (16)$$

Equation (14) states that the flow value of an unemployed worker is equal to her value of leisure z plus the probability she finds a job $\lambda(\theta_{p,\delta})$ times the resulting capital gain $E - U$ plus the probability of an aggregate shock times that capital gain. Equation (15) expresses a similar idea for an employed worker, who receives a wage payment $w_{p,\delta}$ but loses her job at rate δ . Equation (16) 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 (15) and (16) and then subtract equation (14), defining $S_{p,\delta} \equiv J_{p,\delta} + E_{p,\delta} - U_{p,\delta}$:

$$rS_{p,\delta} = p - z - \lambda(\theta_{p,\delta})(E_{p,\delta} - U_{p,\delta}) - \delta S_{p,\delta} + s(\mathbb{E}_{p,\delta}S_{p',\delta'} - S_{p,\delta}). \quad (17)$$

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

$$\frac{E_{p,\delta} - U_{p,\delta}}{\beta} = S_{p,\delta} = \frac{J_{p,\delta}}{1-\beta}. \quad (18)$$

Substituting for $E - U$ in equation (17) yields equation (9).

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 S is unchanged and in particular is not history dependent, regardless of the

frequency of wage renegotiation.

B 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 s . 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 $\gamma \equiv s/n$ and $\sigma \equiv \sqrt{s}\Delta$. 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 $1 - hs$, has increased by Δ with probability $\frac{hs}{2} \left(1 - \frac{y}{n\Delta}\right)$, and has decreased by Δ with probability $\frac{hs}{2} \left(1 + \frac{y}{n\Delta}\right)$. Adding this together shows

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

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

$$\text{Var}(y(t+h) - y(t) | y(t)) = \mathbb{E}((y(t+h) - y(t))^2 | y(t)) - (\mathbb{E}(y(t+h) - y(t) | y(t)))^2.$$

The first term evaluates to $hs\Delta^2$ 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

$$\text{Var}(y(t+h) - y(t) | y(t)) = hs\Delta^2 \equiv h\sigma^2,$$

where $\sigma^2 = s\Delta^2$. 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, s, n) to $(\frac{\Delta}{x}, x^2 s, x^2 n)$ for any $x > 0$. It is easy to verify that this does not change either the autocorrelation parameter $\gamma = s/n$ or the instantaneous variance $\sigma = \sqrt{s}\Delta$. But as $x \rightarrow \infty$, the distribution of the innovation process x converges to a normal by the Central Limit Theorem. Equivalently, y converges to an Ornstein-Uhlenbeck process.²² 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 increasing x without substantially changing the results.

C Derivation of the Wage Equation

Assume that wages are continually renegotiated, so the wage only depends on the current aggregate state (p, δ) . Eliminate current and future values of J from equation (16) using equation (18)

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

Similarly, eliminate current and future values of S using (10):

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

Finally, replace the last two terms using equation (9) to get equation (13).

²²Notably, 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.
- COOLEY, THOMAS, AND EDWARD PRESCOTT (1995): “Economic Growth and Business Cycles,” in *Frontiers of Business Cycle Research*, ed. by Thomas Cooley. Princeton University Press, New Jersey.
- 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.

- 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*.
- (2003): “A Model of Aggregate Fluctuations,” 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.
- JOVANOVIC, BOYAN (1979): “Job Matching and the Theory of Turnover,” *Journal of Political Economy*, 87, 972–990.
- 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.
- 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.
- MOSCARINI, GIUSEPPE (2002): “Skill and Luck in the Theory of Turnover,” Mimeo.
- 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 Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages,” *American Economic Review*, 75, 676–690.
- (2000): *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, second edn.

- PRIES, MICHAEL (2001): “Persistence of Employment Fluctuations: A Model of Recurring Job Loss,” Forthcoming, *Review of Economic Studies*.
- 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.
- SOLON, GARY, ROBERT BARSKY, AND JONATHAN PARKER (1994): “Measuring the Cyclicalities 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.

Summary Statistics, quarterly U.S. data, 1951 to 2001

	p	u	v	$\frac{v}{u}$
Mean	—	0.0567	—	—
Coefficient of Variation	0.018	0.188	0.183	0.349
Autocorrelation (1 Quarter)	0.866	0.918	0.930	0.936
Correlation Matrix	p	1	-0.399	0.399
	u	—	1	-0.896
	v	—	—	1
	$\frac{v}{u}$	—	—	—

Table 1: Average labor productivity p is real average output per hour 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. The unemployment rate u is constructed by the BLS from the Current Population Survey. The help-wanted advertising index v is constructed by the Conference Board. Both u and v are quarterly averages of seasonally adjusted monthly series. Unemployment, vacancies, and productivity are expressed as ratios to an HP filter with smoothing parameter 10^5 . The coefficient of variation is the ratio of the standard deviation to the mean.

Parameter	Source of Shocks	
	Productivity	Job Destruction
productivity p	stochastic	1
job destruction rate δ	0.1	stochastic
discount rate r	0.012	0.012
value of leisure z	0.4	0.4
matching function $m(u, v)$	$1.7\sqrt{uv}$	$1.7\sqrt{uv}$
bargaining power β	0.5	0.5
cost of vacancy c	0.54	0.54
standard deviation σ	0.161	0.076
autoregressive parameter γ	0.080	0.080

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

Productivity Shocks

	p	u	v	θ	
Mean	1.050	0.0568	0.0564	1.071	
Coefficient of Variation	0.260	0.189	0.229	0.428	
Autocorrelation (1 quarter)	0.917	0.959	0.844	0.917	
Correlation Matrix	p	1	-0.904	0.976	1.000
	u	—	1	-0.898	-0.901
	v	—	—	1	0.975
	θ	—	—	—	1

Job Destruction Shocks

	δ	u	v	θ	
Mean	0.102	0.0571	0.0557	0.979	
Coefficient of Variation	0.192	0.186	0.165	0.0219	
Autocorrelation (1 quarter)	0.922	0.958	0.957	0.922	
Correlation Matrix	δ	1	0.978	0.971	-1.000
	u	—	1	0.999	-0.978
	v	—	—	1	-0.972
	θ	—	—	—	1

Table 3: Results from simulating the dynamic stochastic model. The text provides details on the stochastic process for productivity and for the job destruction rate. The coefficient of variation is the ratio of the standard deviation to the mean.

Quarterly U.S. Unemployment Rate and Trend, 1951–2001




Figure 1: The unemployment rate is a quarterly average of the seasonally adjusted monthly series constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey, survey home page <http://www.bls.gov/cps/>. The trend is an HP filter of the quarterly data with smoothing parameter 10^5 .

Quarterly U.S. Help Wanted Advertising Index and Trend, 1951–2001




Figure 2: 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/fred/data/employ/helpwant>. The trend is an HP filter of the quarterly data with smoothing parameter 10^5 .

Quarterly U.S. Beveridge Curve, 1951–2001




Figure 3: The unemployment rate is constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey. The help-wanted advertising index is constructed by the Conference Board. Both are quarterly averages of seasonally adjusted monthly series and are expressed as ratios to an HP filter with smoothing parameter 10^5 .

Quarterly U.S. Average Labor Productivity and Trend, 1951–2001




Figure 4: Real output per hour in the non-farm business sector, constructed by the Bureau of Labor Statistics' (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, 1951–2001




Figure 5: The unemployment rate is constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey. 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 hour in the non-farm business sector, constructed by the BLS from the National Income and Product Accounts and the Current Employment Statistics. The vacancy-unemployment ratio and labor productivity are expressed as ratios to an HP filter with smoothing parameter 10^5 .




Figure 6: $\frac{\bar{p}-z}{c} = \frac{(\beta+(1-\beta)u)v}{(1-\beta)(1-u)u}$ for four different values of β (solid line, left scale) and $\frac{v}{u}$ (dashed line, right scale). The unemployment rate u is constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey. The help-wanted advertising index v is constructed by the Conference Board. Both are quarterly averages of seasonally adjusted monthly series.

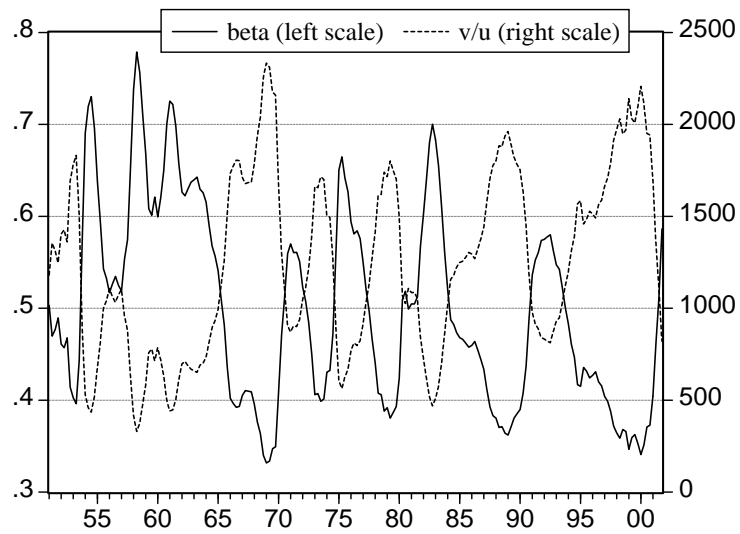


Figure 7: $\beta = 1 - \frac{v}{(1-u)(v+u\frac{\bar{p}-z}{c})}$ with $\frac{\bar{p}-z}{c} = 1282$ (solid line, left scale) and $\frac{v}{u}$ (dashed line, right scale). The unemployment rate u is constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey. The help-wanted advertising index v is constructed by the Conference Board. Both are quarterly averages of seasonally adjusted monthly series.