

Wage Rigidities and Jobless Recoveries

Robert Shimer*

March 9, 2012

Abstract

Real wage rigidities cause jobless recoveries. Suppose that a one-time shock reduces the capital stock below trend. If wages are flexible, they decline and employment increases at the moment of the shock, before both revert back to normal levels as the economy grows back to trend. If wages are completely rigid and the labor market is otherwise frictionless, the shock causes a proportional and permanent decline in employment, capital, output, consumption, and investment relative to trend. In a search model with rigid wages, the shock causes a persistent but not permanent decline in these economic outcomes, a jobless recovery.

*This paper was prepared for the Journal of Monetary Economics/Swiss National Bank/Study Center Gerzensee conference on “Directions for Macroeconomics: What Did We Learn from the Economic Crisis?” I am grateful for comments from Steven Davis, Robert King, Rafael Lavive, Ivan Werning, two anonymous referees, and numerous seminar participants on an earlier version of this paper, and for financial support from the National Science Foundation.

1 Introduction

Wage rigidities can significantly propagate shocks through time. Suppose an economy that is on a balanced growth path is hit by a transitory shock. While the shock is “on,” low investment reduces the level of the capital stock below the balanced growth path. This paper is concerned with the subsequent dynamics of the economy, after the shock is turned “off.” If the labor market clears, wages must fall below the balanced growth path, employment typically rises above the balanced growth level, and the economy reaccumulates capital and eventually returns to the balanced growth path. If instead the path of real wages is unchanged and employment is determined by labor demand, there is a permanent proportional decline in employment and in capital, output, consumption, and investment relative to trend. The economy exhibits a jobless recovery.

The same logic holds if there are search frictions in the labor market. If wages decentralize the social planner’s problem, they are low during the recovery, employment is typically high, and the economy returns to the balanced growth path. If wages are rigid, there is a long-lasting (but not permanent) decline in employment, and in capital, output, consumption, and investment relative to trend. Search frictions are also useful for rationalizing why wages continue to grow on trend after an adverse shock. From the perspective of individual workers and firms in the search economy, expectations of trend wage growth is as plausible as expectations of flexible wages that decentralize the solution to the planner’s problem. In particular, in both cases wages are individually rational in the sense that they lie below the marginal product of labor and above the marginal rate of substitution between consumption and leisure.

The empirical motivation for this paper is the anemic behavior of the U.S. economy during the 2000s and the disastrous experience since the onset of the financial crisis in 2008 (Figure 1). From the first quarter of 2008 until the fourth quarter of 2010—the latest available data at the time this paper was written—the employment-population ratio fell by 7.5 log points and average per capita weekly hours fell by 9.5 log points. Over the same time

period, real GDP per capita grew 7.0 log points slower than trend and real nondurable and services consumption per capita grew 6.5 log points slower than trend.¹ This extraordinary contraction came on the heels of a decade of weak growth, with essentially no recovery from the mild recession at the start of the decade. Many forecasters predict similarly weak growth in the years ahead and few see any prospect for an early return to the employment level in 2007 and certainly not to the level in 2000.

[Figure 1 about here.]

Real labor costs, as measured by the Employment Cost Index for total compensation, deflated by the Consumer Price Index for all urban consumers, and detrended, have been quite stable throughout this episode.² Figure 2 shows little evidence for a decrease in employment costs relative to trend during the 2001 recession, although they did start to fall several years later as the real economy reached its trough. Employment costs started rising again in the middle of 2008, mirroring a decrease in price inflation at this stage of the recession. By the second quarter of 2008, employment costs were essentially back at the same level as in 2007. The model predicts a close link between employment costs and productivity and the figure also confirms this. Both output per worker and output per hour are back at their 2007 levels as well, with time patterns that closely mirror the path of labor costs.³ In short, Figures 1 and 2 do not provide any evidence that real wages are procyclical.

[Figure 2 about here.]

¹These calculations assume a linear trend from 1951 to 2008. Over this time period, the employment-population ratio grew by 0.07 log points per quarter and average weekly hours grew by 0.04 log points, so detrending these variables is quantitatively unimportant. On the other hand, GDP per capita and consumption per capita grew by 0.45 and 0.52 log points per quarterly, respectively, so here detrending is more important. The theory in this paper is consistent with balanced growth. That is, trend productivity growth causes trend increases in output and consumption without affecting employment or hours, offering a theoretical justification for the differential treatment of these time series. But in practice it is not important whether employment and hours are detrended.

²The Employment Cost Index is essentially a price index, measuring the change in the total cost of a hiring a worker to fill a specific, narrowly defined position. This mitigates issues related to changes in the composition of the labor force over the business cycle. The trend for real employment costs is 0.17 log points per quarter.

³The trend for output per hour and output per worker is 0.41 and 0.38 log points.

The basic building blocks of this paper are not new. The underlying framework is a neoclassical growth model, where firms use capital and labor to produce output and individuals consume, save, and work to maximize their utility. I first introduce wage rigidities directly into this model through the extreme assumption that wages grow at the same rate as labor augmenting technology and workers are required to supply whatever labor firms demand. Any excess of labor supply over employment is naturally labeled as unemployment.

While this first model is useful for understanding the consequences of wage rigidities, it is not so useful for understanding why rigid wages may prevail in equilibrium. In particular, this model cannot explain why unemployed workers could not offer to work at less than the prevailing wage (Barro, 1977). Here search frictions are useful. Since Diamond (1971), economists have understood that a small amount of search frictions can lead to an equilibrium that is far from the competitive one. While each worker-firm pair must pay a wage that is close to the wage paid in all other matches, this common wage may potentially take on quite different values, depending on expectations. In particular, there is an interval of wages at which firms are willing to employ workers and workers are willing to work, and this interval may be quite large even if search frictions are relatively unimportant (Hall, 2005). To be clear, however, the model does not explain why wages are rigid, since the flexible wage path is also consistent with individual rationality. Instead, it makes the point that conditional on a path of individually rational wages, the model is consistent with jobless recoveries if and only if wages are rigid.

The literature subsequent to Hall (2005) has focused on whether wage rigidities help to amplify shocks. To my knowledge, this is the first paper that emphasizes how wage rigidities can generate jobless recoveries. Moreover, the previous literature has focused on the response to a particular (aggregate technology) shock. This paper stresses that the behavior of the model does not depend so much on the nature of the shock, but rather on the amount that the capital stock declines and on the subsequent adjustment of wages.

The paper starts by analyzing a neoclassical growth model with a competitive labor

market in which supply and demand are always equal. It considers the transitional dynamics when the capital stock starts off below trend and shows that during the subsequent adjustment dynamics, the real wage always lies below trend. Depending on preferences, employment may rise above trend (because individuals are poorer) or fall below trend (because the wage is below trend). For example, if the period utility function can be expressed as log consumption minus the disutility of work, employment and investment rise on the impact of an adverse shock, while consumption and output fall. In any case, the model never delivers a quantitatively large decline in employment in response to an adverse reallocation shock.

The paper then considers an ad hoc wage rigidity in the neoclassical model: the wage follows the balanced growth path of the neoclassical model and the equilibrium level of employment is determined from firms' labor demand decision, without regard to workers' willingness to supply this labor. In the resulting equilibrium, the levels of output, employment, consumption, and investment are proportional to the capital stock. That is, if a one-time shock destroys one percent of the capital stock, other economic outcomes will also fall by exactly one percent and never recover back to trend.

The paper next turns to the main exercise, showing that the same forces are at work in a model with search frictions. It augments the neoclassical growth model with a recruiting technology through which a firm uses some of its employees to attract unemployed workers (Shimer, 2010). As Rogerson and Shimer (2010) discuss, the natural analog of the (flexible wage) neoclassical growth model is the social planner's solution to that model. The main insights of the frictionless model carry over to the search model. For example, if the period utility function can be expressed as log consumption minus the disutility of work, employment and investment increase following an adverse shock, while consumption and output decline. With strong complementarities between consumption and labor supply, it is possible to generate a decline in all four outcomes; however, employment is still much less volatile than output.

Finally, the paper introduces rigid wages into the search model through an assumption

that the wage path is fixed at the level that decentralizes the balanced growth path in the social planner's problem. Unlike in the frictionless model, the paper does not impose that workers have to supply whatever labor firms demand at that wage. Instead, in equilibrium firms are always willing to employ workers at this fixed wage, although they cut back on recruiting if the wage is too high. Workers are also always willing to work at the wage and indeed would happily take a job at a lower wage if one were available. Thus this wage is consistent with the indeterminacy in equilibrium wages identified by Hall (2005). In this case, even if consumption and leisure are separable, the model generates a decline in output, consumption, investment, and employment relative to trend, with all four detrended variables moving by roughly the same magnitude. The model is therefore broadly consistent with the empirical patterns that we have observed during the recent slump.

The model predicts that employment is low during a recovery because firms cut back on hiring, not because the incidence of unemployment rises. Shimer (2012) argues that this low hiring accounts for the majority of fluctuations in unemployment, a view that was reaffirmed during the recent slump (Elsby, Hobijn, and Şahin, 2010). In particular, this approach accounts for the simultaneous increase in unemployment and decline in vacancies that occurs during most recessions. On the other hand, this approach misses out on some features of the current recession, especially the increase in job vacancies during a period of constant unemployment that started in the second half of 2009. Understanding this breakdown in the Beveridge curve is left for future research.

The paper has three more sections. The next analyzes the frictionless model, covering first the planner's solution, second the decentralized equilibrium, and third the rigid wage model. This develops the intuition for the following section, which analyzes the model with search frictions. It again contains three subsections which go through the same versions of the model. The final section concludes with a brief discussion of the implications of this result for the current recovery and for economic policy.

2 Frictionless Model

An economy consists of a representative household and a representative firm. The household contains a unit measure of individuals with identical preferences. Household members are infinitely-lived and discount the future with factor $\beta \in (0, 1)$. Labor is indivisible and in period t a fraction n_t of household members are employed. The household maximizes the equal-weighted sum of its members' utility. The period utility of an employed household member who consumes c_e is $\frac{c_e^{1-\sigma}(1+(\sigma-1)\gamma)^\sigma}{1-\sigma}$ if $\sigma \neq 1$ and $\log c_e - \gamma$ if $\sigma = 1$. The corresponding value for an unemployed household member⁴ who consumes c_u is $\frac{c_u^{1-\sigma}}{1-\sigma}$ if $\sigma \neq 1$ and $\log c_u$ if $\sigma = 1$. This is the most general specification of preferences consistent with balanced growth and indivisible labor: a proportional change in wages and non-labor income leaves the household's choice of employment unchanged and simply leads to a proportional change in consumption. The parameter $\sigma > 0$ plays several roles. It determines risk-aversion and the intertemporal elasticity of substitution. It is also important for the complementarity between consumption and work and in particular determines the relative consumption of employed and unemployed household members. In particular, if $\sigma > 1$, employed workers optimally consume more than unemployed workers. The parameter $\gamma > 0$ determines the disutility of work or equivalently the utility from leisure. As long as γ and σ are positive, the utility of an unemployed worker is always higher than the utility of an employed worker who consumes the same amount, so working is unpleasant.

Firms use capital k and labor n to produce output with a Cobb-Douglas production technology. The output $k^\alpha(z_t n)^{1-\alpha}$ plus undepreciated capital $(1 - \delta)k$ is then used both for investment and consumption. Labor augmenting technology grows at a constant rate $g \geq 0$, $z_{t+1} = (1 + g)z_t$. The parameter $\alpha \in (0, 1)$ is the capital share of income.

⁴In the frictionless model with a competitive labor market, it would be more correct to refer to household members as nonemployed rather than unemployed, since the activity is in no sense involuntary. The notation introduced here is more consistent with the remainder of the paper.

2.1 Planner's Problem

A planner chooses the time path of consumption for employed workers $\{c_{e,t}\}$, consumption for unemployed workers $\{c_{u,t}\}$, employment $\{n_t\}$, and the capital stock $\{k_{t+1}\}$ to maximize the total utility of the representative household

$$\sum_{t=0}^{\infty} \beta^t \left(n_t \frac{c_{e,t}^{1-\sigma} (1 + (\sigma - 1)\gamma)^\sigma}{1 - \sigma} + (1 - n_t) \frac{c_{u,t}^{1-\sigma}}{1 - \sigma} \right) \quad (1)$$

subject to a law of motion for the capital stock

$$k_{t+1} = k_t^\alpha (z_t n_t)^{1-\alpha} + (1 - \delta)k_t - n_t c_{e,t} - (1 - n_t) c_{u,t}, \quad (2)$$

taking as given the initial capital stock k_0 .

To solve the model, place a multiplier $\beta^t \lambda_t$ on the law of motion for capital. Three sets of conditions describe a solution. The first one gives the allocation of consumption between employed and unemployed workers as

$$c_{e,t} = \frac{c_t (1 + (\sigma - 1)\gamma)}{1 + (\sigma - 1)\gamma n_t} \text{ and } c_{u,t} = \frac{c_t}{1 + (\sigma - 1)\gamma n_t}, \quad (3)$$

where $c_t = n_t c_{e,t} + (1 - n_t) c_{u,t}$ is total consumption. In particular, the ratio of consumption by the employed to consumption by the unemployed is $c_{e,t}/c_{u,t} = 1 + (\sigma - 1)\gamma$. If $\sigma > 1$, the employed consume more than the unemployed because working raises the marginal utility of consumption, so σ indeed represents the complementarity between consumption and employment. In addition,

$$\lambda_t = \left(\frac{c_t}{1 + (\sigma - 1)\gamma n_t} \right)^{-\sigma}, \quad (4)$$

so λ_t is the marginal utility of consumption of a hypothetical representative agent with period utility $c_t^{1-\sigma} (1 + (\sigma - 1)\gamma n_t)^\sigma / (1 - \sigma)$ who consumes c_t and works a fraction n_t of the time.

The second condition is the Euler equation,

$$\lambda_t = \beta \lambda_{t+1} (\alpha \kappa_{t+1}^{\alpha-1} + 1 - \delta), \quad (5)$$

where $\kappa_t \equiv k_t/z_t n_t$ is capital per efficiency unit of labor. This states that the planner must be indifferent between having households consume one additional unit at t , valued at the marginal utility of consumption λ_t , or invest it, with marginal product of capital $\alpha \kappa_{t+1}^{\alpha-1} + 1 - \delta$ at $t + 1$, valued at the discounted marginal utility of consumption at $t + 1$, $\beta \lambda_{t+1}$. These first two conditions hold in every model in this paper.

The final condition is the optimality of employment. For simplicity, assume that some household members work and some are unemployed; this will be the case if the disutility of labor is sufficiently large. At an interior solution for employment, the first order condition for labor is

$$\gamma \sigma \lambda_t^{-\frac{1}{\sigma}} = (1 - \alpha) z_t \kappa_t^\alpha. \quad (6)$$

This equates the marginal rate of substitution between consumption and leisure to the marginal product of labor.

It is straightforward to solve for a balanced growth path along which consumption of employed and unemployed workers, capital, and output grow at rate g , the marginal utility of consumption grows at rate $(1 + g)^{-\sigma} - 1$, and employment is constant. Now suppose that the capital stock is initially below the balanced growth path. To see how the economy behaves, use equation (6) to eliminate κ_{t+1} from equation (5) and equations (3), (4), and (6) to eliminate n_t , $c_{e,t}$, and $c_{u,t}$ from equation (2). Then log-linearize the system of equations in a neighborhood of the balanced growth path. Using $\hat{\lambda}_t$ to denote the log-deviation of λ_t from the balanced growth path, and similarly for \hat{k}_t , gives $\hat{\lambda}_{t+1} = \varepsilon_\lambda \hat{\lambda}_t$ and $\hat{k}_{t+1} = \left(\frac{(\varepsilon_k - \alpha)(1+g) - (1-\alpha)(1-\delta)}{\alpha\sigma(1+g)} \right) \hat{\lambda}_t + \varepsilon_k \hat{k}_t$ where $\varepsilon_\lambda \equiv \frac{\alpha\sigma}{1-\alpha+\alpha\sigma-\beta(1+g)^{-\sigma}(1-\alpha)(1-\delta)} \in (0, 1)$ and $\varepsilon_k \equiv \frac{(1+g)^{\sigma-1}}{\beta\varepsilon_\lambda} > 1$. The eigenvalues of this system are ε_λ and ε_k and so it is saddle-path stable. The eigenvector associated with the stable arm is $(\hat{k}, \hat{\lambda}) = (1, s)$, where $s \equiv -\frac{\alpha\sigma(1+g)(\varepsilon_k - \varepsilon_\lambda)}{(\varepsilon_k - \alpha)(1+g) - (1-\alpha)(1-\delta)}$

is the elasticity of the marginal utility of consumption λ with respect to the capital stock along the saddle path. The slope s is negative, so when the capital stock is below trend, the marginal utility of consumption is above trend.

Now suppose that after some unmodeled adverse shock, the economy starts with one percent less capital than in the balanced growth path, $\hat{k}_0 = -1$. Following this event, the economy stays on the saddle path, and so $\hat{\lambda}_0 = -s$, above its trend value. Eventually $\hat{\lambda}$ falls back to trend as the capital stock is rebuilt. To see how other model outcomes behave along this (log-linear) transition path, start with equation (6), which implies labor productivity $z_t \kappa_t^\alpha$ moves inversely with λ_t , $z_t \kappa_t^\alpha \propto \lambda_t^{-\frac{1}{\sigma}}$. A one log point decrease in the capital stock raises λ by $-s$ log points; therefore, labor productivity must fall upon the impact of the shock, by $-s/\sigma$ log points. It then gradually increases back to its trend value.

Next consider the impact on employment. Equation (6) implies $\hat{n}_0 = \hat{k}_0 + \frac{1}{\alpha\sigma} \hat{\lambda}_0 = -\frac{\alpha\sigma + s}{\alpha\sigma}$. If $\sigma \leq 1$, one can verify that this is negative, so employment rises on the impact of the adverse shock and then gradually falls back to its trend value. This seems like the most natural case. Households are poor and since leisure is a normal good, they cut back on it and so supply more labor. This accelerates the economy's recovery to trend. But this need not be the case. For example, suppose $g = 0$ so there is no growth. In this case, if $\sigma > \frac{(1-\beta(1-\delta))(1-\delta(1-\alpha))}{\alpha\delta} \equiv \bar{\sigma}$, employment initially falls after the adverse shock before rising back to its trend value because substitution effects dominate income effects. Following the loss of capital, the marginal product of labor is low, which discourages households from supplying much labor despite their relative poverty. But while this is a theoretical possibility, I argue later that for any reasonable calibration of the model, employment rises following the impact of the shock.

One can also compute the behavior of consumption, although the expressions are messier. If $\sigma \leq 1$, it is possible to prove analytically that when the capital stock is below trend, consumption is also below trend. If there is no growth, $g = 0$, this is also true when $\sigma \geq \bar{\sigma}$, the case where a reduction in the capital stock lowers employment, although the result may

be reversed at intermediate values.

Finally, turn to investment, $i = z^{1-\alpha}k^\alpha n^{1-\alpha} - c$. In general, the elasticity of investment with respect to the capital stock is also a messy expression, but numerical simulations suggest that when $\sigma \leq 1$, a reduction in the capital stock raises investment as the economy returns to trend. This can again be reversed if σ is sufficiently large. In particular, when $g = 0$ and $\sigma = \bar{\sigma}$, a one percent reduction in the capital stock reduces investment by α percent. Note, however, that the decline in investment is always smaller than the decline in the capital stock and so capital does indeed return to the initial trend.

In summary, at low values of σ , an economy with less capital than trend has high employment and consumption. But for large σ , it is at least theoretically possible to have employment, consumption, and investment all below their trend values, although employment always falls by less than does the capital stock, pushing down labor productivity.

2.2 Decentralized Equilibrium

A competitive equilibrium with complete markets decentralizes the social planner's problem, an application of the welfare theorems. This feature of the neoclassical growth model is well understood, but writing down the decentralized economy formally is a useful prelude to the discussion of a model in which wages do not clear the labor market. Rather than modeling complete markets by specifying employment lotteries (Rogerson, 1988; Hansen, 1985), it is easier to assume that insurance takes place within a large household (Merz, 1995).

A typical household with initial assets a_0 chooses a path for consumption $\{c_{e,t}, c_{u,t}\}$ and employment $\{n_t\}$ to maximize the average utility of its members in equation (1) subject to a lifetime budget constraint

$$\sum_{t=0}^{\infty} q_0^t (n_t c_{e,t} + (1 - n_t) c_{u,t} - w_t n_t) = a_0, \quad (7)$$

where q_0^t is the time 0 price of a unit of consumption at time t and w_t is the wage.

As in the social planner's problem, the household equates the marginal utility of consumption for employed and unemployed individuals, giving equation (3). In addition, the household problem yields the Euler equation

$$\frac{q_0^{t+1}}{q_0^t} = \beta \frac{\lambda_{t+1}}{\lambda_t}, \quad (8)$$

where λ_t is the marginal utility of consumption for a hypothetical representative agent, defined in equation (4). Finally, the household equates the marginal rate of substitution for the representative agent to the wage,

$$\gamma \sigma \lambda_t^{-\frac{1}{\sigma}} = w_t. \quad (9)$$

Next, a firm that initially owns k_0 units of capital chooses a path for capital and employment to maximize the present value of its profits,

$$\sum_{t=0}^{\infty} q_0^t (k_t^\alpha (z_t n_t)^{1-\alpha} + (1-\delta)k_t - k_{t+1} - w_t n_t). \quad (10)$$

The first order condition for capital is

$$q_0^t = q_0^{t+1} (\alpha \kappa_{t+1}^{\alpha-1} + 1 - \delta). \quad (11)$$

Eliminating q_0^{t+1}/q_0^t using the household Euler equation (8) gives the planner's Euler equation (5). In addition, the firm's first order condition for labor equates the marginal product of labor to the wage,

$$(1-\alpha) z_t \kappa_t^\alpha = w_t. \quad (12)$$

Combining this with the household's first order condition (9) recovers the planner's first order condition (6). This establishes the equivalence between the planner's problem and the decentralized equilibrium.

Recall that when the capital stock is below the balanced growth path, labor productivity is below the balanced growth path as well. In the decentralized economy, equation (12) shows that the wage is proportional to labor productivity and so the wage is also low. The next section examines what happens if the wage path cannot adjust in this manner.

2.3 Rigid Wage

This section assumes that the wage grows at the rate of technological progress, $w_t = (1+g)^t w_0$ for some w_0 , and households must supply whatever labor firms demand at that wage. The model is otherwise unchanged. In particular, a household with initial assets a_0 chooses a path for consumption $\{c_{e,t}, c_{u,t}\}$ to maximize the average utility of its members in equation (1) subject to a lifetime budget constraint (7) taking as given a path for intertemporal prices, the wage, and the level of employment. The household Euler equation (8) is unchanged. This, combined with the household budget constraint, completely describes the household's behavior. In particular, equation (9) does not necessarily hold.

The firm's problem is unchanged, since the firm simply maximizes profits taking as given an arbitrary sequence for wages and intertemporal prices. Combining the first order condition for capital, equation (11), with the household Euler equation (8), eliminating employment using equation (12), and using the assumption that $w_t/z_t = w_0/z_0$ gives

$$\lambda_t = \beta \left(\alpha \left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1-\alpha}{\alpha}} + 1 - \delta \right) \lambda_{t+1} = \beta R_0 \lambda_{t+1}, \quad (13)$$

where $R_0 \equiv \alpha \left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1-\alpha}{\alpha}} + 1 - \delta$ is the gross marginal product of capital. In addition, use equations (4) and (12) to eliminate consumption and employment from the resource constraint:

$$k_{t+1} = \left(\left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1-\alpha}{\alpha}} + 1 - \delta \right) k_t - \lambda_t^{-\frac{1}{\sigma}} \left(1 + (\sigma - 1) \gamma \left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1}{\alpha}} \frac{k_t}{z_t} \right). \quad (14)$$

This is again a pair of difference equations for capital and the marginal utility of consumption.

Together with a transversality condition, this completely describes the equilibrium of the economy.

Now suppose that the economy starts from a balanced growth path where $\beta R = (1 + g)^\sigma$. This might, for example, be the decentralized equilibrium of an economy with a competitive labor market that is moving along a balanced growth path. There is a one time, unanticipated, one percent decrease in the capital stock. Thereafter, there are no more shocks but the wage is forced to grow at rate g . I claim that in the unique equilibrium of the economy, employment, consumption, investment, and output all fall by one percent on the impact of the shock and thereafter grow at their usual growth rate: g for consumption, investment, and output, with employment constant at the depressed level. This is a jobless recovery.

To prove this, first observe from equation (13) that if $\beta R = (1 + g)^\sigma$, $\lambda_{t+1} = (1 + g)^{-\sigma} \lambda_t$ along the transition path. Then rewrite (14) as

$$k_{t+1} = \left(\left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1-\alpha}{\alpha}} + 1 - \delta - z_0^{-1} \lambda_0^{-\frac{1}{\sigma}} (\sigma - 1) \gamma \left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1}{\alpha}} \right) k_t - \lambda_0^{-1/\sigma} (1 + g)^t. \quad (15)$$

Now fix

$$\lambda_0^{-\frac{1}{\sigma}} = \frac{\left(\left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1-\alpha}{\alpha}} - \delta - g \right) k_0}{1 + z_0^{-1} (\sigma - 1) \gamma \left(\frac{(1-\alpha)z_0}{w_0} \right)^{\frac{1}{\alpha}} k_0}. \quad (16)$$

Substituting into the difference equation gives $k_{t+1} = (1 + g)k_t$, consistent with the proposed equilibrium. Moreover, at any higher value of $\lambda_0^{-\frac{1}{\sigma}}$, the economy runs out of capital in finite time. At any lower value, the transversality condition is violated. Therefore this value of λ_0 is the only possible equilibrium. Along the subsequent balanced growth path, capital grows at rate g , the marginal utility of consumption shrinks at rate $(1 + g)^{-1/\sigma} - 1$, and employment is constant. This implies that consumption and investment grow at rate g as well.

To summarize, with rigid wages, the model naturally produces a permanent and equal decline in consumption, employment, investment, and output relative to trend in response to a one time shock that reduces the size of the capital stock. Intuitively, employment is

proportional to capital when wages are rigid and so the dynamics of the equilibrium are essentially that of an ‘AK’ model, in which all shocks are permanent (Rebelo, 1991). The question that cannot be answered in the frictionless model is why wages would not adjust following the shock. There are unemployed workers who would like to work at less than the prevailing wage and firms that would be willing to hire them. What prevents these workers from offering to work at that wage? Search frictions provide a natural answer.

3 Search Model

This section extends the neoclassical growth model by assuming that firms have access to two technologies. One uses capital and labor to produce output. The other uses labor to recruit more workers. Each period, every firm divides its workforce between these two tasks. This allocation determines current output and the evolution of the aggregate employment rate, which is now a state variable.

3.1 Planner’s Problem

A planner chooses the time path of consumption $\{c_{e,t}, c_{u,t}\}$, recruiting $\{v_t\}$, employment $\{n_{t+1}\}$, and the capital stock $\{k_{t+1}\}$ to maximize the total utility of the representative household in equation (1) subject to two constraints. First, the economy faces a resource constraint

$$k_{t+1} = k_t^\alpha (z_t(n_t - v_t))^{1-\alpha} + (1 - \delta)k_t - n_t c_{e,t} - (1 - n_t)c_{u,t}. \quad (17)$$

The only modification in this constraint, compared to equation (2) in the frictionless model, is in production. Of the n_t employed workers, v_t are dedicated to recruiting and $n_t - v_t$ to production, reducing the available output. The recruiters are used to attract new workers to the firm, with their success determined by a matching function. In particular, employment evolves as

$$n_{t+1} = (1 - x)n_t + m(v_t, 1 - n_t), \quad (18)$$

where x is the constant fraction of employed workers who lose their job and $m(v_t, 1 - n_t)$ is the number of new matches, an increasing, constant returns to scale function of recruiters and unemployment $1 - n_t$. The matching function follows Pissarides (1985), although its arguments are recruiters and unemployment rather than vacancies and unemployment. The planner takes as given an initial capital stock k_0 and level of employment n_0 .

To solve this model, place a multiplier $\beta^t \lambda_t$ on the law of motion for capital and $\beta^t \xi_t$ on the law of motion for employment. The first order conditions for consumption of the employed and unemployed deliver equations (3) and (4). The first order condition for capital delivers the Euler equation (5), with $\kappa_t \equiv k_t/z_t(n_t - v_t)$ redefined as capital per efficiency units of producers. These results are otherwise unchanged from the model without search frictions.

Next eliminate the Lagrange multiplier ξ between the first order conditions for recruiting and employment. Also replace c_e and c_u with the marginal utility of consumption using equations (3) and (4). This gives

$$\lambda_t(1 - \alpha)z_t\kappa_t^\alpha = \beta m_v(v_t, 1 - n_t) \left(-\gamma\sigma\lambda_{t+1}^{\frac{\sigma-1}{\sigma}} + \lambda_{t+1}(1 - \alpha)z_{t+1}\kappa_{t+1}^\alpha \left(1 + \frac{1 - x - m_u(v_{t+1}, 1 - n_{t+1})}{m_v(v_{t+1}, 1 - n_{t+1})} \right) \right), \quad (19)$$

where m_v and m_u are the derivatives of the matching function with respect to the number of recruiters and unemployed workers, respectively. This expression determines the optimal level of recruiting versus production. An additional producer raises output at t by the marginal product of labor $(1 - \alpha)z_t\kappa_t^\alpha$. This is valued at the marginal utility of consumption, and so the left hand side is the utility cost of recruiting. Each additional recruiter attracts m_v additional workers to the firm next period. This reduces utility directly by the marginal disutility of working. In addition, it allows the planner to shift resources into production next period, while keeping employment at a fixed level two periods hence. More precisely, given the level of employment next period n_{t+1} , the planner sets the number of producers $l_{t+1} \equiv n_{t+1} - v_{t+1}$ to achieve a desired employment level in two periods, $n_{t+2} = (1 - x)n_{t+1} +$

$m(v_{t+1}, 1 - n_{t+1})$. Implicitly differentiating gives $\frac{\partial l_{t+1}}{\partial n_{t+1}} = 1 + \frac{1-x-m_u(v_{t+1}, 1-n_{t+1})}{m_v(v_{t+1}, 1-n_{t+1})}$ when n_{t+2} is held fixed. For each worker who is hired this period, $\frac{\partial l_{t+1}}{\partial n_{t+1}}$ additional workers are available to produce next period, each raising output by the marginal product of labor $(1 - \alpha)z_{t+1}\kappa_{t+1}^\alpha$, which is valued at the marginal utility of consumption λ_{t+1} .

In summary, the model has two endogenous state variables, n and k , two controls, c and v , with κ defined for convenience as the capital-producer ratio and λ as the marginal utility of consumption. The four equations (5), (17), (18), and (19) describe the equilibrium. The consumption of employed and unemployed workers can then be recovered from equation (3).

It does not seem possible to say much more about this model analytically. Instead, I calibrate the model parameters and log-linearize it in a neighborhood of the balanced growth path before studying the behavior of the model economy when it starts off away from the balanced growth path. Since there are two state variables, there are potentially many initial conditions to examine. In practice, however, the economy quickly converges to one-dimensional subset of the (k, n) plane, as explained further below, and so it is natural to focus on that subset of initial conditions.

The calibration fixes a time period as a month and sets the discount factor to $\beta = 0.996$, just under five percent annually. It considers the model both with $\sigma = 1$ and with higher values that allow for consumption-labor complementarity. The parameter γ , governing the taste for leisure, is set to obtain a five percent unemployment rate along the balanced growth path; the exact value depends on the other calibrated parameters. Fix $\alpha = 0.33$ to match the capital share of income in the National Income and Product Accounts and set the growth rate of labor augmenting technology at $g = 0.0018$, or 2.2 percent net growth per year, consistent with the annual measures of multifactor productivity growth in the private business sector constructed by the Bureau of Labor Statistics.⁵ The level of technology z_0 affects the level of output, capital, and consumption, but is otherwise irrelevant. Set the monthly depreciation

⁵See <ftp://ftp.bls.gov/pub/special.requests/opt/mp/prod3.mfptablehis.zip>, Table 4. Between 1948 and 2007, productivity grew by 0.818 log points, or approximately 0.014 log points per year. The model assumes labor-augmenting technical progress, and so multiply g by $1 - \alpha$ to obtain TFP growth.

rate δ to target a trend capital to annual output ratio of 3.2, the average capital-output ratio in the United States since 1948.⁶ The exact value of δ depends on the preference parameter σ , but in the benchmark model this procedure implies $\delta = 0.0028$.⁷

Turn next to the parameters that determine flows between employment and unemployment. The average exit probability from employment to unemployment in the United States is $x = 0.034$ per month Shimer (2005). Following much of the search and matching literature, the matching function is isoelastic function, $m(v, u) = \bar{\mu}v^\eta u^{1-\eta}$, and in particular it is symmetric, $\eta = 0.5$ (see the survey by Petrongolo and Pissarides, 2001). The paper discuss the importance of η more below. Two pieces of evidence pin down the efficiency parameter in the matching function $\bar{\mu}$: the average unemployment rate is about 5 percent during the postwar period; and Hagedorn and Manovskii (2008) and Silva and Toledo (2009) show that recruiting a worker uses approximately 4 percent of one worker's quarterly wage. Shimer (2010) shows that these facts imply that $\bar{\mu} = 2.32$ and that the recruiter-unemployment ratio is $v/u \approx 0.078$ along a balanced growth path. It follows that the share of recruiters in employment is $v/n \approx 0.004$, with 99.6 percent of employees devoted to production.

Log-linearizing around the balanced growth path confirms that the system is saddle-path stable, with two eigenvalues lying inside the unit circle and two lying outside. The stable eigenvalues, associated with the convergent dynamics, are 0.990 and 0.313. The low value of the second eigenvalue implies that when the state variables start away from eigenvector associated with the larger eigenvalue, they quickly converge to that eigenvector. Assume therefore that the economy starts on the eigenvector associated with the larger eigenvalue and ignore the short-lived transitional dynamics coming from the smallest eigenvalue.

In the frictionless model with $\sigma = 1$, a reduction in the capital stock raised employment and investment and reduced consumption. The solid line Figure 3 confirms this result.

⁶More precisely, this is the average ratio of the Bureau of Economic Analysis's Fixed Asset Table 1.1, line 1 measure of the current cost net stock of fixed assets and consumer durable goods, divided by the National Income and Product Accounts' Table 1.1.5, line 1 measure of nominal Gross Domestic Product.

⁷For values of σ in excess of 2.5, the depreciation rate implied by a capital-output ratio of 3.2 is negative, putting some discipline on the complementarity parameter.

Employment and investment increase sharply above trend, while consumption and labor productivity fall below trend during the transitional dynamics. In this case, the employment response is so strong that output is actually above trend when the capital stock is low. All of these effects quickly wear off, closing 1.7 percent of the gap to the steady state in every month. This is faster than the transitional dynamics in a model without an employment margin because the increase in employment accelerates convergence. For example, if the savings rate and employment were constant, the rate of convergence would be $(1 - \alpha)(g - \delta)/(1 + g) = 0.3$ percent per month.

[Figure 3 about here.]

The remaining lines in Figure 3 show various calibrations of the search model. The dashed line shows the benchmark parameterization. The most noticeable impact of the search friction is to dampen the employment response. Intuitively, when employment is high, the efficiency of recruiters is low, discouraging the planner from doing too much recruiting. This offsets, but does not reverse, the incentive to raise employment. Because of the dampened employment response, output of the final good falls in the search model and the response of labor productivity is muted. Finally, search frictions prolong the transitional dynamics, but otherwise do not much affect the dynamics of consumption, investment, labor productivity (final output per worker), or wages.⁸

To see the robustness of this result, consider several variations of the benchmark model. First, raise the complementarity between consumption and labor supply to $\sigma = 2.5$, the highest value consistent with a nonnegative depreciation rate and a capital-output ratio of 3.2. The calibration target for unemployment implies $\gamma = 0.46$, and so employed workers optimally consume $1 + (\sigma - 1)\gamma = 1.69$ times as much as unemployed workers. This is far larger than the reduction in consumption at retirement would suggest is plausible (Aguiar and Hurst, 2005). Still, the model generates an increase in employment and investment

⁸The next section discusses the wage path that decentralizes the planner's problem.

following an adverse shock to the capital stock, with the response simply more muted than in the benchmark model.

The matching function also affects model outcomes. For example, set $\bar{\mu} = 1.61$, which implies that each recruiter attracts only twelve workers per quarter. This moderates the response of employment, but has little effect on the other outcomes. Also consider the elasticity of the matching function. At one extreme, if $m(v, u) = \bar{\mu}u$, employment is fixed at $n = \bar{\mu}/(x + \bar{\mu})$. At the other, if $m(v, u) = \bar{\mu}v$, the transitional dynamics are mathematically identical to those in the frictionless model.⁹ Intuitively, diminishing returns to recruiting limits the willingness of the planner to increase recruiting in response to a shock that would otherwise encourage him to raise the level of employment. This issue disappears in the limit with $\eta = 1$. The bottom line is that no calibration of this model delivers a jobless recovery during the transition back to steady state.

3.2 Decentralization

As a prelude to the rigid wage model, it is useful to show how the social planner's solution can be implemented through Nash bargaining over wages in a decentralized economy.

The problem of a household is unchanged from the frictionless model with rigid wages. A household with initial assets a_0 chooses a path for consumption $\{c_{e,t}, c_{u,t}\}$ to maximize the average utility of its members in equation (1) subject to a lifetime budget constraint (7) taking as given a path for intertemporal prices, the wage, and the level of employment. As in the frictionless model with rigid wages, the household problem yields the Euler equation (8), where λ_t is still the marginal utility of consumption for a hypothetical representative agent. This, combined with the household budget constraint, completely describes the household's behavior.

⁹This result holds if $\sigma = 1$, in which case a constant $\lambda_t z_t k_t^\alpha$ solves equation (19), as in the frictionless first order condition (6); use this to eliminate κ from equation (5) and establish that the dynamics of marginal utility is identical in the two models. Even in this special case, the economy must start with initial conditions consistent with constant $\lambda_t z_t k_t^\alpha$; otherwise the search model has some additional short-lived dynamics coming from the fact that employment is a state variable.

Next write the problem of a firm that starts with k_0 units of capital and n_0 employees. It chooses a time path for recruiting, employment, and capital to maximize its profits

$$\sum_{t=0}^{\infty} q_0^t (k_t^\alpha (z_t(n_t - v_t))^{1-\alpha} + (1 - \delta)k_t - k_{t+1} - w_t n_t), \quad (20)$$

where $n_{t+1} = (1 - x)n_t + v_t \mu(\theta_t)$. The firm places v_t workers into recruiting, each of whom hires $\mu(\theta_t)$ new workers. The function μ is obtained from the matching function as $\mu(v/u) \equiv m(v, u)/v = m(1, u/v)$, decreasing in the recruiter-unemployment ratio. Note that recruiting is constant returns to scale at the firm level, although it generally has diminishing returns at the aggregate level. Each firm takes the aggregate recruiter-unemployment ratio $\theta = v/u$ as given when choosing how much to recruit, although the aggregate outcome will be determined in equilibrium.

To solve this problem, place a multiplier $q_0^t \xi_t$ on the law of motion for employment. The first order condition for capital, combined with the household Euler equation, yields equation (11), where κ_t is again the capital-producer ratio. Eliminating ξ_t between the first order conditions for vacancies and employment and again using the household Euler equation to eliminate intertemporal prices gives

$$(1 - \alpha) z_t \kappa_t^\alpha \lambda_t = \beta \lambda_{t+1} \mu(\theta_t) \left((1 - \alpha) z_{t+1} \kappa_{t+1}^\alpha \left(1 + \frac{1 - x}{\mu(\theta_{t+1})} \right) - w_{t+1} \right). \quad (21)$$

Equation (21) differs from equation (19) because of the presence of the wage. Its logic is basically unchanged, however, still equating the value of a producer and a recruiter. The left hand side is the current marginal product of a producer evaluated at the current marginal utility of consumption. The right hand side is the discounted product of next period's marginal utility of consumption, the number of workers attracted by a recruiter, and the extra output from the newly hired workers and the workers who are freed up from recruiting next period, net of the wage cost of the newly hired workers.

To decentralize the planner's problem, assume that workers bargain with employers to

set wages, with an outcome described by the Nash bargaining solution. A worker's threat point while bargaining is that he becomes unemployed, while a firm's threat point is that it loses the worker. Moreover, let $\phi \in (0, 1)$ denote the worker's bargaining power. Shimer (2010) proves that the wage satisfies

$$w_t = \phi(1 - \alpha)z_t\kappa_t^\alpha(1 + \theta_t) + (1 - \phi)\gamma\sigma\lambda_t^{-\frac{1}{\sigma}}. \quad (22)$$

This is a weighted average of two terms. The first is the marginal product of labor. This accounts both for the output the worker produces $(1 - \alpha)z_t\kappa_t^\alpha$, and for the fact that, if bargaining fails, the firm must place additional workers into recruiting in order to maintain the same size the next period. The second is the marginal rate of substitution between consumption and leisure. In the frictionless model, these two terms are equal and both are equal to the wage. With search frictions, the marginal product of labor generally exceeds the marginal rate of substitution and the wage lies between these two levels. At this wage, firms are happy to hire workers and households are happy to supply labor.

In general, the equilibrium does not decentralize the planner's problem. But suppose that $m(v, u) = \bar{\mu}v^\eta u^{1-\eta}$, so $\mu(\theta) = \bar{\mu}\theta^{\eta-1}$ for some $\eta \in [0, 1]$. If $\phi = 1 - \eta$, one can verify that equations (21) and (22) reduce to the planner's condition in equation (19), and so the decentralized equilibrium is efficient. This is a generalization of the Mortensen (1982)–Hosios (1990) efficiency condition. Conversely, any other wage path that has a different expected present value leads to an inefficient equilibrium.¹⁰

3.3 Rigid Wage Model

The rigid wage model is identical to the decentralization above, except that the wage does not satisfy the Nash bargaining solution, but instead grows at a constant rate g . That is,

¹⁰Other wage paths may have the same expected present value but spread risks differently across workers and firms. For example, an employed worker may earn a constant wage while employed, although the wage depends on the labor market conditions when he is hired. This changes measured wages but does not alter the efficiency of equilibrium (Shimer, 2004).

both the household's problem and firm's problem are unchanged, yielding equation (21).¹¹ All that changes is the wage equation (22); I again impose that the wage that grows with productivity, $w_t = (1 + g)^t w_0$. It is easy to verify that under appropriate conditions, firms do not want to fire their workers at that wage and workers do not want to quit their jobs.

To start, consider the behavior of the model along a balanced growth path. First, the Euler equation (5) reduces to $1 = \beta(1 + g)^{-\sigma} (\alpha z \kappa^{\alpha-1} + 1 - \delta)$, which pins down the capital-producer ratio κ . Then the first order condition for recruiting, equation (21), reduces to $\frac{w_t}{z_t} = (1 - \alpha) \kappa^\alpha \left(1 - \frac{(1+g)^{\sigma-1} + \beta(x-1)}{\beta \mu(\theta)}\right)$, which pins down the recruiter-unemployment ratio θ as a decreasing function of the wage per efficiency unit of labor. Next, the steady state equation for employment, $n = \frac{f(\theta)}{x+f(\theta)}$, pins down n , where $f(\theta) \equiv m(\theta, 1)$ is the job finding probability for unemployed workers as a function of the recruiter-unemployment ratio. Using $\kappa = \frac{k}{z(n-\theta(1-n))}$, one then obtains the ratio of capital to productivity k/z . Finally, the balanced growth equation for capital pins down consumption relative to capital, $c/k = \kappa^{\alpha-1} - \delta - gk$. The wage does not affect κ , but an increase in the wage raises $\mu(\theta)$ and so reduces θ . This in turn lowers employment. In general, the effect on capital and consumption is ambiguous, although it may be natural to focus on parameters for which a reduction in the wage raises the balanced growth number of producers $n - \theta(1 - n)$; this is always the case at the social optimum. But the important point to note is that there is a range of wage paths that are consistent with a well-behaved balanced growth path. In contrast, in the frictionless model there was a unique wage path consistent with balanced growth.

It only remains to check that the wage lies in between the marginal product of labor and the marginal rate of substitution between consumption and leisure, $(1 - \alpha) z_t \kappa^\alpha (1 + \theta) \geq w_t \geq \frac{\gamma \sigma c_t}{1 + (\sigma - 1) \gamma n}$, so firms are willing to hire workers and workers are willing to work for firms. Since analytical results are not available, consider the calibrated model. A wage of $w_t/z_t = 4.022$ decentralizes the social planner's solution. As w_t/z_t increases to 4.040, the balanced growth recruiter-unemployment ratio falls to 0, driving the unemployment rate up to 1, and

¹¹This is in contrast to the frictionless labor market model, in which the household problem changed in the presence of rigid wages.

so at higher wages firms are unwilling to recruit workers. At lower wages, the recruiter-unemployment and employment rate rise, pushing the wage towards the marginal rate of substitution. If $w_t/z_t < 2.995$, households would refuse some employment opportunities because of their capital income is high relative to their potential labor earnings. Still, a large range of wage paths are consistent with workers willingly supplying labor to firms and firms willingly hiring workers.

Why does an equilibrium exist for such a large range of fixed wages? This is essentially the insight in Diamond (1971), extended to a business cycle analysis in Hall (2005). If all firms pay a wage w , no worker would accept a job at a wage much below w because doing so would be dominated by continued job search. On the other hand, if workers anticipate only getting jobs at a lower wage, they will be willing to do so. This strategic complementarity in the wage acceptance decision is enough to leverage small search frictions into a large region of wage indeterminacy. Indeed, with appropriate assumptions about wage setting, the equilibrium of a search economy need not converge to the competitive equilibrium as the search frictions become arbitrarily small.

This analysis illustrates another point: although a balanced growth path exists for a range of fixed wages, the employment rate is extremely sensitive to the wage. A 0.1 percent increase in the wage, from 4.022 to 4.026 raises the unemployment rate from 5 percent to 6.3 percent, while a 0.5 percent increase in the wage shuts down the economy entirely.¹² This occurs because the increase in the wage reduces the profitability of recruiting a worker at a fixed marginal product of labor. With less recruiting, employment falls as well.

To understand why the balanced growth path is so responsive to the wage, it may help to think back to the model without search frictions. In that case, the capital-labor ratio is increasing in the wage relative to productivity and the interest rate is decreasing in the capital-labor ratio. Since a balanced growth path only exists if $\beta R = (1 + g)^\sigma$, it follows that

¹²These are only statements about balanced growth paths. A permanent 0.5 percent increase in the wage from 4.022 to 4.040 will not immediately shut down the economy, but rather causes a gradual increase in the unemployment rate and decline in the capital stock.

it only exists for a single value of the wage relative to productivity. In the search model, this exact logic fails but the balanced growth path remains very sensitive to the wage. Of course, a constant wage in the face of persistent high unemployment may seem unreasonable. A more realistic model would allow the wage to fall, perhaps very slowly, in response to persistent pressure from unemployed workers.

This paper does not attempt to model this type of wage dynamics, but instead considers transitory shocks to the capital stock holding the wage path fixed at some initial level. To do this, use the same parameter values as in the social planner’s problem, with wages that decentralizes the planner’s balanced growth path. Then vary the capital stock holding fixed the wage path. The two stable eigenvalues of the dynamic system are 0.998 and 0.291, so it is saddle-path stable. The low value of the smaller eigenvalue again ensures that the system quickly converges to the eigenvector associated with the larger eigenvalue. And the high value of the larger eigenvalue—it was 0.990 in the flexible wage model—indicates that the convergent dynamics will be slow.

In the rigid wage model, a reduction in the capital stock lowers output, employment, consumption, and investment for a broad range of parameterizations of the model. Moreover, the declines are roughly equal in magnitude. Figure 4 illustrates this both in the benchmark model and in several variants of it. The solid line shows the frictionless model, identical to the model where the elasticity of the matching function is $\eta = 1$. This delivers the exact theoretical result that employment, consumption, investment, and output all decline by an equal magnitude relative to trend.¹³ The dashed line shows the benchmark model, where the decline in employment, output, and investment is slightly smaller. The dotted line shows a higher value of the complementarity parameter, $\sigma = 2.5$; this scarcely affects the results. The dash-dot line shows a much more frictional economy, $\bar{\mu} = 1.61$, so recruiters contact twelve workers per quarter. This slightly reduces the employment response. These figures suggest a jobless recovery, with a low level and anemic growth in employment accompanying

¹³This result holds in the search model with $\eta = 1$ for arbitrary values of the other parameters.

trend growth (at below trend levels) in consumption, investment, and output.

[Figure 4 about here.]

There is one interesting difference between the rigid wage model with a frictionless labor market and the rigid wage model with search frictions: in the first model, the economy never recovers from the adverse shock, while in the second model it does if $\eta < 1$. The reason for the recovery is instructive. Following an adverse shock to the capital stock, suppose employment fell by the same amount. This pushes down the recruiter-unemployment ratio, allowing firms to economize on recruiting. In other words, output of the final good does not fall by as much as the decline in capital and employment. It is then possible to invest a bit more, leading the economy to eventually return to the balanced growth path. The worse is the matching process, the faster is the convergent dynamics, but even with $\bar{\mu} = 1.61$, the larger stable eigenvalue falls only to 0.997. If $m(u, v) \propto v$, this channel is shut down. An adverse shock to the capital stock is permanent, leading to an equal permanent decline in employment and in consumption, investment, and output relative to trend, as in the model without search frictions.

Finally, note that in a flexible wage economy, the wage would fall after this adverse shock. It follows that the rigid wage is always higher than the marginal rate of substitution between consumption and leisure. Put differently, when the unemployment rate is high because wages are too high, individuals are happy to work when they have an opportunity. In addition, the wage is less than the marginal product of labor as long as the shock is not so large that firms stop recruiting new workers. Thus the model can handle a sharp slowdown in recruiting following an adverse reallocation shock without necessitating any change in the wage.

4 Conclusion

Consider an economy that starts along a balanced growth path and then loses some of its capital stock. If the wage does not fall from its trend, the model can explain a persistent

decline in output, employment, consumption, and investment. This is consistent with individually rational behavior in a search model: firms are still willing to hire workers at the high wage, but they simply cut back on recruiting; and workers are happy when they are able to get a job at this high wage. Of course, unemployed workers would like a job, but that is always true in a search model, whether wages are flexible or rigid. The conclusion is robust to the specification of preferences and the nature of the matching technology, arguably the most controversial ingredients in the theory.

I have deliberately tried to avoid a discussion of the nature of the initial adverse shock. For the purpose of exploring the recovery, the nature of the shock is unimportant as long as wages are rigid. That the model is able to deliver a protracted contraction in response to a decline in the capital stock suggests that many different shocks will have a similar effect on economic outcomes. For example, imagine a final good, used both for consumption and investment, that is an aggregate of heterogeneous intermediate goods. When final goods are invested, the capital becomes specific to a particular intermediate variety. Finally, idiosyncratic shocks change the composition of intermediate goods used to produce the final good. When a shock reduces the demand for a particular variety, effectively the associated capital becomes worthless. In a flexible wage model, workers would immediately be reallocated to more productive uses. But with rigid wages this does not happen and the reallocation shock may lead to a protracted contraction.

The quantitative exercises focused on the transitional dynamics of an economy starting with ten percent less capital than along the balanced growth path. This initial condition may come from two sources. First, the protracted contraction in 2008 and 2009 was accompanied by very low investment rates, with gross investment in business fixed capital falling as much as 40 percent below trend. The cumulated effect of this low investment can explain about half the decline in the capital stock. The second source came from poor investment decisions made during the preceding decade, for example over-investment in home construction equipment. The collapse in the housing market reduced the value of that equipment, with little prospect

of reusing it in another industry. Thus the U.S. economy had less effective business capital than trend even at the start of the recession.

I close with brief answers to three questions. First, what is the impact of countercyclical government spending in this type of model? It is straightforward to introduce a government financed by lump-sum taxes. In particular, with lump-sum taxes, Ricardian equivalence holds and so the timing of taxation is irrelevant. In a flexible wage model, wasteful government consumption, or government consumption that enters additively into individuals' utility functions, raises labor supply by making households feel poorer. This carries over to the search model as well. An increase in government consumption lowers wages and labor productivity and raises output, though by less than the increase in spending. With rigid wages, a temporary increase in government consumption crowds out private investment. The decline in the capital stock then reduces employment, consumption, and output.¹⁴ Thus this model, despite its arguably Keynesian rigid wage and noncompetitive labor market, predicts that countercyclical fiscal expansions will have perverse and unintended consequences. On the other hand, policies designed to reduce employment costs, such as countercyclical payroll tax cuts, may have significant, long-lasting, and salutary effects in a depressed economy.

Second, why is the recovery from the 2008–2009 recession so anemic compared to those from past deep recessions, such as the one in 1982–1983? The paper suggests that this reflects an increase in wage rigidities. This may in part be a consequence of the current low inflation environment, combined with a downward nominal wage rigidity. With a higher background inflation rate, firms could achieve a desired reduction in real wages by reducing nominal wage increases, but this process is slow when the inflation rate is below two percent per year. The 2007–2008 financial crisis may also be to blame. To the extent that existing firms are reluctant to cut wages in a depressed labor market, there may be opportunities for new firms to enter. Unfortunately, lingering problems in financial markets may reduce

¹⁴In contrast, an increase in government investment, if more productive than private sector investment, may be expansionary. Still, the model does not explain why this investment should only be performed during recessions.

funding opportunities for these firms, extending the labor market downturn. Both of these hypotheses deserve further scrutiny.

Third, does this paper offer broader lessons for the direction of macroeconomics after the economic crisis? I believe that it does. The crisis did not invalidate the basic amplification and propagation mechanisms stressed by modern macroeconomic models, such as intertemporal substitution in consumption and capital accumulation decisions, but the long slump has made it increasingly difficult to accept that workers are on their labor supply curve. The paper argues that we should put labor market imperfections, and in particular wage rigidities, back into the center of research on both the amplification and propagation of transitory shocks to the economy.

References

- Aguiar, Mark, and Erik Hurst, 2005. “Consumption versus Expenditure.” *Journal of Political Economy*. 113 (5): 919–948.
- Barro, Robert J., 1977. “Long-Term Contracting, Sticky Prices, and Monetary Policy.” *Journal of Monetary Economics*. 3 (3): 305–316.
- Diamond, Peter A., 1971. “A Model of Price Adjustment.” *Journal of Economic Theory*. 3 (2): 156–168.
- Elsby, Michael, Bart Hobijn, and Ayşegül Şahin, 2010. “The Labor Market in the Great Recession.” *Brookings Papers on Economic Activity*. pp. 1–40, Spring.
- Hagedorn, Marcus, and Iourii Manovskii, 2008. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited.” *American Economic Review*. 98 (4): 1692–1706.
- Hall, Robert E., 2005. “Employment Fluctuations with Equilibrium Wage Stickiness.” *American Economic Review*. 95 (1): 50–65.
- Hansen, Gary D., 1985. “Indivisible Labor and the Business Cycle.” *Journal of Monetary Economics*. 16 (3): 309–327.
- Hosios, Arthur J., 1990. “On the Efficiency of Matching and Related Models of Search and Unemployment.” *Review of Economic Studies*. 57 (2): 279–98.
- Merz, Monika, 1995. “Search in the Labor Market and the Real Business Cycle.” *Journal of Monetary Economics*. 36 (2): 269–300.

- Mortensen, Dale T., 1982. "Property Rights and Efficiency in Mating, Racing, and Related Games." *American Economic Review*. 72 (5): 968–979.
- 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): 676–690.
- Rebelo, Sergio, 1991. "Long-Run Policy Analysis and Long-Run Growth." *Journal of Political Economy*. 99 (3): 500–521.
- Rogerson, Richard, 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics*. 21 (1): 3–16.
- Rogerson, Richard, and Robert Shimer, 2010. "Search in Macroeconomic Models of the Labor Market." in Orley Ashenfelter, and David Card (ed.), *Handbook of Labor Economics*, vol. 4A, pp. 619–700. Elsevier.
- Shimer, Robert, 2004. "The Consequences of Rigid Wages in Search Models." *Journal of the European Economic Association*. 2 (2–3): 469–479.
- Shimer, Robert, 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*. 95 (1): 25–49.
- Shimer, Robert, 2010. *Labor Markets and Business Cycles*, Princeton University Press.
- Shimer, Robert, 2012. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics*. forthcoming.
- Silva, José, and Manuel Toledo, 2009. "Labor Turnover Costs and the Cyclical Behavior of Vacancies and Unemployment." *Macroeconomic Dynamics*. 13 (S1): 76–96.

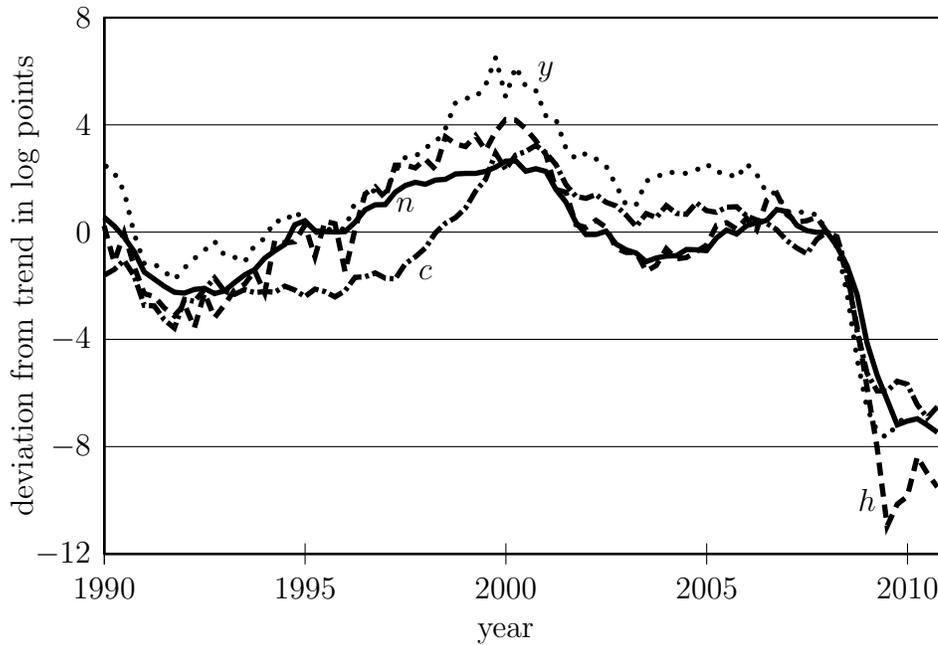


Figure 1: The solid line shows employment per capita n ; the dashed line shows hours per capita h ; the dotted line shows GDP per capita y ; and the dash-dotted line shows consumption per capita c . GDP and consumption are detrended using a linear trend from 1951 to 2008 and expressed as deviations from trend in log points. The end date for all variables is the fourth quarter of 2010.

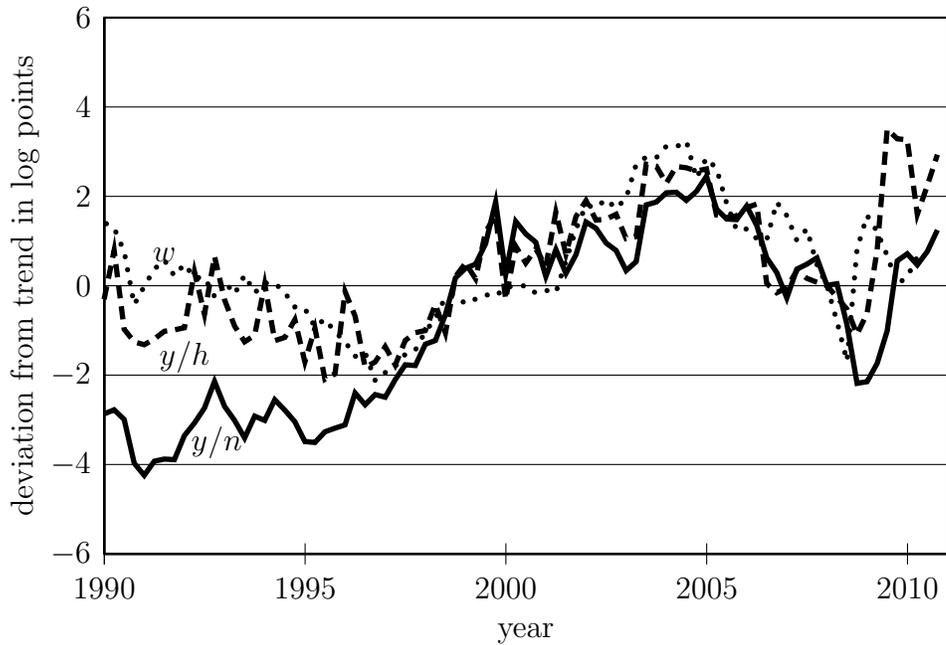


Figure 2: The solid line shows output per capital y/n ; the dashed line shows output per hour y/h ; and the dotted line shows employment costs w . All three variables are detrended using a linear trend from 1951 to 2008 and expressed as deviations from trend in log points. The end date for the first two variables is the fourth quarter of 2010, while the end date for employment costs is the third quarter.

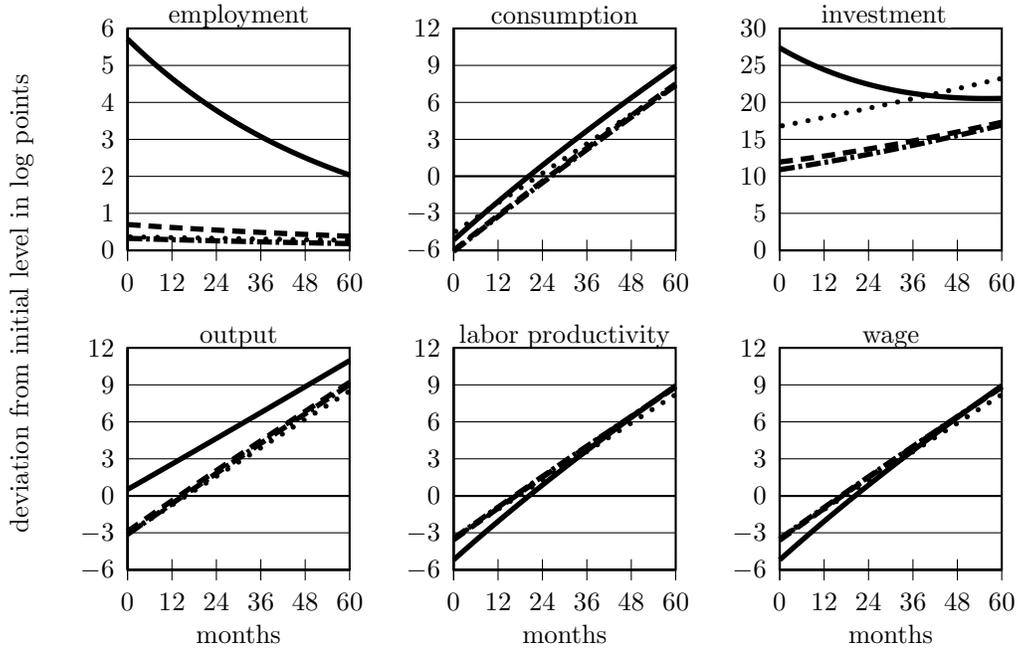


Figure 3: Transitional dynamics starting from ten percent below the trend capital stock, social planner's solution. The solid line shows the frictionless model with $\sigma = 1$. The dashed line shows the benchmark parameterization of the search model. The dotted line shows $\sigma = 2.5$, and the dash-dot line shows $\bar{\mu} = 1.61$. The remaining parameters are fixed at their benchmark values.

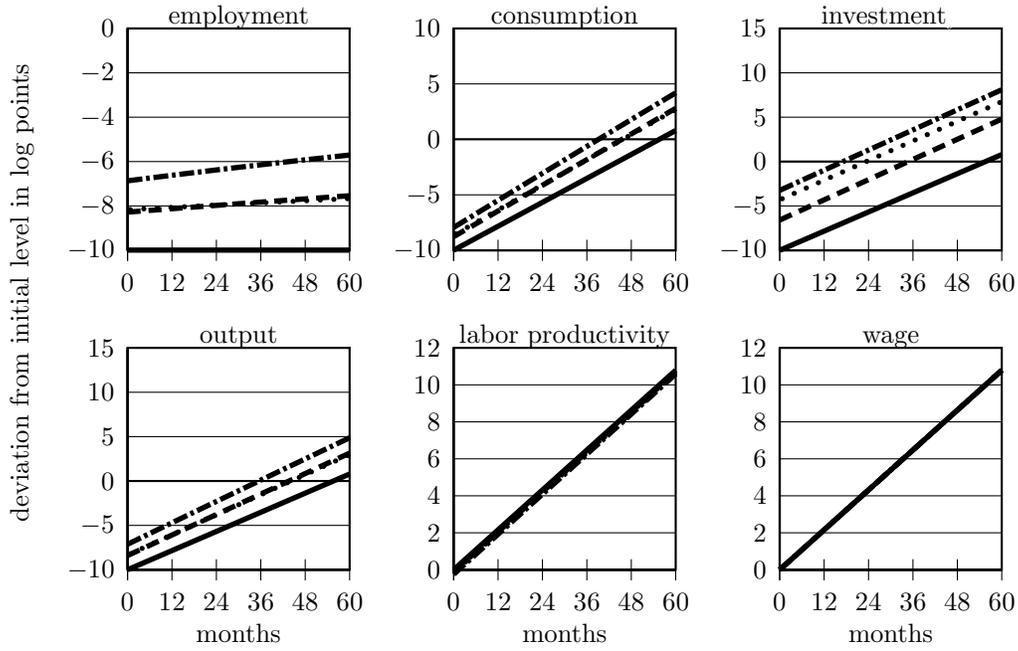


Figure 4: Transitional dynamics starting from ten percent below the trend capital stock, rigid wage model. The solid line shows the frictionless model with $\sigma = 1$. The dashed line shows the benchmark parameterization of the search model. The dotted line shows $\sigma = 2.5$, and the dash-dot line shows $\bar{\mu} = 1.61$. The remaining parameters are fixed at their benchmark values.