The Cyclical Behavior of Labor Markets*

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

August 19, 2003

1 Introduction

A sequence of recent papers (Costain and Reiter 2003, Hall 2003a, Shimer 2003) has argued that the standard theory of equilibrium unemployment, the Mortensen-Pissarides search and matching model (Mortensen and Pissarides 1994, Pissarides 2000) cannot explain the magnitude of the cyclical fluctuations in two of its central elements, unemployment and vacancies. An increase in labor productivity encourages firms to create more vacancies. This reduces the duration of unemployment, which puts upward pressure on wages. In a reasonably calibrated version of the economy, the wage increase absorbs virtually all of the productivity increase, and so the shock has little effect on unemployment and vacancies.

Following a suggestion in Shimer (2003), Hall (2003b) introduces an additional real wage rigidity, a backward-looking social norm, into the Mortensen-Pissarides model. The wage rigidity is socially inefficient: wages are too low in expansions, inducing excessive vacancy creation, and too high in recessions, discouraging most vacancy creation. Nevertheless, Hall (2003b) shows that there are no bilateral gains from renegotiating the wage; every employed worker always prefers to receive a higher wage and every employer always wants to pay a lower wage. The model thereby avoids Barro's (1977) critique of implicit contracting

*Preliminary. This paper will be presented in an invited session at the European Economic Association meetings in Stockholm, Sweden, August 20–24. This material is based upon work supported by the National Science Foundation under grant number 0079345. I am also grateful to the Alfred P. Sloan Foundation for financial support.

models that workers and firms are not exploiting all the potential gains from trade, yet it quantitatively matches the behavior of unemployment and vacancies in the U.S..

This paper first reviews the argument in Shimer (2003) that the Mortensen-Pissarides matching model cannot generate substantial fluctuations in unemployment and vacancies. I then show that with a fixed real wage, the model easily generates large, socially inefficient, fluctuations in these two variables. Finally, I ask a variant of Lucas's (1987) cost-of-business-cycle question: how much would a worker pay to eliminate the real wage rigidity and instead allow wages to vary optimally over the business cycle? The answer is surprisingly little: in a calibrated example, the cost of the real wage rigidity is about 0.1 percent of lifetime consumption, even though the real wage rigidity amplifies fluctuations in unemployment and vacancies by more than a factor of ten. To the extent that policies designed to make real wages more flexible[1] are difficult to implement, unevenly effective, and have unintended costs, this analysis suggests that they are unlikely to be desirable even in an economy in which real wages are otherwise fixed.

2 Flexible Wage Model

The benchmark model extends Pissarides (1985) by making labor productivity p follow a first order Markov process. Because the model has become fairly standard (see Pissarides (2000) for a textbook treatment), I proceed rapidly through the setup.

Time is continuous. The economy consists of a measure 1 of risk-neutral,[2] infinitely-lived workers and a continuum of risk-neutral, infinitely-lived firms. I assume workers and firms discount future payoffs at a common rate $r > 0$, but my numerical results focus on limiting results as $r \to 0$. Workers can either be unemployed or employed. An unemployed worker gets flow utility z from non-market activity ('leisure') and searches for a job. An employed worker earns an endogenous productivity-contingent wage w_p but may not search. Firms have a constant returns to scale production technology that uses only labor; each worker yields profit equal to the difference between labor productivity and the wage, $p - w_p$. Jobs end exogenously at rate $s > 0$, resulting in a separation that leaves the worker unemployed and the firm with a vacancy. In order to hire a worker, a firm must maintain an open vacancy at flow cost c. Free entry drives the expected present value of an open vacancy to zero.

[1] Monetary policy may have such an effect if *nominal* wages are sticky.

[2] Alternatively one can view this as a complete markets model in which labor income risk is insured.

There is a Cobb-Douglas, constant returns to scale matching technology,[3] so that the rate at which unemployed workers find jobs and the rate at which vacancies are filled depends only on the endogenous productivity-contingent vacancy-unemployment ratio θ_p. More precisely, I assume workers find jobs at rate $\mu\theta_p^{1-\alpha}$ and vacancies are filled at rate $\mu\theta_p^{-\alpha}$, where $\alpha \in (0,1)$ is the elasticity of the matching function with respect to the unemployment rate. The unemployment rate $u(t)$ increases with job destruction and decreases when workers find jobs, and so evolves according to

$$\dot{u}(t) = s\big(1 - u(t)\big) - \mu\theta_{p(t)}^{1-\alpha} u(t), \tag{1}$$

where $p(t)$ is the level of labor productivity at time t.

A shock hits the economy according to a Poisson process with arrival rate λ, at which point a new productivity p' is drawn from a distribution that depends on the current productivity level p. Let $\mathbb{E}_p X'_p$ denote the expected value of an arbitrary variable X following the next aggregate shock, conditional on the current state p. I assume that this conditional expectation is finite, which is ensured if the state space is compact. Current productivity and the stochastic process for productivity are common knowledge.

The setup of the model is most easily described through four Bellman equations:

$$rU_p = z + \mu\theta_p^{1-\alpha}\big(E_p - U_p\big) + \lambda\big(\mathbb{E}_p U_{p'} - U_p\big) \tag{2}$$

$$rE_p = w_p + s\big(U_p - E_p\big) + \lambda\big(\mathbb{E}_p E_{p'} - E_p\big) \tag{3}$$

$$rV_p = -c + \mu\theta_p^{-\alpha}\big(J_p - V_p\big) + \lambda\big(\mathbb{E}_p V_{p'} - V_p\big) \tag{4}$$

$$rJ_p = p - w_p + s\big(V_p - J_p\big) + \lambda\big(\mathbb{E}_p J_{p'} - J_p\big) \tag{5}$$

The first pair of equations describe the value of a worker when she is unemployed (U) and employed (E) as a function of the current productivity level p.[4] If she is unemployed, she gets current value from leisure z and finds a job at rate $\mu\theta_p^{1-\alpha}$. There is also an aggregate shock at rate λ, giving a capital gain $\mathbb{E}_p U_{p'} - U_p$. When she is employed, she earns the endogenous wage w_p, loses her job at rate s, and realizes an aggregate shock at rate λ. The

[3] Petrongolo and Pissarides (2001) argue that the matching function exhibits constant returns to scale. There has been less analysis of the Cobb-Douglas assumption, which is central to the interpretation of some of the results that follow. See Blanchard and Diamond (1989) for an estimate of a CES matching function; they cannot reject a unit elasticity of substitution, the Cobb-Douglas case.

[4] These equations implicitly assume that the value functions are independent of the unemployment rate. It is straightforward to show that there is an equilibrium with such a property. In fact, there is no equilibrium in which the value functions depend on the unemployment rate.

second pair of equations similarly describe the value of a job that is vacant (V) or filled (J).

For each productivity level p, there are six endogenous variables within the four equations (2)–(5), four Bellman values, the vacancy-unemployment ratio θ_p, and the wage w_p. To close the model, we need two additional equations. One is the free entry condition, which drives the value of a vacancy to zero,

$$V_p = 0. \tag{6}$$

The other standard assumption, dating back at least to Pissarides (1985), is that wages are set by asymmetric Nash bargaining. I assume here for simplicity that wages are renegotiated following an aggregate shock, which ensures that at any point in time all workers are paid a common wage w_p. Shimer (2003) demonstrates that the behavior of unemployment and vacancies is unchanged if the wage in new jobs is determined by Nash bargaining, while the wage in old jobs may follow an arbitrary contract; for example, the wage in old jobs may remain fixed following an aggregate shock, as long as this does not induce an inefficient separation. In the present context, with linear utility and no on-the-job search, the Nash bargaining assumption amounts to

$$\frac{E_p - U_p}{\beta} = \frac{J_p - V_p}{1-\beta}, \tag{7}$$

where $\beta \in (0,1)$ represents workers' bargaining power.

Since the six equations (2)–(7) are linear in five of the endogenous variables, U_p, E_p, V_p, J_p, and w_p, we can eliminate these variables algebraically to get a forward-looking non-linear difference equation for the vacancy-unemployment ratio:[5]

$$\frac{r+s+\lambda}{\mu\theta_p^{-\alpha}} + \beta\theta_p = \big(1-\beta\big)\frac{p-z}{c} + \lambda\mathbb{E}\frac{1}{\mu\theta_{p'}^{-\alpha}}. \tag{8}$$

Note that only the ratio of market productivity p minus non-market productivity z to the vacancy posting cost c enters this equation. A productivity shock therefore has exactly the same effect as a shock to any of these other variables. It is also possible to express the wage as a function of the contemporaneous vacancy-unemployment ratio and model parameters:

$$w_p = \beta\big(p + c\theta_p\big) + (1-\beta)z. \tag{9}$$

In this case, the ratio of the difference between the wage w_p and non-market productivity

[5]See Shimer (2003) for details.

z to the vacancy posting cost, $\frac{w_p - z}{c}$, is determined by the ratio $\frac{p-z}{c}$. Equations (8) and (9) implicitly define the state-contingent vacancy-unemployment ratio and wage and are easily solved numerically given particular parameter values.

3 Calibration

I calibrate the model to match the quarterly behavior of the U.S. economy from 1951 to 2001.[6] Although it is a continuous time model, I sample the model-generated data at discrete points in time corresponding to the end of each quarter so as to make it comparable with actual data. The parameter choices are summarized in Table 1; the next three paragraphs provide a brief justification.

I assume labor productivity p follows a first order autoregressive process. I normalize its mean to unity and then use data on real output per hour in the non-farm business sector, constructed by the Bureau of Labor Statistics as part of its Major Sector Productivity and Costs program, to quantify the instantaneous standard deviation and persistence of this key variable. In the simplest case, this translates into a three state Markov process for p, with values 0.976, 1 and 1.024. An aggregate shock hits according to a Poisson process with arrival rate 0.16 per quarter. When productivity is in an extreme state and a shock arrives, it moves to the intermediate state. From the intermediate state, productivity is equally likely to move to either of the extreme states.[7]

Next I set the discount rate to zero. Compared with a more standard number, say 0.010 or 0.015 per quarter, this choice scarcely affects the quantitative results; however, it simplifies the normative analysis because welfare only depends on the long-run behavior of the economy and in particular is independent of the current state. The separation rate s is equal to 0.1, based on the Abowd and Zellner's (1985) corrected worker flow data. The value of leisure is set to 0.4, consistent with an unemployment benefit replacement ratio of 40 percent. I normalize the cost of a vacancy to 0.53. This essentially pins down the units of a vacancy; in particular, I target a mean vacancy-unemployment ratio of 1.

Finally, I set the constant in the matching function to $\mu = 1.65$, so as to match the mean unemployment rate in the data, 5.7 percent. I set the elasticity of the matching function

[6] The U.S. economy is an interesting benchmark because, relative to most European economies, it is thought to have flexible wages. The finding that wages in the U.S. economy are significantly more rigid than in the benchmark model is therefore all-the-more-surprising.

[7] The first and second moments reported here are robust to an increase in the number of productivity states. See Shimer (2003) for details.

with respect to unemployment at $\alpha = 0.5$. This is consistent with the evidence summarized by Petrongolo and Pissarides (2001). It also ensures that fluctuations in unemployment and vacancies are of approximately equal magnitude. Finally, I set workers' bargaining power at $\beta = 0.5$, satisfying the so-called Hosios (1990) condition. Shimer (2003) proves that if $\beta = \alpha$, the equilibrium in the economy with Nash bargaining maximizes the expected present value of output net of vacancy costs, even in the presence of productivity shocks. Again, this facilitates the welfare analysis, which asks about the cost of wage rigidity. If wages are set using Nash bargaining with $\beta \neq \alpha$, rigid wages at a more appropriate level may actually raise welfare.

Table 2 shows summary statistics for productivity, wages, unemployment, and vacancies using quarterly data from the U.S. from 1951 to 2001. Productivity and wages are constructed by the Bureau of Labor Statistics (BLS) from the National Income and Product Accounts and the Current Employment Statistics. The former is measured as real average output per hour in the non-farm business sector; the latter is real hourly compensation in the same sector. It is a broad measure of compensation, including wages, salaries, tips, bonuses, and in-kind payments, as well as imputed compensation for proprietors and unpaid family workers. Unemployment is measured by the BLS using the Current Population Survey. Vacancies are crudely measured by the Conference Board help-wanted advertising index, but this variable closely tracks direct measures of vacancies when they are available (Abraham 1987, Shimer 2003). All data are detrended using a very low frequency Hodrick-Prescott filter with smoothing parameter $100,000$.

Table 3 summarizes the model generated data. I chose parameters to generate the correct standard deviation and first-order autocorrelation for labor productivity and the correct mean unemployment rate. The model-generated data can be compared with the actual data along the remaining dimensions. In some cases, the model performs very well. For example, the correlation between unemployment and vacancies in this data set is -0.90, while in the model it is -0.87. In other words, the model can produce a downward sloping 'Beveridge curve' or vacancy-unemployment relationship. On others, the model performs less well. In the data, vacancies are slightly more persistent than unemployment, while in the model vacancies are much less persistent. Introducing planning lags would presumably correct this shortcoming (Fujita 2003).

But the real problem with the model lies in the volatility that it (fails to) generate. The unemployment rate 15.2 times as volatile in the data as in the model, the vacancy rate 11.4 times, and the vacancy-unemployment ratio 12.5 times. The Mortensen-Pissarides model generates only a tiny fraction of the volatility of its two central elements, unemployment and vacancies.

4 Rigid Wage Model

The source of the volatility problem lies in the flexibility of wages, which in the model are perfectly correlated with productivity and are 96 percent as volatile. A productivity increase spurs entry by firms, which reduces unemployment duration and puts upward pressure on wages. In practice, the latter absorbs most of the shock, so there is little change in the vacancy-unemployment ratio. Following the analysis in Hall (2003b), this section demonstrates the improved performance of the model if wages are rigid. I make an extreme assumption here and replace the Nash bargaining solution (7) with a fixed wage, $w_p = \bar{w}$.[8]

It is simplest to view this assumption as a social norm. All workers expect to be paid $\bar{w}$ when employed, and all firms expect to have to pay $\bar{w}$ to any employee. Crucially, a matched worker and firm have no incentive to renegotiate the wage. An employed worker always prefers a higher wage, so long as this does not induce the firm to lay her off (i.e. as long as $J_p > 0$), and an employer always prefers to pay a lower wage, so long as this does not induce the worker to quit (i.e. as long as $E_p > U_p$). I focus throughout on a range of parameter values for which these conditions hold.[9]

An equilibrium is a solution to equations (2)–(6), with a constant and exogenous wage. The vacancy-unemployment ratio satisfies a forward-looking differential equation,

$$\frac{r+s+\lambda}{\mu\theta_p^{-\alpha}} = \frac{p-\bar{w}}{c} + \lambda \mathbb{E}_p \frac{1}{\mu\theta_{p'}^{-\alpha}}, \tag{10}$$

with θ_p truncated at zero. Note that the vacancy-unemployment ratio θ_p may be positive even if the wage exceeds current productivity, $p < \bar{w}$, because of the option value from future productivity shocks.

I fix all the parameters (except workers' bargaining power β) at their values in Table 1 and then characterize the equilibrium with the wage $\bar{w}$ chosen so as to replicate the appropriate

[8] A shortcoming of this assumption is that long-run productivity growth will counterfactually induce a long-run decline in the unemployment rate. Hall (2003b) makes a more complicated assumption on wage setting that ensures a similar behavior of the model in the short-run but a more satisfactory response to long-run trends.

[9] With a constant wage w_p that exceeds the value of leisure, the latter condition always holds. The former condition is equivalent to requiring $\theta_p > 0$; in my numerical example, this is true in all three states as long as $\bar{w} \leq 0.99$.

average unemployment rate, 5.7 percent. Table 4 records the results. There is scarcely any change in the correlation matrix or in the autocorrelation of any of the variables. For example, the correlation between unemployment and vacancies becomes slightly less negative, -0.85 rather than -0.87, while the lack of persistence in vacancies remains a problem. But the variability of unemployment and vacancies rises dramatically. Unemployment is almost exactly as variable in the model as in the U.S. data, while vacancies are somewhat more variable in the model than in the data. This reflects the much greater variability in the vacancy-unemployment ratio when wages do not absorb any of the impact of a shock. From a positive perspective, the naïve assumption that wages are constant significantly improves the performance of the model.

These results are sensitive to the choice of the mean wage. Figure 1 shows the mean unemployment and vacancy rates and the coefficient of variation of these two rates as functions of the fixed wage $\bar{w}$. A higher wage discourages vacancy creation, raising the mean unemployment rate. It also makes the ratio $\frac{p-\bar{w}}{c}$ more volatile, which implies that both vacancies and unemployment are more variable when $\bar{w}$ is larger. Nevertheless, with the wage fixed at approximately the average level in the economy with Nash bargaining, the Mortensen-Pissarides model is capable of matching the cyclical behavior of the key labor market variables.

5 The Cost of Rigid Wages

If rigid wages have a large effect on the equilibrium unemployment and vacancy rates, it seems natural that they would also have a large effect on welfare. During good times, wages are too low and so the vacancy-unemployment ratio θ is too high, while in bad times the opposite is true. This section analyzes the welfare cost of rigid wages in the presence of productivity fluctuations. At least to the authors' surprise, the cost is very small, as low 0.1 percent of net output if the level of wages is approximately correct.

There are several modelling questions to address before demonstrating this result. First, what is the correct benchmark for measuring the cost of rigid wages? It is well-known that the equilibrium of the economy with Nash bargaining does not maximize welfare for generic values of workers' bargaining power β. However, Shimer (2003), extending the results in Hosios (1990), proves that if the matching function is Cobb-Douglas and workers' bargaining power is equal to the elasticity of the matching function with respect to the unemployment rate, α, then the decentralized equilibrium is welfare-maximizing even in the presence of productivity shocks. That is, the state-contingent vacancy-unemployment ratio θ_p is socially optimal. I constructed the calibrated example to have this property, and so the benchmark economy is the social optimum.

Second, the cost of rigid wages might depend on the current state of the economy; it is less costly to have a high rigid wage when productivity is high than when it is low. But the assumption that the interest is zero, $r = 0$, allows us to avoid this complexity, since transitional dynamics do not affect undiscounted average lifetime utility.

Finally, how should welfare be measured? Since workers are risk-neutral, they only care about their average consumption, which is equal to the average output in the economy net of vacancy posting costs. When the productivity level is p and the unemployment rate is u, net output is $(1-u)p + uz - c\theta_p u$, the sum of income from the $1-u$ employed workers and from the u unemployed workers minus the cost of the $v = \theta_p u$ vacancies.

Figure 2 reports the percentage loss in average net income in the economy with fixed wages as a function of the wage $\bar{w}$. If $\bar{w} \approx 0.967$, about 0.12 percent of average income is lost due to fixed wages, a negligible amount. While a fixed real wage dramatically alters the behavior of unemployment and vacancies, it scarcely has any effect on welfare.

There are two important caveats to this result. First, since workers in this economy are risk-neutral, this calculation is based on the cost that rigid wages impose on the mean level of consumption. The variance in consumption also matters if workers are risk-averse and capital markets are incomplete; however, calculations by Lucas (1987) and others indicate that the cost of consumption variability is small even in non-representative agent economies (Imrohoroglu 1989, Atkeson and Phelan 1994, Obstfeld 1994) or in economies with non-standard preferences (Alvarez and Jermann 2000).

Second, if the wage level is fixed at a different level, $\bar{w} \neq 0.967$, the welfare costs may be much more significant. But this is not really a statement about wage rigidity in the presence of productivity fluctuations; it is similarly true that if wages are flexible but workers' bargaining power β is not equal to the elasticity α, welfare may be reduced substantially. I therefore ignore this possibility and leave for future research the question of why a rigid real wage might tend towards the efficient level in the long-run.

6 Conclusion

I have extended the Mortensen-Pissarides search and matching model with the simplest possible model of rigid wages, a constant wage. With wages determined by Nash bargaining, unemployment and vacancies are much less variable in the model than in the U.S. economy. A fixed wage generates approximately the correct variance for these two key variables. At the same time, a fixed real wage need not have any significant welfare cost. This suggests that, to the extent that government policies distort the economy along other dimensions, policies designed to make real wages more flexible are likely to be counterproductive, even if they succeed in moderating employment and vacancy fluctuations.

References

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

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

Alvarez, Fernando, and Urban Jermann (2000): "Using Asset Prices to Measure the Cost of Business Cycles," Mimeo.

Atkeson, Andrew, and Christopher Phelan (1994): "Reconsidering the Costs of Business Cycles with Incomplete Markets," in *NBER Macroeconomics Annual*, ed. by Stanley Fischer, and Julio Rotemberg, vol. 9, pp. 187–207. MIT Press, Cambridge.

Barro, Robert (1977): "Long-Term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics*, 3(3), 305–316.

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

Costain, James, and Michael Reiter (2003): "Business Cycles, Unemployment Insurance, and the Calibration of Matching Models," Mimeo.

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

Hall, Robert (2003a): "Modern Theory of Unemployment Fluctuations: Empirics and Policy Applications," *American Economic Review Papers and Proceedings*, 93(2), 145–150.

——— (2003b): "Wage Determination and Employment Fluctuations," Mimeo.

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

Imrohoroglu, Ayse (1989): "Cost of Business Cycles iwth Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, 97(6), 1364–1383.

Lucas, Robert (1987): *Models of Business Cycles*. Basil Blackwell, Oxford.

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

Obstfeld, Maurice (1994): "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability," *European Economic Review*, 38(7), 1471–1486.

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

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

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

Shimer, Robert (2003): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies: Evidence and Theory," Mimeo.

variable	meaning	value	explanation
$\bar{p}$	labor productivity: mean	1.000	Normalization
σ	labor productivity: inst. std. dev.	0.016	Output per hour in the
γ	labor productivity: mean reversion	0.16	non-farm business sector
r	discount rate	0	Simplify welfare analysis
s	separation rate	0.1	Abowd and Zellner (1985)
z	value of leisure	0.4	Benefit replacement ratio 40%
c	vacancy cost	0.53	Normalization $\theta \approx 1$
μ	matching function constant	1.65	Unemployment rate 5.7%
α	matching function elasticity	0.5	Petrongolo and Pissarides (2001)
β	worker's bargaining power	0.5	Decentralized equilibrium is optimal

Table 1: Parameter choices. Additional details are provided in the text. Note that labor productivity takes on three possible values, 0.976474, 1, and 1.02449. A shock hits at rate $\lambda = 0.16$. If the old productivity level is not equal to 1, it adjusts there immediately. If it is equal to 1, it moves with equal probability to 0.976474 or 1.02449.

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

		p	u	v	θ	w
Mean		—	0.0567	—	—	—
Coefficient of Variation		0.017	0.188	0.183	0.344	0.013
Autocorrelation (1 Quarter)		0.854	0.918	0.930	0.938	0.892
Correlation Matrix	p	1	-0.367	0.362	0.359	0.376
	u	—	1	-0.896	-0.949	-0.167
	v	—	—	1	0.955	0.136
	θ	—	—	—	1	0.175
	w	—	—	—	—	1

Table 2: Average labor productivity p is real average output per hour in the non-farm business sector, constructed by the Bureau of Labor Statistics (BLS) from the National Income and Product Accounts and the Current Employment Statistics. The unemployment rate u is constructed by the BLS from the Current Population Survey. The help-wanted advertising index v is constructed by the Conference Board. The wage w is real hourly compensation in the non-farm business sector, constructed by the Bureau of Labor Statistics (BLS) from the National Income and Product Accounts and the Current Employment Statistics. Both u and v are quarterly averages of seasonally adjusted monthly series. Productivity, unemployment, vacancies, and wages are expressed as ratios to an HP filter with smoothing parameter 10^5. The coefficient of variation is the ratio of the standard deviation to the mean.

Model-Generated Data, Nash-Bargaining

		p	u	v	θ	w
Mean		1.000	0.0569	0.0575	1.011	0.968
Coefficient of Variation		0.017	0.012	0.016	0.027	0.016
Autocorrelation (1 Quarter)		0.852	0.920	0.719	0.852	0.852
Correlation Matrix	p	1	-0.957	0.975	1.000	1.000
	u	—	1	-0.868	-0.957	-0.957
	v	—	—	1	0.975	0.975
	θ	—	—	—	1	1.000
	w	—	—	—	—	1

Table 3: Model-generated data with Nash bargaining. Parameterization given in Table 1.

Model-Generated Data, Fixed Wage

		p	u	v	θ	w
Mean		1.000	0.0572	0.0595	1.112	0.967
Coefficient of Variation		0.017	0.185	0.228	0.381	0
Autocorrelation (1 Quarter)		0.852	0.920	0.715	0.851	—
	p	1	-0.942	0.972	0.996	—
	u	—	1	-0.852	-0.923	—
Correlation Matrix	v	—	—	1	0.967	—
	θ	—	—	—	1	—

Table 4: Model-generated data with the wage fixed at $\bar{w} = 0.967$. Parameterization given in Table 1.

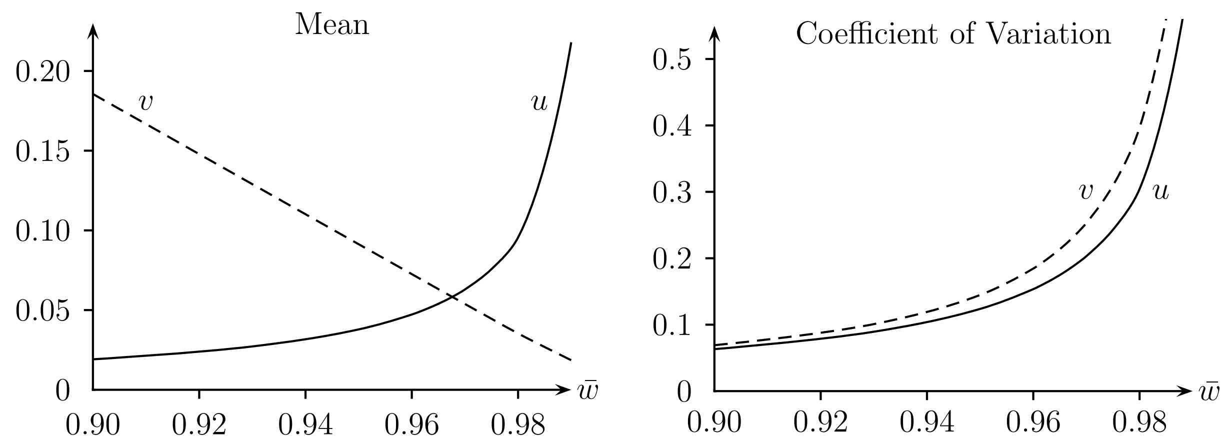

Figure 1: The left panel shows the mean unemployment (solid line) and vacancy (dashed line) rates as a function of the fixed wage $\bar{w}$. The right panel shows the coefficient of variation of the unemployment and vacancy rates.

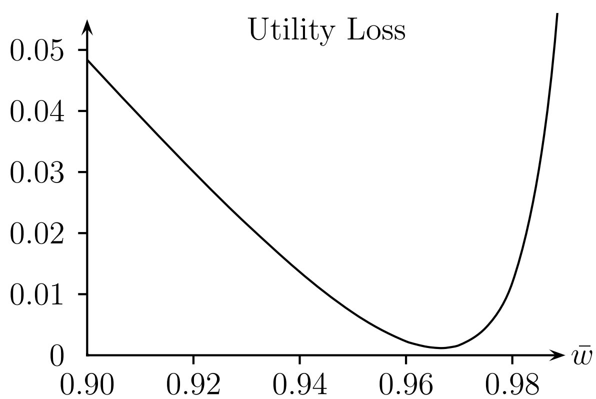

Figure 2: The utility loss from a fixed wage $\bar{w}$, expressed as a percent of the socially optimal, flexible wage benchmark.