

Unemployment and Vacancy Fluctuations in the Matching Model: Inspecting the Mechanism

Andreas Hornstein, Per Krusell, and Giovanni L. Violante

The state of the labor market, employment and unemployment, plays an important role in the deliberations of policymakers, the Federal Reserve Bank included. Over the last 30 years, economic theory has led to substantial progress in understanding the mechanics of business cycles. Much of this progress in macroeconomics has been associated with the use of calibrated dynamic equilibrium models for the quantitative analysis of aggregate fluctuations (Prescott [1986]). These advances have mainly proceeded within the Walrasian framework of frictionless markets. For the labor market, this means that while these theories contribute to our understanding of employment determination, they have nothing to say about unemployment.

Policymakers care about the behavior of unemployment for at least two reasons. First, even if one is mainly interested in the determination of employment, unemployment might represent a necessary transitional state if frictions impede the allocation of labor among production opportunities. Second, job loss and the associated unemployment spell represent a major source of income risk to individuals.

Over the past two decades, the search and matching framework has acquired the status of the standard theory of equilibrium unemployment.¹ This theory is built on the idea that trade in the labor market is costly and takes time. Frictions originating from imperfect information, heterogeneity

■ We wish to thank Kartik Athreya, Sam Malek, Leo Martinez, and Ned Prescott for helpful comments. The views expressed in this article are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

¹ For a textbook survey, see Pissarides (2000).

of firms and workers, and lack of coordination disrupt the ability to form employment relationships. The quantity of idle inputs in the labor market (unemployed workers and vacant jobs) is a measure of such disruption. In its most basic representation, a labor market matching model focuses on the interaction between unemployment and job creation. Higher productivity increases the return to job creation and thereby increases the rate of job creation. In turn, a higher rate of job creation makes it easier for unemployed workers to find jobs and thereby reduces unemployment. This explains the observed counter-cyclical (pro-cyclical) behavior of unemployment (job creation).

Shimer (2005) goes beyond investigating the qualitative features of the basic matching model. He follows the research program on dynamic equilibrium models with Walrasian frictionless markets and explores whether or not a calibrated matching model of the labor market is quantitatively consistent with observed aggregate fluctuations. He surprisingly concludes that a reasonably calibrated matching model does not generate enough volatility in unemployment and cannot explain the strong procyclicality of the job-finding rate. In other words, the matching model stops short of reproducing the cyclical behavior of its two central elements: unemployment and vacancies.²

In this article, we present the basic matching model, also known as the Mortensen-Pissarides model, in detail and, building on Shimer (2005), we explain the reasons for the quantitative problems of the model. Essentially, given the way wages are determined in the (Nash-bargaining) model and the way Nash bargaining is calibrated, wages respond strongly to changes in productivity so that the incentive for firms to create jobs does not change very much. We then discuss two possible ways of reconciling a matching model with the data.

First, as argued by Hall (2005) and Shimer (2004), if wages are essentially rigid, the model performs much better. We contend that rigid wages per se are not sufficient; another necessary requirement is a very large labor share—close to 100 percent of output. Moreover, we show that with rigid wages, the model has implications for the labor share that seem too extreme: the labor share becomes perfectly negatively correlated with—and as volatile as—labor productivity whereas in the data this correlation is -0.5 , and the variation of the share is not nearly as large as that of productivity.

Second, as suggested by Hagedorn and Manovskii (2005), without abandoning Nash bargaining, a different calibration of some key parameters of the

² We should note that Andolfatto (1996) and Merz (1995) were the first to integrate the matching approach to the labor market into an otherwise standard Walrasian model and to evaluate this model quantitatively. Their work, however, was not so much focused on the model's ability to match the behavior of unemployment, but on how the introduction of labor market frictions affects the ability of the otherwise standard Walrasian model to explain movements in employment, hours worked, and other non-labor-market variables. Andolfatto (1996), however, also pointed out the model's inability to generate enough volatility in vacancies.

model also allows one to raise the volatility of unemployment and vacancies in the model. For this calibration to work, however, one again needs a very high wage share. This high share is obtained by “artificially” raising the outside option of the worker through generous unemployment benefits.³ We tentatively conclude, as do Costain and Reiter (2003), that this parameterization has implausible implications for the impact of unemployment benefits on the equilibrium unemployment rate: a 15 percent rise in benefits would *double* the unemployment rate.

Why is a very large (very small) wage share (profit share) so important in order for the model to have a strong amplification mechanism for vacancies and unemployment? The model has a free-entry condition stating that vacancies are created until discounted profits equal the cost of entry. If profits are very small in equilibrium, a positive productivity shock induces a very large percentage increase in profits, and hence a large number of new vacancies must be created—through firm entry—thus lowering the rate of finding workers enough that entry remains an activity with zero net payoff.

We conclude that neither one of the solutions proposed is fully satisfactory, for two reasons. First, they both have first-order counterfactual implications. Second, they both assume a very large value for the labor share. It is hard to assess whether this value is plausible because there is no physical capital in the baseline matching model. We speculate that the addition of physical capital, besides providing a natural way of measuring the labor share of aggregate income, would allow the analysis of another important source of aggregate fluctuations, investment-specific shocks, which have proved successful in Walrasian models.⁴

The present article, which can be read both as an introduction to the matching model of unemployment and as a way of understanding the recent discussions of the model’s quantitative implications, is organized as follows. We first quickly describe the data. Next, we describe in Section 2 the basic model without aggregate shocks. In Section 3, we define and solve for a stationary equilibrium: a steady state. In Section 4, we briefly discuss transition dynamics within the model without shocks. In Section 5, we derive the qualitative comparative statics for a one-time permanent change of the model’s parameters. In Section 6, we present the alternative calibration strategies one could follow to parameterize the matching model, and in Section 7, we show how the quantitative comparative statics results differ according to the model

³ To be precise, a large wage share is also sufficient for a strong amplification mechanism with rigid wages. With flexible wages, the large wage share must be achieved by making unemployment benefits high.

⁴ See Greenwood, Hercowitz, and Krusell (2000) and Fisher (2003) for a quantitative account of the role of this type of shock in U.S. business cycles over the postwar period. Costain and Reiter (2003) illustrate quantitatively that productivity shocks affecting only new jobs improve the performance of the baseline model.

Table 1 Aggregate Statistics: 1951:1–2004:4

HP Smoothing Parameter: 10^5							
	u	v	θ	λ_w	w	s	p
Standard Deviation	0.20	0.23	0.38	0.12	0.02	0.02	0.02
Autocorrelation	0.94	0.95	0.95	0.91	0.95	0.91	0.89
Correlation with p	−0.40	0.31	0.38	0.38	0.69	−0.35	1.00

HP Smoothing Parameter: 1600							
	u	v	θ	λ_w	w	s	p
Standard Deviation	0.13	0.14	0.26	NA	0.01	0.01	0.01
Autocorrelation	0.87	0.90	0.89	NA	0.81	0.77	0.76
Correlation with p	−0.29	0.45	0.38	NA	0.72	−0.61	1.000

Notes: Data are quarterly, and u is the unemployment rate of the civilian population; v is the help-wanted advertising index; $\theta = v/u$ is labor market tightness; p is output per employee in the nonfarm business sector; s is the labor share constructed as the ratio of compensation of employees to output in the nonfarm sector; w is the wage computed as labor share times labor productivity, i.e., $w = s \cdot p$. The statistics for the job-finding rate, λ_w , are those reported in Shimer (2005) for an HP smoothing parameter of 10^5 .

calibration. In Section 8, we introduce explicit stochastic aggregate shocks and discuss how the quantitative comparative statics results for one-time permanent shocks have to be modified to account for persistent but temporary shocks. Section 9 concludes the article.

1. THE DATA

The focus of the analysis is on fluctuations at the business-cycle frequencies, and hence low-frequency movements in the data should be filtered out. For quarterly data, the standard practice (followed by Andolfatto and Merz) is to use a Hodrick-Prescott (HP) filter with a smoothing parameter set to 1600. Shimer (2005) chooses a much smoother trend component, corresponding to an HP smoothing parameter of 10^5 .

Table 1 summarizes the key labor market facts around which this article is centered. We report statistics for the detrended log-levels of each series. When we remove a very smooth trend (smoothing parameter 10^5), we can summarize the data as follows:

- **Unemployment and Vacancies.** First, unemployment, u , and vacancies, v , are about 10 times more volatile than labor productivity, p . Market tightness, θ , defined as the ratio of vacancies to unemployment, is almost 20 times more volatile. Second, market tightness is

positively correlated with labor productivity. Both unemployment and vacancies show strong autocorrelation.

- **Job-Finding Rates.** The job-finding rate, λ_w , is six times more volatile than productivity and is pro-cyclical. It is also strongly autocorrelated.
- **Wages and Labor Share.** Wages and the labor share are roughly as volatile as labor productivity. The correlation between wages and labor productivity is high but significantly less than one, and the labor share is countercyclical.

Using a more volatile trend component (lower smoothing parameter) has almost no effect on the relative volatilities. For the vast majority of the variables, the percentage standard deviation is reduced roughly by one-third. Interestingly, the volatility is cut in half for wages and the labor share. Overall, the autocorrelation of the series is reduced, since some of the persistence is absorbed by the more variable HP trend. Finally, the correlation structure of the series with labor productivity is, in general, unchanged except for the labor share whose negative correlation almost doubles.

We conclude that the choice of smoothing parameter has no impact on the unemployment and vacancy statistics but does affect the labor share statistics somewhat.

2. THE MODEL

We now outline and discuss the basic Mortensen-Pissarides matching model with exogenous separations.⁵ We choose a formulation in continuous time in order to simplify some of the derivations. It is useful to first describe the stationary economy (when aggregate productivity is constant over time) because that model is simple and yet very informative about how the model with random shocks behaves. Later, we will briefly discuss aggregate fluctuations with stochastic productivity shocks that are persistent but not permanent.⁶

Workers and Firms

There is a fixed number of workers in the economy; the model does not consider variations in the labor force or in the effort or amount of time worked by each

⁵ The main reference is Pissarides (1985). Mortensen and Pissarides (1994) extend the model to endogenous separations. Pissarides (2000) contains an excellent survey of the matching models. See also Rogerson, Shimer, and Wright (2005) for a recent survey.

⁶ The view that aggregate fluctuations in output and unemployment are due to fluctuations in productivity is not essential here. For the given environment, one can interpret productivity shocks as actually representing another source of fluctuations (such as “demand shocks,” e.g., shocks to preferences).

worker. For example, think of workers as being uniformly distributed on the interval $[0,1]$ —for any point on this interval, there is one worker—though there is no particular meaning to a worker’s position on the interval.

Workers are all the same from the perspective of both their productivity and their preferences. Workers are infinitely lived and have linear utility over consumption of a homogeneous good, meaning that to the extent that there is uncertainty, workers are risk-neutral. There is constant (exponential) discounting at rate r . One can therefore think of a worker’s expected present value of utility as simply the expected present value of income.⁷

Workers are either employed or unemployed. An employed worker earns wage income, w , but cannot search. Unemployed workers search for jobs. Let $b > 0$ denote the income equivalent of the utility flow that a worker obtains in the nonworking activity when unemployed, e.g., the monetary value of leisure plus unemployment benefits net of search costs.⁸

A firm is a job. The supply of firms (jobs) is potentially infinite. Every firm is equally productive at any point in time. Firms are risk-neutral and they discount future income at the same rate as do workers. Production requires one worker and one firm; firms can really be thought of as another type of labor input, such as an “entrepreneur.” A firm-worker pair produces p units of the homogeneous output per unit of time. We assume that the value of production for a pair always exceeds the value of not working for a worker, i.e., that $p > b > 0$.⁹ There is no cost for a firm to enter the labor market.

The Frictional Labor Market

In a “frictional” labor market, firms and workers do not meet instantaneously. In addition, firms that want to meet workers have to use resources to post a vacancy. In particular, a firm has to pay c units of output per unit of time it posts a “vacancy.” Let the number of idle firms that have an open position be denoted $v(t)$, and let the number of unemployed workers be $u(t)$. Lack of coordination, partial information, and heterogeneity of vacancies and workers are all factors that make trading in the labor market costly.

We do not model these labor market frictions explicitly but use the concept of a matching function as a reduced form representation of the frictions.¹⁰ This

⁷ Alternatively, one could assume that workers are risk-averse but that they can obtain complete insurance against idiosyncratic income risk. In this case, it would also be optimal for workers to maximize the expected present value of income.

⁸ Note that unemployment benefits do not serve an insurance role in this environment since workers are either risk-neutral or they already obtain complete insurance.

⁹ This condition is necessary for ruling out a trivial equilibrium with zero employment: if $b > p$, no worker would be willing to work even if she could extract the entire value of the output produced from the firm.

¹⁰ The concept of an “aggregate matching function” has been around for some time. In their survey of the literature on matching functions, Petrolongo and Pissarides (2001) include a short

formulation specifies that the rate at which new matches, m , are created is given by a time-invariant function, M , of the number of unemployed workers searching for a job and the number of vacant positions: $m = M(u, v)$. At this point, we will assume that M is (1) increasing and strictly concave in each argument separately and (2) constant returns to scale (CRS) in both arguments. Thus, matches are more likely when more workers and firms are searching, but holding constant the size of one of the searching groups, there are decreasing marginal returns in matching.

New matches are formed according to Poisson processes with arrival rates λ_w and λ_f . Given the rate at which new matches are formed, the rate at which an unemployed worker meets a firm is simply $\lambda_w(t) = m(t)/u(t)$, the total number of successful matches per worker searching. Similarly, the rate at which a vacant firm meets a worker is $\lambda_f(t) = m(t)/v(t)$. Since the matching function is CRS, the two meeting rates depend on labor market tightness, $\theta(t) = v(t)/u(t)$, only:

$$\lambda_w(t) = M[1, \theta(t)] \text{ and } \lambda_f(t) = M[1/\theta(t), 1]. \quad (1)$$

As the relative number of vacancies increases, the job-finding rate, λ_w , also increases, but the worker-finding rate, λ_f , decreases. We assume that once a firm and a worker have been matched, they remain matched until “separation” occurs. Separation occurs according to a Poisson process with exogenous arrival rate, σ .

If an unemployed worker meets vacant firms according to a Poisson process with arrival rate, λ_w , then the probability that the worker meets exactly one vacant firm during a time period, Δ , is $\lambda_w \Delta$ if the time period is sufficiently short. Furthermore, the probability that a worker meets two or more vacant firms during this time period is essentially zero.¹¹ Similarly, the probability that a vacant firm meets an unemployed worker is $\lambda_f \Delta$, and the probability that a matched firm-worker pair separates is $\sigma \Delta$. Thus, if we start out with $u(t)$ unemployed workers and $1 - u(t)$ employed workers at time t , after a short time period, Δ , the number of unemployed workers will be

$$u(t + \Delta) = \sigma \Delta [1 - u(t)] + [1 - \lambda_w(t) \Delta] u(t).$$

history of the concept. Lagos (2000) warns against the dangers of such a “reduced-form approach” to frictions when, for example, evaluating the effects of policies. The underlying reason is that policies may affect the search behavior of agents and change the shape of the aggregate matching function.

¹¹ Note that for a Poisson process, the rate λ at which the state changes need not be bounded above by one. Since we are interested in the limiting case when the time interval, Δ , becomes arbitrarily small, the probability of a state change, $\lambda \Delta$, will eventually be less than one for any fixed and finite λ .

Subtracting $u(t)$ from either side of this expression, dividing by Δ , and taking the limit when the length of the time period goes to zero, we obtain

$$\dot{u}(t) = \lim_{\Delta \rightarrow 0} \frac{u(t + \Delta) - u(t)}{\Delta} = \sigma [1 - u(t)] - \lambda_w(t)u(t). \quad (2)$$

Here $\dot{u}(t)$ denotes the time derivative (change per unit of time) of $u(t)$: $\dot{u}(t) = \partial u(t)/\partial t$. This equation captures that the change in unemployment is the flow into unemployment (the number of employed workers times the rate at which they separate) minus the flow from unemployment (the number of unemployed workers times the rate at which they find a job).

The dynamic evolution of unemployment is one of the key concerns in this model. Notice, however, that the job-finding rate for workers, $\lambda_w(t)$, in equation (2) depends on vacancies through labor market tightness, $\theta(t)$. What determines vacancies, $v(t)$? In order to answer this question, we need to describe what determines profits for entering firms, which in turn requires us to discuss what wages workers receive.

With matching frictions, both workers and firms have some bargaining power since neither party can be replaced instantaneously, as is commonly assumed in competitive settings. There is a variety of theories that describe how bargaining allocates output between firms and workers under these circumstances. Below we will determine wages according to the widely used Nash-bargaining solution. For simplicity, from now on we will mainly consider steady states, situations in which all aggregate variables are stationary over time. Thus, $u(t)$, $v(t)$, $\lambda_w(t)$, and $\lambda_f(t)$ are all constant, even though individual workers and firms face uncertainty in their particular experiences.

3. STATIONARY EQUILIBRIUM

Values

Denote the net present value of a matched firm J (which in general would depend on time but in a steady state does not). Given output, p , and the wage, w , paid to its worker, J must satisfy

$$rJ = p - w - \sigma(J - V), \quad (3)$$

where V is the value of the firm when unmatched. This equation is written in flow form and is interpreted as follows: the flow return of being matched—the capital value of being matched times the rate of return on that value—equals the flow profits minus the expected capital loss resulting from match separation—the rate at which the firm is separated, σ , times the latter capital loss equals $J - V$.¹²

¹² This equation is written in flow form but can be derived from a discrete-time formulation analogous to the derivation of equation (2). Suppose that the value of being vacant is constant

Similarly, the value of a vacant firm satisfies

$$rV = -c + \lambda_f(J - V). \quad (4)$$

Here, there is a flow loss due to the vacancy posting cost and an expected capital gain from the chance of meeting a worker.

Turning to the net present value of a matched worker, W , and an unemployed worker, U , we similarly have

$$rW = w - \sigma(W - U), \text{ and} \quad (5)$$

$$rU = b + \lambda_w(W - U). \quad (6)$$

The flow return from not working, b , could be a monetary unemployment benefit collected from the government, a monetary benefit from working in an informal market activity, or the monetary equivalent of not working in any market activity (the value of being at home). We will discuss the role and interpretation of b more extensively below, because it turns out that it matters how one thinks of this parameter.

Wage Determination

The values of (un)matched workers and firms depend on the wages—yet to be determined—paid in a match. Obviously, for a match to be maintained it must be beneficial for both the worker, $W - U \geq 0$, and the firm, $J - V \geq 0$. We define the total surplus of a match, $S \equiv (J - V) + (W - U)$, as the sum of the gain of the firm and worker being in a match relative to not being in a match. We assume that the wage is set such that the total match surplus is shared between the worker and firm according to the Nash-bargaining solution with share parameter β :¹³

$$W - U = \beta S \text{ and } J - V = (1 - \beta)S. \quad (7)$$

over time from the perspective of a matched firm and that we are looking at one period being of length Δ . During this period, there is production, and wages are paid, the net amount being $(p - w)\Delta$ since p and w are measured per unit of time. At the end of the period, the match separates with probability $\sigma\Delta$ and remains intact with probability $1 - \sigma\Delta$. So it must be that $J(t) = (p - w)\Delta + (1 - \sigma\Delta)e^{-r\Delta}J(t + \Delta) + \sigma\Delta e^{-r\Delta}V$. Here, $e^{-r\Delta} \equiv \delta(\Delta)$ is a discount factor; it gives a *percentage* decline in utility as a function of the length of time, $-(d\delta(\Delta)/d\Delta)/\delta(\Delta)$, which is constant and equal to r . Subtract $J(t + \Delta)e^{-r\Delta}$ on both sides and divide by Δ . That delivers $\frac{J(t) - J(t + \Delta)}{\Delta} + \frac{(1 - e^{-r\Delta})}{\Delta}J(t + \Delta) = p - w - \sigma e^{-r\Delta}(J(t + \Delta) - V)$. Take limits as $\Delta \rightarrow 0$. Then the left-hand side becomes $\dot{J}(t) + rJ(t)$, the second term coming from an application of l'Hôpital's rule and the value being a continuous function of time. The right-hand side gives $p - w - \sigma(J(t) - V)$. In a steady state, $J(t)$ is constant and equal to J , satisfying the equation in the text.

¹³ The Nash-bargaining solution does not describe the outcome of an explicit bargaining process; rather, it describes the unique outcome among the set of all bargaining processes whose outcomes satisfy certain axioms (Nash [1950]). Also, one can derive the Nash-bargaining solution as the outcome of a bargaining process where participants make alternating offers until they reach agreement. For a survey of the bargaining problem, see Osborne and Rubinstein (1990).

Summing the value equations for matched pairs and subtracting the values of unmatched firms and workers, using the Nash-bargaining rule, we therefore obtain

$$rS = p - \sigma S + c - \lambda_f(1 - \beta)S - b - \lambda_w\beta S, \quad (8)$$

which implies that

$$S = \frac{p + c - b}{r + \sigma + (1 - \beta)\lambda_f + \beta\lambda_w}. \quad (9)$$

That is, we can express the surplus as a function of the primitives and the matching rates, which are endogenous and will be determined by the free entry of firms as shown below. We see that the surplus from being in a match is

- decreasing in the interest rate (a higher interest rate reduces the present value of remaining in the match),
- decreasing in the separation rate (a higher separation rate lowers the expected value of remaining together),
- decreasing in the bargaining share of workers times the rate at which they meet vacant firms (the higher the chance that unemployed workers meet vacant firms and the higher the share that workers receive in that case, the less valuable it is to be matched now), and
- decreasing in the bargaining share of firms times the rate at which vacant firms meet unemployed workers (the higher the chance that vacant firms meet unemployed workers and the higher the share that firms receive in that case, the less valuable it is to be matched now).

To derive a useful expression for the wage, subtract rV from the value equation for matched firms, (3), and use the Nash-bargaining rule to obtain

$$r(1 - \beta)S = p - w - \sigma(1 - \beta)S - rV. \quad (10)$$

Also, notice that given the surplus sharing rule, (7), and the expressions for the vacancy and unemployment values, (4) and (6), the surplus in (8) can be written as

$$rS = p - \sigma S - rV - rU. \quad (11)$$

Now multiply equation (11) by $1 - \beta$, subtract it from equation (10), and solve for the wage:

$$w = \beta(p - rV) + (1 - \beta)rU. \quad (12)$$

Thus, the wage is a weighted average of productivity minus the flow value of a vacancy and the flow value of unemployment with the weights being β and $1 - \beta$, respectively. Intuitively, one can understand this equation as follows:

$w - rU$, the flow advantage of being matched for a worker, is just its share, β , of the overall advantage of being matched for the worker and the firm together, $\beta(p - rV - rU)$.

Firm Entry

There is an infinite supply of firms that can post vacancies, and entry is costless. Therefore, in an equilibrium with a finite number of firms posting vacancies, the value of a posted vacancy is zero:

$$V = 0. \quad (13)$$

If $V < 0$, no firm would enter, and if $V > 0$, an infinite number of firms would enter. This means that the number of vacancies, $v(t)$, adjusts at each point in time so that there are zero profits from entering, given the matching rate with workers, λ_f , which depends on $u(t)$ and on $v(t)$.

The free-entry condition (13), together with the definition of the vacancy value (4) and the surplus sharing rule (7) then determine the surplus value:

$$S = \frac{c}{(1 - \beta)\lambda_f}. \quad (14)$$

Moreover, we can use the free-entry condition to simplify the expression for the surplus in (9); the surplus can now be expressed as

$$S = \frac{p - b}{r + \sigma + \beta\lambda_w}. \quad (15)$$

These two expressions for the surplus can be combined to write

$$\frac{p - b}{r + \sigma + \beta\lambda_w} = \frac{c}{(1 - \beta)\lambda_f}. \quad (16)$$

This is an equation in one unknown, labor market tightness (θ), since both meeting rates (λ_w and λ_f) depend only on the number of vacancies relative to the unemployment rate (see equation (1)).

We also see that free entry implies that the wage expression (12) simplifies to

$$w = \beta p + (1 - \beta)rU. \quad (17)$$

Equilibrium Unemployment

In a steady state, $\dot{u}(t) = 0$, so the evolution for unemployment as given by equation (2) becomes

$$\sigma(1 - u) = \lambda_w u. \quad (18)$$

Thus, in a steady state, the flow into unemployment—the separation rate in existing matches times the number of matches—must equal the flow out of unemployment—the job-finding rate times the number of unemployed.

The steady state expression for unemployment can, on the one hand, be used to express unemployment as a simple function of the separation rate and the job-finding rate. On the other hand, it can be used to write the job-finding rate in terms of the unemployment rate and the separation rate. If we know, for example, that the unemployment rate is 10 percent and that the monthly separation rate is 5 percent, then the chance of finding a job within a month must be $\sigma \frac{1-u}{u} = 0.05 \cdot \frac{0.9}{0.1} = 0.45$; that is, just under one-half.

Solving the Model

Solving the model is now straightforward. We have derived (16) and (18) in two unknowns, θ and u . Furthermore we can solve the two equations sequentially. First, from (1) it follows that λ_w (λ_f) is increasing (decreasing) in θ . This, in turn, implies that the left-hand side (LHS) of (16) is decreasing in θ and that the right-hand side (RHS) is increasing in θ . Thus, if a solution, θ , to (16) exists, it is unique. Second, conditional on θ , we can solve (18) for the equilibrium unemployment rate.

One can show that a solution to (16) exists if we assume that the matching function satisfies the Inada conditions.¹⁴ We assume a particular functional form for the matching function that meets these conditions and that is the most common one in the literature, the Cobb-Douglas (CD) matching function,

$$M(u, v) = Au^\alpha v^{1-\alpha}. \quad (19)$$

The CD matching function has convenient properties in terms of how the matching rates change with changes in labor market tightness,

$$\lambda_w = A\theta^{1-\alpha} \text{ and } \lambda_f = A\theta^{-\alpha}. \quad (20)$$

Independent of the level of unemployment, if the labor market tightness increases by 1 percent, the rate at which a worker (firm) finds a firm (worker) goes up (down) by $1 - \alpha$ (α) percent.¹⁵

Using the CD matching function, our equilibrium condition, (16), becomes

$$\frac{p - b}{r + \sigma + \beta A\theta^{1-\alpha}} = \frac{c}{(1 - \beta)A\theta^{-\alpha}}. \quad (21)$$

For $\theta = 0$, the LHS of (21) is finite and positive, and the RHS is zero. As θ becomes arbitrarily large, the LHS converges to zero and the RHS becomes arbitrarily large. Thus there exists a positive θ that solves (21). The unem-

¹⁴ Let $f(\theta) = M(\theta, 1)$. Then the Inada conditions are $f(0) = 0$, $f(\infty) = \infty$, and $f'(0) = \infty$.

¹⁵ Shimer (2005) argues that the constant elasticity CD matching function describes the data for the U.S. labor market well. See also Section 7 on calibration.

ployment rate then can be solved for in a second step, using (18), as

$$u = \frac{\sigma}{\sigma + A\theta^{1-\alpha}}. \quad (22)$$

We obtain the wage by using the definitions of the matching rates, (20), and substituting the expressions for rU and the value of S , (6), and (15) in wage equation (17):

$$w = \beta p + (1 - \beta) \left(b + \frac{\beta c}{1 - \beta} \theta \right) = \beta(p + c\theta) + (1 - \beta)b. \quad (23)$$

A Digression: The Frictionless Model

We now show that as search frictions become small, the equilibrium of the economy with matching frictions converges to the equilibrium of the corresponding economy without matching frictions. Search frictions can become small either because the cost of searching for vacant firms, c , becomes small or because the efficiency of the matching process, A , improves.

The frictionless economy is identical to the one outlined so far, except that matching between vacant firms and unemployed workers is instantaneous and costless. The resource allocation problem in the frictionless economy, which can be studied from the perspective of a benevolent social planner, is trivial. There will always be the same number of firms as workers operating because there is no cost in creating vacancies, and the matching process is instantaneous. Leaving workers idle would therefore be inefficient since $p > b$. There are no vacancies since matching is instantaneous. There is a competitive equilibrium that supports this allocation given some wage rates, $w(t)$, specified at all points in time. It is clear that for these wages, $w(t)$ must equal p for all t because workers are in short supply, and firms are not. That is, firm entry bids down profits to zero, and workers obtain the entire output.

Now suppose that the vacancy-posting costs become arbitrarily small: $c \rightarrow 0$. Then for any finite θ , the LHS of (21) is strictly positive, but the RHS converges to zero. Therefore, it must be that $\theta \rightarrow \infty$. To find the wage, some care must be taken, since the wage expression contains $c\theta$, i.e., $0 \cdot \infty$. Since workers meet firms at an ever-increasing rate, $\lambda_w \rightarrow \infty$, the unemployment rate becomes arbitrarily small, $u \rightarrow 0$, and from equation (9) it follows that the surplus from being in a match becomes arbitrarily small: $S \rightarrow 0$. Then simply inspect (10), which implies that $w \rightarrow p$, as expected: workers obtain the whole production value.

The same kind of result is obtained if the matching efficiency becomes arbitrarily large, $A \rightarrow \infty$. Now, however, there will be no vacancies, and θ will remain finite. To see this formally, multiply (21) with $A\theta^{-\alpha}$, divide the numerator and denominator of the LHS by A , and take the limit as $A \rightarrow \infty$:

$$\frac{p - b}{\beta\theta_\infty} = \frac{c}{1 - \beta}.$$

Since $\theta_\infty = \lim_{A \rightarrow \infty} \theta(A)$ is finite, the limits of both λ_f and λ_w are infinite. Thus from equation (9) it follows that the limit of the surplus is zero; from (22) it follows that the limit of the unemployment rate is zero; and from (10) it follows that the limit of the wage again must equal p . Since θ_∞ is positive and finite, v_∞ must equal 0 since u_∞ equals 0. There is no unemployment, and there are no vacancies.

4. TRANSITION DYNAMICS

So far, we have discussed how the key endogenous variables—unemployment, vacancies, job-finding rates, and wages—are determined in steady state. But how does the economy behave out of a steady state? To answer this question, one needs to find out what the economy's state variables are. A state variable is a variable that is predetermined at time t and that matters to outcomes. Here, unemployment is clearly a state variable, because it is a variable that moves slowly over time according to (2). In fact, it is the only state variable. No other variable is predetermined. This means that, in general, allocations at t depend on $u(t)$ but not on anything else.

So what is a dynamic equilibrium path of the economy if it starts with an arbitrary $u(0)$ at time zero? It turns out that the equilibrium is very easy to characterize. All variables except $u(t)$ and $v(t)$ will be constant over time from the very beginning.¹⁶ To show that this is indeed an equilibrium, simply assume that θ is constant from the beginning of time and equal to its steady state value and then verify that all equilibrium conditions are satisfied. Since θ is constant, all job-finding rates— $\lambda_w(t)$ and $\lambda_f(t)$ —will be constant and equal to their steady state values because they depend on θ and on nothing else. Since the λ s are the only determinants of the values J , V , W , and U , the solution for the values will be the same as the steady state solution. It then also follows that the wage must be the steady state wage. To find $u(t)$ and $v(t)$, we conclude that $u(t)$ will simply follow

$$\dot{u}(t) = \sigma [1 - u(t)] - \lambda_w u(t), \quad (24)$$

which differs from (2) only in that λ_w is now constant. Once we have solved for $u(t)$, we can find the path for $v(t)$ residually from $v(t) = \theta u(t)$. Moreover, note that if $u(0)$ is above the steady state, u , the RHS of equation (24) is negative, which means that $\dot{u}(0)$ is negative. Unemployment falls, and as long as it is still above u , it continues falling until it reaches (converges to) u . Similarly, if it starts below u , it rises monotonically over time toward u .¹⁷

¹⁶ Pissarides (1985; 2000, Chapter 1) shows that this is the unique equilibrium path.

¹⁷ Formally, the solution for $u(t)$ is the solution to the linear differential equation (24): $u(t) = u + e^{-(\sigma + \lambda_w)t} (u(0) - u)$, where $u = \frac{\sigma}{\sigma + \lambda_w}$.

The fundamental insight here is that there are no frictions involved in firm entry, but there are frictions in movement of workers in and out of jobs.¹⁸ Therefore, $u(t)$ is restricted to follow a differential equation which is “slow-moving,” whereas $v(t)$ does not have to satisfy such an equation. It can jump instantaneously to whatever it has to be so that θ is equal to its steady state value from the beginning of time.

5. COMPARATIVE STATICS

We now analyze how different parameters influence the endogenous variables. In particular, how does unemployment respond to changes in productivity? Here, we emphasize that these are steady state comparisons. We find the long-run effect of the permanent change in the parameter. For most variables—all except $u(t)$ and $v(t)$ —the impact of a permanent change in the parameter is instantaneous because θ immediately moves to its new, long-run value (see the discussion in the previous section). Of course, in the section below where some of the primitives are stochastic, their changes need not be permanent, and slightly different results apply.

For example, if we are looking at a 1 percent permanent increase in productivity, p , the comparative statics analysis in this section will correctly describe the effect on θ both in the long and in the short run, whereas the effect on unemployment recorded here only pertains to how it will change in the long run. The short-run effect on unemployment of a permanent change in a parameter is straightforward to derive, nevertheless: It simply involves tracing out the new dynamics implied by the linear differential equation (24) evaluated at the new permanent value for λ_w (which instantaneously adopts its new value because θ does). In particular, one sees from the differential equation that an increase in θ will increase λ_w and thus increase the speed of adjustment to the new steady state rate of unemployment.

We are mainly interested in how the economy responds to changes in p , but we will also record the responses to b , σ , and c . We compute elasticities, i.e., we use percentage changes and ask by what percent θ and u will change when p , b , σ , and, c change by 1 percent. We derive the relevant expressions by employing standard comparative statics differentiation of (21) and (22). Using $\hat{x}$ to denote $d \log(x) = dx/x$, it is straightforward to derive

¹⁸ The speed of movements from unemployment into employment is regulated by the hiring rate, λ_w , which, in turn, depends on the endogenous market tightness, θ . Separations instead are exogenous, and, hence, the speed of movements from employment to unemployment is simply determined by the parameter, σ .

$$\hat{\theta} = \frac{r + \sigma + \beta\lambda_w}{\alpha(r + \sigma) + \beta\lambda_w} \left[\frac{p}{p - b} \hat{p} - \frac{b}{p - b} \hat{b} - \frac{\sigma}{r + \sigma + \beta\lambda_w} \hat{\sigma} - \hat{c} \right], \text{ and} \quad (25)$$

$$\hat{u} = (1 - u) \left[\hat{\sigma} - (1 - \alpha) \hat{\theta} \right]. \quad (26)$$

The Effect of an Increase in Productivity

From equation (25), we see that an increase in p of 1 percent leads to more than a 1 percent increase in θ since $\alpha < 1$, and $p > b > 0$. Intuitively, p increases the value of matches, and given that firms capture some of the benefits of this increase in value, there will be an increase in the number of firms per worker seeking to match. The larger the fraction of the surplus going to the firm (β small), the more vacancies and market tightness will respond to a change in labor productivity. We also see that to the extent that b is close to p , the effect can be large, since $p/(p - b)$ can be arbitrarily large. Why is this effect larger the closer b is to p ? When $(p - b) \simeq 0$, the profit from creating vacancies is small, and $\theta \simeq 0$. Hence, even a small change in p induces very large changes in firms' profits and market tightness, θ , in percentage terms, through the free-entry condition (21).

Because the job-finding rate, λ_w , equals $A\theta^{1-\alpha}$, we obtain that $\hat{\lambda}_w = (1 - \alpha)\hat{\theta}$, so the effect of p on θ is higher than that on job-finding rates by a constant factor, $1/(1 - \alpha)$. If we look at the effect on unemployment, note from (26) that a 1 percent increase in θ lowers unemployment by $(1 - u)(1 - \alpha)$ percent.

The Effects of Changing b , σ , and c

Changes in income when unemployed, b , have a very similar effect to productivity changes, p , but with an opposite sign. Increasing b , in particular, lowers θ significantly if b is near p , but it has very little effect on θ if b is close to zero. An increase in the match separation rate, σ , decreases labor market tightness. More frequent separations reduce the expected profits from creating a vacancy, and, thus, θ falls. The effects on labor market tightness of higher vacancy-posting costs, c , are negative as well. A 1 percent increase in the vacancy cost lowers the labor market tightness (by less than 1 percent) because it requires firms' finding rates to go up in order to preserve zero profits, and, hence, there must be fewer vacant firms relative to unemployed workers. There is less than a one-for-one decrease because the surplus, once matched, increases as well, as is clear from equations (14) and (15).

The effects on the job-finding rate of all the above changes in primitives are all one minus α times the effect on $\hat{\theta}$. Similarly, the effects on unemployment are $-(1-u)(1-\alpha)$ times those on $\hat{\theta}$, with the exception of a change in σ because from (22), the total effect on unemployment of a rise in σ by 1 percent is twofold. The first effect is an indirect decrease through the impact on θ (a higher σ leads to a higher θ), which lowers unemployment. The second effect is a direct increase of unemployment due to the higher rate at which matches separate. The total effect cannot be signed without more detailed assumptions; for example, if $\alpha \geq 1/2$, the net effect is to increase unemployment.

An Additional Friction: Rigid Wages

In the model just described, productivity changes arguably have such a small impact on labor market tightness and unemployment that they cannot account for the observed fluctuations in the data. Hall (2005) and Shimer (2004) suggest that one way to address this shortcoming is to change the wage-setting assumption. We now describe a very simple model that captures this idea.

The values for workers and entrepreneurs continue to be defined by equations (3), (4), (5), and (6). Now, assume that wages are fixed at some exogenous level, $\bar{w}$, such that the implied capital values for entrepreneurs and workers satisfy $J > 0$ and $W > U$. Hall (2005) justifies this assumption on wage determination as a possible sustainable outcome of a bargaining game. The new equilibrium zero-profit condition from a vacancy creation is

$$\frac{p - \bar{w}}{r + \sigma} = \frac{c}{\lambda_f} = \frac{c}{A\theta^{-\alpha}}. \quad (27)$$

It follows that the impact of a change in labor productivity on market tightness is given by

$$\hat{\theta} = \frac{p}{\alpha(p - \bar{w})} \hat{p}. \quad (28)$$

Comparing this last expression to that in equation (25), we see that the rigid-wage model gives a stronger response. In particular, independent of b , if the average wage, $\bar{w}$, is large as a fraction of output (i.e., if the labor share is large), then market tightness will be very sensitive to small changes in productivity.

The effect on unemployment, given the changes in θ , is the same whether or not wages are rigid, as given by equation (26). Finally, a comparison of equations (21) and (27), reveals that by choosing a value for the worker's bargaining power, β , close to zero in the model with Nash bargaining, one achieves essentially rigid wages, since w is then almost the same as b .

Table 2 Parameters and Steady States for Calibrations

Common across Calibrations			
$r = 0.012, \alpha = 0.72, p = 1,$ $A = 1.35, \lambda = 1.35, \theta = 1, u = 0.07$			
Specific to Calibrations			
	Shimer	Hagedorn & Manovskii	Hall
β	0.72	0.05	NA
b	0.40	0.95	0.40
w/p	0.98	0.97	0.98
b/w	0.41	0.98	0.41
η_{wp}	1.00	0.50	0.00

6. CALIBRATION

In the previous section on comparative statics we demonstrated how steady states change when primitives change. In particular, we have analyzed qualitatively how a permanent productivity change affects labor market tightness (recall that the effect is the same in the short as well as in the long run) and how it influences unemployment in the long run. However, what are the magnitudes of these effects? In order to answer this question we need to assign values to the parameters, and we will do this using “calibration.” We will, to the extent possible, select parameter values based on long-run or micro-economic data. Hence, we will not necessarily select those parameters that give the best fit for the time series of vacancies and unemployment, since we restrict the parameters to match other facts.

The parameters of the model are seven: β , b , p , σ , c , A , and α . The steady state equations that one can use for the calibration are three: (21), (22), and (23). Some aspects of the calibration are relatively uncontroversial, but as we will see below, some other aspects are not. Therefore, we organize our discussion in two parts. We first describe how to assign values to the subset of parameters that allows relatively little choice. We then discuss the remaining parameters and show how, depending on what data one uses to calibrate these, different parameter selections may be reasonable. We also explain why this is a crucial issue—the effect of productivity changes for vacancies, and unemployment may differ greatly across calibrations. We summarize the different calibration procedures in Table 2.

Basic Calibration...

In this section, we follow the calibration in Shimer (2005). We think of a unit of time as representing one quarter. Therefore, it is natural to select $r = 0.012$,

given that the annual real interest rates have been around 5 percent. We choose the separation rate, $\sigma = 0.10$, based on the observation that jobs last about two and a half years on average.¹⁹

Job-finding rates in the data are estimated by Shimer to be 0.45 per month. Thus, a target for λ_w of 1.35 per quarter seems reasonable. Notice from equations (25), (26), and (28) that the response of labor market tightness and the unemployment rate to changes in productivity and other parameters does not depend on the worker-finding rate, λ_f . We therefore follow Shimer and simply normalize labor market tightness, $\theta = 1$, so that the worker-finding rate is equal to the job-finding rate.²⁰

Next, consider the elasticity of the matching function: what should α be? Shimer plots the logarithm of job-finding rates against $\log(v/u)$ and observes something close to a straight line with a slope coefficient of about 0.28, which the theory's formulation, $\lambda_w = A\theta^{1-\alpha}$, says it should be. Therefore, we set $\alpha = 0.72$. Since we have set θ equal to one and λ_w equal to $A\theta^{1-\alpha}$, it follows that $A = 1.35$. From the condition determining steady state unemployment, (22), we now obtain that $0.1(1.35 - u) = u$, so that u is 6.9 percent, which is roughly consistent with the data. Notice also that the system of equilibrium conditions is homogeneous of degree one in c , p , and b . Therefore, we normalize $p = 1$ in steady state.

It remains to select c , b , and β . We have two equations left: the wage equation, (23), and the free-entry equilibrium condition, (21), which is the one that solves for θ in terms of primitives. We can think of this latter equation as residually determining c once b and β have been selected. Two more aspects of the data therefore need to be used in order to pin down b and β .

...but what are b and β ?

We now turn to the more contentious part of the calibration.

Completing Shimer's Calibration

It is common to regard b as being the monetary compensation for the unemployed. The OECD (1996) computes average "replacement rates" across countries, i.e., the ratio of benefits to average wages, and concludes that, whereas typical European replacement rates can be up to 0.70, replacement

¹⁹ For a Poisson process with arrival rate σ , the average time to the arrival of the state change is $1/\sigma$. Thus, the average time from forming a match to separation is $1/\sigma = 10$ quarters.

²⁰ Alternatively, we could have followed Hall (2005) and set the monthly worker-finding rate to one so that $\lambda_f = 3$, implying that $\theta = 1/3$. The value chosen for θ does not influence our results.

rates are at most 0.20 in the United States.²¹ Shimer (2005) sets b equal to 0.4, which is even beyond this upper bound for the replacement rate since it turns out that the wage is close to one in his calibration. One reason why b should be higher than 0.2 is that it also includes the value of leisure associated with unemployment. We will discuss some alternative ways to calibrate b below.

Regarding β , it is common to appeal to the Hosios condition for an efficient search.²² This condition says that in an economy like the present one, firm entry is socially efficient when the surplus sharing parameter, β , is equal to the elasticity parameter of the matching function, α . Thus, Shimer (2005) assumes that $\beta = \alpha$. This is one possible choice, though it is not clear why one should necessarily regard the real-world search outcome as efficient. In conclusion, if $\beta = 0.72$ and $b = 0.4$, from the free-entry condition we obtain $c = 0.324$, and the calibration in Shimer (2005) is completed. Note that Shimer does not use the wage equation in his calibration.

Alternative: Use the Wage Equation

Let us now look at an alternative way of calibrating the model that exploits the wage equation. Hagedorn and Manovskii (2005) point to two observations that arguably can be used to replace those used by Shimer to calibrate b and β .

First, they argue that one can look at the size of profits in the data. Referring to empirical studies, Hagedorn and Manovskii argue that the profit share, which they identify as $(p - w)/p$ in the model, is about 0.03.²³ That is, this calibration strategy is equivalent to selecting a wage share a few percentage points below one. Second, they argue that one can look at how much wages respond to productivity. Using microeconomic data, Hagedorn and Manovskii conclude that a 1 percent productivity increase raises wages by half a per-

²¹ In the United States, unemployment insurance replaces around 60 percent of past earnings, but in the data, unemployed workers earn much less than the average wage.

²² See Hosios (1990). Free entry of firms involves an externality since individual vacant firms do not take into account that variations in the vacancy rate affect the rate at which they meet unemployed workers and the rate at which unemployed workers meet them.

²³ A pure aggregate profit measure should probably take the cost of vacancies into account, and, as such, it should be computed somewhat differently:

$$((1 - u)(p - w) - vc)/(p(1 - u)) = 1 - (w/p) - \theta(c/p)(u/(1 - u)).$$

If this expression equals 0.03, one obtains a smaller wage share, but since c must be less than one for zero profits to be feasible, w/p cannot be below $0.97 - 1 \cdot 1 \cdot 0.05/0.95 \sim 0.92$. Thus, both computations lead to a wage share close to one.

cent.²⁴ We now show how one can use these two observations to determine b and β .

The wage share. The wage income share in the model is obtained by dividing the wage equation (23) by productivity:

$$\frac{w}{p} = \beta \left(1 + \frac{c\theta}{p} \right) + (1 - \beta) \frac{b}{p}. \quad (29)$$

Rearranging the equilibrium condition (21) yields

$$\frac{c\theta}{p} = \frac{(1 - \beta) \lambda_w}{r + \sigma + \beta \lambda_w} \left(1 - \frac{b}{p} \right). \quad (30)$$

It is informative to calculate the wage share implied by Shimer's calculations. Now, given Shimer's preferred parameter values,

$$\frac{c\theta}{p} \approx \frac{1 - \beta}{\beta} \left(1 - \frac{b}{p} \right)$$

since $(r + \sigma)$ is small relative to λ_w . Therefore, with this expression inserted into (29), we conclude that

$$\frac{w}{p} \approx 1,$$

meaning that calibration of the wage share to 0.97 will not *by itself* be a large departure from Shimer's parameterization. Indeed, Shimer obtains a wage share of $w/p = 0.973$.

However, there are several different choices of the pair $(b/p, \beta)$ that can achieve this value of the labor share. To see this, combine equations (29) and (30) by eliminating $c\theta/p$:

$$\frac{w}{p} = \frac{(r + \sigma) [\beta + (1 - \beta) b/p] + \beta \lambda_w}{(r + \sigma) + \beta \lambda_w}. \quad (31)$$

Shimer chooses a relatively large value of β , which makes the wage share in (31) close to one without imposing constraints on b/p . Alternatively, β can be set close to zero, in which case a value for b/p needs to be around one. Recall that with b close to p , the dynamic properties of the model change dramatically. The model has a much stronger amplification mechanism, but how can one justify this choice of β ?

The wage elasticity with respect to productivity. We differentiate (31) in order to derive a relation between $\eta_{\theta p} \equiv d \log \theta / d \log p$, the percentage change in θ in response to a 1 percent increase in p , and $\eta_{wp} \equiv d \log w / d \log p$ (the

²⁴ When we regress the cyclical component of wages on labor productivity (see Table 1 for a description of the data), we obtain an elasticity of 0.57 with the low smoothing parameter and 0.72 with Shimer's smoothing parameter. The first number is higher than, but not too distant from, Hagerdorn and Manovskii's preferred estimate of 0.5. In particular, it is not statistically different from 0.5.

corresponding measure for how wages respond to productivity). We obtain

$$\eta_{wp} = \frac{\beta [1 + (c\theta/p) \eta_{\theta p}]}{w/p}. \quad (32)$$

When $r + \sigma$ is small relative to $\beta\lambda_w$, as in Shimer's calibration, the elasticity of labor market tightness with respect to productivity satisfies

$$\eta_{\theta p} \approx \frac{1}{1 - b/p},$$

demonstrating that the wage elasticity must be

$$\eta_{wp} \approx \frac{1}{w/p}$$

(it must be close to one if the labor share is near one). That is, Shimer's calibration generates a one-for-one wage increase in response to productivity, measured in percentage terms, which is twice as large as the estimates cited by Hagedorn and Manovskii.

To obtain such a low elasticity, one needs to decrease β , so that $r + \sigma$ is no longer small relative to $\beta\lambda_w$, and this is how Hagedorn and Manovskii accomplish the task. A combination of (32) and the exact expression for $\eta_{\theta p}$ from (25) allows us, after some simplifications, to solve for η_{wp} as

$$\eta_{wp} = \left(\frac{\beta}{w/p} \right) \left[\frac{\alpha(r + \sigma) + \lambda_w}{\alpha(r + \sigma) + \beta\lambda_w} \right]. \quad (33)$$

It is now easy to see that using the baseline (uncontroversial) calibration together with $w/p \approx 1$ and $\beta = 0.13$ takes us to a number for η_{wp} that is closer to one-half.²⁵ Notice also that when β is close to zero, the approximation that $\eta_{\theta p} \approx p/(p - b)$ is no longer so good; rather, $\eta_{\theta p}$ is significantly higher than $p/(p - b)$, thus further strengthening the amplification of shocks in the model.

Put differently, if we restrict the model so that it generates a weaker response of wages to productivity, then expression (33) tells us that β has to be significantly smaller. And as we saw before, that (together with a wage share sufficiently close to one) totally changes the dynamics of this model.

How does the calibration influence the amplification from productivity to unemployment? As seen in (26), the transmission from θ to u depends only on α and on u itself, so there is little disagreement here. The contentious parts of the calibration do not influence this channel. That is, the differences in the amplification of unemployment between the alternative calibrations are

²⁵ Again, we need to remind the reader that our wage elasticity is defined for a one-time permanent change of productivity. Hagedorn and Manovskii (2005) base their analysis on an economy with recurrent and persistent, but not permanent, shocks. Therefore, our calibration results for various parameters can differ somewhat.

inherited from the differences in how these calibrations amplify labor market tightness.

Some Further Remarks on Calibration

What is the value of the labor share? Apparently, relatively minor differences in how close the wage share is to one make a significant difference in the results. It seems to us, however, that wage shares are very difficult to calibrate properly without having the other major input in the model, namely capital. Of course, some search/matching models do allow an explicit role for capital. Pissarides (2000), for example, discusses a matching model where firms, once they have matched, rent capital in a frictionless market for capital. Thus, a neoclassical (or other) production function can be used, and the wage share can be calibrated to the ratio of wage income to total income using the national income accounts. The relevant wage income share, however, is then net of capital income, and the same applies to the definition of output. Alternatively, in Hornstein, Krusell, and Violante (2005), we assume that capital is purchased in competitive markets but that an entrepreneur has to purchase capital first in order to be able to search for a worker—in order to qualify as a “vacant firm.” It is an open question as to whether models with capital will also embody a sensitivity of the amplification mechanism to the calibration of the labor share.

What is the value for the wage elasticity? If one insists that wages are less responsive to the cycle than what is implied by Shimer’s calibration, then there is more amplification from productivity shocks, and the model’s implications are closer to the data. Hall (2005) maintains an even more extreme assumption that wages are entirely rigid; this is why we considered a version of the model with rigid wages. Going back to equation (28), we see that a rather extreme outcome is produced, provided that we still calibrate so that the wage share is close to one. Now inelastic wages and a high wage share interact to boost the amplification mechanism. However, the model has the counterfactual implication that the labor share, w/p , is perfectly negatively correlated with output while only mildly countercyclical in the data.

What is the value for the elasticity of unemployment to benefits? Finally, a possible third clue for calibrating this model can be obtained if one has information about how the economy responds to changes in unemployment compensation.²⁶ Of course, the absence of controlled experiments makes it difficult to ascertain the magnitude of such effects. The upshot, however, is that if the response of θ to p is large (because b is close to p), then the response of an increase in unemployment compensation would be a very sharp decrease in θ (and increase in unemployment). In particular, as explained in Section 5, the elasticity of the exit rate from unemployment with respect to b equals

²⁶ This way of assessing matching models was proposed in Costain and Reiter (2003).

$(1 - \alpha)$ times $\eta_{\theta b}$. Given $\alpha = 0.72$, the Hagedorn-Manovskii calibration implies that this elasticity equals -6.3 (see Table 3). Thus, a 10 percent rise in unemployment benefits would increase expected unemployment duration ($1/\lambda_w$) by roughly 60 percent.

The existing estimates of the elasticity of unemployment duration with respect to the generosity of benefits, which are based on “quasi-natural” experiments, are much smaller. Bover, Arellano, and Bentolila (2002) find for Spain that not receiving benefits increases the hazard rate at most by 10 percent, implying a local elasticity of 0.1. For Canada, Fortin, Lacroix, and Drolet (2004) exploit a change in the legislation that led to a rise in benefits by 145 percent for singles below age 30 and estimate an elasticity of the hazard rate around 0.3. For Slovenia, van Ours and Vodopivec (2004) conclude that the 1998 reform which cut benefits by 50 percent was associated with a rise in the unemployment hazard by 30 percent at most, implying an elasticity of 0.6. Finally, an earlier survey by Atkinson and Micklewright (1991) argues that reasonable estimates lie between 0.1 and 1.0.

In sum, these estimates mean that the elasticity implied by the Hagedorn-Manovskii parametrization is between six and sixty times larger than the available estimates.

7. QUANTITATIVE RESULTS FOR THE DIFFERENT CALIBRATIONS

In this section, we show that the three alternative calibrations discussed in Section 6 have very different quantitative implications for the comparative statics discussed in Section 5. Note that, although the values for certain key parameters— β and b in particular—are different, the steady state values of the key aggregate variables are the same across parameterizations. The reason, as explained, is that certain parameters are not uniquely identified in steady state.

Implications for θ , λ_w , and u

Table 3 summarizes the results for the preferred calibrations of Shimer, Hagedorn and Manovskii, and Hall. Recall that Hall’s parameterization has a constant wage.

With Shimer’s calibration, the model has a very poor amplification mechanism.²⁷ A 1 percent permanent rise in productivity leads only to a 1.7 percent

²⁷ Though Table 3 contains information about the comparative statics of separation rates, we focus the discussion on the effects of productivity. Shimer (2005) shows that in terms of equation (1), most unemployment volatility in the U.S. economy is accounted for by variations in job creation (the job-finding rate), as opposed to job destruction (the job-separation rate). Furthermore,

Table 3 Steady State Elasticities

Response of to change in	θ			λ_w			u		
	p	b	σ	p	b	σ	p	b	σ
Shimer	1.72	-0.69	-0.07	0.48	-0.19	-0.02	-0.45	0.18	0.95
Hagedorn & Manovskii	23.72	-22.51	-0.08	6.64	-6.30	-0.02	-6.18	5.87	0.95
Hall	81.70	0.00	-8.17	22.88	0.00	-2.29	-21.30	0.00	3.06

rise in market tightness, a response that is below that in the data by a factor of 16. Similarly, unemployment and the job-finding rate move very little in the wake of a productivity change. Shimer attributes the failure of the model to the fact that, with Nash bargaining, the wage is too closely linked to productivity and absorbs too large a fraction of the productivity fluctuations. As a result, profits do not rise enough to give firms the incentive to create many additional vacancies.

Hall's calibration imposes a constant wage.²⁸ The consequences of this assumption are striking: Market tightness and unemployment respond almost 50 times more than in Shimer's baseline model. Since wages are fixed, a rise in productivity translates entirely into profits. Firms post many more vacancies, which also boost the volatility of the job-finding rate, λ_w .

Hagedorn and Manovskii's calibration, finally, leads to the best results for the volatility of market tightness and for the job-finding rate with respect to productivity shocks: A 1 percent productivity increase leads to a 20 percent increase of market tightness and a 7 percent increase of the job-finding rate. The main problem, however, is that this calibration induces what seems to be excessive sensitivity of u to unemployment benefits b . The elasticity is about six—almost 20 times larger than the number resulting from Shimer's calibration. To interpret what this magnitude means, consider a policy experiment where unemployment benefits are raised by 15 percent; the unemployment rate would then double under Hagedorn and Manovskii's calibration.²⁹

as Table 3 demonstrates, variations in the job-separation rate have a negligible effect on the job-finding rate.

²⁸ For the calibration of Hall's sticky-wage model, we match the wage income share and the unemployment benefits from the Shimer calibration. In all other respects, the calibration is the same as for the Shimer calibration.

²⁹ The fact that the Hagedorn and Manovskii parameters are chosen such that wages do not respond strongly to changes in productivity implies that wages respond strongly to changes in benefits. For the Hall calibration, wages are simply assumed to be fixed, which imposes no additional restrictions on calibration. Thus, even though wages are less responsive than under Hagedorn and Manovskii, changes in b have no impact on the equilibrium. When wages are fixed exogenously, the level of benefits is irrelevant.

Quantitative Implications for the Cyclical-ity of the Labor Share

From equation (32), it is straightforward to rewrite the elasticity of wages, w , and the labor share, s , with respect to a productivity shock, p , as

$$\begin{aligned}\hat{w} &= \frac{\beta}{s} \hat{p} + \beta \frac{c\theta/p}{s} \eta_{\theta p}, \text{ and} \\ \hat{s} &= \hat{w} - \hat{p},\end{aligned}$$

where $\eta_{\theta p}$ denotes the elasticity of θ with respect to p .

For Hall's calibration, the implications are immediate—the model has the counterfactual implication that the volatility of wages is zero and that the correlation between the labor share and labor productivity is minus one. With Shimer's calibration, $\hat{w} \approx 1.15$, and, hence, wages respond one-for-one to labor productivity, absorbing most of their impact, as explained above. Compared to the data, wages are too volatile. The labor share is essentially acyclical, in contrast with the data. Thus, the baseline calibration of the matching model with a low b also fails along these two dimensions.

Hagedorn and Manovskii's parameter choice is constructed to match $\hat{w} = 0.5$, and therefore $\hat{s} = -0.5$. Under this parameterization, the model is quite successful in matching the elasticity of the labor share, since in the data, the labor share is about as volatile as labor productivity and is countercyclical. Here, it is evident that the choice made by Hagedorn and Manovskii of setting β near zero is useful since one can reconcile a large value for $\eta_{\theta p}$ with small fluctuations in the wage and a countercyclical labor share.

8. THE MATCHING MODEL WITH AGGREGATE RISK

In the comparative statics exercise above we have studied how long-run outcomes in our model economy respond to one-time permanent changes in parameters. Yet we want to evaluate how well the model matches the business cycle facts of the labor market, and the business cycle is arguably better described by recurrent stochastic changes to parameters. For this reason we now modify the model and include stochastic productivity shocks that are persistent but not permanent.

One might conjecture that the difference between the effects of one-time permanent shocks and persistent—but not permanent—shocks will be smaller, the more persistent the shocks are. In this case the difference between the comparative statics exercise and the analysis of the explicit stochastic model might be small since labor productivity is quite persistent. The autocorrelation coefficient is around 0.8 (see Table 1). It turns out that the difference between the two approaches is noticeable, but it does not overturn the basic conclusion from the comparative statics analysis. If the calibration is such that wages

respond strongly to changes in productivity, then productivity shocks cannot account for the volatility of the labor market.

The modified model can be analyzed in almost closed form—again because free entry makes vacancies adjust immediately to any shock. Thus, as before, unemployment is a state variable, but it will only influence its own dynamics (and, residually, that of vacancies), whereas all other variables will depend only on the exogenous stochastic shocks in the economy. Again, the argument that backs this logic up proceeds by construction: specify an equilibrium of this sort, and show that it satisfies all the equilibrium conditions.

We will focus on a simple case in which the economy switches between a low-productivity state, $p_1 = p(1 - \mu)$, and a high-productivity state, $p_2 = p(1 + \mu)$, with $\mu > 0$. The switching takes place according to a Poisson process with arrival rate τ .³⁰ The capital values of (un)matched firms and workers, (3) to (6), are easily modified to incorporate the dependence on the aggregate state of the economy:

$$rJ_i = p_i - w_i - \sigma(J_i - V_i) + \tau(J_{-i} - J_i), \quad (34)$$

$$rV_i = -c + \lambda_f(\theta_i)(J_i - V_i) + \tau(V_{-i} - V_i), \quad (35)$$

$$rW_i = w_i - \sigma(W_i - U_i) + \tau(W_{-i} - W_i), \text{ and} \quad (36)$$

$$rU_i = b + \lambda_w(\theta_i)(W_i - U_i) + \tau(U_{-i} - U_i), \quad (37)$$

for $i = 1, 2$, where $-i$ denotes 1 if $i = 2$ and vice versa. Each value equation now includes an additional capital gain/loss term associated with a change in the aggregate state. We continue to assume that wages are determined to implement the Nash-bargaining solution for the state-contingent surplus, $S_i = J_i - V_i + W_i - U_i$, and that there is free entry: $V_i = 0$.

We now apply the surplus value definition and the free-entry condition to equations (34) to (37) in the same way as for the steady state analysis in the previous sections. The equilibrium can then be characterized by the following equations:

$$(r + \sigma + \tau) \frac{c}{(1 - \beta)\lambda_f(\theta_i)} = p_i - rU_i + \tau \frac{c}{(1 - \beta)\lambda_f(\theta_{-i})}, \text{ and} \quad (38)$$

$$(r + \tau)U_i = b + \frac{\beta\lambda_w(\theta_i)c}{(1 - \beta)\lambda_f(\theta_i)} + \tau U_{-i}, \quad (39)$$

for $i = 1, 2$. The idea is to see how an increase in μ from zero—when $\mu = 0$, we are formally in the previous model without aggregate shocks—will influence labor market tightness: If p goes up by 1 percent, that is, μ increases by 0.01, by how many percentage points does θ_1 go down and θ_2 go up? And how does the answer depend on τ ? We will find answers with two different methods. First we will use a local approximation around $\mu = 0$,

³⁰ The model can easily be extended to include a large but finite number of exogenous aggregate states.

which allows us to derive an elasticity analytically. Then we will look at a particular value of $\tau > 0$ and compute exact values for θ_1 and θ_2 .

Local Approximations

For a local approximation at a point where the two states are identical ($\mu = 0$), the equilibrium is symmetric such that θ_1 goes down by the same percentage amount by which θ_2 goes up. For this case, we can show explicitly how the equilibrium elasticity depends on the persistence parameter, τ . We solve for the elasticity in two steps.

First, taking the total derivative of expression (38) with respect to a change in productivity, μ yields

$$\left\{ (r + \sigma + \tau) \frac{\alpha c}{\lambda_{f,i}} \right\} \eta_i = (1 - \beta) \left(\left[(-1)^i p - \frac{\partial r U_i}{\partial \mu} \right] \right) + \left\{ \tau \frac{\alpha c}{\lambda_{f,-i}} \right\} \eta_{-i}, \quad (40)$$

where $\eta_i \equiv (\partial \theta_i / \partial \mu) (1 / \theta_i)$ denotes the elasticity of tightness in state i with respect to a change in productivity. Since we consider only a small productivity difference across states, we approximate the terms in curly brackets by the non-state-contingent steady state values for $\mu = 0$. Furthermore, since everything is symmetric, the solution is such that

$$\eta = \eta_2 > 0 > \eta_1 = -\eta. \quad (41)$$

Inspecting the results, it is easy to see that *keeping the effect of productivity on the flow value of unemployment constant*, the absolute value of the elasticity is higher, the lower τ is. The response of labor market tightness to productivity is stronger, the more persistent the shock is. However, higher productivity also raises the flow value of unemployment, which hurts firms, and this effect goes in the opposite direction.

In order to understand the latter effect, we need to solve expression (39) for the flow return on unemployment as a function of labor market tightness:

$$r U_i = b + \frac{\beta}{1 - \beta} c \frac{(r + \tau) \theta_i + \tau \theta_{-i}}{r + 2\tau} \text{ for } i = 1, 2.$$

We see here that if there is no discounting ($r = 0$), productivity would not affect the flow value of unemployment since it would raise θ_2 and lower θ_1 , but the two effects are symmetric and cancel each other out. However, discounting results in a larger weight on the current aggregate state. To analyze the effect in detail, take the total derivative with respect to the change in productivity,

Table 4 Elasticity of Tightness with Respect to Productivity, $\eta_{\theta p}$; Local Approximation

τ	0.00	0.01	0.02	0.05	0.10	0.50
Average Duration (in years)	∞	25.00	12.50	5.00	2.50	0.50
Shimer	1.72	3.94	5.42	7.04	6.47	2.23
Hagedorn & Manovskii	23.67	29.28	28.24	22.12	15.49	4.42
Hall	81.70	69.32	60.20	43.16	29.33	8.23

μ , to deliver³¹

$$\frac{\partial r U_i}{\partial \mu} = c\theta \frac{\beta}{1-\beta} \frac{[(r+\tau)\eta_i + \tau\eta_{-i}]}{r+2\tau} \text{ for } i = 1, 2. \quad (42)$$

Again using symmetry (41) in (42) we obtain

$$\frac{\partial r U_2}{\partial \mu} = -\frac{\partial r U_1}{\partial \mu} = c\theta \frac{\beta}{1-\beta} \frac{r}{r+2\tau} \eta. \quad (43)$$

It is apparent that with discounting ($r > 0$), the elasticity of the flow return on unemployment is positively influenced by productivity (it goes up in state two relative to state one), and as shocks become more persistent (as τ falls), the effect is stronger (for a given value of η). We also see that changes in the persistence parameter have a bigger impact on how the flow unemployment value responds to productivity when the persistence parameter is large: τ appears in the denominator so that when it is large, the effects are close to zero. Intuitively, when there is almost no persistence, the flow value of unemployment almost does not react to productivity because it is so short-lived, and small changes in persistence become unimportant too.

Inserting expression (43) into (40) and using symmetry again, one obtains the following expression for the elasticity of labor market tightness:

$$\left\{ (r + \sigma + 2\tau) \frac{\alpha c}{\lambda_f} + \theta \beta \frac{r}{r+2\tau} \right\} \eta = \frac{(1-\beta)p}{c}. \quad (44)$$

In Table 4, we display the elasticity of labor market tightness with respect to labor productivity for our three different calibrations and how the elasticity depends on the persistence of the aggregate state. For the purpose of business cycle analysis, an average duration of the state between 2.5 ($\tau = 0.1$) and 5 years ($\tau = 0.05$) appears to be appropriate. We see that for business-cycle durations the results differ from the $\tau = 0$ case, which reproduces the numbers from the comparative statics analysis for a one-time permanent shock. In particular, for less-than-permanent shocks, the different calibrations produce

³¹ We have again approximated the state-contingent values of θ_i with the non-state-contingent steady state value, θ .

Table 5 Elasticity of Tightness with Respect to Productivity; Exact Solution for $\mu = 0.005$ and $\tau = 0.05$

	1 to 2	2 to 1
Shimer	7.30	−6.80
Hagedorn & Manovskii	24.98	−19.93
Hall	54.52	−35.28

results that are more similar. The amplification under Shimer's calibration increases relative to the amplification under the alternative calibrations/models. It remains true that with the Shimer calibration labor productivity fluctuations cannot account for the volatility of labor market tightness, whereas the Hagedorn and Manovskii and Hall calibrations come close. Recall that in the U.S. economy labor market, tightness is about 20 times as volatile as labor productivity (see Table 1). For an arrival rate consistent with the persistence of business cycles, $\tau \in [0.05, 0.1]$, labor market tightness for the Shimer calibration is only seven times more volatile than productivity, whereas for the Hagedorn and Manovskii calibrations, tightness is 20 times more volatile than productivity. It is 30 to 40 times as volatile for the Hall calibration.

We also confirm the theoretical analysis above regarding the effects of persistence. The table reveals that the elasticity of labor market tightness with respect to changes in productivity is not necessarily monotone with respect to the arrival rate of the aggregate state change. On the one hand, the more persistent shocks are, the more a productivity increase influences the present value revenue of the firm. As a consequence, more firms enter and labor market tightness goes up. However, as persistence increases, so do the costs of the firms—they are determined by the workers' outside options—and this effect works in the opposite direction. Moreover, as shown above, this latter effect is a particularly important one when persistence is large but a relatively unimportant one when shocks are very short-lived. Thus, as Table 4 shows, the response of labor market tightness to productivity first increases as persistence goes up and then decreases when shocks become close to permanent, and the effect on workers' outside options dominates.

Exact Solution

In Table 5 we display exact results for a case in which a switch from the low productivity to the high productivity represents a 1 percent change in productivity. For Shimer's calibration, this results in a 7.3 percent increase or a 6.8 percent reduction of labor market tightness. The approximation in Table 4 for $\tau = 0.05$ is roughly the average of the elasticities reported in Table 5; thus, the accuracy of the approximations is reasonable.

9. CONCLUSION: WHERE NEXT?

We have reviewed recent literature that assesses the ability of the search/matching model of the labor market to match some key characteristics of labor markets, namely, the large fluctuations in vacancies and in unemployment. We have, in particular, discussed what features of a calibration seem necessary for matching the data within the context of the standard model or of one augmented with an assumption that real wages are rigid. In this discussion, we have tentatively concluded that there is no wholly satisfactory calibration of the basic setup or a simple alteration thereof that allows the key characteristics of the data to be roughly reproduced. On the one hand, one can assume that the value of being at home is almost as large as that of having a job, but that seems somewhat implausible on a priori grounds, and it implies that there must also be strong sensitivity of unemployment to unemployment benefits, which arguably we do not observe. On the other hand, one can assume rigid wages, but we show that rigid wages necessitate a wage share close to one in order to be powerful in creating large fluctuations in labor market variables, and this route moreover produces an excessively volatile labor share.

It is an open question as to where one might go next. In our view, it seems important to first examine a model with capital, because the results we report above are very sensitive to the value of the labor share. In a model with capital, there is no ambiguity about how one should interpret the labor share. Moreover, a model with capital offers another natural source of fluctuations in vacancies and unemployment, namely, fluctuations in the price of investment goods. Such fluctuations will directly influence the incentives for firms to enter/open new vacancies, and, hence, seem a promising avenue for further inquiry.

REFERENCES

- Andolfatto, David. 1996. "Business Cycles and Labor Market Search." *American Economic Review* 86 (1): 112–32.
- Atkinson, Anthony B., and John Micklewright. 1991. "Unemployment Compensation and Labor Market Transitions: A Critical Review." *Journal of Economic Literature* 29 (4): 1679–1727.
- Bover, Olympia, Manuel Arellano, and Samuel Bentolila. 2002. "Unemployment Duration, Benefit Duration and the Business Cycle." *Economic Journal*, Royal Economic Society 112 (127): 223–65.
- Costain, James S. and Michael Reiter. 2003. "Business Cycles,

Unemployment Insurance, and the Calibration of Matching Models.” CESifo Working Paper 1008.

Fisher, Jonas D.M. 2003. “Technology Shocks Matter.” Federal Reserve Bank of Chicago Working Paper 2002-14, rev. 2003.

Fortin, Bernard, Guy Lacroix, and Simon Drolet. 2004. “Welfare Benefits and the Duration of Welfare Spells: Evidence from a Natural Experiment in Canada.” *Journal of Public Economics* 88 (7–8): 1495–1520.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. 2000. “The Role of Investment-Specific Technological Change in the Business Cycle.” *European Economic Review* 44 (1): 91–115.

Hagedorn, Marcus, and Iourii Manovskii. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited.” Mimeo, University of Pennsylvania.

Hall, Robert E. 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American Economic Review* 95 (March): 50–65.

Hornstein, Andreas, Giovanni L. Violante, and Per Krusell. Forthcoming. “The Replacement Problem in Frictional Economies: A Near Equivalence Result.” *Journal of the European Economic Association*.

Hosios, Arthur J. 1990. “On the Efficiency of Matching and Related Models of Search and Unemployment.” *Review of Economic Studies* 57(2): 279–98.

Lagos, Ricardo. 2000. “An Alternative Approach to Search Frictions.” *Journal of Political Economy* 108 (October): 851–73.

Merz, Monika. 1995. “Search in the Labor Market and the Real Business Cycle.” *Journal of Monetary Economics* 36 (2): 269–300.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies* 61 (3): 397–416.

Nash, John F. 1950. “The Bargaining Problem.” *Econometrica* 18 (2): 155–62.

Organisation for Economic Co-operation and Development. 1996. *OECD Employment Outlook*. OECD: Paris.

Osborne, Martin J., and Ariel Rubinstein. 1990. *Bargaining and Markets*. San Diego: Academic Press.

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

- Pissarides, Christopher A. 1985. "Short Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages." *American Economic Review* 75 (4): 676–90.
- _____. 2000. *Equilibrium Unemployment Theory*, 2nd edition. Cambridge: MIT Press.
- Prescott, Edward C. 1986. "Theory Ahead of Business Cycle Measurement." Federal Reserve Bank of Minneapolis *Quarterly Review* 10 (Fall): 9–22.
- Rogerson, Richard, Robert Shimer, and Randall Wright. Forthcoming. "Search-Theoretic Models of the Labor Market." *Journal of Economic Literature*.
- Shimer, Robert (2004). "The Consequences of Rigid Wages in Search Models." *Journal of the European Economic Association (Papers and Proceedings)* 2 (May–April), 469–79.
- _____. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (March): 24–49.
- van Ours, Jan C., and Milan Vodopivec. 2004. "How Changes in Benefits Entitlement Affect Job-Finding: Lessons from the Slovenian 'Experiment.'" Institute for the Study of Labor Discussion Paper 1181.