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

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

October 29, 2002

Abstract

In the United States, the unemployment rate is strongly countercyclical, vacancies are strongly procyclical, and wages are very mildly procyclical. This paper argues that these facts cannot be reconciled with standard models of search unemployment. Using both an analytical solution to a deterministic model and simulations of a stochastic model, the paper shows that search models predict a strong positive correlation between wages and the vacancy-unemployment ratio. Productivity shocks generate strongly procyclical vacancies and hence strongly procyclical wages; job destruction shocks generate weakly procyclical wages because vacancies are strongly countercyclical. Implicit contracts do not resolve this shortcoming. The paper discusses alternative means of wage determination that may be more capable of matching the business cycle facts.

*This is a major revision of a paper previously entitled 'Equilibrium Unemployment Fluctuations'. I thank seminar participants at the Minneapolis Fed, NBER Summer Institute, Princeton, St. Louis Fed, Universidad Di Tella, Wharton, and the CEP Macroeconomics of the Labour Market conference for comments on that paper, and especially Daron Acemoglu, Olivier Blanchard, V. V. Chari, Joao Gomes, Dale Mortensen, Christopher Pissarides, and Richard Rogerson. This material is based upon work supported by the National Science Foundation under Grant No. 0079345. I am grateful to the Alfred P. Sloan Foundation for financial support, to the Minneapolis Fed for its hospitality while I revised a previous version of this paper, and to 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 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 reconcile two central cyclical characteristics of the labor market: vacancies are volatile and strongly negatively correlated with unemployment, while real wages are rigid and only weakly correlated with the vacancy-unemployment ratio.

I make this argument in several steps, beginning with the empirical evidence and then considering the theoretical predictions of the model both analytically and numerically. In Section 2, I present the two business cycle facts. Both have been noted before, and so I attempt to keep the discussion of the data brief. I focus my attention on attempting to convince the reader that these facts 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 no aggregate uncertainty. The model I consider is a generalization of those in the literature. For example, I allow individual productivity to follow an arbitrary stochastic process and I do not impose constant returns to scale on the matching function. I show that if agents do not discount future income, the average wage can be expressed in closed form as a function of the vacancy-unemployment ratio and three parameters that are typically not thought of as business cycle impulses: the value of leisure, the cost of advertising a job vacancy, and workers' bargaining power. Moreover, the difference between wages and the value of leisure is proportional to the vacancy-unemployment ratio, so any change in the stochastic productivity process that raises the vacancy-unemployment ratio will raise the wage rate proportionately. This is the fundamental reason why the model cannot reconcile the weak procyclicality of real wages with the strong procyclicality of the vacancy-unemployment ratio.

I then examine the impact of different shocks on the vacancy-unemployment ratio. I show that any change in the stochastic productivity process that leaves the average level of productivity unchanged will have little effect on the vacancy-unemployment ratio and therefore on the average wage. If such shocks affect the unemployment rate, for example by raising the job destruction rate, they will induce a positive correlation between unemployment and vacancies, i.e. will counterfactually predict that vacancies are countercyclical. Conversely, I show that any shock that does not affect the fraction of meetings that result in matches or the average job destruction rate will induce a strong negative correlation between vacancies and unemployment and hence highly cyclical wages. For example, an increase in average productivity induces firms to create vacancies, resulting in a sharp increase in wages.

Section 4 shows that the comparative static results carry over to a stochastic model with aggregate shocks and discounting. I allow for two types of shocks: aggregate productivity shocks that raise the productivity of all matches and have no effect on the rate at which employed workers lose their job; and job destruction shocks that raise the rate at which employed workers become unemployed but have no effect on the productivity in surviving matches. I also specialize the model, 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, this is a stochastic version of the Pissarides (1985) model.[1]

I first derive a simple expression for the wage rate as a function of the current vacancy-unemployment ratio and other parameters and then derive a forward-looking equation for the vacancy-unemployment ratio in terms of model parameters. Next, I calibrate the model to match the data along as many dimensions as possible. The calibration confirms the intuition provided by the comparative statics in the deterministic model. If the economy is hit only by productivity shocks, it is easy to generate a strong negative correlation between unemployment and vacancies, but wages are about an order of magnitude more volatile than in the data. On the other hand, if the economy is hit only by job destruction shocks, wages do not vary very much over the business cycle, but this is because vacancies are countercyclical. In other words, the model cannot reconcile the strongly procyclicality of the vacancy-unemployment ratio with the weak procyclicality of wages.

One possible explanation for this finding is that wages are not continually renegotiated, as the benchmark model presumes. For example, while an aggregate shock may affect the wage in new employment relationships, implicit contracts smooth wages across these shocks in existing relationships. Section 5 shows that while this reduces the measured volatility of wages considerably, it does not resolve the basic problem. In the model with continual wage renegotiation, workers capture most of a productivity shock in the form of higher wages. Firms' profit, the difference between the present

[1] To my knowledge, this is the first paper that deals with a stochastic search and matching model with many aggregate states; for example, Mortensen and Pissarides (1994) and Cole and Rogerson (1999) assume that there are only two states.

value of productivity and of wages, changes very little, and so small productivity shocks have little effect on vacancy creation and hence on unemployment. I show that implicit contracts do not affect the present value of wages in new employment relationships, and so the relationship between productivity and vacancy creation is unchanged. As a result, either the model with or without implicit contracts requires extremely large productivity shocks, again approximately an order of magnitude larger than in the data, in order to generate the observed fluctuations in the vacancy-unemployment ratio.

A number of existing papers argue that, contrary to the conclusions of this study, standard search and matching models are consistent with the business cycle behavior of labor markets. Some look at the implications of the model for the behavior of various stocks and flows, for example the unemployment rate, the job destruction rate, and gross worker flows, but do not deal with wages. Another group of papers examines the implications of the model for wages, but ignore the empirical evidence that vacancies are procyclical, or equivalently, that the ease of gaining employment is procyclical. A third group of papers has tried to reconcile both facts, but has failed on one count or the other. None of these papers appears to have recognized the basic tension in the model that makes it impossible to reconcile the two facts.

Papers by Abraham and Katz (1986), Blanchard and Diamond (1989), Mortensen and Pissarides (1994), and Cole and Rogerson (1999) fit into the first category, matching the behavior of labor market stocks and flows by sidestepping the behavior of wages. For example, Abraham and Katz (1986) argue that the negative correlation between unemployment and vacancies is inconsistent with models in which unemployment is driven by fluctuations in the job destruction rate, notably the Sectoral Shifts model of Lilien (1982). That leads them to advocate an alternative model in which unemployment fluctuations are driven by aggregate disturbances, e.g. productivity shocks, although they never consider the wage implications of that model. Similarly, Blanchard and Diamond focus only on the dynamic equation for the evolution of unemployment. They do not model the supply of jobs, but rather assume that that it follows a slow-moving Markov process driven by aggregate shocks.[2] Their conclusion that such a model can generate a strong negative correlation between unemployment and vacancies is correct, but their modelling strategy prevents them from analyzing the behavior of wages in response to the substantial time-variation in the vacancy-unemployment ratio. On the other hand, the Mortensen and Pissarides (1994) model has implications

[2]More precisely, they assume the total supply of jobs is fixed, but an exogenous fraction of jobs is idle. The probability of an idle job becoming active and an active job becoming idle follows a two-state Markov process. Thus in response to an aggregate shock, the stock of active and idle jobs responds only gradually.

for wages, but in their numerical work, they focus exclusively on the cyclical behavior of job creation and destruction. Likewise, Cole and Rogerson (1999) argue that the Mortensen-Pissarides model can match a variety of business cycle facts, but they do so in a reduced form model that sidesteps the role of wages entirely.

The second group of papers, including work by Hall (1995), Pries (2001), Ramey and Watson (1997), and Gomes, Greenwood, and Rebelo (2001), argues that variants of the standard search and matching model can explain the observed real wage rigidity, but these papers ignore the procyclicality of vacancies. 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 and hence wages remain constant, since aggregate economic conditions, e.g. the productivity level and exogenous job destruction rate, have returned to normal. Equivalently, vacancies decline with unemployment 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 creation rate is exogenous and acyclic, which is equivalent to assuming that vacancies are proportional to unemployment. My analysis suggests that this is probably an important part of the explanation for why their model produces rigid wages. Similarly, Gomes, Greenwood, and Rebelo (2001) sidestep the vacancy-unemployment issue completely by looking at a model in which the job creation rate is exogenous and constant, i.e. vacancies are proportional to unemployment. This helps keep wages relatively rigid in their model.

Finally, there are two papers that have placed fairly standard search models into a real business cycle framework. Merz (1995) is unable to generate a statistically significant contemporaneous relationship between unemployment and vacancies but replicates the cyclicality of real wages (see her Table 3). Andalfatto (1996) also has trouble matching the negative correlation between unemployment and vacancies. Moreover, in his model the correlation between the real wage and output is 0.95, compared with 0.04 in the data. Thus these papers appear to encounter the problem I highlight in this paper, although they do not emphasize this shortcoming of the model.

Section 2 presents the main empirical evidence. Section 3 analyzes the deterministic model analytically. Section 4 solves the stochastic model numerically. Section 5 argues that implicit contract theory does not resolve the gap between theory and evidence.

Section 6 concludes by discussing alternative models of wage determination that might do a better job of matching the empirical evidence.

2 U.S. Labor Market Facts

This section discusses the time series behavior of unemployment, vacancies, and real wages in the United States. Table 1 summarizes the detrended data.

2.1 Unemployment

The unemployment rate is the most commonly used cyclical indicator of the amount of job search activity. In an average month from 1964 to 2001,[3] 5.9 percent of the U.S. labor force was out of work. This time series exhibits considerable temporal variation, falling as low as 3.4 percent in 1968 and 1969 and reaching as high as 10.8 percent in 1982 and 1983 (Figure 1). Some of these fluctuations are almost certainly unrelated to business cycles, the focus of this paper. To highlight short-term fluctuations, I take the ratio of the unemployment rate to an extremely low frequency trend, a Hodrick-Prescott (HP) filter with smoothing parameter 10^7. The ratio of the unemployment rate to its trend has a standard deviation of 0.17, so the unemployment rate is often as much as 34 percent above or below trend. Detrended unemployment also exhibits considerable persistence, with first order autocorrelation 0.97 in monthly data.

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 would imply a more cyclical measure of job search activity, worsening the problems that I

[3] I focus on data since 1964 because this is the longest time period during which a consistent monthly series for wages is available. Unemployment data are available back to 1948 and vacancy data to 1951.

highlight 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 decline in the number of job search methods used or switch towards towards less time-intensive methods. A preliminary examination of data using the Current Population Survey indicates no cyclical variation in the number or type of job search methods utilized. Indirect evidence comes from estimates of matching functions, recently reviewed by Petrongolo and Pissarides (2001), which universally find that an increase in the unemployment rate is associated with an increase in the number of matches. 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 the aggregate amount of job search is positively correlated with the unemployment rate and hereafter use the unemployment rate to proxy for the amount of job search activity.

2.2 Vacancies

The flip side of unemployment is job vacancies. The Job Openings and Labor Turnover Survey (JOLTS), a new establishment survey, provides an ideal empirical definition of vacancies: "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."[4] Unfortunately, JOLTS only began in December 2000. Comparable data has never previously been collected in the U.S..[5]

I instead use the best available proxy, the Conference Board help-wanted advertising index, measured as the number of help-wanted advertisements in 51 major news-

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

[5]JOLTS looks at the number of job openings on the last business day of the month. Although there is too little data to look at this time series systematically, its behavior during the 2001 recession is instructive. During the first eight months of the survey, December 2000 to July 2001, firms on average had 4.18 million job openings. A year later, this number declined by 24.2 percent to 3.17 million. Over the same period, the help-wanted advertising index fell slightly more, by 30.9 percent.

papers.[6] A potential shortcoming is that help-wanted advertising is subject to low frequency fluctuations that are only tangentially related to the labor market: the Internet may have reduced firms' reliance on newspapers as a source of job advertising; newspaper consolidation may have increased advertising in surviving newspapers; and Equal Employment Opportunity laws may have encouraged firms to advertise job openings more extensively. Fortunately, a low frequency trend should remove the effect of those and other secular shifts.

Figure 2 shows that the correlation between detrended unemployment and help-wanted advertising is -0.875 during this 35 year period. Moreover, the magnitude of variation in unemployment and vacancies is very similar, with both detrended series having a standard deviation of 0.17. In other words, when the cyclical 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 0.31 around its trend.

The negative correlation between unemployment and vacancies has been noted before. A scatter diagram of the inverse relationship between vacancies and unemployment is often called the 'Beveridge curve'. Abraham and Katz (1986) argued that the correlation is inconsistent with Lilien's (1982) sectoral shifts hypothesis, and instead indicates that business cycles are driven by aggregate fluctuations. In their article entitled 'The Beveridge Curve', Blanchard and Diamond (1989) conclude that at business cycle frequencies, shocks generally drive the unemployment and vacancy rates in the opposite direction.

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 unemployment to employment transition rate $\lambda(\theta)$. This variable can be measured directly using gross labor market flows data and is indeed procyclical; see Blanchard and Diamond (1990), Bleakley, Ferris, and Fuhrer (1999), or Abraham and Shimer (2001). The nonparticipation to employment transition rate

[6] See Abraham (1987) for a detailed discussion of this measure. 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.

is also procyclical.

The procyclicality of transition rates into employment 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, further support for the direct evidence on the cyclicality of the two variables.

2.3 Real Wages

The modest cyclicality of the real wage is a central fact in macroeconomics.[7] In this paper, I focus on the behavior of detrended average hourly earnings for production workers in private industry (hereafter 'AHE'), measured by the Bureau of Labor Statistics (BLS) establishment survey, the Current Employment Statistics. This is the only wage measure that is available at monthly frequencies. I also consider two quarterly measures of wages, average hourly compensation in the non-farm business sector constructed from the National Income and Product Accounts; and the Employment Cost Index, wages and salaries only, for private industry workers from the National Compensation Survey. I deflate all three wage series using both a measure of consumer prices, the Consumer Price Index for all urban consumers (CPI) and a measure of producer prices, the Producer Price Index for finished goods (PPI). The monthly data are detrended using an HP filter with smoothing parameter 10^7 and the quarterly data uses a comparable smoothing parameter, 10^5.

Figure 3 plots monthly AHE deflated by the CPI and PPI against the vacancy-unemployment ratio. There are two important results. First, real wages are substantially less variable than the vacancy-unemployment ratio. One measure of this is the coefficient of variation, the ratio of the standard deviation to the mean. For the vacancy-unemployment ratio, this is 0.31, while for the consumption wage, it is 0.02 and for the production wage it is 0.04. Second, the correlation between the vacancy-unemployment ratio and the wage measures is weak but positive, 0.33 and 0.32 for the

[7] Abraham and Haltiwanger (1995) recently reviewed the relevant empirical evidence.

consumption and production wage, respectively. This is often expressed in terms of a regression coefficient; after correcting for serial correlation, for example by running the regression in first differences, I find that a ten percent increase in the vacancy-unemployment ratio is associated with a 0.1 percent increase in both the consumption and production wage. Although statistically significant at conventional levels, this is economically a very small number.[8]

Average hourly compensation (AHC) is a significantly broader measure of wages than AHE. It includes the wage and salary information in AHE, as well as tips, bonuses, and in-kind payments. Moreover, this broad measure of compensation includes the imputed cost of proprietors' and unpaid family workers' labor services. An additional advantage of AHC data is that it is available at a relatively high frequency, quarterly, since 1951, the year in which the help-wanted advertising index was first constructed.[9] AHC is both less variable than AHE (coefficient of variation 0.02 for the consumption wage and 0.03 for the production wage) and less correlated with the vacancy-unemployment ratio (correlation coefficient 0.22 and 0.28). If anything, looking at this broader measure of compensation over a longer time period reduces the measured cyclicality of wages.

A common explanation for the near acyclicality of real wages is that this is due to a composition bias as the structure of employment naturally changes over the business cycle.[10] In part to address this concern, the BLS has since 1975 produced the employment cost index (ECI), essentially a Laspeyres price index with a fixed job composition.[11] If wages do not change in surviving jobs, the ECI will also not change, even if the composition of job loss is unequally distributed across income groups. Unfortunately, the ECI is only available since 1975. Despite the short time period and distinct method of producing the data, the results are familiar. The coefficient of variation is less than 0.03 and the correlation coefficient of wages and the vacancy-unemployment ratio is below 0.4. All of these data series together suggest that a ten percent increase in the vacancy-unemployment ratio is associated with at most a 0.2 percent increase in wages. Real wages are very modestly procyclical.

[8]Since all variables are expressed as ratios to their trend, i.e. have a mean approximately equal to 1, it is unimportant whether I run this regression in logs, levels, or some combination of the two.

[9]Average hourly earnings are available annually since 1951. Annual and monthly data yield similar results during the periods in which the two samples overlap.

[10]Solon, Barsky, and Parker (1994), probably the best known paper on composition bias, concludes that real wages are modestly procyclical. Translating their findings into the terms used here, a ten percent increase in the vacancy-unemployment ratio is associated with a quarter percent increase in real wages for workers who were employed during the downturn. Bils (1985) finds a smaller effect of composition on wages.

[11]See Ruser (2001) for an overview of the ECI and especially pages 11–12 for a comparison of the ECI and AHE measures over the business cycle.

3 Deterministic Model

I now examine whether search theory can reconcile the strong procyclicality of the vacancy-unemployment ratio with the weak procyclicality of real wages. 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 all agents discount future payoffs at rate $r > 0$. I focus throughout on steady states.

Workers can either be unemployed (unmatched) or employed in a match with flow productivity $p \in [0, \hat{p}]$. Unemployed workers get flow utility from leisure z and search for a job. Employed workers earn a flow wage $w(p)$ that depends on the productivity of their match, but cannot search. An employed worker may quit her job at will. She will typically do so if the match quality is sufficiently low, not only because she misses out on leisure, but also because she foregoes the opportunity to search for a better match.

Firms have a constant returns to scale production technology that uses only labor. If the firm employs a worker in a match with productivity p, the firm earns (expected) flow revenue p and pays the wage $w(p)$. The firm may fire the worker at will. In order to hire a worker, a firm must maintain an open vacancy at flow cost c. Firms create vacancies as long as there are profits to be made from doing so.

Let u denote the unemployment rate and v denote the measure of vacancies in the economy. The flow of matches in the economy is given by a matching function $m(u, v)$, increasing in both arguments. The matching 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 an unemployed worker and a vacancy meet, they realize the initial productivity of their match $p \sim F(p)$. Thereafter, match 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 $p \in [0, \hat{p}]$, $G_p(0) = 0$ and $G_p(\hat{p}) = 1$, so productivity always remains within these bounds.[12] 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 produc-

[12]Note that although this stochastic process is first-order Markov, it can approximate a higher order Markov process arbitrarily closely by embedding information about history in the remote decimal places of p. Indeed, the entire analysis carries over to the case of a higher-order Markov process only at the cost of notational complexity.

tivity 0. Mortensen and Pissarides (1994) assume new matches have productivity $\hat{p}$, while subsequent shocks occur at rate δ, with the new productivity level 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 the match. This process can easily be generalized to allow for more gradual learning.

There are generally bilateral gains from matching. I follow the bulk of the search literature and assume the gains from trade are divided according to a Nash bargaining solution. At any point in time, 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". The Nash bargaining assumption is central to my findings. I consider alternative assumptions in the conclusion.

3.1 Definition of Equilibrium

I characterize the equilibrium of the economy using Bellman equations. Let U denote the expected present value of an unemployed worker, $E(p)$ denote the expected present value of an employed worker in a match with productivity p, and $\mathbb{I}(p)$ be the probability that a match with productivity p is mutually acceptable. Then

$$rU = z + \frac{m(u,v)}{u} \int_0^{\hat{p}} \mathbb{I}(p)(E(p) - U)dF(p) \tag{1}$$

$$\text{and } rE(p) = w(p) + \delta(p) \int_0^{\hat{p}} \Big(\mathbb{I}(p')E(p') + \big(1 - \mathbb{I}(p')\big)U - E(p)\Big)dG_p(p'). \tag{2}$$

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 the increase in present value $E(p) - U$ over the density of productivity in new, mutually acceptable matches. An employed worker earns a wage $w(p)$ and faces a shock at rate $\delta(p)$, at which time she either becomes unemployed or accepts employment at the new productivity level p', drawn from $G_p(p')$. The match may be destroyed immediately following the shock, leaving the worker unemployed, or it may be maintained, leaving her with value $E(p')$. Similarly, let V denote the expected present value of a vacancy and $J(p)$ denote the expected present value of a job with productivity p:

$$rV = -c + \frac{m(u,v)}{v} \int_0^{\hat{p}} \mathbb{I}(p)(J(p) - V) dF(p) \tag{3}$$

$$\text{and } rJ(p) = p - w(p) + \delta(p) \int_0^{\hat{p}} \Big(\mathbb{I}(p')J(p') + \big(1 - \mathbb{I}(p')\big)V - J(p)\Big) dG_p(p'). \tag{4}$$

The logic behind these equations is similar, except a vacancy imposes a flow cost c while a filled job earns profits $p - w(p)$.

Four additional conditions close the model. First, free entry drives the value of a vacancy to zero:

$$V = 0. \tag{5}$$

Second, Nash bargaining implies that the wages are set to maximize the Nash product:

$$w(p) = \arg\max_{w(p)} (E(p) - U)^{\beta} (J(p) - V)^{1-\beta} \tag{6}$$

for all p. Third, matches are mutually acceptable if both $E(p) > U$ and $J(p) > V$ and unacceptable if one inequality is reversed:

$$\mathbb{I}(p) = \begin{cases} 1 & \text{if} \quad E(p) > U \text{ and } J(p) > V \\ 0 & \phantom{\text{if}} \quad E(p) < U \text{ or } J(p) < V \end{cases} \tag{7}$$

If $E(p) = U$ and $J(p) = V$, the acceptance probability may lie freely between 0 and 1. It is straightforward to verify that the Nash bargaining assumption (6) aligns the worker's and firm's preferences, i.e. $E(p) \gtreqless U$ if and only if $J(p) \gtreqless V$.

Finally, the unemployment rate is at its steady state value. Let $M(p)$ denote the cumulative distribution of productivity across workers, with $M(0) = u$, the unemployment rate, and $M(\hat{p}) = 1$. Then in steady state, the destruction of existing matches with productivity $p' < p$ must equal the creation of such matches:

$$\int_0^p \delta(p') dM(p') = \int_0^{\hat{p}} \delta(p'') \left(\int_0^p \mathbb{I}(p') dG_{p''}(p') \right) dM(p'') + m(u,v) \int_0^p \mathbb{I}(p') dF(p') \tag{8}$$

for all p. The left hand side is equal to the flow of existing matches with productivity $p' \in [0, p]$ that are hit by a shock. The first term on the right hand side is the flow of existing matches with arbitrary productivity $p'' \in [0, \hat{p}]$ that are hit by a shock and survive at a new productivity level $p' \in [0, p]$. The second term is the flow of new matches that begin with productivity $p' < p$. A special case of this equation gives the familiar 'job destruction equals job creation' relationship. Evaluate (8) at $p = \hat{p}$ to get

$$\int_0^{\hat{p}} \delta(p') \left(\int_0^{\hat{p}} (1 - \mathbb{I}(p''))dG_{p'}(p'') \right) dM(p') = m(u,v) \int_0^{\hat{p}} \mathbb{I}(p')dF(p')$$

The left hand side is the rate at which existing matches with productivity p' are hit by a shock and the new productivity level p'' leads to the termination of the match (job destruction). The right hand side is the rate at which meetings occur and generate a match that survives the initial shock (job creation).

Definition 1. *A steady state equilibrium is a set of Bellman values U, $E(p)$, V, and $J(p)$, an acceptance decision $\mathbb{I}(p)$, a wage function $w(p)$, and measures of vacancies v, unemployment u, and distribution of productivity $M(p)$ satisfying equations (1)–(8).*

3.2 Wages and the Vacancy-Unemployment Ratio

There is no analytic solution to the model at this level of generality. This section instead develops a key relationship between wages and the vacancy-unemployment ratio in a limiting case, $r = 0$. 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. Perhaps more surprisingly, the analytical comparative static results are nearly the same as the numerical results from simulations of the stochastic generalization of the model analyzed in Section 4 and thus provide insight into the mechanism driving the behavior of the stochastic model.

To begin, differentiate the logarithm of the Nash product (6) with respect to the wage $w(p)$ to get the first order condition

$$\frac{\beta}{E(p) - U} \frac{dE(p)}{dw(p)} + \frac{1-\beta}{J(p) - V} \frac{dJ(p)}{dw(p)} = 0.$$

Equations (2) and (4) imply $\frac{\partial E(p)}{\partial w(p)} = \frac{1}{r + \delta(p)} = \frac{-\partial J(p)}{\partial w(p)}$ and so

$$(1 - \beta)(E(p) - U) = \beta(J(p) - V). \tag{9}$$

Next, multiply equation (1) by $1-\beta$ and (3) by β and combine using equation (9):

$$\begin{aligned}(1-\beta)(rU-z)u &= m(u,v)\int_0^{\hat{p}} \mathbb{I}(p)(1-\beta)(E(p)-U)dF(p) \\ &= m(u,v)\int_0^{\hat{p}} \mathbb{I}(p)\beta(J(p)-V)dF(p) = \beta(rV+c)v\end{aligned}$$

Using the free entry condition (5), the outer set of equalities gives

$$rU = z + \frac{\beta cv}{(1-\beta)u} \tag{10}$$

Equation (10) 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 an unemployed worker who is offered a payoff y forever, in return for staying out of the labor market. If the worker accepts the offer, the present value of her income is y/r, and so equation (10) implies the offer is accepted if $y > \bar{y} \equiv z + \frac{\beta cv}{(1-\beta)u}$, 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.[13] This implies that if $r = 0$, a worker's average flow income is

$$\bar{y} = z + \frac{\beta cv}{(1-\beta)u}.$$

Her average flow income can also be expressed as a weighted average of her value of leisure z and the average wage she earns while employed $\bar{w}$. Since she is unemployed a fraction u of her lifetime, this gives

$$\bar{y} = uz + (1-u)\bar{w}. \tag{11}$$

Eliminating average flow income $\bar{y}$ between the previous equations gives an expression for the average wage in an economy without discounting:

$$\bar{w} = z + \frac{\beta cv}{(1-\beta)u(1-u)}. \tag{12}$$

Consider two economies that have different stochastic processes for productivity but the same value of leisure z, the same bargaining parameter β, and the same vacancy

[13]Stronger 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.

cost c. In general, the economies are likely to have different vacancy and unemployment rates v and u and hence, according to equation (12), a different average wage level.

Equation (12) has quantitative implications as well. Suppose that one economy has a four percent unemployment rate while the other has a five percent unemployment rate. The empirical evidence on vacancies in Section 2.2 suggests that the vacancy rate in the former economy should be 20 or 25 percent higher than in the latter economy, and so roughly speaking, the ratio $\frac{v}{u(1-u)}$ is approximately 50 percent higher in the low unemployment economy. Equation (12) then tells us that the difference between the average wage and the value of leisure, $\bar{w} - z$, must be 50 percent larger in the low unemployment economy as well. Unless average wages are very close to the value of leisure, i.e. workers are largely indifferent between working and staying home, this implies that there must be very large differences in average wages associated with modest differences in the unemployment rate generated by different stochastic processes for productivity.

3.3 Comparative Statics

Equation (12) provides a tight link between the behavior of average wages and the vacancy-unemployment ratio. If a change in parameters results in a proportional change in vacancies and unemployment, average wages will not change very much, while if it generates opposing movements in vacancies and unemployment, average wages will move strongly in the same direction as vacancies. Which comparative static one observes depends on which parameter changes. Making this statement precise requires a bit more algebra. Combine equations (2), (4), and (9):

$$\frac{w(p) - rU}{\beta} = \frac{r(E(p) - U)}{\beta} - \delta(p) \int_0^{\hat{p}} \frac{\mathbb{I}(p')(E(p') - U) - (E(p) - U)}{\beta} dG_p(p') =$$
$$\frac{r(J(p) - V)}{1 - \beta} - \delta(p) \int_0^{\hat{p}} \frac{\mathbb{I}(p')(J(p') - V) - (J(p) - V)}{1 - \beta} dG_p(p') = \frac{p - w(p) - rV}{1 - \beta}.$$

The outer equalities and the free entry condition (5) give a simple linear equation for the wage as a function of productivity:

$$w(p) = \beta p + (1 - \beta) rU. \tag{13}$$

Averaging across all jobs in the economy gives

$$\bar{w} = \beta \bar{p} + (1 - \beta) rU, \tag{14}$$

where $\bar{p}$ is the average level of productivity. Since by definition $\bar{y} \equiv rU$, eliminate $\bar{y}$ between equations (11) and (14) to get:

$$\bar{w} = \frac{\beta\bar{p} + (1-\beta)uz}{\beta + (1-\beta)u}, \tag{15}$$

i.e. in an economy without discounting, the average wage is a weighted average of the average level of productivity and the value of leisure. Moreover, at low unemployment rates, almost all of the weight falls on the first term. I can also obtain a relationship between average productivity and the vacancy-unemployment ratio. Eliminate $\bar{w}$ from (15) using (12):

$$\bar{p} = z + \frac{\big(\beta + (1-\beta)u\big)cv}{(1-\beta)u(1-u)}. \tag{16}$$

Now consider a change in the stochastic productivity process that does not alter average productivity. An example is a 'sectoral shift' (Lilien 1982), resulting in an increase in the dispersion of productivity without altering the mean. Equation (15) implies that the wage rate will be a modestly decreasing function of the unemployment rate. More precisely, the semielasticity of the average wage with respect to the unemployment rate when the unemployment rate is zero is

$$\left.\frac{d\log\bar{w}}{du}\right|_{u=0} = -\left(\frac{1-\beta}{\beta}\right)\left(1-\frac{z}{\bar{p}}\right),$$

assuming the average level of productivity $\bar{p}$ does not change. Unless workers' bargaining power β is very small, so workers gain little from participating in the market, this semielasticity is small. Equation (16) explains why wages are stable in response to such shocks. If average productivity is constant, the vacancy-unemployment ratio cannot vary much in response to the unemployment rate, counter to the empirical evidence.

On the other hand, equation (15) indicates that at low unemployment rates, the average wage and average productivity are nearly the same. In combination with equation (12), this implies that a change in the stochastic productivity process that raises average productivity raises both average wages and the vacancy-unemployment ratio. What happens to the components of this ratio is less obvious. Define $\rho \equiv \int_0^{\hat{p}} \mathbb{I}(p)dF(p)$, the steady state probability that a worker and firm choose to match when they meet. Also let δ denote the steady state rate of match destruction. If a change in the stochastic productivity process does not affect ρ and δ, then the steady state unemployment relationship $\rho m(u,v) = \delta(1-u)$ is also unaffected, which in turn implies that vacancies and unemployment must move in opposite directions. In the special case $m(u,v) = \mu(uv)^\alpha$, the product uv is proportional to $(1-u)^{1/\alpha}$, and so is nearly unaffected by average productivity. This accords with the empirical evidence on vacancies, that during an expansion, unemployment falls and vacancies rise by roughly equal amounts. But it is inconsistent with the empirical evidence on wages, that real wages are only mildly procyclical.

In summary, comparative statics suggest that a search and matching model can deliver a weak correlation between wages and the unemployment in response to changes in the stochastic productivity process, but only if average productivity is constant so vacancies are positive correlated with unemployment. It can deliver a negative correlation between vacancies and unemployment, but can generate volatility in this ratio only if average productivity, and hence average wages, are equally volatile.

4 Stochastic Model

This section generalizes the model to allow for aggregate shocks and discounting but specializes it to keep individual behavior simple. I assume that at any point in time, all jobs have a common productivity $p > z$ and end at rate δ, the job destruction rate. In 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 puts all its weight on $0 < z$, so old jobs are endogenously destroyed following the first shock. Both parameters of the stochastic process for productivity themselves follow a first order Markov processes. In particular, 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 the cumulative distribution function $H_{p,\delta}(p', \delta')$.[14] I assume that the set of aggregate states (p, δ) is contained in $[z, \hat{p}] \times [0, \bar{\delta}]$, which guarantees that meetings result in matches in every state of the world. At every point in time, the current values of productivity and the job destruction rate are common knowledge, and wages adjust so as to satisfy the Nash bargaining solution.

I further simplify the analysis by imposing constant returns to scale on the matching function $m(u, v)$.[15] The rate at which workers contact firms is $m(u, v)/u = m(1, v/u) \equiv \lambda(v/u)$ and the rate at which firms contact workers is $m(u, v)/v \equiv q(v/u)$, both functions of the vacancy-unemployment ratio $\theta \equiv v/u$. I look for an equilibrium in which

[14]As with idiosyncratic shocks in the deterministic model, this structure allows us to approximate a higher-order Markov process by embedding information about history in the remote decimal places of p. The restriction to first-order processes is done both for notational convenience and because the simulations suggest that there is no need to incorporate additional history dependence into the model.

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

the vacancy-unemployment ratio depends only on the current value of p and δ, $\theta_{p,\delta}$, and conjecture that there is no other equilibrium. In particular, θ is independent of the unemployment rate.

From the comparative statics exercise, one would expect productivity shocks to induce a negative correlation between vacancies and unemployment and high variability in wages, while job destruction shocks should cause a positive correlation between vacancies and unemployment and little variation in wages.

4.1 Definition of Equilibrium

As in the deterministic model, I characterize the equilibrium using Bellman equations. Now $U_{p,\delta}$ is the expected present value of an unemployed worker as a function of the current aggregate state and $E_{p,\delta}$ is the expected present value of an employed worker (in a job with productivity p) as a function of the aggregate state.

$$rU_{p,\delta} = z + s\big(\mathbb{E}_{p,\delta}U_{p',\delta'} - U_{p,\delta}\big) + \lambda(\theta_{p,\delta})\big(E_{p,\delta} - U_{p,\delta}\big) \tag{17}$$

An unemployed worker gets leisure z. An aggregate shock hits at rate s, changing the state to (p', δ'), drawn randomly from the distribution $H_{p,\delta}$. This results in an expected capital gain of $\mathbb{E}_{p,\delta}U_{p',\delta'} - U_{p,\delta}$, where the first term denotes the expected value of unemployment following the shock conditional on the current state (p, δ). Finally, the worker finds a job at rate $\lambda(\theta_{p,\delta})$, getting capital gain $E_{p,\delta} - U_{p,\delta}$. Note that the probability of the worker finding a job and an aggregate shock hitting at the same instant is zero.

Similar equations describe the value of an employed worker, a vacant job, and a filled job:

$$rE_{p,\delta} = w_{p,\delta} + s\big(\mathbb{E}_{p,\delta}E_{p',\delta'} - E_{p,\delta}\big) + \delta\big(U_{p,\delta} - E_{p,\delta}\big) \tag{18}$$

$$rV_{p,\delta} = -c + s\big(\mathbb{E}_{p,\delta}V_{p',\delta'} - V_{p,\delta}\big) + q(\theta_{p,\delta})\big(J_{p,\delta} - V_{p,\delta}\big) \tag{19}$$

$$rJ_{p,\delta} = p - w_{p,\delta} + s\big(\mathbb{E}_{p,\delta}J_{p',\delta'} - J_{p,\delta}\big) + \delta\big(V_{p,\delta} - J_{p,\delta}\big) \tag{20}$$

$w_{p,\delta}$ denotes the wage conditional on the current aggregate state. In each case, the flow value is the sum of three terms reflecting the current payoff, the probability of an aggregate shock times the resulting expected capital gain, and the probability of an idiosyncratic shock times that capital gain.

Next, free entry drives the value of a vacancy to zero in every state:

$$V_{p,\delta} = 0 \tag{21}$$

Finally, wages adjust so as to maximize the Nash product:

$$w_{p,\delta} = \arg\max_{w_{p,\delta}} (E_{p,\delta} - U_{p,\delta})^{\beta} (J_{p,\delta} - V_{p,\delta})^{1-\beta} \tag{22}$$

Definition 2. *An equilibrium is a set of Bellman values U, E, V, and J, a wage w, and a vacancy-unemployment ratio θ in each state (p, δ) that satisfy equations (17)–(22).*

In equilibrium, the unemployment rate evolves according to

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

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

4.2 Characterization of Equilibrium

As in the deterministic model, Bellman equations (18) and (20) and the Nash bargaining assumption (22) imply workers and firms split the match surplus so that

$$(1 - \beta)(E_{p,\delta} - U_{p,\delta}) = \beta(J_{p,\delta} - V_{p,\delta}). \tag{23}$$

Rewrite (18), use the Nash bargaining solution (23) evaluated both at (p, δ) and an arbitrary future (p', δ'), and then rewrite (20) to get

$$\begin{aligned}
\frac{w_{p,\delta} - (r+s)U_{p,\delta} + s\mathbb{E}_{p,\delta}U_{p',\delta'}}{\beta} &= (r + \delta + s)\left(\frac{E_{p,\delta} - U_{p,\delta}}{\beta}\right) - s\mathbb{E}_{p,\delta}\left(\frac{E_{p',\delta'} - U_{p',\delta'}}{\beta}\right) \\
&= (r + \delta + s)\left(\frac{J_{p,\delta} - V_{p,\delta}}{1-\beta}\right) - s\mathbb{E}_{p,\delta}\left(\frac{J_{p',\delta'} - V_{p',\delta'}}{1-\beta}\right) \\
&= \frac{p - w_{p,\delta} - (r+s)V_{p,\delta} + s\mathbb{E}_{p,\delta}V_{p',\delta'}}{1-\beta}.
\end{aligned}$$

Solve the outer equalities for $w_{p,\delta}$, simplifying with the free entry condition (21), to get an expression for the wage as a linear combination of current productivity and current and expected future flow values of unemployment:

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

This is a natural generalization of equation (13) in the deterministic model. Eliminate $\mathbb{E}_{p,\delta}U_{p',\delta'}$ between this expression and (17):

$$w_{p,\delta} = \beta p + (1-\beta)\big(\lambda(\theta_{p,\delta})\big(E_{p,\delta} - U_{p,\delta}\big) + z\big).$$

The Nash bargaining solution (23) then implies

$$w_{p,\delta} = \beta\big(p + \lambda(\theta_{p,\delta})(J_{p,\delta} - V_{p,\delta})\big) + (1-\beta)z.$$

Eliminate $J_{p,\delta}$ between this equation and Bellman equation (19) and simplify with the free entry condition (21) to get an equation for the wage as a function of the current productivity level and the current vacancy-unemployment ratio.

$$w_{p,\delta} = \beta\big(p + \theta_{p,\delta}c\big) + (1-\beta)z, \tag{24}$$

where the final simplification comes from $\lambda(\theta) \equiv \theta q(\theta)$. This generalizes equation (1.20) in Pissarides (2000) to a stochastic environment. Note that if shocks to the job destruction rate do not affect the vacancy-unemployment ratio very much, wages will also be unresponsive to job destruction shocks. On the other hand, productivity shocks will affect both the wage rate both directly and through the vacancy-unemployment ratio.

To solve for the vacancy-unemployment ratio, use equation (24) to eliminate $w_{p,\delta}$ from Bellman equation (20):

$$(r + \delta + s)J_{p,\delta} = (1-\beta)(p - z) - \beta\theta_{p,\delta}c + s\mathbb{E}_{p,\delta}J_{p',\delta'}.$$

Bellman equation (19) and the free entry condition (21) imply $J_{p,\delta} = c/q(\theta_{p,\delta})$, hence

$$\mathbb{E}_{p,\delta}J_{p',\delta'} = c\mathbb{E}_{p,\delta}\left(\frac{1}{q(\theta_{p',\delta'})}\right).$$

Substituting these into the preceding equation gives a forward looking equation for θ.

$$\frac{r+\delta+s}{q(\theta_{p,\delta})} + \beta\theta_{p,\delta} - (1-\beta)\frac{p-z}{c} = s\mathbb{E}_{p,\delta}\left(\frac{1}{q(\theta_{p',\delta'})}\right). \tag{25}$$

Note that the left hand side of this equation is an increasing function of the current vacancy-unemployment ratio θ, while the right hand side depends on expected future values of the same variable. This equation can easily be solved numerically. Once the state-contingent vacancy-unemployment ratio is determined, equation (24) defines the state-contingent wage. All that remains is to calculate the ergodic distribution of the aggregate state vector (p, δ) and the unemployment rate.

In a special case of some independent economic interest, equation (25) can be solved analytically. Suppose that each vacancy contacts an unemployed worker at a constant Poisson rate μ, independent of the unemployment rate. That is, $q(\theta) = \mu$ and $\lambda(\theta) = \mu\theta$. Given the risk-neutrality assumptions, this is equivalent to assuming that firms must pay a fixed cost $k \equiv c/\mu$ in order to hire a worker. The vacancy-unemployment ratio, or equivalently workers' hiring rate, adjusts so that the bargained wage leaves the revenue from contacting a worker equal to the cost. If the vacancy-unemployment ratio is too high, workers demand a high wage (equation 24), leaving the present value of a new job lower than the hiring cost, and conversely if the vacancy-unemployment ratio is too low. With this restriction on the matching function, equation (25) reduces to

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

In other words, the vacancy-unemployment ratio, and hence workers' job contact rate, is determined entirely by static considerations. Moreover, if the interest rate is equal to zero, $r = 0$, this equation can be written as

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

where $u^*_{p,\delta}$ is the stationary unemployment rate associated with a particular value of p and δ, i.e. the solution to $\delta(1 - u^*_{p,\delta}) \equiv \mu\theta_{p,\delta}u^*_{p,\delta}$. This is essentially the same as equation (16) in the deterministic model. In other words, with this matching function, the comparative statics in the deterministic model are virtually identical to the nonstationary behavior of the stochastic model. I do not use this restriction in the numerical work below, however, because it generates too little cyclical volatility in vacancies.

4.3 Parameterization

In this section, I choose parameter values to match the time series behavior of the U.S. unemployment rate. The most important question is the choice of the stochastic process for productivity and job destruction. I model each as a discrete state space approximation to a first order autoregressive process with Brownian motion innovations. More precisely, consider a random variable $y(t)$ that is hit with shocks according to a Poisson process with arrival rate s. When a shock hits, the new value y' is determined from y according to:

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

Here $\Delta > 0$ represents the step size and $n > 0$ the number of steps before $y(t)$ hits the boundary of admissible values. In addition, I assume that the initial value $y(0)$ is one of a discrete number of points, $y(0) \in Y \equiv \{-n\Delta, -(n-1)\Delta, \ldots, 0, \ldots, (n-1)\Delta, n\Delta\}$.

This stochastic process has some convenient properties (see Appendix A for proofs). First, at any date t, $y(t) \in Y$, confirming the state space is discrete. Second, define $\gamma \equiv s/n$ and $\sigma \equiv \sqrt{s}\Delta$. At any date t and for any initial condition $y_0 \in Y$,

$$y(t) = e^{-\gamma t} y_0 + \varepsilon(t),$$

where $\varepsilon(t)$ is distributed with mean zero and variance $\frac{(1-e^{-2\gamma t})\sigma^2}{2\gamma}$. In other words, γ measures mean reversion in this stochastic process, with $\gamma = 0$ corresponding to a random walk, while σ^2 is the instantaneous variance. The unconditional distribution of $y(t)$, i.e. the limit as t converges to infinity, has mean zero and variance $\frac{\sigma^2}{2\gamma} = \frac{n\Delta^2}{2}$.

Third, 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 $\left(\frac{1}{x}\Delta, x^2 s, x^2 n\right)$ for any $x > 0$. It is easy to verify that this does not change either the autocorrelation parameter γ or the instantaneous variance σ^2. But as $x \to \infty$, the distribution of innovations $\varepsilon(t)$ converges to a normal. Equivalently,

$$y(t) = e^{-\gamma t} y_0 + \sigma \int_0^t e^{-\gamma(t-\tau)} dZ(\tau),$$

where Z is a standard Brownian motion. 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. For sufficiently large values of x, $y(t)$ is very nearly a first order autoregressive process with Brownian motion innovations.[16]

I consider two different simulations. In the first, productivity is stochastic and job destruction is deterministic. More precisely, $\tilde{p}(t) \equiv \log \frac{p(t)-z}{\bar{p}-z}$ follows a discrete state space autoregressive process. This ensures that $p(t) > z$ for all t while allowing for the

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

possibility that $p(t)$ rises substantially above that level. In the second, job destruction is stochastic and productivity is deterministic. In this case, I assume $\tilde{\delta}(t) \equiv \log \frac{\delta(t)}{\delta}$ follows a discrete state space autoregressive process. Again, the logarithmic formulation ensures that the job destruction rate remains nonnegative while allowing for substantial intertemporal variability in the parameter.

Consider first 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 mean level of productivity $\bar{p}$, the value of leisure z, the 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 two parameters of the stochastic process, the standard deviation σ and the mean reversion parameter γ. I use empirical evidence to motivate the choice of these parameters.

Without loss of generality, I normalize the mean productivity level to $\bar{p} = 1$.[17] 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, since the mean wage rate in the model is 1.02. I set the bargaining parameter to $\beta = 1/2$, a standard assumption in a literature that lacks any evidence to the contrary. The deterministic model suggests that much lower values of β or much higher values of z (i.e. $z \to \bar{p}$) will significantly reduce wage cyclicality.

I normalize a time period to be one month, and therefore set the discount rate to $r = 0.004$, 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, a new Bureau of Labor Statistics program, the Job Openings and Labor Turnover Survey, constructs an employer-based measure of labor turnover. From December 2000 to July 2002, 3.36 percent of employed workers left their current employer in a typical month.[18] Although some of these undoubtedly moved to another employer,[19] it is reassuring how similar this number is to Abowd and Zellner's (1985) estimates. In the model with productivity

[17] That is, $\tilde{p} = \log \frac{p-z}{1-z}$ has mean zero. The mean of p is approximately 1.05.

[18] This number is not seasonally adjusted due to the small amount of available data.

[19] In an average month, more than half of the turnover, 1.88 percent of employed workers, quit their current job, while only 1.21 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.

shocks, I fix the job destruction rate at $\delta = 0.034$ per month.

I set the elasticity of the matching function with respect to vacancies at $\alpha = 1/2$. 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 vacancy rate. The data indicates 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 (25) 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 $\lambda(\theta)$. I choose to target a mean vacancy-unemployment ratio of 1, which requires setting $\mu = 0.566$ and $c = 0.57$.

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 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 $\sigma = 0.106$ and $\gamma = 0.04$. Note that the quarterly autocorrelation of productivity is 0.88, lower than in standard real business cycle models (Cooley and Prescott 1995). Finally, I set the number of grid points on each side of the mean, n, to 1000, so as to closely approximate Gaussian innovations. This implies that Poisson arrival rate of shocks is $s = n\gamma = 40$ times per month. I also consider the effects of using a coarser grid, $n = 1$, i.e. three grid points, one above the mean, one below the mean, and one equal to the mean, so there is on average one large shock every $1/s = 25$ months.

In the case of shocks to the job destruction rate, I change only the standard deviation of the stochastic process, reducing it to $\sigma = 0.050$. Productivity is constant and equal to 1, while the mean job destruction rate $\bar{\delta} = 0.034$. 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 these parameter choices in the two simulations.

4.4 Results

This section examines the joint behavior of wages, vacancies, and unemployment 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.

I use equation (25) 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 stochastic mean reverting process described above, and repeat. At the end of each period (month), 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 'months' of data encompassing approximately 400 million Poisson shocks with $n = 1000$. The top panel depicts the effect of productivity shocks, while the bottom panel shows the effect of job destruction shocks. In Table 4, I reduce the grid size to $n = 1$, with virtually no effect on the results.[20]

The response of unemployment and vacancies to productivity shocks, shown in the top panels of Tables 3 and 4, is consistent with the empirical evidence summarized in Table 1. In particular, 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 α in the matching function. A more serious concern lies with the persistence of the two variables. The model predicts unemployment is very persistent, with a monthly first order autocorrelation of 0.99, while vacancies are much less persistent, with first order autocorrelation 0.90. In the data, both have an intermediate level of persistence, with vacancies slightly less persistent than unemployment. It seems likely that anything that increases the persistence of vacancies, 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.

[20] The grid size affects the higher moments of the variables of interest. For example, reducing n from 1000 to 1 lowers both the skewness and kurtosis of the time series distribution of unemployment. In response to productivity shocks, the unemployment rate is skewed right (skewness 0.405) and has fat tails (kurtosis 3.25) when $n = 1000$. It is less skewed (0.211) and has thin tails (2.04) when $n = 1$. The data indicates that the unemployment rate is more skewed (0.621) and has approximately normal tails (2.94). The response of the higher moments to job destruction shocks is similar.

The real problem with the model lies in the response of wages to productivity shocks. According to the data, wages are significantly less volatile than unemployment or vacancies and only weakly correlated with those cyclical indicators. The stochastic model predicts that in response to productivity shocks, wages are more volatile than either variable, highly correlated with both, and almost perfectly correlated with the vacancy-unemployment ratio. The behavior of wages in the stochastic model is also quantitatively consistent with the predictions of the deterministic model. Equation (12) suggests that a one percent increase in the vacancy-unemployment ratio should raise the difference between the wage and the value of leisure by about 1 percent. In the calibration, the value of leisure $z = 0.4$ is fixed at just below 40 percent of the average wage $w = 1.01$, so the wage should increase by 0.6 percent in response to a one percent increase in the vacancy-unemployment ratio. In the simulations of the stochastic model, the coefficient of variation on wages is indeed 60 percent of the coefficient of variation on the vacancy-unemployment ratio.

Turning this around, suppose the largest increase in the real wage during the past 37 years was due to a productivity shock. This event occurred in 1972, when the production wage rose 12 percent above trend.[21] A regression of log wages on the log vacancy-unemployment ratio in the model implies a semi-elasticity of approximately 1, and so the productivity shock should also have raised the vacancy-unemployment ratio by about 12 percent above trend. It follows that if the model is correct, almost all of the fluctuations in the vacancy-unemployment ratio during this period were not a consequence of productivity shocks.[22]

The bottom panels of Tables 3 and 4 show that job destruction shocks induce an almost perfectly positive correlation between unemployment and vacancies, an event that is essentially never observed (see Figure 2). This is a dynamic reflection of equation (16) in the deterministic model, which shows that at a given level of productivity, vacancies are almost perfectly correlated with unemployment when the unemployment rate is low. The strong positive correlation between unemployment and vacancies implies the vacancy-unemployment ratio fluctuates very little, with a coefficient of variation in the

[21]Even this event appears to be a consequence of the detrending procedure. The rise in commodity prices during the subsequent years sharply increased the producer price index and lowered real producer wages. The detrending procedure measures this as an increase in the real producer wage in advance of this shock. Note that the consumption wage never rose more than seven percent above trend.

[22]This argument does not preclude the possibility that there are coincident productivity and labor supply shocks, e.g. a strong negative correlation between productivity p and the value of leisure z. A decline in p reduces employment and lowers wages while an increase in z reduces employment and raises wages, potentially making wages relatively stable. The problem with this argument is that there is no reason to think there should be a negative correlation between these two shocks.

model equal to 6 percent of its value in the data. And this implies that wages also fluctuate very little in the model. For example, equation (24) implies that if there are no fluctuations in productivity, the coefficient of variation of wages is $\beta c = 0.27$ times the coefficient of variation of the vacancy-unemployment ratio.

To summarize, the stochastic model confirms that wages fluctuate little in response to job destruction shocks, but this is because the model generates counterfactually countercyclical vacancies. It also confirms that wages fluctuate too much in response to productivity shocks, although those shocks generate strongly procyclical vacancies. These results are insensitive to the choice of grid size.

5 Implicit Contracts

One explanation for why it might be difficult to reconcile the model with the data is that implicit contracts smooth the timing of wage payments across cyclical fluctuations (Baily 1974, Azariadis 1975). In the model, I assume workers are risk-neutral and so have no desire to smooth the timing of wage payments. If in reality workers are risk averse and capital markets are imperfect, firms may partially ensure workers against aggregate fluctuations by decreasing the link between wages and the state of the aggregate economy. At its most extreme, wages might be fixed at the beginning of an employment relationship and cannot change even if one party wishes to terminate the match; in this case, matches are no longer mutual tenant-at-will.

Although such implicit contracts complicate the algebra of the model, it is straightforward to show that they have no allocational effects. I distinguish an employment relationship both based upon its initial productivity when the wage was set, (p, δ), and on its current productivity, (p', δ'). With this new notation, the recursive equations become:

$$rU_{p,\delta} = z + s\big(\mathbb{E}_{p,\delta} U_{p',\delta'} - U_{p,\delta}\big) + \lambda(\theta_{p,\delta})\big(E^{p,\delta}_{p,\delta} - U_{p,\delta}\big) \tag{26}$$

$$rE^{p,\delta}_{p',\delta'} = w_{p,\delta} + s\big(\mathbb{E}_{p',\delta'} E^{p,\delta}_{p'',\delta''} - E^{p,\delta}_{p',\delta'}\big) + \delta\big(U_{p',\delta'} - E^{p,\delta}_{p',\delta'}\big) \tag{27}$$

$$rV_{p,\delta} = -c + s\big(\mathbb{E}_{p,\delta} V_{p',\delta'} - V_{p,\delta}\big) + q(\theta_{p,\delta})\big(J^{p,\delta}_{p,\delta} - V_{p,\delta}\big) \tag{28}$$

$$rJ^{p,\delta}_{p',\delta'} = p' - w_{p,\delta} + s\big(\mathbb{E}_{p',\delta'} J^{p,\delta}_{p'',\delta''} - J^{p,\delta}_{p',\delta'}\big) + \delta\big(V_{p',\delta'} - J^{p,\delta}_{p',\delta'}\big). \tag{29}$$

As usual, I close the model with the free entry condition $V_{p,\delta} = 0$ and the Nash bargaining solution. Now, however, I require only that the Nash product is maximized in new matches. It will be convenient to define the surplus function $S^{p,\delta}_{p',\delta'} \equiv E^{p,\delta}_{p',\delta'} + J^{p,\delta}_{p',\delta'} - U_{p',\delta'} - V_{p',\delta'}$, the sum of the gains the worker and firm get from trade as a function of the initial productivity (p, δ) and the current productivity (p', δ'). Then Nash bargaining implies

$$\frac{E_{p,\delta}^{p,\delta} - U_{p,\delta}}{\beta} = S_{p,\delta}^{p,\delta} = \frac{J_{p,\delta}^{p,\delta} - V_{p,\delta}}{1-\beta}. \tag{30}$$

Moreover, by adding together Bellman equations (27) and (29), one finds that the surplus function depends only on the current productivity level. Suppressing the redundant superscript:

$$(r + \delta + s)S_{p,\delta} = p - (r+s)U_{p,\delta} + s\mathbb{E}_{p,\delta}U_{p',\delta'} + s\mathbb{E}_{p,\delta}S_{p',\delta'}, \tag{31}$$

since $V_{p,\delta} = \mathbb{E}_{p,\delta}V_{p',\delta'} = 0$.

To solve this system of equations, eliminate the value of employment from (26) using the Nash bargaining solution (30) and eliminate the value of a job from (28) using (30) and the free entry condition:

$$\begin{aligned}(r+s)U_{p,\delta} - s\mathbb{E}_{p,\delta}U_{p',\delta'} &= z + \lambda(\theta_{p,\delta})\beta S_{p,\delta} \\ \text{and } c &= q(\theta_{p,\delta})(1-\beta)S_{p,\delta}.\end{aligned}$$

Use the first equation to eliminate the current and future value of unemployment U from (31) and the second equation to eliminate the current and future surplus S from (31), yielding

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

which is identical to equation (25). In other words, implicit contracts do not affect the vacancy-unemployment ratio conditional on the current state of the economy. They do, however, alter the wage setting rule, which is now forward-looking, unlike equation (24). The reason the change in wage setting does not affect the vacancy-unemployment ratio is because implicit contracts do not alter the expected present value of wages in new employment relationships. Together with the expected present value of productivity, this variable determines firms' incentive to create new jobs.

If implicit contracts make wages more rigid without changing the behavior of vacancies and unemployment, they might appear to solve the shortcomings of the model. Unfortunately, this is not the case. With or without implicit contracts, it takes a big productivity shock in order to generate a moderate movement in the vacancy-unemployment ratio because most of a productivity gain is extracted by workers in the form of wages and does not result in additional job creation. Qualitatively, we can see this in equation (16), which shows that in the deterministic model, the difference between productivity and the value of leisure is very nearly proportional to the vacancy-unemployment ratio. The numerical analysis of the stochastic model bears this intuition out. Table 3 indicates that in order to generate fluctuations in the unemployment rate of the magnitude observed in the U.S., the coefficient of variation of productivity must equal 0.20, i.e. productivity often lies forty percent above or below its trend.

Compare this with the empirical variability of productivity. I measure the average product of labor as output per hour in the non-farm business sector (APL). This is collected by the Bureau of Labor Statistics as part of the National Income and Product Accounts and is available quarterly since before 1951, when vacancy data first became available. I compare the average product of labor with a comparable measure of wages, average hourly compensation, discussed previously in Section 2.3. Table 5 shows that the average product of labor is slightly more variable than average hourly compensation and slightly more strongly correlated with the vacancy-unemployment ratio, but it is not nearly as variable in the data as in the model with productivity shocks. I conclude that implicit contracts do not reconcile the theory and evidence.

6 Conclusion

I have argued in this paper that a search and matching model in which wages are determined by Nash bargaining cannot reconcile the strong procyclicality of the vacancy-unemployment ratio with the weak procyclicality of wages and productivity. Productivity shocks generate strongly procyclical wages, while job destruction shocks generate countercyclical vacancies. Both of these findings are counterfactual. 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 job creation, helping to reconcile the evidence and theory.

The relevant question is therefore 'why are wages so rigid'? Unfortunately, I do not have a satisfactory answer at this time. One thought is that perhaps a rigid wage structure is socially optimal. A search equilibrium with Nash bargaining over wages is generally not Pareto optimal because agents do not take into account the positive and negative externalities from their search process. For example, by creating a vacancy, firms make it easier for workers to find jobs and harder for other firms to find workers. Perhaps the difficulty in matching theory with evidence is a manifestation of these externalities.

An examination of a 'social planner's' problem rejects that hypothesis for empirically plausible parameterizations of the model. The social planner chooses a state-contingent vacancy level in order to maximize the present discounted value of output net of vacancy creation costs. In the stochastic model, the planner's problem can be written recursively as

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

Instantaneous output is equal to z times the unemployment rate u plus p times the employment rate minus c times the number of vacancies $v \equiv u\theta$. The value changes gradually as the unemployment rate adjusts, with $\dot{u}(t) = \delta(1-u(t)) - u(t)\lambda(\theta)$, 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

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

independent of the unemployment rate. In particular, with a Cobb-Douglas matching function $m(u,v) = \mu u^{\alpha} v^{1-\alpha}$, this implicit definition of $\theta_{p,\delta}$ reduces to

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

This is a special case of equation (25), with workers' bargaining power β equal to 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 4.4 assumed a Cobb-Douglas matching function with $\alpha = \beta$, the equilibrium allocation described in that section solves the social planner's problem. Conversely, if those parameter values describe the U.S. economy, the observed degree of wage rigidity is inconsistent with output maximization.

With other matching functions, the link between the equilibrium with wage bargaining and the solution to the planner's problem is broken. At one extreme, if unemployment and vacancies are perfect substitutes, i.e. $\lambda(\theta) = \alpha_u + \alpha_v \theta$, 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. More generally, if unemployment and vacancies are nearly perfect substitutes, the output-maximizing vacancy-unemployment ratio is very sensitive to current productivity. On the other hand, if unemployment and vacancies are less substitutable than in the Cobb-Douglas case, the impact of productivity shocks is even more muffled. Which case is relevant is an empirical issue. Blanchard and Diamond (1989) use nonlinear least squares to estimate a Constant Elasticity of Substitution 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. From this I conclude that for empirically plausible parameter values, the observed behavior of unemployment, vacancies, and productivity are inconsistent with output-maximizing behavior.

A number of papers examine a 'competitive' search economy, in which firms can commit to wages before hiring workers and can increase their hiring rate by promising higher wages (Peters 1991, Montgomery 1991, Moen 1997, Shimer 1996, Burdett, Shi, and Wright 2001). It is by now well-known that a competitive search equilibrium maximizes output, essentially by creating a market for job applications. The discussion of output maximizing search behavior therefore also implies that the observed behavior of unemployment, vacancies, and productivity is inconsistent with a competitive search equilibrium if the matching function is Cobb-Douglas.[23]

Dropping some of the informational assumptions in the standard search model appears to be a more promising line of research.[24] 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. Job creation, on the other hand, depends on the expected value of productivity in excess of workers' optimal wage demand. In such a model, a

[23]In the previous version of this paper, "Equilibrium Unemployment Fluctuations", I analyze a deterministic competitive search model with a CES matching function more thoroughly. I also present additional evidence that matching functions are Cobb-Douglas. Details are available on request.

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

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 (1998) 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. 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 future wage offers. Both effects may help keep wages low in expansions and high in recessions, although it will take further research to see whether this mechanism is quantitatively significant.

A Properties of the Stochastic Process

Consider a random variable $y(t)$ that is hit with shocks according to a Poisson process with arrival rate s. When a shock hits, the new value y' is determined from y according to:

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

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 of two Poisson shocks arriving is negligible, and so $y(t+h)$ is equal to y 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

$$\mathbb{E}\left(y(t+h)|y(t)\right) = \left(1 - \frac{hs}{n}\right) y(t) \equiv (1 - h\gamma) y(t),$$

where $\gamma \equiv s/n$. Next, note that the variance of $y(t+h) - y(t)$ given $y(t)$ is

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

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

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

where $\sigma^2 = s\Delta^2$.

Equivalently, it is straightforward to verify that for any t and initial condition $y_0 \in Y$, one can represent $y(t)$ as

$$y(t) = e^{-\gamma t} y_0 + \sigma \int_0^t e^{-\gamma(t-\tau)} d\nu(\tau),$$

where $\nu(\tau)$ is a random variable with mean zero and variance τ for all $\tau \in [0,t]$. This in turn implies that the conditional expected value of $y(t)$ is $e^{-\gamma t} y_0$ and that the conditional variance is $\sigma^2 \int_0^t e^{-2\gamma(t-\tau)} d\tau \equiv \frac{(1-e^{-2\gamma t})\sigma^2}{2\gamma}$. Finally, convergence to Brownian motion as the arrival rate of shocks increases follows from the Central Limit Theorem.

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.

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

Azariadis, Costas (1975): "Implicit Contracts and Underemployment Equilibria," *Journal of Political Economy*, 83, 1183–1202.

Baily, Martin (1974): "Wages and Employment under Uncertain Demand," *Review of Economic Studies*, 41, 37–50.

Bils, Mark (1985): "Real Wages over the Business Cycle: Evidence from Panel Data," *Journal of Political Economy*, 93, 666–689.

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.

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.

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.

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

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

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

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

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

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

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

MORTENSEN, DALE, AND CHRISTOPHER PISSARIDES (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61, 397–415.

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

PETRONGOLO, BARBARA, AND CHRISTOPHER PISSARIDES (2001): "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, 39(2), 390–431.

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

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

PRIES, MICHAEL (2001): "Persistence of Employment Fluctuations: A Model of Recurring Job Loss," Mimeo.

RAMEY, GARY, AND JOEL WATSON (1997): "Contractual Fragility, Job Destruction, and Business Cycles," *Quarterly Journal of Economics*, 112(3), 873–911.

RUSER, JOHN (2001): "The Employment Cost Index: What is it?," *Monthly Labor Review*, pp. 3–21.

SHIMER, ROBERT (1996): "Contracts in a Frictional Labor Market," Mimeo.

SOLON, GARY, ROBERT BARSKY, AND JONATHAN PARKER (1994): "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?," *Quarterly Journal of Economics*, 109, 1–25.

Summary Statistics, monthly U.S. data, 1964 to 2001

		u	v	$\frac{v}{u}$	AHE/CPI	AHE/PPI
	Mean	0.0594	—	—	—	—
	Coefficient of Variation	0.171	0.174	0.312	0.024	0.037
	Autocorrelation (1 month)	0.970	0.976	0.981	0.979	0.985
Correlation Matrix	u	1	-0.875	-0.948	-0.349	-0.383
	v	—	1	0.954	0.298	0.257
	$\frac{v}{u}$	—	—	1	0.330	0.317
	AHE/CPI	—	—	—	1	0.802
	AHE/PPI	—	—	—	—	1

Table 1: 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. Nominal average hourly earnings for production workers in private industry (AHE) is constructed by the BLS from the Current Employment Statistics. It is deflated either by the consumer price index for all urban consumers (CPI) or the producer price index for finished goods (PPI). Unemployment, vacancies, and real wages are expressed as ratios to an HP filter with smoothing parameter 10^7. The coefficient of variation is the ratio of the standard deviation to the mean. The autocorrelation is the first order autocorrelation.

Parameter	Source of Shocks Productivity	Job Destruction
productivity p	stochastic	1
job destruction rate δ	0.034	stochastic
discount rate r	0.004	0.004
value of leisure z	0.4	0.4
matching function $m(u,v)$	$0.566\sqrt{uv}$	$0.566\sqrt{uv}$
bargaining power β	0.5	0.5
cost of vacancy c	0.570	0.570
standard deviation σ	0.106	0.050
autoregressive parameter γ	0.040	0.040
grid size n	1000 or 1	1000 or 1

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

		u	v	$\frac{v}{u}$	w	p
Mean		0.0593	0.0557	0.998	1.006	1.044
Coefficient of Variation		0.169	0.215	0.388	0.234	0.240
Autocorrelation (1 month)		0.991	0.898	0.958	0.958	0.958
Correlation Matrix	u	1	-0.873	-0.905	-0.906	-0.908
	v	—	1	0.970	0.971	0.971
	$\frac{v}{u}$	—	—	1	1.000	1.000
	w	—	—	—	1	1.000
	p	—	—	—	—	1

Job Destruction Shocks

		u	v	$\frac{v}{u}$	w	δ
Mean		0.0597	0.0549	0.922	0.963	0.035
Coefficient of Variation		0.171	0.151	0.020	0.005	0.179
Autocorrelation (1 month)		0.990	0.990	0.961	0.961	0.960
Correlation Matrix	u	1	0.999	-0.967	-0.967	0.967
	v	—	1	-0.957	-0.957	0.957
	$\frac{v}{u}$	—	—	1	1	-1.000
	w	—	—	—	1	-1.000
	δ	—	—	—	—	1

Table 3: Results from simulating the dynamic stochastic model with $n = 1000$. The text provides details on the stochastic process for productivity and for the job destruction rate.

Productivity Shocks

	u	v	$\frac{v}{u}$	w	p
Mean	0.0594	0.0557	0.997	1.006	1.043
Coefficient of Variation	0.168	0.212	0.371	0.224	0.230
Autocorrelation (1 month)	0.991	0.896	0.960	0.960	0.960
Correlation Matrix u	1	-0.882	-0.937	-0.938	-0.938
v	—	1	0.974	0.974	0.974
$\frac{v}{u}$	—	—	1	1.000	1.000
w	—	—	—	1	1.000
p	—	—	—	—	1

Job Destruction Shocks

	u	v	$\frac{v}{u}$	w	δ
Mean	0.0597	0.0549	0.922	0.963	0.035
Coefficient of Variation	0.169	0.150	0.020	0.005	0.176
Autocorrelation (1 month)	0.990	0.990	0.960	0.960	0.960
Correlation Matrix u	1	0.999	-0.967	-0.967	0.967
v	—	1	-0.958	-0.958	0.957
$\frac{v}{u}$	—	—	1	1	-1.000
w	—	—	—	1	-1.000
δ	—	—	—	—	1

Table 4: Results from simulating the dynamic stochastic model with $n = 1$. The text provides details on the stochastic process for productivity and for the job destruction rate.

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

		$\frac{v}{u}$	Compensation AHC/CPI	Compensation AHC/PPI	Productivity APL/CPI	Productivity APL/PPI
Coefficient of Variation		0.349	0.017	0.029	0.020	0.032
Correlation Matrix	$\frac{v}{u}$	1	0.219	0.284	0.353	0.380
	AHC/CPI	—	1	0.710	0.654	0.533
	AHC/PPI	—	—	1	0.523	0.866
	APL/CPI	—	—	—	1	0.769
	APL/PPI	—	—	—	—	1

Table 5: 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 monthly series. Average Hourly Compensation (AHC) and Average Output per Hour (APL) in the non-farm business sector are constructed by the BLS for the National Income and Product Accounts. Both measures are deflated both by the consumer price index for all urban consumers (CPI) or the producer price index for finished goods (PPI). Unemployment, vacancies, real wages, and real 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.

Figure 1: The unemployment rate is constructed by the Bureau of Labor Statistics from the Current Population Survey. The trend is an HP filter with smoothing parameter 10^7.

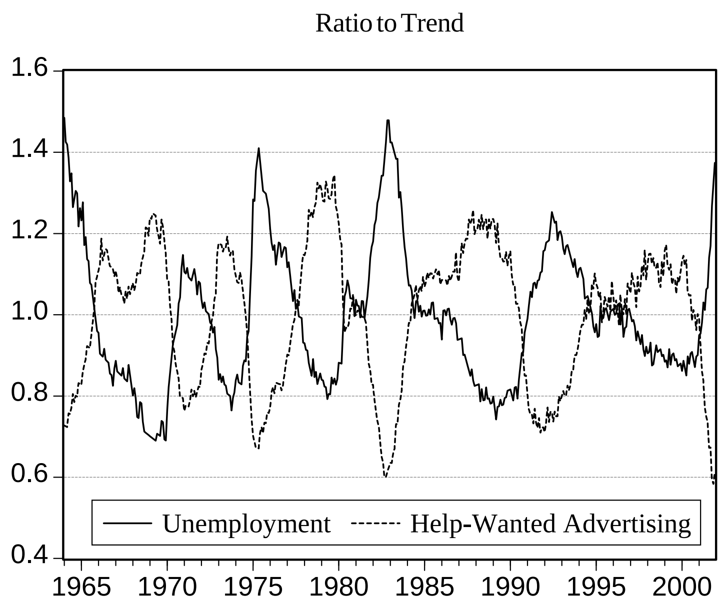

Figure 2: The unemployment rate is constructed by the Bureau of Labor Statistics from the Current Population Survey and the help-wanted advertising index is constructed by the Conference Board. Both variables are expressed as ratios to an HP filter with smoothing parameter 10^7.

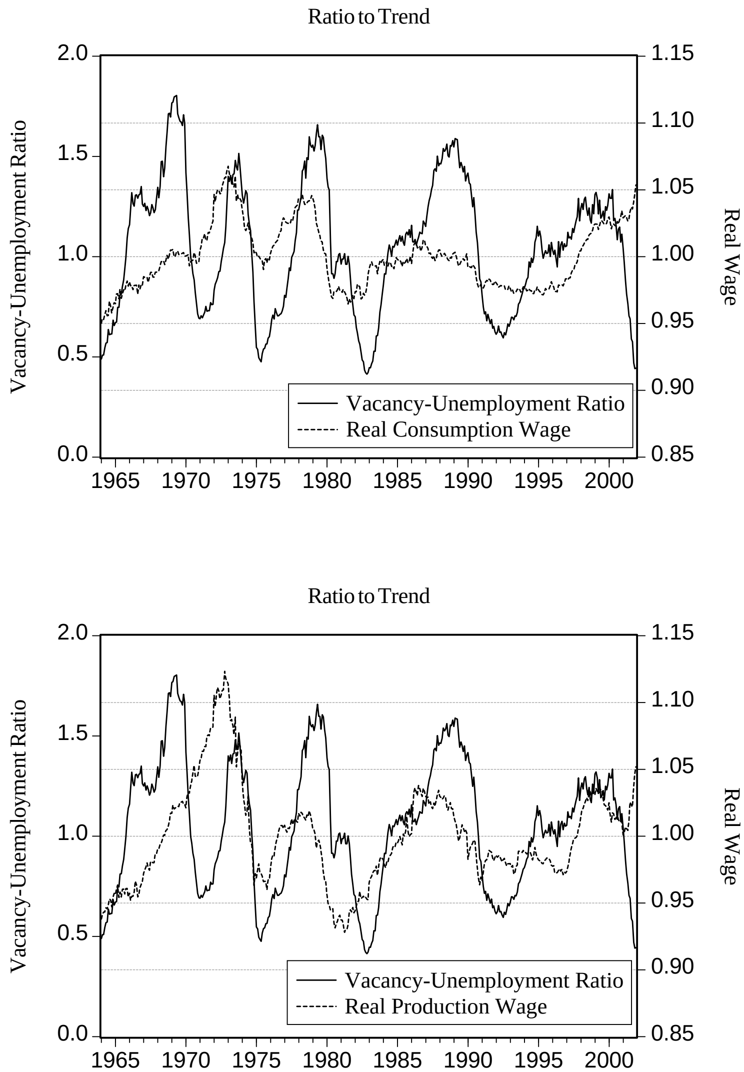

Figure 3: The unemployment rate u is constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey. The help-wanted advertising index v is constructed by the Conference Board. Nominal average hourly earnings for production workers in private industry (AHE) is constructed by the BLS from the Current Employment Statistics. It is deflated by the consumer price index for all urban consumers (CPI) in the top panel and by the producer price index for finished goods (PPI) in the bottom panel. Unemployment, vacancies, and real wages are expressed as ratios to an HP filter with smoothing parameter 10^7.