EFRW 7/25/02
SHIMER 4:15 PM

Equilibrium Unemployment Fluctuations*

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

June 6, 2002

Abstract

This paper develops three facts about the cyclical behavior of the labor market in the U.S. and explores whether existing models are consistent with those facts. I find that a standard model of search unemployment with wage bargaining either counterfactually predicts a positive correlation between unemployment and vacancies over the business cycle or predicts an order of magnitude more wage cyclicality than is observed in the data. I then consider an alternative environment with endogenous wage determination and show that while theoretically this can improve the performance of the model, it does not help under the most empirically plausible parameter restrictions. I conclude that some other source of real wage rigidities amplifies cyclical fluctuations.

*I thank seminar participants at the Minneapolis Fed, St. Louis Fed, Wharton, and the CEP Macroeconomics of the Labour Market conference for comments on previous versions of this 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 this paper, and to Thijs van Rens for excellent research assistance.

1 Introduction

In recent years, the search and matching model has become the standard theory of equilibrium unemployment (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 existing literature on the cyclical behavior of equilibrium search models, reviewed briefly in Section 5, has misread some important facts and as a result has claimed unwarranted success. This paper documents three facts that a model of equilibrium unemployment fluctuations should be able to explain — at least one of which is contrary to the conventional wisdom — and then explores whether the search and matching model can explain those facts. My main conclusion is that the standard model of equilibrium unemployment lacks sufficient wage rigidity.

Section 2 describes the three facts. First, I document the near acyclicality or modest procyclicality of real wages. A one percentage point increase in the unemployment rate is associated with at most a half a percent decline in wages on average. By using several measures of wages, I show that that this is true after including non-wage compensation and after controlling for the various "composition biases" that alter the pool of employed workers during the business cycle. On the other hand, I find that nonlabor payments are more strongly procyclical. Second, I document the strong negative correlation between unemployment and vacancies at business cycle frequencies. Indeed, I show that the product of unemployment and vacancies is nearly acyclic. As long as the difficulty of finding a job is decreasing in the number of vacancies and increasing in the number of unemployed workers, this implies it is more difficult for workers to find jobs during a recession. Finally, I show that in a well-defined sense, the labor market is always very nearly in steady state, although the steady state moves around during the business cycle. More precisely, I use labor market transition rates to construct a measure of what the unemployment rate and employment ratio would have been in each month if the labor market stocks were in steady state. I show that this construct, the 'flow unemployment rate', very closely tracks the usual 'stock' measure of unemployment.

Section 3 examines whether the standard model of equilibrium unemployment (Pissarides 2000) is consistent with these facts. I show that the model makes predictions about the relationship between real wages and labor market transition rates and between vacancies and labor market transition rates. Using the third fact, that the labor market is very nearly in steady state, I transform these into relationships between real wages and the unemployment rate and between vacancies and the unemployment rate. I then compare these theoretical relationships with the associated empirical relationships, the first two facts. I show that if cyclical unemployment fluctuations are not driven by movements in the job destruction rate or shifts in the matching function, the standard search model can easily explain why vacancies, and hence the difficulty of finding a job, is procyclical. On the other hand, if fluctuations were driven by such shocks, there would be a strong and counterfactual positive correlation between unemployment and vacancies over the business cycle, so that the difficulty of finding a job would be acyclic. Conversely, I show that the cyclicality of the difficulty of finding a job is inconsistent with the near acyclicality of the real wage. The model predicts that a decline in the unemployment rate from 5 percent to 4 percent should be associated with approximately a 30 percent decline in wages. I conclude that the standard model of equilibrium unemployment cannot reconcile the first two facts.

Section 4 considers an alternative wage setting environment, the 'competitive search equilibrium' (Montgomery 1991, Peters 1991). Firms commit to wages before hiring workers and are willing to offer high wages because doing so increases their hiring rate. I show that under some additional assumptions, the competitive search model predicts that real wages are much less procyclical or even mildly countercyclical, even in the face of procyclical vacancies. Effectively, it is as if in the standard search model, workers' bargaining power is countercyclical. However, I offer some evidence that suggests these additional assumptions are not satisfied. Instead, wages are as flexible in the competitive search equilibrium as in the standard model. This means that I am left without an explanation for the observed rigidity of real wages. Moreover, since the equilibrium of the competitive search model is socially optimal, the rigidity of real wages is socially inefficient.

Section 5 reviews the existing literature on search models of unemployment. Section 6 concludes by conjecturing about why wages might be less cyclical than standard models predict.

2 Three Labor Market Facts

2.1 Real Wage Cyclicality

The modest cyclicality of the real wage is a central fact in macroeconomics.[1] Since a substantial contribution to this literature goes beyond the scope of this paper, this section simply replicates some standard evidence. I consider three Bureau of Labor Statistics (BLS) measures of real wages: Average Hourly Earnings as measured by the BLS establishment survey, the Current Employment Statistics; Average Hourly Compensation, constructed from the National Income and Product Accounts; and the Employment Cost Index, measured using the BLS National Compensation Survey. All wage series are deflated by the Consumer Price Index for all Urban Consumers.[2] Despite the differences in these three measures of real wages, discussed in more detail below, the results are quite consistent across measures: a one percentage point increase in the unemployment rate reduces real wages by at most half a percent.

Average Hourly Earnings (AHE) of production and non-supervisory workers in private non-farm employment, constructed from employers' monthly payroll data, is a relatively narrow measure of wages. While AHE measures income before any deductions and includes paid vacations, sick days, and overtime, it does not include employee benefits, irregular bonuses, tips, and in-kind payments. Nevertheless, it is probably the best-known measure of wages, and so I begin with it. Figure 1 shows that there is a weak negative correlation between the real AHE of production workers and the unemployment rate, a convenient measure of the state of the business cycle. This is commonly quantified via a regression of the *log* deviation of the wage from its trend on the *absolute* deviation of the unemployment rate from its trend, so the coefficient estimates are semi-elasticities. Column 1 in Table 1 shows that a one percentage point increase in the unemployment rate is associated with a half a percent increase in the wage rate using annual data from 1947 to 2001. The low Durbin-Watson statistic indicates autocorrelation in the residuals, and so next column re-estimates the relationship in first differences. The point estimate of the response of wages declines insignificantly to about 0.4 percent. Since 1964, AHE estimates have been available on a monthly basis. Column 4 shows that adjusting the frequency of the data has some effect on these estimates; estimating the same equation on monthly data yields a still-lower responsiveness of AHE to the unemployment rate. Moreover, column 3 indicates that

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

[2] Using other price deflators, e.g. the producer price index, which measures the firms' labor costs, makes wages appear more countercyclical (Abraham and Haltiwanger 1995).

this is not due to the time period considered. If anything, AHE are more cyclical in recent years. Finally, in all cases the low R^2 shows that little of the variation in AHE is correlated with movements in the unemployment rate.

The National Income and Product Accounts (NIPA) provide a broader measure of compensation, including both 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. Dividing by hours data from the Current Employment Statistics yields the NIPA measure of Average Hourly Compensation (AHC) in the non-farm business sector. Despite the different methodology for constructing AHE and AHC, the correlation between the detrended series is 0.80, as Figure 1 demonstrates graphically. Column 5 in Table 1 shows that there is essentially no correlation between changes in AHC and changes in the unemployment rate in annual data from 1947 to 2001, while column 6, which uses quarterly data over the same time period, finds statistically insignificant evidence for a slight negative correlation between AHC and the unemployment rate.

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.[3] In part to address this concern, the BLS has since 1975 produced a third measure of wages, the Employment Cost Index (ECI), wages and salaries only, for private industry workers. Essentially, the ECI is a Laspeyres price index, meaning that in calculating ECI inflation from quarter to quarter, the composition of jobs is held fixed at the beginning-of-period value and used as weights in the computation of the overall employment cost inflation. 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. Figure 2 shows that again, despite the very different method by which they are constructed, the ECI and AHE behave similarly over their common sample period. Still, a regression of the detrended log ECI on the detrended unemployment rate in first differences yields an insignificantly *positive* coefficient estimate and has almost no explanatory power (Table 1, column 7).[4] During the same sample period and at a quarterly frequency, there is a modest negative relationship between the unemployment rate and the AHE measure (column 8) and a slight positive correlation

[3]Solon, Barsky and Parker (1994) is probably the best known paper on composition bias. Appendix A discusses the relevant composition biases and the extent to which they are addressed by the ECI.

[4]The National Compensation Survey has two measures of compensation costs, wages and salaries, which I have used here, and benefits. Benefit cost data are available quarterly since 1980. Although real benefit costs have increased sharply, they show no significant cyclical volatility. As a result, regressing the change in detrended log total compensation costs on the change in the detrended unemployment rate yields very similar results: a coefficient estimate of 0.17, smaller than the standard error 0.20.

between the unemployment rate and AHC (column 9). I conclude that no matter how we measure them, real wages are nearly acyclical. A one percentage point increase in the unemployment rate is associated with at most a half a percent drop in wages, although some measures indicate that there is no change at all.

Explicit or implicit contract theory (Baily 1974, Azariadis 1975) argues that firms smooth risk-averse workers' wages over the business cycle in the presence of incomplete markets. When the unemployment rate rises, the expected present value of wages may decline sharply but the current flow of labor income scarcely moves at all. Since in a standard competitive economy with a Cobb-Douglas aggregate production function, wages are equal to the marginal product of labor, which is in turn proportional to the average product of labor (APL), we can quantify the importance of this phenomenon by looking at the cyclical behavior of the APL rather than at wages. The NIPA provides the standard measure of the APL, output per hour in the non-farm business sector. Like AHC, this is available quarterly since 1947. Column 10 in Table 1 shows that a one percentage point increase in the unemployment rate is associated with a one-third of a percent decline in the APL, just at the borderline of statistical significance. Although the greater cyclicality of the APL relative to AHC is qualitatively consistent with the predictions of implicit contract theory, it is still not quantitatively very large.

Another way to look at this is through an accounting identity. Output either accrues to labor in the form of compensation or else is measured as nonlabor payments. The latter category includes payments to capital, corporate income taxes, and pure profits. If an increase in the unemployment rate is associated with a larger decline in the APL (output per hour) than in AHC (labor compensation per hour), nonlabor payments normalized by hours worked (NLP) must decline by even more. Column 11 in Table 1 confirms this logic. The semi-elasticity of NLP with respect to the unemployment rate is -0.8, and is statistically negative at conventional levels of significance. In other words, during recessions there is essentially no change in wages and a modest decline in profits, taxes, and payments to capital. From the perspective of implicit contract theory, the puzzle is not to explain why wages are acyclical but instead why measured nonlabor payments are not more strongly procyclical.

2.2 Vacancies and Unemployment

This section examines the cyclical behavior of the Conference Board help-wanted advertising index, the best available proxy for vacancies (Abraham 1987).[5] Panel a in

[5] "Job vacancies are defined as current unfilled job openings in your establishment which are immediately available for occupancy by workers from outside your firm and for which your firm is actively seeking such

Figure 3 shows that after accounting only for extremely low frequency fluctuations — an HP filter with smoothing parameter 10^7 — the correlation between help-wanted advertising and unemployment from 1951 to 2001 is -0.876. This relationship 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. Less well known is that the standard deviation of the cyclical fluctuations in the two series is almost identical, 0.192 for unemployment and 0.195 for help-wanted advertising. Put differently, the product of unemployment and vacancies is much less cyclical that either of the underlying series (panel b). Over the full sample period, the correlation between the unemployment rate and the product of unemployment and vacancies is 0.22, so the product is mildly countercyclical. However, in more recent years the product has become slightly procyclical, e.g. a correlation of -0.1 with the unemployment rate since 1970.[6]

An implication of this finding is that it should be harder to find jobs during a recession. Assume that the number of newly hired workers is given by an increasing and constant returns to scale matching function $m(u, v)$, depending on the number of unemployed workers and the number of vacancies. Then the probability that any individual unemployed worker finds a job, the average transition rate from unemployment to employment, is $\lambda^{ue} \equiv \frac{m(u,v)}{u} = m(1, v/u)$, increasing in v and decreasing in u. Since v is procyclical and u is countercyclical, it follows immediately that λ^{ue} is strongly procyclical. In the next section, I measure this transition rate directly and confirm that it is strongly procyclical (Figure 4, panel d). The transition rate from out of the labor force directly into employment is also procyclical (panel g).

The finding that the transition rates into employment are procyclical might appear to contradict Blanchard and Diamond's (1990) conclusion that "the amplitude

workers. Included are full-time, part-time, permanent, temporary, seasonal, and short-term job openings." (from Subcommittee on Economic Statistics of the Joint Economic Committee 1966, p. 166). Abraham (1987) shows that vacancies and help-wanted advertising in Minnesota track each other very closely through two business cycles from 1972 to 1981. Vacancy data are otherwise unavailable in the United States.

[6] A more robust result is that the product of unemployment and vacancies lags the unemployment rate by about two years. This may reflects a better known result, that vacancies lead unemployment by about two years. Pissarides (1987) argues that the ability of a simple search model to match this fact is an empirical success. Since the loops are much less pronounced than the inverse relationship between unemployment and vacancies at business cycle frequencies, I do not emphasize them here.

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 $m(u, v)$, that is only possible if procyclicality of vacancies offsets countercyclicality of unemployment, which is further support for the direct evidence.

2.3 Labor Market Stocks and Flows

Consider the evolution of the stock of workers in different employment states, e_t (employment), u_t (unemployment), and n_t (not in the labor force). Standard models of equilibrium unemployment, including the models analyzed in Sections 3 and 4 below, describe a two-step link between the exogenous state of the economy, e.g. productivity, and the state of the labor market $s_t \equiv \{e_t, u_t, n_t\}$. First, the state of the economy determines the flow of workers between labor market states, a Markov transition matrix Λ_t, independent of s_{t-1}. Second, the transition matrix determines the evolution of the state of the labor market via the transition equation $s_t = \Lambda_t s_{t-1}$. The goal of this section is to show that, empirically, the dynamic behavior of the state of the labor market is governed almost entirely by the first link.[7] For all practical purposes, the state of the labor market is given by the stationary distribution of Λ_t, $s_t^* \equiv \Lambda_t s_t^*$. This finding is convenient because it means that, although economic theory provides a direct link only between the state of the economy and transition rates, we can usefully discuss the relationship between the state of the economy and the state of the labor market by imposing the restriction that the state of the labor market is given by the stationary distribution of Λ.

I first use matched Current Population Survey (CPS) data sets to measure labor market transition rates Λ_t.[8] In a typical month, I observe the current and lagged

[7]Pissarides (1986) makes a similar point using annual data for Britain. He writes "...the rates in and out of unemployment in Britain (and much more so in countries with higher turnover in the labour force like the US) are sufficiently high that for given constant rates, the adjustment in the unemployment relationship implicit in the stock-flow identity is completed within a year. Hence for year-to-year changes in unemployment ... we may solve for unemployment from the stock-flow identity by setting [the change in the unemployment rate] to zero" (p. 505). Abraham and Shimer (2001) demonstrate this using annual data for the U.S.. I believe this is the first paper to document the monthly behavior of the flow unemployment rate and employment ratio.

[8]The CPS is a rotating panel. Individuals are in the survey for four months, out for eight months, and

employment status of at least 70,000 individuals. Weighting these observations by the average of the individual's current and lagged basic CPS weights and summing yields a measure of the nine labor market transition rates, e.g. the employment-to-unemployment transition rate λ^{eu}, etc..[9] Figure 4 shows the nine time series plotted against the unemployment rate for reference.

Next, observe that if the transition matrix Λ were constant over time, the state vector s would eventually converge to a stationary distribution $s^* \equiv \{e^*, u^*, n^*\}$, the normalized eigenvector associated with Λ's unit eigenvalue. Even out of steady state, the time-varying eigenvector s_t^*, the solution to $s_t^* \equiv \Lambda_t s_t^*$, measures the current state of the labor market. In particular, call $\frac{u_t^*}{e_t^*+u_t^*}$ the 'flow unemployment rate', since it uses labor market flows (Λ_t) to measure something akin to the unemployment rate. Similarly, call e_t^* the 'flow employment ratio'.

Panel a in Figure 5 compares the flow unemployment rate to the standard 'stock' measure of unemployment from 1967 to 2001, while panel b compares the flow and stock employment ratios. Both flow series track the associated stock series very closely, although the flow series exhibit more high frequency volatility and have a somewhat lower mean value than the stock series. This implies that the additional dynamics generated by the transition equation $s_t = \Lambda_t s_{t-1}$ are not quantitatively important. The dynamic properties of the state of the labor market simply reflect the dynamic properties of the

then in again for four months, with approximately $\frac{1}{8}$ of the sample in each of the eight 'rotation groups' in each month. The structure of the rotating panel implies that an individual who is the first or fifth rotation group in one month was not interviewed the previous month, and so $\frac{1}{4}$ of the sample cannot be matched. I attempt to match the remaining individuals using household identifiers, line numbers, rotation group, race, sex, and age. The values for each of these variables must be the same in month $t-1$ and t, except for rotation group (which must increase by 1) and age (which may increase by 1 or remain constant). I also drop all observations that are not uniquely identified by these measures within a given survey. This matching procedure is imperfect both because of survey attrition and because of mistakes in recording identifying information (Madrian and Lefgren 2000). In practice, I am able to match approximately 90 percent of the individuals in rotation groups 2 to 4 and 6 to 8 with records in the previous month's survey. I defer a discussion of the shortcomings of matching CPS files until Section 2.3.

[9]Hoyt Bleakley provided me with time series from June 1967 to December 1975, which were originally constructed by Joe Ritter. Because of survey redesigns and privacy restrictions, I could not construct transition rates in January 1976, January 1978, July 1985, October 1985, January 1994, and June to October 1995. I seasonally adjust the transition rate data using a ratio-to-moving average technique that allows for missing observations and for time-variation in the seasonal factors. I first take the ratio of the actual time series to a centered 12-month moving average of the time series. I only construct this ratio when all of the relevant 13 months' data are available. I then calculate a nine-year centered moving average of the ratio for each month. I construct this variable, the seasonal factor, even if some of the relevant data is unavailable. Finally, I take the ratio of the unadjusted time series to the seasonal factor whenever the unadjusted data is available. In practice, I obtain very similar results with the Census X-12 procedure. The disadvantage to that approach is that it is harder to understand exactly what is happening and moreover it is necessary to interpolate the missing observations.

transition matrix. This result should be of some independent interest.[10]

There is a large literature documenting deficiencies in the measurement of labor market transition rates. Some of these might have implications for the close relationship between the stock and flow measures of the state of the labor market. I discuss three of them here.

First, there is a strong rotation group effect that reduces the measured flow unemployment rate and employment ratio. From 1976 to 2001, the period of time for which I have access to the CPS micro data, the reported unemployment rate in the first rotation group averaged 7.6 percent (0.47 percentage points) higher than the unemployment rate in the remaining sample. The comparable figure for the employment ratio was 1.5 percent (0.76 percentage points) higher. Pairwise tests strongly reject equality of the unemployment rate and employment ratio across rotation groups (Table 2), although the bias in other rotation groups is much less significant.[11] Because of this rotation group effect, labor market transition estimates will overestimate the number of people exiting employment and to a lesser extent unemployment. This reduces my estimate of the steady state fraction of the population in that state, i.e. it reduces the measured flow unemployment rate and employment ratio. To get around this problem, I drop matched records linking the first and second rotation groups.

Second, there is non-random attrition from the survey, with unemployed workers more likely to exit the sample. While the CPS weights compensate for this in measuring the actual unemployment rate and employment ratio, this may still bias estimated transition rates. Suppose, for example, that workers tend to drop out of the sample when they lose their job. Then I will underestimate the employment-to-unemployment transition rate λ^{eu}, reducing the measured flow unemployment rate. In addition, since the BLS time series for the unemployment rate and employment survey come from the full CPS sample, not the subset of matchable files, my stock and flow measures refer to slightly different populations. I take two measures to correct for non-random attrition. First, I measure the stock data directly from the matched CPS files. Second,

[10]For example, Cole and Rogerson (1999) explain the persistence of unemployment fluctuations, e.g. the high serial correlation of the change in the unemployment rate, by arguing that the unemployment rate only slowly adjusts to its steady state value, the flow unemployment rate. On the contrary, I have shown that the flow unemployment rate exhibits as much persistence as the stock unemployment rate, and so to explain the persistence of the unemployment rate, one must explain why transition rates adjust very slowly.

[11]There is a small literature on rotation group biases in the CPS. Solon (1986) writes "Why responses vary systematically with time in the sample is unknown, but possible factors include conditioning of respondents or interviewers by repeated contacts, differences among rotation groups in the length and content of the questionnaire and differences in nonresponse, which household member is interviewed, and whether the interview is conducted by telephone or in person" (p. 105). Curiously, the rotation group bias was essentially unchanged by the major redesign of the CPS in 1994.

I introduce a fourth state to the vector s_t, 'missing from the CPS'.[12] I measure the transition rate in and out of this state using all workers who cannot be matched from month to month (except those who cannot be matched because their rotation group enters or exits the sample).[13] Together with the correction for rotation group biases, this eliminates the average gap between the stock and flow unemployment rates and employment ratios.

Finally, Abowd and Zellner (1985) and Poterba and Summers (1986) have documented that classification errors, where an individual's employment status is mistakenly recorded, significantly increase the estimated labor market transition rates. Fortunately, this should not be a significant problem for this exercise. Assuming that workers are on average classified correctly, i.e. that the stock unemployment rate and employment ratios are correctly measured in every month, classification errors will not affect the measured flow rates. The proof of this claim is simple. Let $\tilde{\Lambda}_t$ denote the measured transition matrix and Λ_t denote the true transition matrix, with no classification errors. If workers are on average classified correctly, the evolution of the state vector satisfies both $s_t = \tilde{\Lambda}_t s_{t-1}$ and $s_t = \Lambda_t s_{t-1}$. I have already shown that using the measured transition matrix $\tilde{\Lambda}_t$, the state vector is always essentially in steady state, $s_t = s_t^*$ satisfying $s_t^* = \tilde{\Lambda}_t s_t^*$. This implies $s_t^* = \Lambda_t s_t^*$ as well, so the current state is also the eigenvector associated with the unit eigenvalue of the true transition matrix Λ_t. As further confirmation of this reasoning, I have replicated my analysis using Abowd

[12] This appears to be a novel way to address non-random attrition. Other methods, e.g. estimating the likelihood of attrition for workers with different characteristics and adjusting the measured transition rates accordingly, have little effect on the measured transition rates (Bleakley, Ferris and Fuhrer 1999).

[13] A potential issue with this procedure is that I cannot measure the "missing-to-missing" transition rate, since such workers never appear in the CPS. Fortunately, the number of such transitions does not affect the distribution of workers across the three 'real' employment states, and hence does not affect my estimate of the employment ratio or the unemployment rate. To see why, denote the state vector by $s = \{x, m\}$, where $x = \{e, u, n\}$ are the three usual states. Similarly divide the transition matrix Λ into blocks, a 3×3 matrix Λ^{xx} containing the usual transition rates, a 3×1 vector Λ^{mx} containing the likelihood of exiting 'missing', a 1×3 vector Λ^{xm} containing the likelihood of entering missing, and λ^{mm}, the unmeasured missing-missing transition rate. Steady state imposes two restrictions: $x = \Lambda^{xx}x + \Lambda^{mx}m$ and $m = \Lambda^{xm}x + \lambda^{mm}m$. Eliminate m from the first equation using the second to get

$$x = \left(\Lambda^{xx} + \frac{\Lambda^{mx}}{1 - \lambda^{mm}} \cdot \Lambda^{xm}\right) x$$

Since the elements of the last column of Λ sum to 1, $\mathbf{1} \cdot \Lambda^{mx} = 1 - \lambda^{mm}$, where $\mathbf{1}$ is a 1×3 vector of 1s. This means that the fraction $\frac{\Lambda^{mx}}{1-\lambda^{mm}}$ is independent of λ^{mm}. By normalizing λ^{mm} to zero, we find that accounting for non-random attrition effectively modifies the transition matrix for x from Λ^{xx} to $\Lambda^{xx} + \Lambda^{mx} \cdot \Lambda^{xm}$. Since $\mathbf{1} \cdot \Lambda^{mx} = 1$, it is straightforward to verify that this new 3×3 matrix is still Markov. I conclude that the stationary distribution of workers across the three real employment states x is given by the normalized eigenvector associated with the unit eigenvalue of $\Lambda^{xx} + \Lambda^{mx} \cdot \Lambda^{xm}$, given the normalization $\lambda^{mm} = 0$.

and Zellner's adjusted labor market transition rate series from 1968 to 1986,[14] with no significant effect on the relative behavior of the stock and flow rates.

In any case, a useful way to think about classification errors are that they increase the measured duration dependence in the data. It is more probable that a worker who is measured unemployed in two consecutive months is correctly classified than a worker who just became unemployed. Thus thinking about classification error naturally leads one to think about the labor market as having many more employment states. By matching CPS files across three months, I can begin to capture this idea. Define a worker's employment state by her employment status in the previous and current month. Thus EE denotes a worker who was employed last month and this month, EU a worker who was employed last month and unemployed this month, etc. Treating 'missing from the sample' as an employment state, this expands the state vector to 16 possibilities. Using matched CPS files across three months, I can measure the transition rate between these 16 states. For example, a worker who is employed in January, unemployed in February, and employed again in March counts as an EU to UE transition from February to March. Obviously many of the 256 transitions are impossible, since a worker who is unemployed in February can only transition to four states in March: UE, UU, UN, or UM. Nevertheless, the resulting Markov transition matrix is regular and so the eigenvector associated with its (unique) unit eigenvalue gives the steady state distribution of workers across employment states. This flow distribution can be compared to the stock distributions obtained by matching individual records across only two months.

In any given month, the necessary information exists for at most half the CPS sample (e.g. in March, rotation groups 3, 4, 7, and 8 can potentially be matched with records from January and February). Moreover, I throw away matches between the first three rotation groups to avoid the rotation group effects discussed above, leaving me with $\frac{3}{8}$ of the original CPS sample. Because I cannot locate public use micro data from the CPS before 1976, I restrict the sample period to 1976 to 2001. Figure 6 confirms that performing all three corrections only tightens the relationship between the stock and flow measures of employment and unemployment.[15] On average, the flow unemployment rate is only 2.1 percent below the actual unemployment rate (0.11 percentage points). The standard deviation of the difference between the two series

[14]I am grateful to Olivier Blanchard for providing me with this data. This series attempts to reduce classification errors by using data from reinterviews.

[15]The (flow or stock) unemployment rate, for example, is defined by $\frac{EU+UU+NU}{(EU+UU+NU)+(EE+UE+NE)}$, where XY is the (flow or stock) measure of the number of individuals in state XY in the appropriate month. Thus both measures ignore individuals who were missing from the survey in either month.

is also small, 0.38 percentage points. The correlation between the flow and stock employment ratio is less impressive but still quite strong (Panel b). In particular, the flow series captures the trends and cyclical fluctuations in the employment ratio, although with considerable noise.

3 Search Model of Equilibrium Unemployment

This section explores whether a standard search model of equilibrium unemployment (Pissarides 1985, Pissarides 2000) is consistent with the three facts outlined in the previous section. I use both theoretical arguments and numerical examples to show that if cyclical fluctuations were driven by shocks to the rate of job destruction, the *ratio* of vacancies to unemployment would be nearly acyclic. This sharply contrasts with the empirical evidence that the *product* of vacancies and unemployment is acyclic. Next I use a similar combination of theory and numerical examples to examine the behavior of the model if cyclical fluctuations are driven by shocks that affect the real marginal product of labor, e.g. productivity shocks. I find that while the model is consistent with the Beveridge curve evidence, it predicts that wages would be strongly procyclical. For example, suppose wages are approximately twice unemployment income. Then the model predicts that at a five percent unemployment rate, the semi-elasticity of wages with respect to the unemployment rate would be approximately -40 (see equation 17 below). The data indicates that the semi-elasticity likely lies somewhere between -0.5 and 0. In response to both types of shocks, the model is unable to replicate key employment fluctuation facts.

The model deliberately abstracts from other potential explanations for this evidence, e.g. a high intertemporal elasticity of substitution in labor supply Lucas and Rapping (1969). In part this represents my skepticism about these explanations. But in addition, by considering one aspect of the labor market in isolation, I hope to clarify whether search frictions alone reduce real wage cyclicality and increase the size of employment fluctuations.

3.1 Benchmark Model

The economy consists of a continuum of measure 1 of infinitely-lived, risk-neutral workers and a continuum of infinitely-lived, risk-neutral firms. Time is discrete and all agents discount future payoffs with factor $\beta < 1$. Workers can be either employed or unemployed. Employed workers inelastically supply a unit of labor and earn a wage w_H, where H is the current state of the economy, defined more precisely below. Unemployed workers search for a job and earn an unemployment benefit b. Firms have a constant returns to scale production technology that uses only labor. Each employee produces p_H units of output. Although I will call p_H firms' productivity, aggregate demand shocks would also enter the model via their influence on the value of firms' output. In order to hire workers, the firms must open vacancies, which they do if there are profits to be made. Each open vacancy costs c, but may give the firm the opportunity to hire a worker.

More precisely, let u_H denote the beginning-of-period measure of unemployed workers (hereafter the unemployment rate) and v_H the measure of vacancies. The number of matches in the economy is given by the constant returns to scale matching function $m(u_H, v_H)$, increasing in both arguments. The probability with which an unemployed worker contacts a firm is $\lambda_H^{ue} \equiv m(1, \theta_H)$ and the probability with which each vacancy contacts a worker is $q_H \equiv m(1/\theta_H, 1)$, where $\theta_H \equiv v_H/u_H$ measures the tightness of the labor market. Finally, employed workers lose their job and become unemployed with exogenous probability λ_H^{eu}, the same probability with which firms lose employees.

There are generally bilateral gains from matching. In this section, I follow the bulk of the 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 $\gamma \in (0, 1)$ of the surplus, her "bargaining power". The next section considers an alternative wage-setting arrangement.

The state of the economy H consists of the entire history of productivity p and of the job destruction probability, λ^{eu}. The history H is common knowledge and the joint distribution of future productivity and job destruction shocks is measurable with respect to H. I assume that both p and λ^{eu} have bounded support. Moreover, for all histories H, $p_H > b$ with probability 1.

The key endogenous variables are the wage w_H, the vacancy-unemployment ratio θ_H, and the unemployment rate u_t. The first two variables are pinned down by the Nash bargaining solution and the free entry condition for vacancies. The third variable is a state variable, and it evolves according to the difference equation

$$u_t = \lambda_H^{eu}(1 - u_{t-1}) + (1 - \lambda_H^{ue})u_{t-1} \tag{1}$$

where H is the state of the economy at the start of period t. The unemployment rate this period is the sum of the number of employed workers, of measure $1 - u_{t-1}$, who lose their job, with probability λ_H^{eu}; and the number of unemployed workers, of measure u_{t-1}, who do not find a job, with probability $1 - \lambda_H^{ue}$. Note that the unemployment rate u_{t-1} is not part of the state of the economy. I show below that it is possible to characterize the equilibrium without reference to this variable and then to go back and solve for the dynamics of the unemployment rate using (1).

3.2 Equilibrium Analysis

I solve for the equilibrium of the economy using Bellman equations. Let U_H denote the expected present value of income for an unemployed worker and E_H denote the expected present value of income for an employed worker, each expressed as a function of the history H. These satisfy

$$U_H = b + \beta \mathbb{E}_H \big(\lambda_H^{ue} E_{H'} + (1 - \lambda_H^{ue}) U_{H'} \big) \tag{2}$$

$$\text{and } E_H = w_H + \beta \mathbb{E}_H \big((1 - \lambda_H^{eu}) E_{H'} + \lambda_H^{eu} U_{H'} \big), \tag{3}$$

where $\mathbb{E}_H$ is the expectations operator conditional on the current history H and H' is the history in the next period. For notational simplicity I express the transition rates λ^{ue} and λ^{eu} directly as functions of H. Similarly, the value of a vacancy V_H and the value of a filled job F_H satisfy:

$$V_H = -c + \beta \mathbb{E}_H \big(q_H F_{H'} + (1 - q_H) V_{H'} \big) \tag{4}$$

$$\text{and } F_H = p_H - w_H + \beta \mathbb{E}_H \big((1 - \lambda_H^{eu}) F_{H'} + \lambda_H^{eu} V_{H'} \big) \tag{5}$$

There are also two key equilibrium conditions. The Nash bargaining solution implies

$$(1 - \gamma)(E_H - U_H) = \gamma (F_H - V_H) \tag{6}$$

for all H, while the free entry condition on vacancies implies $V_H = 0$ for all H.

Rewrite equations (3) and (5) as

$$E_H - U_H - \beta(1 - \lambda_H^{eu}) \mathbb{E}_H \big(E_{H'} - U_{H'} \big) = w_H - (U_H - \beta \mathbb{E}_H U_{H'})$$

$$\text{and } F_H - \beta(1 - \lambda_H^{eu}) \mathbb{E}_H F_{H'} = p_H - w_H,$$

where the last equation is simplified using the free entry condition. Then subtract $1-\gamma$ times the first equation from γ times the second equation, simplifying with the Nash bargaining solution (6):

$$w_H = \gamma p_H + (1-\gamma)(U_H - \beta \mathbb{E}_H U_{H'}). \tag{7}$$

The wage is a weighted average of current productivity and the difference between the value of unemployment today and the expected discounted value of unemployment tomorrow.

To solve for $U_H - \beta \mathbb{E}_H U_{H'}$, eliminate $\mathbb{E}_H E_{H'}$ between equations (2) and (3) and $\mathbb{E}_H F_{H'}$ between equations (4), (5), and the free entry condition $V_H = \mathbb{E}_H V_{H'} = 0$:

$$E_H - U_H = w_H - \frac{1-\lambda_H^{eu}}{\lambda_H^{ue}} b + \frac{1-\lambda_H^{eu}-\lambda_H^{ue}}{\lambda_H^{ue}}(U_H - \beta \mathbb{E}_H U_{H'}) \tag{8}$$

$$\text{and } F_H = p_H - w_H + \frac{1-\lambda_H^{eu}}{q_H} c, \tag{9}$$

where the second equation is further simplified using the free entry condition. Now eliminate $E_H - U_H$, F_H, and $U_H - \beta \mathbb{E}_H U_{H'}$ between the Nash bargaining solution (6), the wage equation (7), and equations (8) and (9):

$$w_H = \gamma \left(p_H + \frac{\lambda_H^{ue}}{q_H} c \right) + (1-\gamma) b = \gamma \left(p_H + \theta_H c \right) + (1-\gamma) b, \tag{10}$$

where the second inequality follows because $\frac{\lambda_H^{ue}}{q_H}$ is the ratio of matches per unemployed worker to matches per vacancy, which is necessarily the vacancy-unemployment ratio θ_H. This states that the wage is a weighted average of productivity plus vacancy costs per unemployed worker and the unemployment benefit. An increase in the vacancy-unemployment ratio (a tighter labor market) raises wages, as does an increase in productivity or the unemployment benefit.

To calculate the vacancy-unemployment ratio, use (10) to eliminate the wage from equation (9):

$$F_H = (1-\gamma)(p_H - b) - \frac{1-\lambda_H^{eu} - \gamma\lambda_H^{ue}}{q_H} c$$

Lead this equation by one period and take expectations. Then substitute it back into equation (4) and combine with the free entry condition:

$$c = \beta q_H \left((1-\gamma)(\mathbb{E}_H p_{H'} - b) + \mathbb{E}_H \left(\frac{1-\lambda_{H'}^{eu}}{q_{H'}} - \gamma \theta_{H'} \right) c \right). \tag{11}$$

Note that firms' hiring probability q is decreasing in the vacancy-unemployment ratio θ_H, and so this is a forward-looking nonlinear difference equation for θ_H. There are several ways to identify the particular solution of interest. Perhaps most intuitively, we are interested in a solution in which θ depends only on the payoff-relevant portion of the history. For example, if the aggregate state follows a first-order Markov process, so the values of p and λ^{eu} next period depend only the current values of those two variables, then θ should only depend on the current value of p and λ^{eu}. Any other solution to these difference equations has the property that θ_H explodes, tending towards plus or minus infinity with positive probability.

3.3 Steady State

Before looking at the behavior of the full stochastic model, it is useful to consider the steady steady of this economy, in which p and λ^{eu} are time-invariant. Since I argued in Section 2.3 that the unemployment rate is essentially always in steady state, the resulting comparative statics may not be a poor approximation to the behavior of the stochastic model.

In the fundamental solution, all variables are independent of the (trivial) history of the economy. Then equation (11) yields simple solution for the steady state hiring rate q:

$$q = \frac{\big(1 - \beta(1 - \lambda^{eu} - \gamma\lambda^{ue})\big)c}{\beta(1-\gamma)(p-b)}. \tag{12}$$

For simplicity, I focus on the limiting behavior of the economy as β approaches 1:

$$q = \frac{(\lambda^{eu} + \gamma\lambda^{ue})c}{(1-\gamma)(p-b)}. \tag{13}$$

Later numerical simulations show that this does not qualitatively affect the results because the relevant discount factor is $\beta(1-\lambda^{eu})$, accounting both for the rate of time preference and for the temporary nature of matches. In practice, $1-\lambda^{eu}$ is significantly smaller than β and so is the dominant determinant of the rate of discounting.

Beveridge Curve

The first question to ask is whether the model is consistent with the Beveridge curve evidence that the product of vacancies and unemployment is nearly acyclical. The answer depends on the nature of shocks to the economy. In steady state, the number of matches $m(u, v)$ is equal to the number of workers who lose their job, $\lambda^{eu}(1-u)$:

$$m(u, v) = \lambda^{eu}(1-u). \tag{14}$$

Suppose there is an increase in the unemployment rate caused by fluctuations in any variable that does not appear in this equation. Such variables include productivity p, unemployment income b, or bargaining power γ, but not the match destruction rate λ^{eu} or the matching function m. Then the economy will move along this Beveridge curve locus. In particular, an increase in the unemployment rate raises the left hand side of equation (14) at any level of vacancies and reduces the right hand side. To restore the steady state relationship, vacancies must fall, causing an inverse relationship between vacancies and unemployment.

Indeed, suppose $m(u,v) \equiv m(\kappa u, v/\kappa)$ for any $\kappa > 0$ and all (u,v), as is the case if $m(u,v) = \mu u^{\frac{1}{2}} v^{\frac{1}{2}}$. Consider an increase in the unemployment rate from 0.05 to 0.10, a large 'recession'. The right hand side declines by about 5.3 percent $(.90/.95-1)$. This implies that the vacancy rate must decline by slightly over fifty percent, so that the left hand side falls by 5.3 percent as well. The product of unemployment and vacancies is therefore nearly acyclic, as the data indicates. I conclude that productivity, unemployment income, or bargaining power fluctuations are consistent with the empirical Beveridge curve evidence in Section 2.2.

What about fluctuations in the job destruction rate λ^{eu}? Since the variable appears directly in equation (14), that equation is not very useful in answering the question. Instead, we must delve deeper into the model. Assuming the unemployment rate is at its steady state value, i.e. that the flow and stock unemployment rates are identical, equation (1) implies

$$\lambda^{eu} = \frac{u}{1-u}\lambda^{ue}.$$

Substituting this into equation (13) gives

$$\frac{\lambda^{ue}}{q} = \frac{(1-\gamma)(p-b)}{\left(\frac{u}{1-u}+\gamma\right)c}.$$

The left hand side is identical to the vacancy-unemployment ratio, making it easy to solve for the relationship between vacancies and unemployment:

$$v = \frac{(1-\gamma)(p-b)u(1-u)}{(u+\gamma(1-u))c}. \tag{15}$$

Shifts in the job destruction rate move vacancies and unemployment along this locus. The right hand side is an increasing function of u for $u < \frac{\sqrt{\gamma}}{1+\sqrt{\gamma}}$. Indeed, a linear approximation around $u = 0$ implies

$$v \approx \frac{(1-\gamma)(p-b)}{\gamma c}u,$$

so unemployment and vacancies are roughly proportional in response to job destruction shocks. This contradicts the empirical Beveridge curve evidence.

Intuitively, an increase in the job destruction rate makes both firms and workers worse off, since it decreases the duration of a match. From the workers' perspective, this results in a reduction in the value of unemployment, the threat point in wage negotiations. Wages fall, which partially compensates firms for the shorter match duration. The vacancy-unemployment ratio, which measures the tightness of the labor market, need not respond much at all. Hence the increase in unemployment directly caused by an increase in the job destruction rate is nearly matched by an increase in vacancies.

Real Wages

The fundamental tension in matching this model with the empirical evidence can be understood by eliminating productivity p from the wage equation (10) using the steady state relationship (12) and simplifying:

$$w = b + \frac{\big(1 - \beta(1 - \lambda^{eu} - \lambda^{ue})\big)\gamma c}{\beta(1-\gamma)q}. \tag{16}$$

Again, we further simplify this equation in three steps: focus on the limit as β approaches 1; use the steady state relationship $\lambda^{eu} + \lambda^{ue} = \frac{\lambda^{ue}}{1-u}$; and replace $\frac{\lambda^{ue}}{q} = \frac{v}{u}$. This gives:

$$w = b + \frac{\gamma c v}{(1-\gamma)(1-u)u}. \tag{17}$$

Since the data indicates that v is procyclical and u is countercyclical, the fraction, and hence the wage rate, is strongly procyclical in response to movements in any parameter that satisfies two conditions: it must not appear on the right hand side of equation (17) and it must induce a negative correlation between vacancies and unemployment. This includes, most notably, productivity p.

To quantify how procyclical wages are in response to such shocks, recall that the empirical Beveridge curve evidence in Section 2.2 indicates that the product of vacancies and unemployment is acyclical, so $v \approx k/u$ for some constant k. Taking this approximation as an exact statement would imply

$$w = b + \frac{\gamma c k}{(1-\gamma)(1-u)u^2} \Rightarrow \frac{d\log(w-b)}{du} = -\frac{2-3u}{u(1-u)}.$$

As long as the unemployment rate is less than 2/3, this is semi-elasticity is negative.

Indeed, it explodes as u approaches zero. For example, if a productivity shock reduces the unemployment rate from 5 percent to 4 percent, the difference between wages and unemployment income increases by about 55 percent. In contrast, assuming wages are normally about twice unemployment income, the empirical evidence in this paper suggests this semi-elasticity should be between -1 and 0. Even Solon et al. (1994) predict a semi-elasticity of wages in excess of unemployment benefits with respect to the unemployment rate in the range of -2.

Conversely, job destruction shocks induce a positive correlation between unemployment and vacancies, and so this argument does not apply. Equation (17) therefore implies that real wages will scarcely respond to such shocks. To quantify this observation, replace $\theta = \frac{v}{u}$ in the wage equation (10) using (15) and simplify:

$$w = b + \frac{\gamma(p-b)}{u + \gamma(1-u)} \Rightarrow \frac{d\log(w-b)}{du} = -\frac{1-\gamma}{u+\gamma(1-u)} \geq -\frac{1-\gamma}{\gamma}.$$

Unless workers' bargaining power γ is very small, this semi-elasticity is roughly consistent with the empirical evidence described in Section 2.1. I conclude that it is easy to explain either the Beveridge curve using productivity shocks or real wage cyclicality using job destruction shocks, but that the model is inconsistent with both facts obtaining simultaneously.

Implicit Contracts

A problem with connecting the theory to the data on wage cyclicality is that implicit contracts might smooth the timing of wage payments across cyclical fluctuations. The benchmark model assumes that wages are continually renegotiated, while alternative assumptions might be more realistic, for example that wages are fixed at the beginning of an employment relationship and then change only in the event that one party would otherwise wish to terminate the arrangement. It is relatively easy to address this possibility because such wage smoothing will not have any allocational effects on the economy. In particular, in deciding whether to create a job, firms care about the expected present value of wages, as summarized by the workers' Bellman value E, not about the exact timing of wages. This means that implicit contracting will not affect the responsiveness of the unemployment rate to shocks, and so we can still look at such elasticities to evaluate the success of the model.

In particular, suppose cyclical fluctuations are driven by productivity shocks. Solve equation (15) for p and replace v with the approximation k/u to get

$$p = b + \frac{(u + \gamma(1-u))ck}{(1-\gamma)(1-u)u^2}.$$

Now the semi-elasticity of the difference between productivity and unemployment income with respect to the unemployment rate is a decreasing function of workers' bargaining power γ. It is maximized, i.e. closest to zero, when $\gamma = 0$, in which case

$$\frac{d\log(p-b)}{du} = -\frac{1-2u}{u(1-u)}.$$

This is positive for $u < 1/2$ and it explodes as u approaches zero. A decrease in the unemployment rate from 5 percent to 4 percent requires a 24 percent increase in the difference between productivity and unemployment income. For larger values of γ, the semi-elasticity is even more negative, and is the same as the semi-elasticity of wages, $-\frac{2-3u}{u(1-u)}$, at the other extreme, $\gamma = 1$. In contrast, the empirical evidence in Section 2.1 suggested that a one percentage point reduction in the unemployment rate is associated with a 0.3 percent increase in the average product of labor.

Regardless of whether wages are continuously renegotiated or the timing of wage payments is determined by implicit contracts, wage flexibility is central to this finding. A negative productivity shock reduces the present value of firm profits and thus lowers the amount of job creation. This increases the duration of unemployment, which reduces the value of an unemployed worker, workers' threat point in bargaining. The present value of wages falls, mitigating the decline in the present value of firm profits. In practice, the decline in the present value of wages is almost as large as the original productivity shock, leaving the present value of the difference between productivity and wages almost unchanged. This means job creation scarcely falls, and the effect of the productivity shock on the unemployment rate is small.

Suppose instead that wages are exogenously fixed, with $b < \bar{w} < p$. Combining the Bellman equations (4) and (5) with the free entry condition $V = 0$ gives

$$c = \beta q \frac{p - \bar{w}}{1 - \beta(1 - \lambda^{eu})}$$

in steady state. As usual, take the limit as β approaches 1, replace $\lambda^{eu} = \frac{u}{1-u}\lambda^{ue}$, and approximate $v = k/u$ to get

$$p = \bar{w} + \frac{ck}{u(1-u)} \Rightarrow \frac{d\log(p-\bar{w})}{du} = -\frac{1-2u}{u(1-u)}.$$

Note that $p - \bar{w}$ is equal to firms' gross profit per employee. In an economy in which firms must pay for capital as well as labor services, this would be considerably less than the capital share, and it is also less than net profits because firms with vacancies lose money in each period.[16] Thus this equation says that it takes a 24 percent increase in firms' gross profit margin in order to reduce the unemployment rate by from 5 percent to 4 percent. If, for example, firms' gross profit margin averages about 4 percentage points, i.e. $p \approx 1.04\bar{w}$, an increase in p to $1.05\bar{w}$ raises the gross profit margin $p - \bar{w}$ by the necessary 25 percent. Such shocks are within the realm of plausibility. I conclude that the problem with the benchmark model lies in the extent of wage flexibility implicit in the Nash bargaining solution (6). This is true whether wages are continuously renegotiated or wages are fixed according to this equation once at the beginning of the employment relationship.

3.4 Dynamic Simulations

To be done.

4 Competitive Search Model of Unemployment

In the previous section I argued that the flexibility of real wages in the standard search model of unemployment makes it difficult to reconcile the model with the data. This section therefore considers an alternative wage setting arrangement. I assume that firms can commit to wages before hiring workers and that by promising higher wages, a firm can increase its hiring rate. This is a 'competitive search' or 'directed search' model (Peters 1991, Montgomery 1991, Moen 1997, Shimer 1996, Burdett, Shi and Wright 2001).

I show that if the matching function $m(u, v)$ is Cobb-Douglas, the competitive search model is indistinguishable from the standard bargaining model. However, if the elasticity of substitution between unemployment and vacancies is larger than 1, its value in the Cobb-Douglas case, wages are less flexible in the competitive search model than in the standard model, improving the model's performance along the dimensions discussed above. Indeed, if $m(u, v)$ is linear in u and v, the competitive search model predicts that wages should be mildly countercyclical. Although this appears to be a promising explanation for real wage rigidities, I also provide some empirical evidence

[16] In an economy without discounting, the flow of net profits is zero in steady state. With discounting, the discounted value of profits is zero, but since firms lose money on vacancy costs up front, they must expect to earn more profit from filled jobs later on. This means that the flow of net profits is positive.

that suggests the elasticity of substitution is close to unity. Thus the competitive search model does not explain the relative inflexibility of real wages.

This finding is relevant for two reasons. It rules out one plausible explanation for real wage rigidity and, perhaps more importantly, the allocation in the competitive search equilibrium maximizes output in the economy given the search frictions (Moen 1997, Shimer 1996). In particular, the vacancy-unemployment ratio θ, and hence the matching rates q and λ^{ue}, are optimal. To the extent that wages are more rigid than the competitive search equilibrium predicts, they are either too high in recessions or too low in expansions or both. This tends to amplify shocks to the economy.

4.1 Wage Setting and Search

In the competitive search equilibrium, when a firm opens a vacancy, it announces a complete wage contract, specifying a wage in each history H for any worker whom it hires. Unemployed workers observe all the available contracts and decide where to apply for a job. In equilibrium, the expected number of unemployed workers applying for each job adjusts so that a worker gets the same expected utility U_H regardless of which job she applies for. According to equation (2), this implies that the probability the worker is hired, an increasing function of the vacancy-unemployment ratio θ associated with a particular job, $\lambda^{ue}(\theta_H)$, adjusts so as to offset any difference in the expected capital gain $\mathbb{E}_H(E_{H'} - U_{H'})$ across wage contracts. Since a firm's probability of hiring a worker is a decreasing function $q(\theta_H)$, firm profits are higher when θ_H and hence $\lambda^{ue}(\theta_H)$ are lower. Conversely, for workers to get a bigger capital gain $\mathbb{E}_H(E_{H'} - U_{H'})$, wages must be higher on average, which reduces firm profits. Firms recognize this tradeoff when choosing their wage contract. Finally, free entry drives the maximum possible value of a vacancy to zero.

Although in principle the wage contract space is complicated, the risk-neutrality assumption implies that workers only care about the expected present value of wages and that likewise the expected cost of the contract to the firm is the expected present value of wages. In other words, contracts are effectively distinguished only according to how they divide the surplus from a match $S_H \equiv F_H + E_H - U_H$ between the worker and firm. As before, we can eliminate $\mathbb{E}_H E_{H'}$ between equations (2) and (3) to get equation (8) and eliminate $\mathbb{E}_H F_{H'}$ between equations (4), (5), and the free entry condition to get (9). Summing these equations gives

$$S_H = p_H + \frac{1-\lambda_H^{eu}}{q_H}c - \frac{1-\lambda_H^{eu}}{\lambda_H^{ue}}b + \frac{1-\lambda_H^{eu}-\lambda_H^{ue}}{\lambda_H^{ue}}(U_H - \beta\mathbb{E}_H U_{H'}), \tag{18}$$

confirming that match surplus is independent of the wage. If the firm takes surplus F_H from the match, the worker gets surplus $E_H - U_H = S_H - F_H$.

To summarize, in a competitive search equilibrium, firms optimally set wages to maximize the value of a vacancy. This replaces the Nash bargaining assumption (6). Otherwise, the model is unchanged from Section 3.

4.2 Equilibrium Analysis

A firm with an open vacancy chooses a vacancy-unemployment ratio θ and an expected capital gain $\mathbb{E}_H F_{H'}$ to maximize the value of the vacancy subject to the constraint that the value of an unemployed worker satisfies equation (2), where $\mathbb{E}_H(E_{H'} - U_{H'}) = \mathbb{E}_H(S_{H'} - F_{H'})$. Moreover, the maximized value of a vacancy is driven to zero by free entry. This can be expressed as a constrained maximization problem:

$$c = \max_{\theta, \mathbb{E}_H F_{H'}} \beta q(\theta) \mathbb{E}_H F_{H'}$$

$$\text{subject to } U_H - \beta \mathbb{E}_H U_{H'} = b + \beta \lambda^{ue}(\theta) \mathbb{E}_H (S_{H'} - F_{H'})$$

Use the constraint to eliminate $\mathbb{E}_H F_{H'}$ from the objective function:

$$c = \max_{\theta} \frac{1}{\theta} \big(\beta \lambda^{ue}(\theta) \mathbb{E}_H S_{H'} + b - (U_H - \beta \mathbb{E}_H U_{H'})\big), \tag{19}$$

where the expression is further simplified with the identity $q(\theta) \equiv \lambda^{ue}(\theta)/\theta$. The first order condition from this maximization problem is

$$\frac{1}{\theta} \big(\beta \lambda^{ue}(\theta) \mathbb{E}_H S_{H'} + b - (U_H - \beta \mathbb{E}_H U_{H'})\big) = \beta \lambda^{ue\prime}(\theta) \mathbb{E}_H S_{H'} \tag{20}$$

Assuming $m(\frac{u}{v}, 1)$ is a concave function of $\frac{u}{v} = 1/\theta$, the objective is a concave function of $1/\theta$, so the first order conditions are both necessary and sufficient for a maximum. In particular, all firms with vacancies choose the same value of θ_H conditional on the history H.

Now eliminate $\mathbb{E}_H S_{H'}$ between (19) and (20) and solve for $U_H - \beta \mathbb{E}_H U_{H'}$:

$$U_H - \beta \mathbb{E}_H U_{H'} = b + \left(\frac{\lambda^{ue}(\theta_H)}{\lambda^{ue\prime}(\theta_H)} - \theta_H \right) c.$$

Then substitute this into (18) to get

$$S_H = p_H - b + \left(\frac{1 - \lambda_H^{eu} - \lambda^{ue}(\theta_H)}{\lambda^{ue\prime}(\theta_H)} + \theta_H \right) c$$

Lead this equation by one period and take expectations to get an expression for $\mathbb{E}_H S_{H'}$. Since (19) and (20) imply $c = \beta\lambda^{ue\prime}(\theta_H)\mathbb{E}_H S_{H'}$, combining this with the previous equations gives a difference equation for θ_H:

$$c = \beta\lambda^{ue\prime}(\theta_H)\mathbb{E}_H\left(p_{H'} - b + \left(\frac{1-\lambda^{eu}_{H'}-\lambda^{ue}_{H'}}{\lambda^{ue\prime}(\theta_{H'})} + \theta_{H'}\right)c\right)$$

In order to compare this with (11), it is useful to define $\varepsilon(\theta) \equiv \frac{d\log m(u,v)}{du} = 1 - \frac{\lambda^{ue\prime}(\theta)\theta}{\lambda^{ue}(\theta)}$, the elasticity of the number of matches with respect to the unemployment rate. Use this to eliminate $\lambda^{ue\prime}$ from the previous equation, simplifying with $\lambda^{ue}(\theta_H)/\theta_H = q_H$:

$$c = \beta q_H(1-\varepsilon(\theta_H))\mathbb{E}_H\left(p_{H'} - b + \left(\frac{1-\lambda^{eu}_{H'}-\lambda^{ue}_{H'}}{q_{H'}(1-\varepsilon(\theta_{H'}))} + \theta_{H'}\right)c\right). \tag{21}$$

If the elasticity of the matching function with respect to the unemployment rate $\varepsilon(\theta)$ is equal to workers' 'bargaining power' γ for all θ, this is identical to equation (11). In this case, the competitive search model effectively endogenizes workers bargaining power and there is no observable difference between the behavior of the benchmark and competitive search models. But to the extent that the elasticity of the matching function $\varepsilon(\theta)$ is decreasing in θ, the competitive search model offers an explanation for why workers' 'bargaining power' might appear to decrease in expansions, reducing real wage cyclicality.

To see when this is the case, consider the constant elasticity of substitution (CES) matching function $m(u,v) = \mu(\alpha u^\rho + (1-\alpha)v^\rho)^{1/\rho}$, with Cobb-Douglas limit $\mu u^\alpha v^{1-\alpha}$ when ρ converges to zero, and with parameter restrictions $\rho < 1$, $\alpha \in (0,1)$, and $\mu > 0$.[17] Higher values of ρ are associated with greater substitutability between unemployment and vacancies. One can verify that this function is concave and so the first order conditions are necessary and sufficient. It is straightforward to show that in this case $\varepsilon(\theta) = \frac{\alpha}{\alpha+(1-\alpha)\theta^\rho}$. When ρ is positive, ε is decreasing in labor market tightness θ, and so is countercyclical.

Since the model does not pin down the timing of wage payments, I require an additional assumption to discuss the cyclicality of real wages. One that makes it relatively easy to link the empirical evidence with the model is that all firms pay the same wage w_H after any payoff relevant history H. In this case, we can eliminate $\mathbb{E}_H F_{H'}$ between (4) and (5) yielding equation (9). Lead that equation by one period

[17]The matching function should satisfy $m(u,v) < \min\langle u,v\rangle$ as well, so this is best viewed as a local approximation.

and take expectations. Then use this expression to eliminate $\mathbb{E}_H F_{H'}$ from (4):

$$c = \beta q(\theta_H)\mathbb{E}_H \left(p_{H'} - w_{H'} + \frac{1-\lambda^{eu}_{H'}}{q(\theta_{H'})} c \right). \tag{22}$$

Since there is one such equation for each payoff-relevant history H, this can be inverted to solve for the wage. Unfortunately, I cannot find an explicit expression for w_H. Of course, even without this assumption, it is possible to discuss the required magnitude of cyclical variation in profits.

4.3 Steady State

In a deterministic steady state, equation (21) reduces to

$$q = \frac{\big(1-\beta(1-\lambda^{eu}-\varepsilon(\theta)\lambda^{ue})\big)c}{\beta(1-\varepsilon(\theta))(p-b)}, \tag{23}$$

analogous to equation (12). As usual, focus on the no-discounting limit and substitute the steady state condition $\lambda^{eu} = \frac{u}{1-u}\lambda^{ue}$:

$$\theta = \frac{(1-\varepsilon(\theta))(p-b)(1-u)}{(u+\varepsilon(\theta)(1-u))c}.$$

In the case of a CES matching function, this becomes

$$\theta = \frac{(1-\alpha)\theta^{\rho}(p-b)(1-u)}{(u(1-\alpha)\theta^{\rho}+\alpha)c}.$$

With a Cobb-Douglas matching function, $\rho = 0$, this reduces to the Beveridge curve expression (15). But in the limit as ρ approaches 1, it implies

$$v = \frac{(p-b)(1-u)}{c} - \frac{\alpha}{1-\alpha}.$$

Even if fluctuations are driven by the job destruction rate λ^{eu}, there will be a weak negative correlation between vacancies and unemployment. This contrasts with the result that with job destruction shocks, unemployment and vacancies are proportional in the standard model.

This is a consequence of the behavior of wages, which no longer fall when the job destruction rate increases. In steady state, equation (22) is easily solved to give

$$w = p + \left(1 - \lambda^{eu} - \frac{1}{\beta}\right)\frac{c}{q}.$$

Eliminate p using equation (23):

$$w = b + \frac{\big(1 - \beta(1 - \lambda^{eu} - \lambda^{ue})\big)\varepsilon(\theta)c}{\beta(1 - \varepsilon(\theta))q}.$$

This is equivalent to (16) with workers' bargaining power γ equal to the elasticity ε. Take the limit as β converges to one, replace $\lambda^{eu} + \lambda^{ue} = \frac{\lambda^{ue}}{1-u}$, and focus on the case of a CES matching function:

$$w = b + \frac{\alpha c \theta^{1-\rho}}{(1-\alpha)(1-u)}.$$

Once again, this reduces to equation (17) if $v \equiv k/u$ and $\rho = 0$, and so wages are strongly procyclical. On the other hand, if ρ is sufficiently close to 1, this equation implies that wages are mildly countercyclical.

Thus a high elasticity of substitution between unemployment and vacancies improves the behavior of the model on two counts: first, there may be a negative correlation between unemployment and vacancies even if fluctuations are driven primarily by job destruction shocks; and second wages are less procyclical and possibly even countercyclical. The two phenomena are related. In the standard model, an increase in job destruction shocks leads to a decline in wages. From the firm's perspective, the two effects nearly cancel, and so the incentive to create a job, summarized by the vacancy-unemployment ratio, declines only slightly. But in the competitive search model, this decline in the vacancy-unemployment ratio effectively raises workers' bargaining power, increasing wages and further reducing the vacancy-unemployment ratio. With an infinite elasticity of substitution, wages actually increase when the job destruction rate increases. This implies that vacancies decline in the face of rising unemployment. Similarly, with a high elasticity of substitution, the productivity shock required for a given movement in the unemployment rate is much smaller, since lower wages no longer offset lower productivity in the firms' profit function.

A matching function with a high elasticity of substitution between unemployment and vacancies is theoretically appealing. Suppose each unemployed worker can contact a vacancy at rate α_u, while each vacancy can contact an unemployed worker at rate α_v. Then the total flow of matches is $m(u, v) = \alpha_u u + \alpha_v v$, the CES matching function in the limit with $\rho = 1$. Moreover, the fraction of matches initiated by workers is $\varepsilon(\theta) = \frac{\alpha}{\alpha + (1-\alpha)\theta}$, where $\alpha = \frac{\alpha_u}{\alpha_u + \alpha_v}$ is the relative efficiency of search by workers compared to firms. Thus another way to think of wage setting in this limit of the competitive search equilibrium is that the party who initiates a match can make a take-it-or-leave it offer to the other agent. Since workers initiate relatively more matches in recessions, this effectively gives workers countercyclical bargaining power $\varepsilon(\theta)$. Such a matching function and wage setting arrangement has an excellent pedigree in the theoretical search literature (Mortensen 1982).

4.4 Empirical Evidence on Matching Functions

Empirically, there is less evidence on whether this is a reasonable matching function. To my knowledge, the only previous attempt to estimate the elasticity of substitution between unemployment and vacancies is Blanchard and Diamond (1989), who estimate $\rho = -0.35$, i.e. unemployed workers and vacancies are somewhat less substitutable than in the Cobb-Douglas case.[18] This section updates Blanchard and Diamond's (1989) finding by estimating the coefficients of a CES matching function

$$\log m_{t+1} = \log \mu + \frac{\log(\alpha u_t^\rho + (1-\alpha) v_t^\rho)}{\rho} + \nu_t \tag{24}$$

using nonlinear least squares. I measure m_{t+1} by the number of workers who move from unemployment to employment between month t and $t+1$, u_t by the number of unemployed workers, and v_t by the help-wanted index. ν_t is an independent and identically distributed error term.

Table 3 shows the results using U.S. data from June 1967 to December 2001. The column (1) estimates are unrestricted, and indicate that ρ is insignificantly negative. In column (2), I run the regression with the restriction $\rho = 0$, the Cobb-Douglas case. Finally, column (3) shows the results with the restriction $\rho = 1$, perfect substitutability. Most noticeably, I cannot reject the null hypothesis that $\rho = 0$, i.e. that the competitive search and standard search models are indistinguishable. On the other hand, the data does not speak very strongly on this issue. Even when I impose $\rho = 1$, time variation in u and v explains almost 88 percent of the variation in the number of matches, 1 percent less than in the unrestricted regression.

Some of this explanatory power comes from the increasing size of the U.S. economy during the sample period, which raises both the number of unemployed workers and the number of job finders. To compensate for that, I subtract $\log u_t$ from both sides of equation (24). Since workers' job finding rate λ_{t+1}^{ue} is the ratio of the number of

[18] Petrongolo and Pissarides (2001) review the large literature on matching function estimation. Most of the literature estimates a Cobb-Douglas matching function, i.e. imposes the restriction $\rho = 0$, and focuses on whether matching functions exhibit constant returns to scale. I impose constant returns throughout, as the competitive search model only makes sense under this restriction.

matches to the number of unemployed workers and θ_t is the vacancy-unemployment ratio, this implies that I estimate

$$\log \lambda_{t+1}^{ue} = \log \mu + \frac{\log(\alpha + (1-\alpha)\theta_t^{\rho})}{\rho} + \nu_t.$$

Although the coefficient estimates are unchanged by this manipulation, it reduces the variation in the dependent variable and so lowers the R^2 of the regression. Nevertheless, I continue to explain 69 percent of the variation in job finding rates through the vacancy-unemployment ratio alone in the unrestricted regression and 66 percent in the regression with the restriction $\rho = 1$.

I have performed a number of other robustness checks, for example instrumenting the right hand side variables using their lagged values to compensate for measurement error, aggregating the data up to annual averages, estimating on sub-samples, and allowing for time trends. Although sometimes the sign of ρ switches, it is never significantly different from zero and the point estimates are usually very small. Indeed, the magnitude of the standard errors is robust to corrections for autocorrelation in the residual or to the inclusion of lagged dependent variables in the regression. I conclude that although competitive search equilibrium and a high elasticity of substitution between unemployment and vacancies in the matching function could theoretically provide a reason why workers' bargaining share appears to be countercyclical, there is no empirical support for this explanation.

5 Literature Review

A growing body of papers has attempted to use search models to think quantitatively about employment fluctuations; however, none has been able to reconcile the first two labor market facts, that wages are nearly acyclical and that vacancies and unemployment are negatively correlated over the business cycle. In part this appears to be due to a misconception that job finding rates do not fluctuate cyclically, which is only possible if the vacancy-unemployment ratio is acyclical.

A number of recent papers 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. Gomes, Greenwood and Rebelo (2001) sidestep the vacancy-unemployment issue completely by looking at a model without a notion of vacancies. Still, they report that average labor productivity, equal to the average wage in their economy, has a correlation of 0.87 with output and a standard deviation that is nearly as large as the standard deviation in employment. Ramey and Watson (1997) emphasize the importance of contractual fragility and job destruction for unemployment fluctuations. Indeed, in order to focus on job destruction, they assume that the job creation rate is exogenous and acyclic, which is equivalent to assuming that vacancies are proportional to unemployment. It is likely that this is why they find that shocks to the economy are amplified and wages are rigid.

Hall (1995) argues that "...brief, sharp episodes of primary job loss are followed by long periods of slowly rebuilding employment relationships over the business cycle. Although the case is far from complete, I believe that these events in the labor market play an important part in the persistence of high unemployment and low output long after the initial shock that triggers a recession" (p. 221). Building on this idea, Pries (2001) shows that a brief adverse shock that destroys some old employment relationships can generate a long 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. Pries emphasizes that this model generates persistent employment fluctuations in response to a one-time shock. This framework succeeds in generating relatively acyclical wages, since wages are constant during the long recovery period. However, this only happens because the vacancy-unemployment ratio remains counterfactually constant though the business cycle.

Abraham and Katz (1986) argue that the negative correlation between unemployment and vacancies is inconsistent with Lilien's (1982) sectoral shifts hypothesis: periods of high unemployment represent episodes of high frictional unemployment as workers move across sectors of the economy. This is quite closely related to the possibility that high unemployment is due to an increase in job destruction not a decrease in productivity. This paper confirms that if wages do not fluctuate in response to job destruction shocks, as is the case in the standard bargaining model of search unemployment, unemployment and vacancies will be positive correlated. But it also shows that Abraham and Katz's (1986) preferred aggregate disturbances imply implausible wage dynamics in standard models of equilibrium unemployment. Abraham and Katz do not discuss the wage implications of these models.

Likewise, papers by Blanchard and Diamond (1989), Mortensen and Pissarides (1994), and Cole and Rogerson (1999) match the behavior of labor market quantities, e.g. unemployment and vacancies, but sidestep wage issues entirely. For example, Blanchard and Diamond assume that the total supply of active jobs follows a slow-moving and exogenous Markov process in response to aggregate shocks, gradually increasing during expansions and then declining during recessions.[19] While this provides some guidance as to how to generate realistic Beveridge curve dynamics, it is an unsatisfactory explanation for labor market fluctuations because it does not explain the nature of the aggregate shocks nor does it explain why firms cannot create more vacancies when the stock of active jobs is unusually low. In a model with wages, recessions would likely be an attractive time to create jobs. While I concur with their conclusion that fluctuations are due to gradual movements in the supply of active jobs, I am unable to reconcile this in a satisfactory manner with the weak cyclicality of real wages. The Mortensen and Pissarides (1994) model obviously has implications for wages, but in their numerical work, they focus exclusively on the cyclical behavior of job creation and job destruction.

6 Conclusion

Standard models of equilibrium unemployment imply that because vacancies and unemployment are negatively correlated over the business cycle, real wages must be extremely procyclical. Moreover, such responsiveness of real wages to shocks implies that in order to generate the observed fluctuations in employment, productivity must also be orders of magnitude more procyclical than the data suggests. The competitive search equilibrium offers one explanation for why this might not be the case. If the elasticity of substitution between unemployment and vacancies in the matching function is larger than one, workers' bargaining power effectively increases in recessions, reducing or even eliminating the cyclicality of real wages and amplifying shocks to the economy. But since there is no evidence that the elasticity of substitution is larger than one, one must look for other explanations for the weak cyclicality of real wages.

One possibility is that within the worker-firm relationship, 'fairness' is as important as outside options for the determination of wages. Careful studies of wage determination offer some support for this hypothesis (Bewley 1999). In the context of search models, perhaps this should not be too surprising. Workers and firms are generally in a bilateral monopoly situation, so there is some possibility of altering the wage

[19] 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.

without affecting the incentive of worker to take a job or the firm to employ her. Fairness might dictate that wages are unaffected by business cycle conditions unless such rigidity would generate an inefficient layoff or quit. By construction, this behavior is inconsequential at the level of the individual employment relationship. But at the aggregate level, this real wage rigidity matters a lot. Firms will be reluctant to hire workers when productivity is low if they perceive that the 'fair' wage has not declined. The decline in job creation thus amplifies the otherwise small effects of a typical negative productivity shock. While it is clear that such behavior can reconcile the labor market facts described in this paper, one would like the fairness hypothesis to have testable predictions — along the lines of the predictions of the bargaining and competitive search models analyzed here — and for those predictions to receive empirical support, before concluding that observed real wage rigidities are a privately efficient but socially inefficient response to fairness concerns.

A Composition Biases in Wage Measures

It is well-known that compositional effects create a countercyclical bias in both AHE and AHC. Less skilled and less experienced workers earn lower wages and are more likely to lose their job during downturns. As this occurs, the composition of the labor force shifts towards more skilled and more experienced workers, raising measured AHE and AHC even if no individual's wage actually changes. It is less frequently recognized that compositional effects also introduce a procyclical bias in these measures, because high wage blue collar industries like durable goods manufacturing and construction are highly procyclical.[20] The ECI is designed in part to address the latter composition bias issues.[21] The important difference between the ECI and other measures of average wages is that the ECI is a Laspeyres price index. In calculating ECI inflation from quarter to quarter, the composition of jobs is held fixed at the beginning-of-period value and used as weights in the computation of the overall employment cost inflation. 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.

The ECI has some drawbacks. Most obvious is the short sample period available, from the third quarter of 1975 to the end of 2001. Since the period includes four recessions and some dramatic fluctuations in unemployment, however, this problem is

[20] Barsky and Solon (1989) show that wages are less cyclical within industries than in the aggregate data.

[21] 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.

not as significant as it once was. Another disadvantage of the ECI is its treatment of overtime hours. The ECI assumes that the usage rate of overtime work and the associated additional labor costs are fixed at an initial rate even as aggregate hours fluctuate. Since overtime hours are strongly procyclical, this mutes the volatility of the ECI. Moreover, to the extent that firms use overtime hours as an (imperfect) means of increasing wage cyclicality, the ECI will miss out on that practice. Finally, if firms link promotions to wages, the ECI will understate the cyclicality of firm's wage policies. As an extreme example, suppose firms associate a time-invariant real wage with every job and give workers raises only by giving them promotions. Then regardless of whether promotions are strongly procyclical, the ECI would always measure a constant employment cost because it tracks jobs not individuals.

The finding in the text that the ECI is less procyclical than AHE might appear to contradict a well-known paper by Solon et al. (1994), which uses longitudinal data from the Panel Study of Income Dynamics (PSID) to examine the cyclical behavior of the change in wages for individuals who keep their job from one year to the next. They conclude that accounting for composition bias by following individuals over time doubles the estimated effect of the unemployment rate on average wage cyclicality, so that in their sample, a one percentage point decline in the unemployment rate is associated with more than a one percent increase in wages.[22] It is worth noting that for the purpose of this paper, the difference between these results is quantitatively irrelevant. Section 3 shows that a standard model of equilibrium unemployment predicts a semi-elasticity of wages with respect to unemployment of about 30 percent, so the model is inconsistent with the data whether the data indicates a semi-elasticity of half a percent or 1 percent. In any case, the different results in part reflect attempts to answer a different question. I am interested in whether firms reduce wages during cyclical downturns, while Solon et al. (1994) are interested in whether workers get lower wages during downturns. If workers move from high-wage to low-wage sectors of the economy but firms do not change wages, I would conclude that firms do not reduce wages while Solon et al. would conclude that workers get lower wages.[23]

[22] Other papers that have looked at longitudinal data find that composition bias is less important. See especially Bils (1985).

[23] See the example on pages 26–27 of Barsky and Solon (1989) for similar logic.

References

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

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

___ **and John Haltiwanger**, "Real Wages and the Business Cycle," *Journal of Economic Literature*, September 1995, *33*, 1215–1264.

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

___ **and Robert Shimer**, "Changes in Unemployment Duration and Labor-Force Attachment," in Alan Krueger and Robert Solow, eds., *The Roaring Nineties: Can Full Employment be Sustained?*, New York: Russell Sage Foundation, 2001, chapter 8, pp. 367–420.

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

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

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

Barsky, Robert and Gary Solon, "Real Wages over the Business Cycle," 1989. NBER Working Paper 2888.

Bewley, Truman, *Why Wages Don't Fall in a Recession*, Cambridge, MA: Harvard University Press, 1999.

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

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

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

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

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

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

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

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

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

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

Madrian, Brigitte and Lars Lefgren, "An Approach to Longitudinally Matching Current Population Survey (CPS) Respondents," *Journal of Economic and Social Measurement*, 2000, *26* (1), 31–62.

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

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

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

Mortensen, Dale, "Property Rights and Efficiency in Mating, Racing, and Related Games," *American Economic Review*, Dec 1982, *72* (5), 968–979.

___ **and Christopher Pissarides**, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 1994, *61*, 397–415.

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

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

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

___ , "Unemployment and Vacancies in Britain," *Economic Policy*, 1986, *3*, 499–559.

___ , "Search, Wage Bargains and Cycles," *Review of Economic Studies*, 1987, *54*, 473–483.

___ , *Equilibrium Unemployment Theory*, second ed., Cambridge, MA: MIT Press, 2000.

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

Pries, Michael, "Persistence of Employment Fluctuations: A Model of Recurring Job Loss," 2001. Mimeo.

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

Ruser, John, "The Employment Cost Index: What is it?," *Monthly Labor Review*, September 2001, pp. 3–21.

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

Solon, Gary, "Effects of Rotation Group Bias on Estimation of Unemployment," *Journal of Business and Economic Statistics*, January 1986, *4* (1), 105–109.

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

Subcommittee on Economic Statistics of the Joint Economic Committee, *Job Vacancy Statistics* 89th Congress, 2nd. Session, Washington: USGPO, 1966.

	(1)	(2)	(3)	(4)
sample period	1947–2001	1947–2001	1964–2001	'64M1–'01M12
measure	AHE	AHE	AHE	AHE
estimation method	level	first diff	first diff	first diff
unemployment rate	-0.53	-0.36	-0.59	-0.19
	(0.23)	(0.18)	(0.24)	(0.09)
R^2	0.087	0.067	0.142	0.010
Durbin-Watson	0.61	1.55	1.26	1.99
Nobs	55	54	38	455

	(5)	(6)	(7)	(8)
sample period	1947–2001	'48Q1–'01Q4	'75Q3–'01Q4	'75Q3–'01Q4
measure	AHC	AHC	ECI	AHE
estimation method	first diff	first diff	first diff	first diff
unemployment rate	-0.03	-0.13	0.10	-0.25
	(0.19)	(0.12)	(0.20)	(0.17)
R^2	0.000	0.005	0.002	0.022
Durbin-Watson	1.56	1.97	1.68	1.06
Nobs	54	215	105	105

	(9)	(10)	(11)
sample period	'75Q3–'01Q4	'48Q1–'01Q4	'48Q1–'01Q4
measure	AHC	APL	NLP
estimation method	first diff	first diff	first diff
unemployment rate	0.15	-0.34	-0.82
	(0.20)	(0.18)	(0.40)
R^2	0.005	0.005	0.019
Durbin-Watson	1.34	2.25	2.04
Nobs	105	215	215

Table 1: The Unemployment Rate is from the Current Population Survey. Average Hourly Earnings of production and non-supervisory workers in private non-farm employment (AHE) is from the Current Employment Statistics. Average Hourly Compensation (AHC), Average Product of Labor (APL), and Non-Labor Payments per Hour (NLP) in the non-farm business sector are from the National Income and Product Accounts. Employment Cost Index for private industry workers (ECI) is from the National Compensation Survey. AHE, AHC, APL, NLP, and ECI are deflated using the consumer price index for all urban consumers. The unemployment rate is expressed as an absolute deviation from trend and AHE, AHC, APL, NLP, and ECI are log deviations from trend, so the coefficient estimates are semi-elasticities. To obtain similar trends using data of different frequencies, the trend of monthly series is given by an HP filter with smoothing parameter 10^7, quarterly series by smoothing parameter 10^5, and annual series by smoothing parameter 450.

		Employment-Population Ratio							
		1	2	3	4	5	6	7	8
	mean	0.6195	0.6140	0.6128	0.6130	0.6122	0.6099	0.6097	0.6115
Unemployment Rate 1	0.0682	—	231.7	311.0	210.9	402.2	709.6	635.4	441.3
2	0.0646	198.9	—	10.5	0.1	23.1	124.1	133.7	39.9
3	0.0635	319.6	17.9	—	9.6	2.2	67.9	75.2	11.2
4	0.0648	155.2	0.4	23.1	—	19.2	116.3	130.9	41.9
5	0.0643	223.8	1.4	10.3	2.9	—	44.3	46.8	3.6
6	0.0625	437.0	72.2	13.2	76.4	52.9	—	0.4	19.5
7	0.0616	701.1	167.4	54.0	145.9	127.3	13.0	—	22.7
8	0.0633	336.5	25.5	0.6	34.7	18.0	9.8	47.5	—

Table 2: The first row shows the mean employment ratio in each of the eight CPS rotation groups, 1976 to 2001. The upper triangle shows F(1,311) statistics for pairwise tests of equality of the underlying series. The first column shows the corresponding figure for the unemployment rate and the lower triangle shows F(1,311) statistics for pairwise tests of equality. The five percent critical value is 3.9 and the one percent critical value is 6.7. The underlying time series were extracted by the author from the CPS and are not seasonally adjusted.

Dependent Variable: matches m_{t+1}

	(1)	(2)	(3)
$\log \mu$	-0.271 (0.280)	0.039 (0.044)	1.929 (0.035)
α	0.833 (0.107)	0.701 (0.010)	0.028 (0.001)
ρ	-0.168 (0.171)	0*	1*
R^2	0.888	0.887	0.875
D-W	1.749	1.745	1.574
Nobs	405	405	405

Table 3: The dependent variable is the log of the number of unemployed workers moving to employment in a month, seasonally adjusted. Independent variables include the unemployment rate and the Conference board help-wanted index, both seasonally adjusted. The specification is given by equation (24). Column (1) is unrestricted and is estimated using nonlinear least squares. Column (2) imposes $\rho = 0$, i.e. $\log m_{t+1} = \log \mu + \alpha \log u_t + (1-\alpha) \log v_t + \nu_t$ and is estimated using ordinary least squares. Column (3) imposes $\rho = 1$ and is estimated using NLS. Sample period is June 1967 to December 2001. Standard errors are in parentheses.

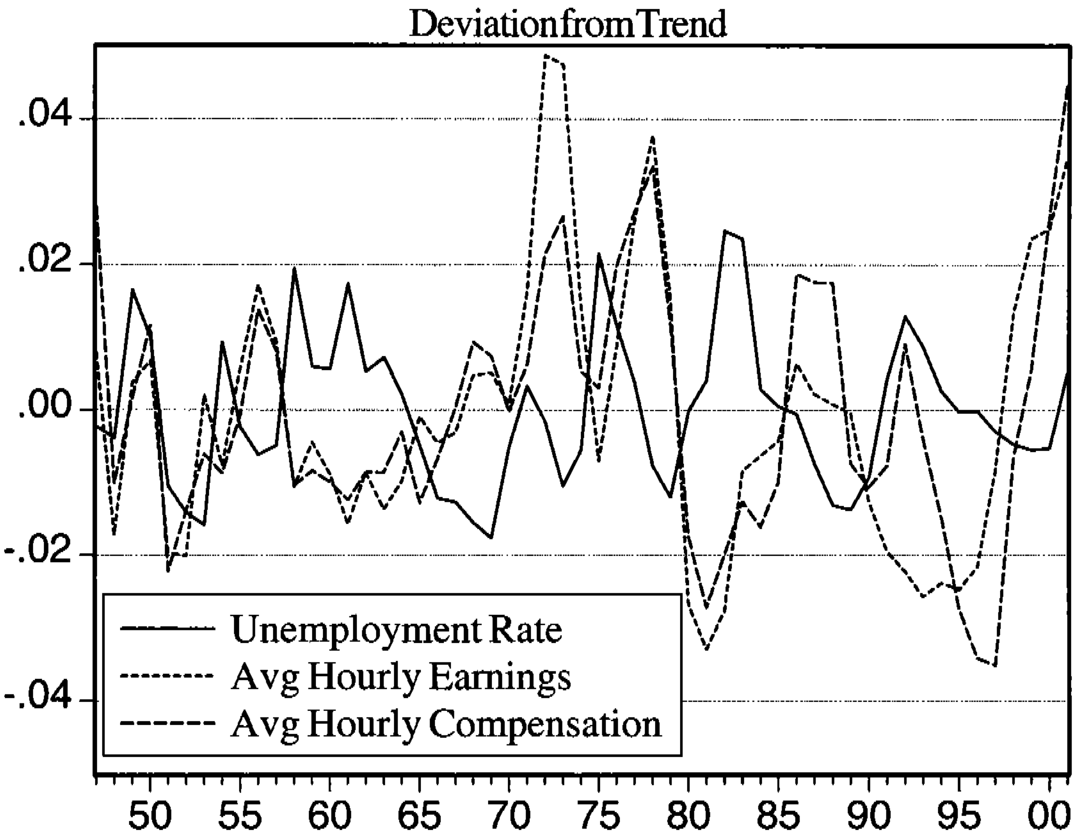

Figure 1: The Unemployment Rate from the Current Population Survey, Average Hourly Earnings of production and non-supervisory workers in private non-farm employment from the Current Employment Statistics, and Average Hourly Compensation in the non-farm business sector from the National Income and Product Accounts. AHE and AHC are deflated using the consumer price index for all urban consumers. The unemployment rate is expressed as an absolute deviation from trend and AHE and AHC are log deviations from trend. In all cases, the trend is given by an HP filter of the annual data with smoothing parameter 450, a very low frequency filter.

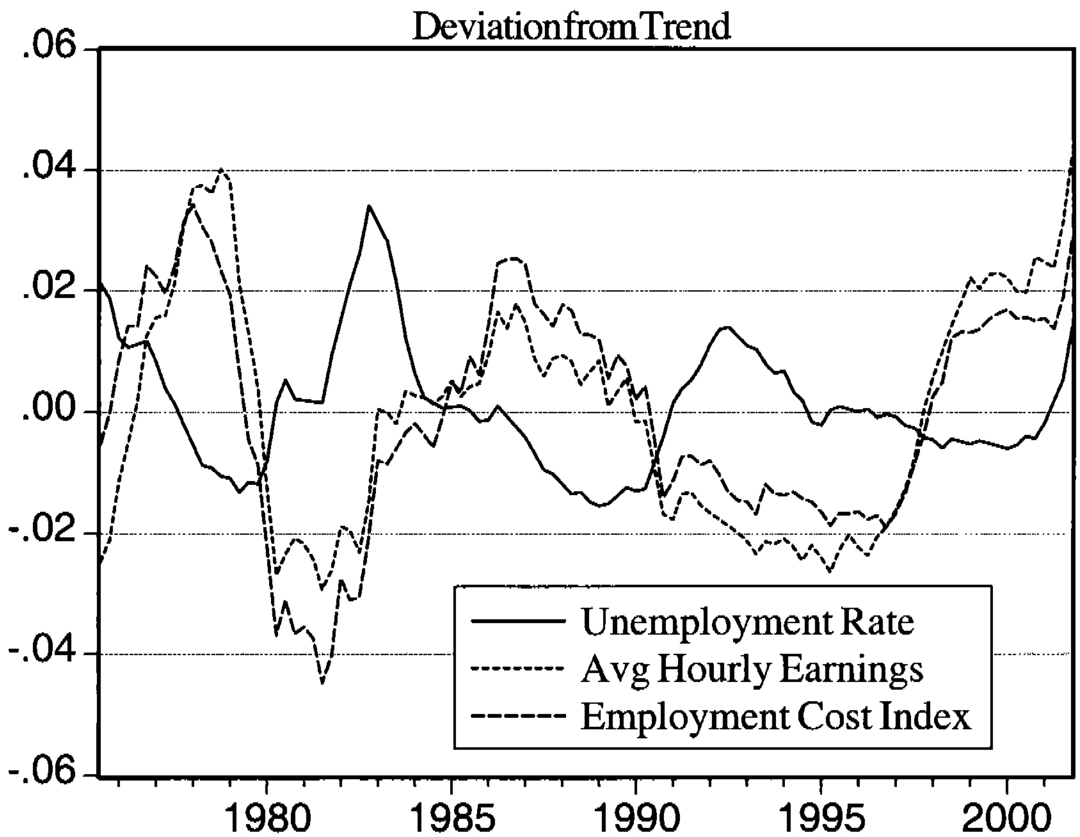

Figure 2: The Unemployment Rate from the Current Population Survey, Average Hourly Earnings of production and non-supervisory workers in private non-farm employment from the Current Employment Statistics, and the Employment Cost Index for private industry workers. AHE and ECI are deflated using the consumer price index for all urban consumers. The unemployment rate is expressed as an absolute deviation from trend and AHE and AHC are log deviations from trend. In all cases, the trend is given by an HP filter of the quarterly data with smoothing parameter 10^5.

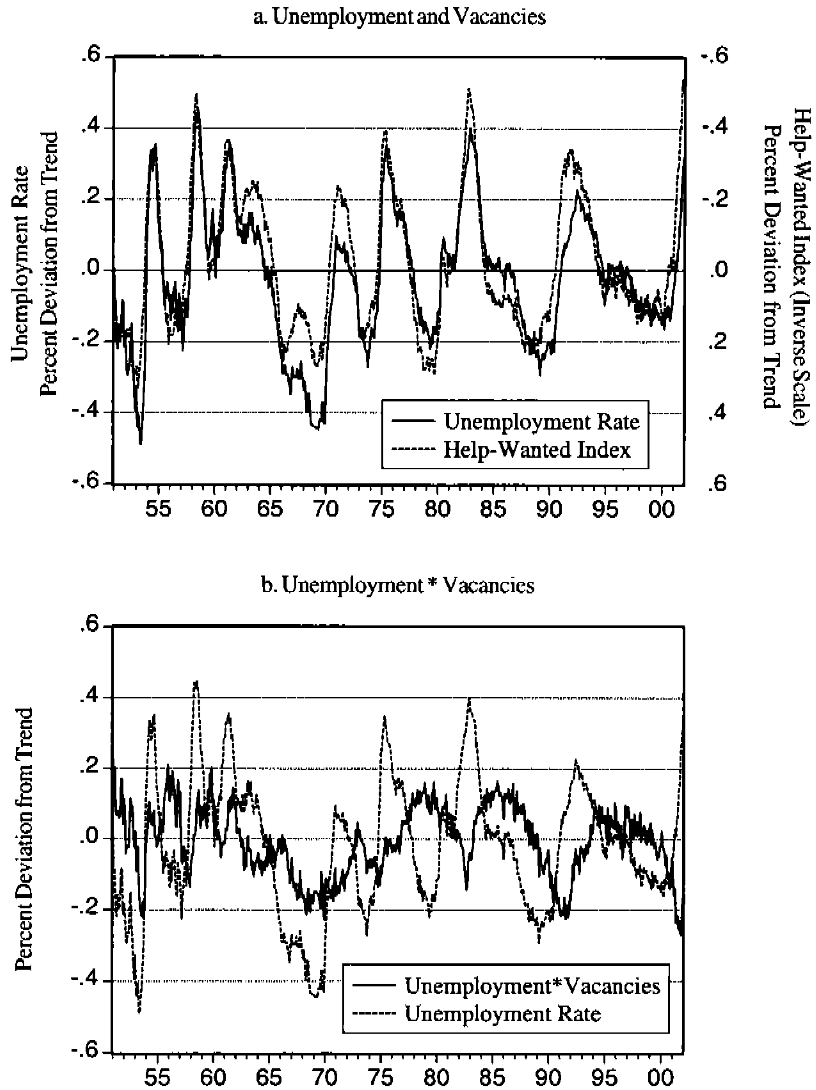

Figure 3: The unemployment rate is measured by the Bureau of Labor Statistics and the Help-Wanted Index by the Conference Board. All three series (unemployment, help-wanted, and their product) are detrended by taking the log deviation from an HP filter with smoothing parameter 10^7, i.e. a very low frequency filter. In panel a, the help-wanted index is plotted on an inverted scale.

Labor Market Transition Rates

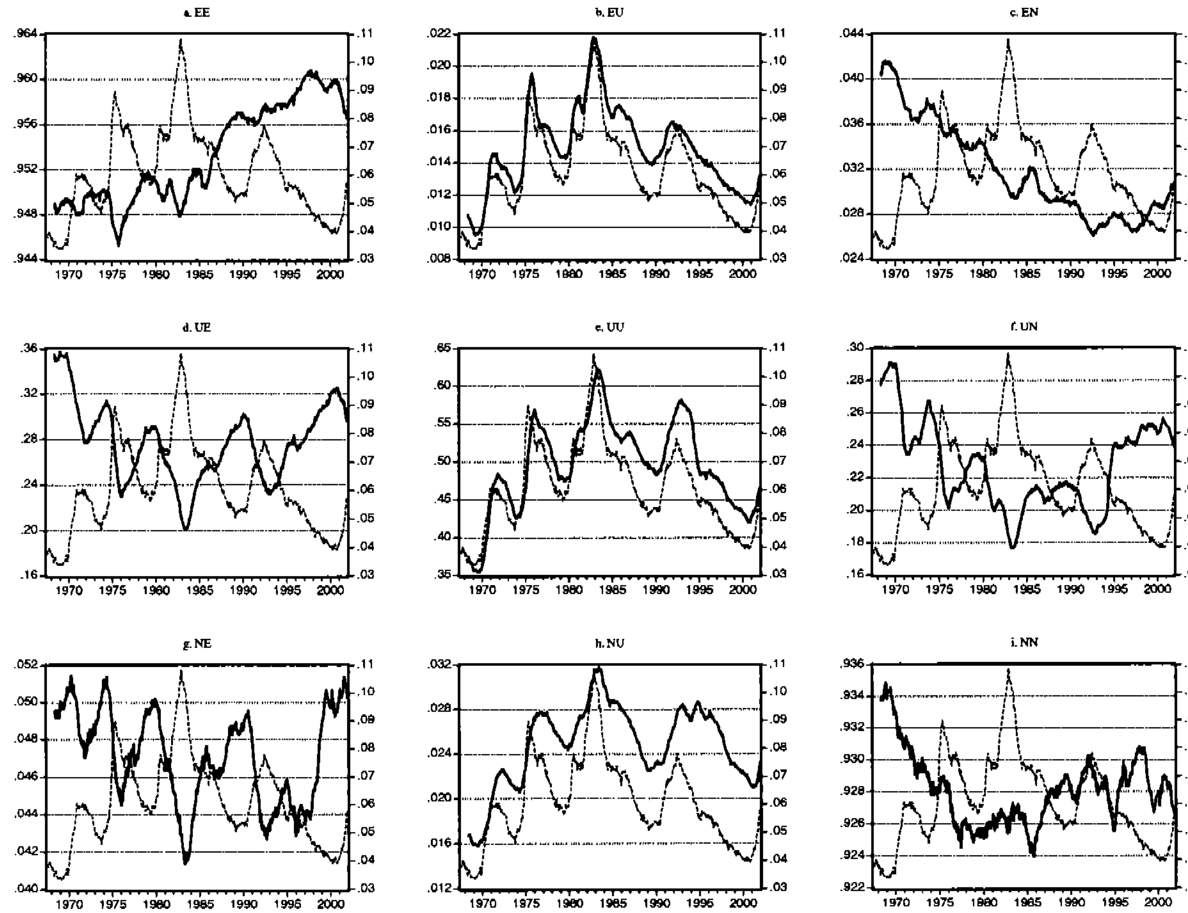

Figure 4: The transition rate between employment, unemployment, and not-in-the-labor force, 12 month moving average (solid line, left scale) and the unemployment rate (dashed line, right scale). Hoyt Bleakley provided me with the transition from June 1967 to December 1975, which were originally constructed by Joe Ritter. I calculated the transition rates from January 1976 to December 2001 from the CPS.

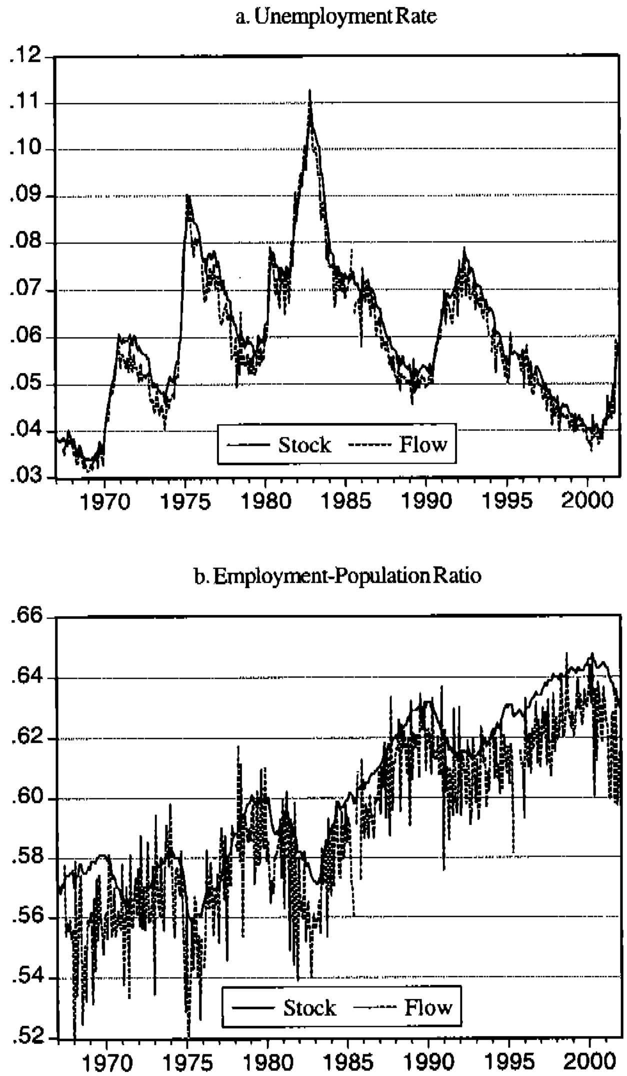

Figure 5: The stock series are monthly Bureau of Labor Statistics measures from the Current Population Survey. The flow series are measured using a 3-state transition matrix constructed from the Current Population Survey. Hoyt Bleakley provided me with the transition rates from 1967 to 1975. The data are seasonally adjusted using a modified ratio-to-moving average technique (see footnote 9).

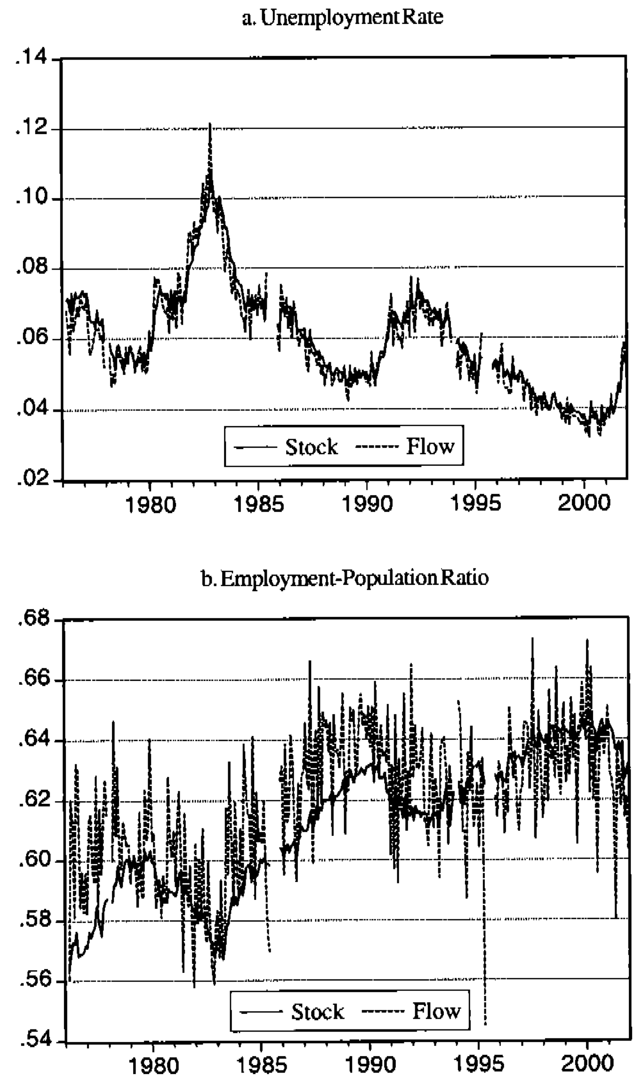

Figure 6: All time series are constructed from the Current Population Survey. The 'stock' series are measured using individuals in rotation groups 4, 7 and 8 who can be matched with records from the previous month. The flow series are measured using a 16-state transition matrix constructed by matching rotation groups 2, 3, and 4; 5, 6, and 7; and 6, 7, and 8, as described in the text. The data are seasonally adjusted using a modified ratio-to-moving average technique (see footnote 9).