EFRSW
7/17/06
12:40 PM

Heterogeneity and the Amplification of Shocks in Matching Models

Michael J. Pries*

June 2006

Abstract

Shimer (2005) demonstrated that in a standard matching model aggregate productivity shocks lead to fluctuations in key labor market statistics—such as the vacancy/unemployment ratio, the job-finding rate, and the unemployment rate—that are too small by an order of magnitude. This paper shows, however, that when the standard model is extended to allow for worker heterogeneity and match heterogeneity, it can generate considerably more realistic amplification. In the model, there are low-productivity and high-productivity workers. Low-productivity workers account for a disproportionate share of unemployment and their share rises during downturns. Fluctuations in the vacancy/unemployment rate are amplified because the value to firms from matching with those low-productivity workers is very sensitive to aggregate conditions. Moreover, fluctuations in the job-finding rate are amplified by the fact that worker-firm meetings are less likely to yield successful matches during downturns.

Keywords: Amplification; Matching Models; Worker Heterogeneity
JEL Classifications: E24, E32, J63, J64

1 Introduction

Search and matching models such as the Mortensen and Pissarides (1994) model have become a primary tool for analyzing labor market fluctuations. The frictions associated with job search have natural appeal as a source of amplification and propagation of small shocks to the economy. In an important recent paper, however, Shimer (2005) showed persuasively that a carefully calibrated version of a standard matching model[1] exhibits fluctuations in

*Dept. of Economics, 3105 Tydings Hall, University of Maryland, College Park, MD 20742; email: pries@econ.umd.edu.

[1]Specifically, Shimer's model is the Pissarides (1985) model modified to include aggregate productivity shocks. In contrast with the Mortensen and Pissarides (1994) model (and the model considered here), there is no endogenous destruction of matches.

vacancies, job-finding rates, and unemployment that are too small, relative to US data, by an order of magnitude. For example, in response to aggregate labor productivity shocks whose standard deviation and persistence are calibrated to US data, the standard deviation of the model's (log) job-finding rate is 0.01, while its empirical counterpart is 0.118. The standard deviation of the model's (log) unemployment rate is 0.009, compared with an empirical value of 0.19. Put starkly, the model fails miserably as an amplification mechanism.

As Shimer (2005) explains, the flexibility of wages, which are determined by Nash bargaining, plays an important role reducing the amplification of shocks in the standard model. When aggregate conditions improve, workers' threat point in the bargain (i.e. the value they can obtain by searching for a new job) increases roughly in proportion with the aggregate shock, which allows workers to obtain most of the additional output in the form of higher wages. Because workers absorb most of the additional output, firms have little additional incentive to open more vacancies. The very small increase in the vacancy/unemployment ratio translates into only a slightly higher job-finding rate and thus only a slightly lower unemployment rate.

Shimer (2005) also shows very clearly that the aspect of the model's calibration that is most important for the lack of amplification is the gap between a match's productivity and the worker's payoff from unemployment. For a broad range of values for the model's other parameters, the elasticity of the vacancy/unemployment ratio with respect to that gap is in the neighborhood of one. Thus, only if that gap is small, so that a small increase in labor productivity results in a large percentage increase in the gap, can empirically plausible labor productivity shocks generate vacancy/unemployment ratio movements of realistic amplitude. Thus, in Shimer's calibration, with average productivity normalized to 1.0 and income while unemployed set to 0.4, a one standard deviation increase in (log) labor productivity (i.e. 0.02) merely generates about a 3.3% increase in the vacancy/unemployment ratio. In the very standard model that Shimer considers, these results are highly robust to changes in parameter values.

However, the model presented here shows that extending the standard model to allow for worker heterogeneity and match heterogeneity can generate significant amplification. The model highlights the fact that the unemployed are not a representative sample of the population. Instead, the pool of unemployed consists disproportionately of low-productivity workers who chronically straddle the employment–unemployment margin. While small productivity shocks have a very limited impact on the 'average' worker, as in Shimer (2005), the impact on the 'marginal' workers is significant. A small adverse productivity shock can push many of the marginal low-productivity workers toward unemployment, as it becomes more difficult for those workers to find viable matches and as firms, which recognize the increased proportion of marginal workers in the pool of unemployed, reduce their vacancy creation.

While it is reasonable to assume, as Shimer does, that *on average* there is a sizeable gap between workers' productivity and their payoff from unemployment, it is also reasonable to assume that for some workers the gap indeed is small. Of course, the gap could be small due either to low productivity or to a high payoff from leisure or unemployment. For simplicity, the model presented below assumes that all heterogeneity is in productivity. More specifically, there are two types of workers: those with high average productivity and those with low average productivity. In addition, and in contrast with Shimer, there is a match-specific component of productivity that is revealed when a worker and firm meet, and which thereafter changes stochastically. As a result, not all meetings result in matches—only those whose match-specific productivity exceeds a threshold—and some matches are endogenously terminated when their match-specific productivity falls below the threshold.

This heterogeneity increases the responsiveness of vacancy creation to small labor productivity shocks in two ways. The first relates to the average composition of the unemployed. Note that low-productivity workers, whose productivity on average only slightly exceeds their unemployment income (or leisure value), are less likely than high-productivity workers to receive a draw of match-specific productivity that is high enough for their match to be viable. As a result, their job finding rates are lower and their separation rates are higher. It follows immediately that their unemployment rate is higher and their share of overall unemployment exceeds their share in the population. Because low-productivity workers constitute a disproportionate share of unemployment, and because the value to a firm of meeting with a low productivity worker is more sensitive to aggregate productivity shocks (compared with high-productivity workers), the value from posting a vacancy is also more sensitive to aggregate shocks than in a model without heterogeneity.

The second way in which heterogeneity increases amplification relates to the cyclical changes in the composition of the unemployed. An adverse labor productivity shock increases the share of low-productivity workers among the unemployed. Because the value to a firm from meeting with a low-productivity worker is considerably lower than the value from meeting with a high-productivity worker, this change in the composition of the unemployed reduces the profitability from posting vacancies during downturns.

In addition to those two source of amplification for vacancy creation, heterogeneity contributes to overall amplification in two other ways. First, Because the composition of the employed workers changes with aggregate conditions, measured average labor productivity will be less pro-cyclical than the actual aggregate component of matches' productivity. During booms, employment of low-productivity workers rises and dampens the increase in measured average labor productivity. These composition effects are important because they suggest that the standard deviation of the aggregate component of a match's productivity is in fact greater than the standard deviation of measured average labor productivity. In other words, the misleadingly small standard deviation of measured labor productivity overstates the amount of amplification that is needed.

Second, the endogeneity of match creation and destruction decisions contributes to the overall amplification of shocks. Because workers and firms do not accept all matches, but rather only those with a match-specific component of productivity that surpasses some optimal threshold, matching rates (rates of match formation) differ from meeting rates. During downturns, when meeting rates decline due to a lower vacancy/unemployment rate, the matching rate declines by even more since the threshold for the match-specific component of productivity increases, so that fewer meetings lead to successful matches. Thus, while meeting rates exhibit greater cyclical fluctuations due to the bigger fluctuations in vacancy creation discussed above, the procyclical success rates of worker-firm meetings means that matching rates are even more volatile than meeting rates.

The next section briefly discusses the related literature. Section 3 describes the model's assumptions. Section 4 defines a recursive equilibrium of the model and explains how the curse of dimensionality renders standard numerical solution techniques implausible. An alternative equilibrium concept—an 'approximate equilibrium' with a more parsimonious state space—is introduced and the numerical procedure for computing that equilibrium is explained. Section 5 calibrates the model and quantitatively assesses the model's amplification by way of comparison with Shimer (2005) and with the data.

2 Related Literature

The 'flows approach' to analyzing labor markets—that is, understanding the level and fluctuations of labor market aggregates like employment, unemployment, and non-participation by analyzing the factors affecting flows between labor market states—has dominated labor market research during the last two decades. Empirically, data on job creation and job destruction rates—in particular, the data of Davis, Haltiwanger and Schuh (1996)—showed that during recessions increases in job destruction are as important as decreases in job creation for understanding the cyclical behavior of the unemployment rate. Supplementing the data on job flows with evidence on worker flows and vacancy rates, Blanchard and Diamond (1989 and 1990) offered a comprehensive view of the various contributors to unemployment rate fluctuations. While confirming the importance of countercyclical fluctuations in separations, they also showed that vacancy rates and job-finding rates are procyclical.

Search and matching models such as Pissarides (1985) and Mortensen and Pissarides (1994) form the core of the theoretical dimension of the 'flows approach.' Cole and Rogerson (1999) rigorously examined whether a carefully calibrated version of the Mortensen and Pissarides (1994) model could quantitatively match the volatilities, auto-correlations, and cross-correlations of job creation, job destruction, and employment. They found that to account for the business cycle facts the model required that average unemployment spell durations be longer (equivalently, average job-finding rates lower) than observed in the data. Moreover, they work with a clever reduced form of the model, which has the by-product of sidestepping the need to calibrate the underlying shock process that drives fluctuations. In other words, they do not have to grapple with the central issue explored here and in Shimer (2005).

Pries (2004) explores whether the frictions emphasized by matching models are capable of serving as a significant propagation mechanism that translates short-lived shocks into long-lived consequences. More specifically, it examines whether frictions can account for the substantial persistence in fluctuations in the unemployment rate. High average job-finding rates suggest limited propagation of shocks since unemployment should adjust quickly to a short-lived burst of job loss as newly unemployed workers rapidly find new jobs. However, Pries (2004) shows that if workers displaced during the initial burst of job loss typically experience several short-lived jobs before finding longer lasting jobs, the elevated separation rate can account for the unemployment rate's persistent response.

As the above discussion indicates, Shimer (2005) focuses primarily on the amplification of shocks, rather than on the propagation of shocks. The persistence of unemployment fluctuations in his model derives almost entirely from the persistence in the labor productivity process assumed in his calibration. In terms of the amplification of shocks, as mentioned above, the key question—to which Shimer (2005) provides a negative answer—is whether a standard matching model can deliver the observed amplitude of fluctuations in vacancy/unemployment ratios and job-finding rates—and thus, unemployment—given an exogenous labor productivity process with fluctuations of realistic amplitude. Shimer (2005) identifies the flexibility with which wages are renegotiated, due to the Nash Bargaining assumption of standard matching models, as the main reason for the models' inability to generate realistic fluctuations.

In a paper highly complementary to Shimer (2005), Hall (2005) shows that with an alternative assumption about wage determination, matching models can indeed deliver realistic fluctuations in job-finding rates and in unemployment. More specifically, in his model wages are inflexible in that they only adjust if the current wage falls outside the bargaining set defined by workers and firms's outside options. This type of inflexibility is immune to the criticism often made of other sticky-wage models; namely, that the worker and firm fail to exploit opportunities for mutual gain. Nevertheless, while decisions by individual workers and firms are efficient in that narrow sense, this type of wage inflexibility generates aggregate inefficiencies precisely because it distorts the incentives to open vacancies at different points in the business cycle.

Gertler and Trigari (2005) consider another form of wage rigidity. Specifically, they consider a matching model with time-dependent, staggered wage-setting of the sort that is widespread in New Keynesian macro models. This kind of wage rigidity induces a spillover that enhances the impact of aggregate shocks on labor market outcomes.

Hagedorn and Manovskii (2005), like the present paper, examine whether matching models can be reconciled with the key facts about labor market fluctuations without resorting to wage rigidities. Their model is exactly Shimer's. As noted above, Shimer observed that vacancy creation and job-finding rates are indeed sensitive to aggregate conditions if the gap between matches' labor productivity and the income that can be earned while unemployed is small. Whereas Shimer rejected this as empirically implausible, Hagedorn and Manovskii (2005) argue for its plausibility. Costain and Reiter (2005) point out, however, that such a small gap implies that the model's equilibrium is implausibly sensitive to labor market policies. Mortensen and Éva Nagypál (2006) and Éva Nagypál (2005) review and extend this literature, emphasizing that unemployment fluctuations can be significant, even without a small gap, when hiring costs are taken into account.

The current paper takes the position that while the gap between productivity and workers' opportunity cost is unlikely to be small on average, it is likely to be small for some portion of the population. Moreover, by accounting for heterogeneity, we can account for the composition effects emphasized above, which are absent in Hagedorn and Manovskii (2005). Another important distinction between the present paper and both Shimer (2005) and Hagedorn and Manovskii (2005) is that the present model allows for endogenous match creation and destruction decisions, which means that matching rates differ from meeting rates and in fact are more volatile.

3 The Model

There are two types of risk-neutral agents: firms and workers. Both types discount future values at the same rate; β denotes the discount factor. There is a continuum of ex ante homogeneous potential firms. These firms can costlessly create vacancies but they incur a recruiting cost c each period that they search to fill them.

There is a continuum of workers of measure one. They are heterogeneous. For tractability, we assume that there are only two types of workers: a fraction ν_l of the population has low average productivity and the remainder, $1-\nu_l$, has high average productivity. When a high-productivity worker meets a vacancy, the productivity of the match is given by $y^a + y_h + \epsilon$, where y^a is an aggregate component of productivity, y_h is the type-specific component of productivity, and ϵ is the match-specific component of productivity. For low-productivity workers, match productivity is given by $y^a + y_l + \epsilon$, with $y_l < y_h$. For both types of workers, the initial match-specific component is drawn from a mean-zero distribution on $[-\infty, \infty]$ with cdf $G(\epsilon)$. The aggregate component of productivity is assumed to follow a discrete state Markov process. The vector of values is given by $\bar{y}^a$ and the elements of the transition matrix, Π, are given by $\pi_{ij} = prob\{y^{a\prime} = \bar{y}^a_j | y^a = \bar{y}^a_i\}$.

When searching, workers receive a flow utility of z. It is worth noting that the key distinction between the two types of workers is the differing values of $y_i - z$. Although for simplicity we have assumed that all of the heterogeneity stems from differences in average output, in reality workers are also heterogeneous in z.

The process by which workers and firms meet each other is described by a constant-returns-to-scale function of the number (measure) of vacancies and the total number of unemployed workers of either type (they are assumed to be perfect substitutes in the search process). Specifically, the number of meetings is determined by the function $m(v, u)$, where v is the number of vacancies and u is the unemployment rate. It follows from the constant-returns-to-scale assumption that the rate at which workers meet firms depends only on the ratio of vacancies to unemployed workers, $\theta = \frac{v}{u}$. We let $f(\theta)$ denote the workers' meeting rate and $q(\theta)$ the firms' meeting rate.

When a worker and firm meet, they observe their match-specific productivity ϵ and then decide whether to establish a match. If ϵ is too low, they will choose not to establish the match. Accordingly, matching rates differ from meeting rates. After a worker and firm establish a match, in subsequent periods they receive, with probability λ, a new draw of ϵ from $G(\epsilon)$. If the new ϵ is too low, then the firm and worker will agree to separate. Matches are also destroyed exogenously with probability s.

Wages are determined by generalized Nash Bargaining. That is, each period the worker and firm agree on a wage that gives to the worker a fraction ϕ of the match surplus that arises from the model's search frictions.

The timing of events within a period is fairly straightforward. At the beginning of the period, y^a is revealed. Next, matched workers and firms choose whether to produce that period. Workers and firms in matches with sufficiently low ϵ will choose to separate. Those workers enter the pool of unemployed and immediately begin to search for a new employer. Among the workers and firms who remain matched, the wage is negotiated and paid, and then production takes place. After production, exogenous separations occur and with probability λ matches draw a new ϵ. Matches with a new draw of ϵ do not decide whether to remain matched until they observe y^a at the beginning of the subsequent period. Likewise, among the unemployed workers who meet a vacancy and receive a draw of ϵ, the decision to form a match is only made at the beginning of the next period, after observing y^a.

4 Equilibrium with Aggregate Fluctuations

The combination of heterogeneous workers and aggregate fluctuations substantially complicates the state space that is relevant to workers and firms' decisions. In the standard model without worker heterogeneity, the free-entry condition assures that the vacancy-unemployment ratio correlates perfectly with the aggregate component of productivity. Thus, the aggregate component of productivity is the only state variable. In the model with heterogeneity, this is not the case. To understand why, note that firms considering whether to open a vacancy would prefer to match with a high-productivity worker and thus they care about the fraction of unemployed who are low-type workers, denoted by μ. As that fraction changes, in response to a change in aggregate conditions, the equilibrium vacancy-unemployment also changes (even if the aggregate component of labor productivity does not change further). Thus, workers and firms must form expectations over values of μ and to do so they must know how currently employed workers are distributed over the different values of ϵ. If there are a lot of low-type workers near the threshold $\bar{\epsilon}_l$, a decline in aggregate productivity will lead to a larger increase in μ than if there are few low-type workers near the threshold.

It follows that the state space includes the entire distribution of employed workers over values of ϵ and consequently the curse of dimensionality renders the computational problem intractable. Nevertheless, it is possible to compute an 'approximate equilibrium,' as in Krusell and Smith (1998), by tracking a more parsimonious state space that still captures nearly all of the information that is relevant to workers and firms. This section proceeds by first outlining the true recursive equilibrium in which the state space includes the distribution of workers. It then defines an 'approximate equilibrium' and describes how it is computed.

The match-specific state variable is ϵ. Aggregate state variables are $X = \{y^a, \psi_l(\epsilon), \psi_h(\epsilon)\}$, where $\psi_l(\epsilon)$ and $\psi_h(\epsilon)$ are the measures of low- and high-productivity workers in matches with match-specific productivity ϵ at the beginning of the period (after the exogenous separations and new draws of ϵ that occur at the end of the previous period, but before the current y^a is observed).

The value functions for unemployed workers and for workers with a match opportunity with productivity ϵ are given (for $i = l, h$) by

$$U_i(X) = z + \beta\big(f(\theta(X))\mathbb{E}_{\tilde{X}|X}\int_{-\infty}^{\infty} E_i(z, \tilde{X})dG(z) + \big(1 - f(\theta(X))\big)\mathbb{E}_{\tilde{X}|X}U_i(\tilde{X})\big) \quad (1)$$

and

$$E_i(\epsilon, X) = \max\Big[U_i(X), w_i(\epsilon, X) + \beta(1-s)\mathbb{E}_{\tilde{X}|X}\Big((1-\lambda)E_i(\epsilon, \tilde{X}) + \lambda\int_{-\infty}^{\infty} E_i(z, \tilde{X})dG(z)\Big) + \beta s\mathbb{E}_{\tilde{X}|X}U_i(\tilde{X})\Big] \quad (2)$$

where tildes indicate values in the subsequent period. The expectation operator $\mathbb{E}_{\tilde{X}|X}$ depends on the transition matrix Π and on the laws of motion for the measures ψ_l and ψ_h. These laws of motion, denoted by $\tilde{\psi}_i = H_i(X)$, are explicitly defined below. The equilibrium vacancy/unemployment ratio $X(\theta)$ is determined by a free-entry condition, also discussed below.

The interpretation of (1) and (2) is straightforward. For example, in (1), the payoff in the current period is z. With probability $f(\theta(X))$ search in the current period results in a match opportunity in the following period. The value of that match opportunity is the expected value of $E_i(\epsilon, \tilde{X})$. As (2) shows, the value of that opportunity is $U_i(\tilde{X})$ when it is rejected and the worker remains unemployed. With probability $1 - f(\theta(X))$, there is no match opportunity and the continuation value is $\mathbb{E}_{\tilde{X}|X} U_i(\tilde{X})$.

The value to a firm of opening a vacancy is given by

$$\begin{aligned} V(X) = -c + \beta \mathbb{E}_{\tilde{X}|X} \Big[q(\theta(X)) \Big(\mu(X) \int_{-\infty}^{\infty} J_l(z, \tilde{X}) dG(z) + (1 - \mu(X)) \int_{-\infty}^{\infty} J_h(z, \tilde{X}) dG(z) \Big) \\ + \big(1 - q(\theta(X))\big) V(\tilde{X}) \Big] \qquad (3) \end{aligned}$$

where $\mu(X)$ is the faction of the unemployed who are low-productivity workers (determined endogenously, as discussed below). Because firms can freely open vacancies, in equilibrium $V(X)$ must always equal zero. $J_i(\epsilon, X)$ is the firm's value from a match opportunity with a worker of type i and match-specific productivity ϵ:

$$\begin{aligned} J_i(\epsilon, X) = \max \Big[V(X),\ y^a + y_i + \epsilon - w_i(\epsilon, X) \\ + \beta(1 - s) \mathbb{E}_{\tilde{X}|X} \Big((1 - \lambda) J_i(\epsilon, \tilde{X}) + \lambda \int_{-\infty}^{\infty} J_i(z, \tilde{X}) dG(z) \Big) \Big] \qquad (4) \end{aligned}$$

We denote the surplus of a match with a worker of type i with match-specific productivity ϵ as

$$S_i(\epsilon, X) = J_i(\epsilon, X) + E_i(\epsilon, X) - U_i(X) - V(X) \qquad (5)$$

As explained above, the generalized Nash bargain is such that a matched worker receives the fraction ϕ of the surplus: $E_i(\epsilon, X) - U_i(X) = \phi S_i(\epsilon, X)$ and $J_i(\epsilon, X) - V(X) = (1 - \phi) S_i(\epsilon, X)$. From the above five equations, together with the free-entry condition $V(X) = 0$, it follows that the surplus of a match with a worker of type i in state $\{\epsilon, X\}$ can be expressed as

$$\begin{aligned} S_i(\epsilon, X) = \max \Big[0, y^a + y_i + \epsilon - z + \beta(1 - s) \mathbb{E}_{\tilde{X}|X} \Big((1 - \lambda) S_i(\epsilon, \tilde{X}) + \lambda \int_{-\infty}^{\infty} S_i(z, \tilde{X}) dG(z) \Big) \\ - \beta f(\theta(X)) \phi \mathbb{E}_{\tilde{X}|X} \int_{-\infty}^{\infty} S_i(z, \tilde{X}) dG(z) \Big] \qquad (6) \end{aligned}$$

For a given aggregate state X and vacancy-unemployment ratio $\theta(X)$, the second term of the max operator is monotonically increasing in ϵ, which guarantees a unique threshold value $\bar{\epsilon}_i(X)$ that satisfies $S_i(z, X) = 0$ for all $z \leq \bar{\epsilon}_i(X)$. Moreover, it is clear that $y_l < y_h$ guarantees that $\bar{\epsilon}_l(X) > \bar{\epsilon}_h(X)$. It follows that the endogenous separation rate, $\lambda G(\bar{\epsilon}_i(X))$, is higher for low-productivity workers, while the job-finding rate $f(\theta(X))\big(1 - G(\bar{\epsilon}_i(X))\big)$ is lower.

An recursive equilibrium of the model can now be defined as follows:

Definition 1. A *recursive equilibrium* is a pair (for $i = \{l, h\}$) of surplus functions $S_i(\epsilon, X)$, a pair of threshold rules $\bar{\epsilon}_i(X)$, a pair of laws of motion $H_i(X)$, and a function $\theta(X)$ relating the vacancy-unemployment to the aggregate state, such that:

1. Given $\theta(X)$, $S_l(\epsilon, X)$ and $S_h(\epsilon, X)$ solve the functional equations in (6).
2. Given the surplus functions $S_l(\epsilon, X)$ and $S_h(\epsilon, X)$, the threshold rules satisfy $\bar{\epsilon}_l(X) = \sup\{\epsilon \in \mathbb{R} : S_l(\epsilon, X) = 0\}$ and $\bar{\epsilon}_h(X) = \sup\{\epsilon \in \mathbb{R} : S_h(\epsilon, X) = 0\}$.
3. The laws of motion $\tilde{\psi}_i = H_i(X)$ are given by
$$\tilde{\psi}_i(\epsilon) = \psi_i(\epsilon)(1-\lambda)(1-s)I(\epsilon \geq \bar{\epsilon}_i) + \big(u_i f(\theta(X)) + \lambda(1-s)(1-u_i)\big)G'(\epsilon)$$
where $u_i = (\nu_i - \int_{\bar{\epsilon}_i}^{\infty} \psi_i(z)dz)/\nu_i$ and the indicator function $I()$ indicates the matches that actually produce once y^a is revealed.
4. $\theta(X)$ satisfies the free-entry condition $\frac{c}{q(\theta(X))} = \beta(1-\phi)\mathbb{E}_{\tilde{X}|X}\Big(\mu(X)\int_{-\infty}^{\infty} S_l(z, \tilde{X})dG(z) + (1-\mu(X))\int_{-\infty}^{\infty} S_h(z, \tilde{X})dG(z)\Big)$, where $\mu(X) = \frac{\nu_l u_l}{\nu_l u_l + (1-\nu_l)u_h}$ gives low-productivity workers' share of unemployment.

The first two conditions are straightforward. In the laws of motion in the third condition, the first term represents the continuing matches that retain their value of ϵ. These workers only remain matched, and survive into the next period, if $\epsilon \geq \bar{\epsilon}_i(X)$. The second term represents newly matched workers and workers who draw a new value of ϵ. Whether those workers will ultimately establish the match in the subsequent period will depend on the threshold value in that period.

An important feature of the equilibrium is that the measures ψ_l and ψ_h only directly affect the equilibrium objects through their impact on $\theta(X)$. More specifically, ψ_l and ψ_h are relevant to workers and firms solely because they impact the share of the unemployed who are low-productivity workers, $\mu(X)$, which affects the vacancy/unemployment ratio $\theta(X)$ that satisfies the free-entry condition (condition 4 of the equilibrium).

This feature is significant because it suggests a means of bypassing the curse of dimensionality introduced by the presence of those two measures in the state space. That is, if some set of information smaller than ψ_l and ψ_h can allow workers and firms to anticipate μ with little error, then that more limited information can serve as the state space of an 'approximate equilibrium' that is very close to the true equilibrium defined above.[2]

This is the approach that I pursue. More specifically, the more parsimonious state space for the 'approximate equilibrium' is assumed to be $\chi = \{y^a, y^a_{-1}, \mu_{-1}\}$. Moreover, workers and firms assume a law of motion $\mu = \hat{H}(\chi)$ that takes the form

$$\mu = \alpha_0 + \alpha_1 y^a + \alpha_2 y^a_{-1} + \alpha_3(y^a - y^a_{-1})I(y^a > y^a_{-1}) + \alpha_4\mu_{-1} \tag{7}$$

To understand why χ is a reasonable approximation of the true state space X that workers and firms need to anticipate μ, consider how μ evolves. When y^a declines, $\bar{\epsilon}_l$ increases and a mass of low-productivity workers becomes unemployed. While a few high-productivity workers may also become unemployed, due to the increase in $\bar{\epsilon}_h$, more low-productivity workers are affected so that the net effect is an upward spike in μ. The law of motion in (7) takes account of the impact of changes in y^a by including its contemporaneous value and its lag. Moreover, the fourth term in the law of motion $\hat{H}$ accounts for the asymmetric impact of changes in y^a (increases in y^a do not induce a downward spike in μ, but rather a gradual decline). The final term in the law of motion, the lagged value μ_{-1}, is useful information for anticipating the current μ because following a change in y^a there is a gradual adjustment of μ toward the value that it would have in the long run, absent another change in y^a. The coefficient on the lagged value of μ relates to the speed of this adjustment.

Given this more parsimonious aggregate state space, we can define an approximate recursive equilibrium:

Definition 2. An *approximate recursive equilibrium* is a pair (for $i = \{l, h\}$) of surplus functions $S_i(\epsilon, \chi)$, a pair of threshold rules $\bar{\epsilon}_i(\chi)$, a pair of laws of motion $\hat{H}_i(\chi)$, and a function $\theta(\chi)$ relating the vacancy-unemployment to the aggregate state, such that:

1. Given $\theta(\chi)$, $S_l(\epsilon, \chi)$ and $S_h(\epsilon, chi)$ solve the functional equations (for $i = \{l, h\}$)

$$S_i(\epsilon, \chi) = \max\Big[0,\ y^a + y_i + \epsilon - z + \beta(1-s)\mathbb{E}_{\tilde{\chi}|\chi}\Big((1-\lambda)S_i(\epsilon, \tilde{\chi}) + \lambda\int_{-\infty}^{\infty} S_i(z, \tilde{\chi})dG(z)\Big) - \beta f(\theta(\chi))\phi\mathbb{E}_{\tilde{\chi}|\chi}\int_{-\infty}^{\infty} S_i(z, \tilde{\chi})dG(z)\Big]. \tag{8}$$

[2]Krusell and Smith (1998) introduced the idea of solution by 'approximate equilibrium' for a general model of uninsured idiosyncratic risk in which the state space included the entire wealth distribution.

2. Given the surplus functions $S_l(\epsilon,\chi)$ and $S_h(\epsilon,\chi)$, the threshold rules satisfy $\bar{\epsilon}_l(\chi) = \sup\{\epsilon \in \mathbb{R} : S_l(\epsilon,\chi) = 0\}$ and $\bar{\epsilon}_h(\chi) = \sup\{\epsilon \in \mathbb{R} : S_h(\epsilon,\chi) = 0\}$.
3. The law of motion $\mu = \hat{H}_i(\chi)$ is given by (7).
4. $\theta(\chi)$ satisfies $\frac{c}{q(\theta(\chi))} = \beta(1-\phi)\mathbb{E}_{\tilde{\chi}|\chi}\Big(\mu\int_{-\infty}^{\infty} S_l(z,\tilde{\chi})dG(z) + (1-\mu)\int_{-\infty}^{\infty} S_h(z,\tilde{\chi})dG(z)\Big)$.

For the expectation operator $\mathbb{E}_{\tilde{\chi}|\chi}$, note that the first element of $\tilde{\chi}$, i.e. $\tilde{y}^a$, evolves in accordance with the transition matrix Π. The second element is simply the next period's lagged value of y^a, which of course is just the current value of y^a. The third element is next period's lagged value of μ, which is the current value of μ and which is determined by the law of motion in the third condition of the above definition.

The computational algorithm used to solve for the 'approximate equilibrium' has six basic steps (appendix A provides details). First, given parameters for the law of motion in (7) and given an initial guess for $\theta(\chi)$, the surplus functions in (8) are solved by value function iteration after discretizing the continuous state variables. Second, given those surplus functions, the free-entry condition can be solved for $\theta(\chi)$. Third, $\theta(\chi)$ is updated using the result from step two, the surplus function is re-solved, and the first two steps are repeated until the surplus functions and $\theta(\chi)$ jointly converge.

Fourth, given the optimal thresholds $\bar{\epsilon}_l(\chi)$, $\bar{\epsilon}_h(\chi)$ and the vacancy-unemployment ratio $\theta(\chi)$, a random sequence of y^a can be used to simulate a sequence of distributions ψ_l and ψ_h using

$$\tilde{\psi}_i(\epsilon) = \psi_i(\epsilon)(1-\lambda)(1-s)I(\epsilon \geq \bar{\epsilon}_i) + \big(u_i f(\theta(\chi)) + \lambda(1-s)(1-u_i)\big)G'(\epsilon). \tag{9}$$

From these distributions, sequences for u_l, u_h, and μ are calculated. Fifth, the simulated sequences of y^a and μ are used to estimate the parameters in the law of motion in (7). Sixth, using these parameters to update the law of motion, the first five steps are repeated until a fixed point for the law of motion is found. If the law of motion provides a satisfactory representation of the simulated sequences, in the sense of goodness of fit, then the 'approximate equilibrium' is deemed a good approximation of the computationally intractable true equilibrium and the algorithm is finished. If not, a different form is considered for the law of motion. In practice, the law of motion in (7) provided a good fit; appendix A provides information on the goodness of fit for the simulations discussed below.

Given an equilibrium, many other variables of interest can be easily determined. The aggregate unemployment rate is given by $u = \nu_l u_l + (1-\nu_l)u_h$ and the measure of vacancies is given by $v = u\theta$. Furthermore, the wage schedules are given by:

$$w_i(\epsilon, \chi) = z + \phi(y^a + y_i + \epsilon - z) + \frac{(1-\phi)\beta f(\theta(\chi))\phi}{1-\beta(1-s)(1-\lambda)} \int_{\bar{\epsilon}_i}^{\infty} 1 - G(z)dz. \tag{10}$$

Aggregate labor productivity can be calculated as

$$\int_{\bar{\epsilon}_l}^{\infty} (y^a + y_l + z)\psi_l(z)dz + \int_{\bar{\epsilon}_h}^{\infty} (y^a + y_h + z)\psi_h(z)dz. \tag{11}$$

5 Quantitative Results

We now turn to a quantitative assessment of the significance of worker heterogeneity for the amplification of aggregate productivity shocks. The calibration of the model's parameters follows Shimer (2005) as closely as possible, in order to allow direct comparison.

Table 1 provides parameter values used in the quantitative results discussed below. The parameters can be separated into two types. For the first type (β, ϕ, η, and z), values are chosen directly to match the values chosen by Shimer. The values of these parameters are held constant across the various quantitative exercises presented below. For the second set of parameters, the values are chosen so as to match target values for moments from simulations of the model. For some of the moments, there is a single targeted value, which again is chosen to be consistent with Shimer. However, for two moments, both related to the worker heterogeneity in the model, Shimer's calibration can offer no guide. Moreover, because the unavailability of reliable evidence on those two moments makes it difficult to stipulate a specific parameter value, a set of plausible values is considered. For each combination of values for those two moments, the entire second set of parameters must be recalibrated in order to match all of the moments. For the parameters that take on different values in the different simulations, table 1 lists the range of values considered, while the table in appendix B details the exact values used in each of the various parameterizations

Parameter	Description	Value
β	Monthly discount factor	0.996
ϕ	Worker's bargaining share	0.72
γ	Meeting function coef.	0.493 – 0.690
η	Meeting function elasticity	0.28
c	Recruiting cost	0.136 – 0.193
ν_l	Proportion low-type workers	0.2
z	Value of leisure	0.4
y_l	Mean productivity for low-type	0.472 – .699
y_h	Mean productivity for high-type	1.08 – 1.10
s	Exogenous job destruction rate	0.019 – 0.029
λ	Probability of new ϵ draw	0.08
$\sigma(\epsilon)$	Standard dev. of ϵ	0.058 – 0.100
$\sigma(y^a)$	Standard dev. of y^a	0.034 – 0.037
$\rho(y^a)$	Autocorr. of y^a	0.974

Table 1: Parameter values. For the parameters whose values change across the six calibrations, the table gives the range of values. Appendix B details the parameter values for each calibration.

presented here.

I normalize the length of each time period to be one month. The discount factor is thus set to $\beta = 0.996$, which corresponds to an annual discount rate of nearly 5%. As in Shimer (2005), workers' bargaining share is set to $\phi = 0.72$.

The markov process for y^a is assumed to have 20 states. The vector of values and the transition matrix are set so that y^a approximates an AR(1) process with a zero mean. The autocorrelation and the standard deviation of the process are chosen so that the model's measured average labor productivity (aggregated to a quarterly frequency), which differs from y^a due to composition effects, matches the autocorrelation (0.88) and standard deviation (0.02) of measured labor productivity in US quarterly data (taken from Shimer (2005)).

The meeting function has a Cobb-Douglas specification: $m(v, u) = \gamma v^{\eta} u^{1-\eta}$. In Shimer (2005), all meetings result in a match so that there is no distinction between meeting rates and matching rates. Thus, Shimer correctly calibrates his matching function by using the available evidence on the relationship between job-finding rates and vacancy/unemployment ratios. In the present model, the technology requiring calibration is the meeting function, which cannot be estimated given the absence of data on the rate at which workers and firms meet. As a starting point, we therefore adopt the same elasticity as Shimer and set $\eta = 0.28$. To the extent that the success rates of meetings—i.e. $1 - G(\bar{\epsilon}_l)$ and $1 - G(\bar{\epsilon}_h)$—rise as aggregate conditions (y^a) improve, the observed elasticity of the matching rate with respect to the vacancy-unemployment rate will be greater than the elasticity of the meeting technology. We will return to this issue below.

The parameters γ and c are the key determinants of the (average) equilibrium values of θ and of the worker's matching rate. Thus, the values for γ and c are guided by targets for those two statistics. As in Shimer, I target an average value for θ of 1.0 and an average monthly matching rate for workers of 0.45.

Because of the assumed coarseness of the heterogeneity, with just two worker types, there is no straightforward empirical counterpart that allows us to directly calibrate the share of low-productivity workers, ν_l. We set ν_l at a relatively low value, 0.2, for all the quantitative exercises considered below. The mean productivities of the two types of workers, y_l and y_h are similarly difficult to parameterize. However, given a value for y_l, the value of y_h can be chosen by normalizing average labor productivity to one. Given this normalization, the flow utility while unemployed, z, is set to 40% of average labor productivity, just as in Shimer (2005).

$G(\epsilon)$ is assumed to be a normal distribution, with mean zero and standard deviation $\sigma(\epsilon)$. Thus, the parameters that still must be tied down are the exogenous separation rate s, the productivity parameter y_l, the probability λ of a new draw of ϵ, and $\sigma(\epsilon)$. These four parameters jointly determine three important moments: the average separation rate, the average value of μ, and the standard deviation of μ. The target for the average (monthly) separation rate is 0.033, as in Shimer (2005). For the average and standard deviation of μ, direct evidence is not available, so a range of target values is considered. For the average value of μ, values of 0.35, 0.45, and 0.55 are considered. These values for the share of unemployment accounted for by low-productivity workers should be understood relative to the share of low-productivity workers in the population as a whole, $\nu_l = 0.2$. For the standard deviation of μ, two values are considered: 0.04 and 0.08. In the simulations, these are calculated as the standard deviation of the log of μ (detrended by the Hodrick-Prescott

	US data	Shimer	Model with heterogeneity					
std. dev. μ:			0.04			0.08		
ave. μ:			0.35	0.45	0.55	0.35	0.45	0.55
ave. labor prod.	0.020	0.020	0.020	0.020	0.020	0.020	0.020	0.020
			(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
job-finding rate	0.118	0.010	0.029	0.046	0.069	0.044	0.069	0.105
			(0.004)	(0.006)	(0.008)	(0.005)	(0.009)	(0.014)
$\theta = v/u$	0.382	0.035	0.056	0.074	0.086	0.081	0.100	0.121
			(0.008)	(0.013)	(0.013)	(0.009)	(0.013)	(0.013)
separation rate	0.075	0	0.028	0.036	0.044	0.034	0.043	0.054
			(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
u	0.190	0.009	0.046	0.065	0.086	0.063	0.089	0.123
			(0.005)	(0.007)	(0.009)	(0.007)	(0.010)	(0.014)

Table 2: Standard deviations of key statistics for the US data, Shimer (2005), and the present model. The numbers in parenthesis indicate the bootsrapped standard errors for the model's simulated statistics.

filter), and should thus be interpreted as percentage deviations.

While there are only three moments that guide the four parameters, in practice there is a normalization that allows us to fix the value for λ. That is, when different values of λ in a reasonable range are considered, with the remaining three parameters recalibrated to match the three target statistics, the results do not change in any significant way. We set $\lambda = 0.08$ in all of the parameterizations discussed below.

Table 2 presents a three-way comparison of the standard deviations of the main statistics of interest. The first column shows the standard deviations in US data (these statistics are taken from table 1 in Shimer (2005)). The second column shows the same statistics for data simulated from Shimer's model (taken from his table 3, which gives results for the simulations his model with only labor productivity shocks and no exogenous separation rate shocks). Statistics for simulated data from the present model are given in the third column. To create those statistics, 1,000 samples of 736 observations were generated. The first 100 observations of each sample were discarded to eliminate sensitivity to initial conditions. For each sample, the remaining 636 observations were aggregated up to the quarterly averages, yielding 212 quarters of 'data' per sample. As in Shimer, the natural log of the quarterly data were detrended using a Hodrick-Prescott filter with smoothing parameter 10^5. The table presents the mean (and, in parenthesis, the standard error) of the samples' standard deviations.

As the first row of table 2 indicates, the standard deviation of labor productivity in both models matches the standard deviation of measured aggregate labor productivity in the US (0.02). This equalization of the driving force behind the models' fluctuations forms the basis for comparisons of other variables. It is worth noting, however, that in fact the standard deviation of shocks to y^a in the model with heterogeneity is somewhat larger (ranging between 0.034 and 0.037 in the different calibrations) than the standard deviation of measured labor productivity. To understand why, note that following a positive shock the measured aggregate labor productivity increases by less than the increase in y^a due to the changing composition of workers (low-skilled workers constitute a greater fraction of employment in good times). The significance of these composition effects is consistent with the similar role that they play in explaining the cyclicality of wages, as highlighted for example by Barsky, Parker and Solon (1994).

The dramatic lack of amplification in the standard model with homogeneous workers and matches, highlighted by Shimer, is apparent from a comparison of the standard deviation of his model's job-finding rate with its empirical counterpart. Despite shocks of similar magnitude, the Shimer model has a job-finding rate whose standard deviation is not even 10% of the empirical value. The same is true of the vacancy/unemployment ratio. For the unemployment rate, the discrepancy is even larger.

The results for the current model's job-finding rate show substantially more amplification than in Shimer's model without heterogeneity. Even for the first calibration, with low variation in the composition of the unemployed ($\sigma(\mu) = 0.04$) and a relatively small average share of low-productivity workers among the unemployed ($\bar{\mu} = 0.35$), the model generates nearly three times as much amplification as the Shimer model. For the calibration in the last column, with higher variation in composition of the unemployed ($\sigma(\mu) = .08$), and a higher average share of low-productivity workers among the unemployed ($\bar{\mu} = 0.55$), the model can explain nearly all of the variation in the job-finding rate.

The third row of table 2 shows that while heterogeneity in workers and in match quality does generate considerably more volatile fluctuations in the vacancy/unemployment ratio (relative to Shimer (2005)), the amplification is not as great as for the job-finding rate.

For example, the calibration in the final column nearly matches the standard deviation of the job-finding rate, but can account for slightly less than a third of the variation in θ. This, of course, suggests that a substantial fraction of the amplification in the model with heterogeneity can be attributed to the pro-cyclical success rate of meetings between firms and workers. With an elasticity of 0.28 on vacancies in the matching function, the volatility of the meeting rate should be about 28% of the volatility of θ. The volatility of the job-finding rate, however, is never less that 50% of the volatility of θ, due to the fact that worker-firm meetings are more likely to be successful in good times, when θ is high, than in bad times, when it is low. This effect is strongest for the calibrations with the highest average μ, since those calibrations have more low-productivity workers near the destruction margin, and thus greater movements in the success rate of matches.

Of course, these results imply that in the calibrated model's simulations the elasticity of the job-finding rate with respect to θ is much greater than the elasticity of 0.28 estimated by Shimer. This need not be troubling, however, because there are at least two reasons to be skeptical of estimates of this elasticity. First, for the primary measure of vacancies, which is based on help-wanted listings, measured vacancies will be more pro-cyclical than actual vacancies to the extent that firms are more likely to advertise vacancies during booms when labor markets are tighter and workers are harder to hire, as seems reasonable. Thus, the elasticity of the job-finding rate with respect to the correctly measured v/u ratio would be higher than the one estimated using the help-wanted measure. Second, in reality the total number of searching workers includes workers who search on-the-job. To the extent that workers are more likely to search on-the-job during booms, when the prospects of meeting a new employer are greater, then the ratio of vacancies to searchers is much less pro-cyclical than v/u. Consequently, the elasticity of the job-finding rate with respect to the vacancies/searchers ratio would be higher than the elasticity of the job-finding rate with respect to the v/u ratio that Shimer estimated.

Nevertheless, in simulations not reported here (but available upon request) I consider a smaller elasticity of the meeting function—specifically $\eta = 0.10$—to try to reduce the elasticity of the job-finding rate with respect to θ in the model's simulations. However, even with $\eta = 0.10$, that elasticity still remains considerably higher than Shimer's estimate (the standard deviation of the job-finding rate is still over 50% of the standard deviation of θ). Moreover, the alternative value for η does not affect the other results in any significant way.

The fourth row of table 2 reports standard deviations for the separation rate. While the behavior of the separation rate is not the primary focus of the paper, it is significant that the model can generate important endogenous fluctuations in the separation rate. The calibration in the final column comes closest to matching the standard deviation observed in US data, but still only accounts for about 75% of the variation. In Shimer (2005), by contrast, all separations are exogenous and there are no movements in the separation rate in response to labor productivity shocks. Moreover, whereas the exogenous stochastic changes in the separation rate considered by Shimer (2005) induce a (counterfactually) positive correlation between v and u, the endogenous movements in separation rates in the model presented here are consistent with the empirical evidence of a negatively sloped Beveridge Curve. Moreover, the average (across samples) correlation of u and v falls between -0.93 and -1.0 for the different calibrations, which is consistent with the empirical evidence on the slope of the Beveridge curve.

The final row of table 2 shows the standard deviation of the unemployment rate. As expected, given the greater variation—relative to Shimer (2005)—in the job-finding rate and separation rate, the unemployment rate also exhibits substantially more variation than in the Shimer model without heterogeneity. Nevertheless, the calibration that exhibits the greatest variation in job-finding rates and separations rates (in the far right column) still only accounts for about 65% of the variation observed in US data. The failure to account for 100% of the variation is not surprising given the fact that neither the job-finding rate nor the separation rate fully match the variation in their empirical counterparts.

In addition to evaluating the volatility of the model's simulated data, it is also worthwhile to evaluate the serial correlation of the data. Table 3 provides a comparison of the autocorrelations of the variables of interest. Recall that the autocorrelation of aggregate labor productivity in the present model and in the Shimer model are equated with the corresponding value for the US data (0.88). As the table shows, there is little variation, in terms of the autocorrelations, across the different calibrations of the model. In each of

	US data	Shimer	Model with heterogeneity					
std. dev. μ:			0.04			0.08		
ave. μ:			0.35	0.45	0.55	0.35	0.45	0.55
job-finding rate	0.908	0.878	0.865	0.863	0.867	0.858	0.860	0.860
			(0.029)	(0.030)	(0.030)	(0.031)	(0.032)	(0.032)
$\theta = v/u$	0.941	0.878	0.864	0.855	0.851	0.804	0.847	0.758
			(0.031)	(0.035)	(0.035)	(0.044)	(0.040)	(0.051)
separation rate	0.733	0	0.387	0.310	0.179	0.359	0.288	0.213
			(0.073)	(0.080)	(0.078)	(0.071)	(0.078)	(0.071)
u	0.936	0.939	0.820	0.822	0.086	0.831	0.836	0.843
			(0.034)	(0.036)	(0.033)	(0.034)	(0.035)	(0.033)

Table 3: Autocorrelations of key statistics for the US data, Shimer (2005), and the present model. The numbers in parenthesis indicate the bootstrapped standard errors for the model's simulated statistics.

the calibrations, the present model falls a bit short in explaining the persistence of several variables. Of particular note, the separation rate in the model is much less persistent than in the data. One possible explanation for this shortcoming is the absence of the kind of repeated job loss emphasized in Pries (2004). It is worth pointing out that the insufficient persistence in the separation rate is likely responsible for the slightly insufficient persistence of the other variables, such as the unemployment rate.

While tables 2 and 3 provide results for calibrations with a range of values for the average and standard deviation of μ, to gain a better assessment of the importance of heterogeneity for amplification would require better evidence on those two moments. At least two issues complicate this. First, the heterogeneity that matters for worker's employment experience is not just their productivity, but also their payoff from unemployment z. Although we have assumed away heterogeneity in the latter, dispersion in z is likely significant. Moreover, it is even more difficult to measure than individual productivity. Second, in reality neither a worker's productivity nor their payoff from unemployment are fixed, though for simplicity the model has assumed that they are. For example, the evidence on job displacement, such as Jacobson, LaLonde and Sullivan (1993), has shown that individuals who have attained a high level of job-specific human capital frequently suffer significant reductions in earnings potential following a job loss. Accordingly, workers who at one time may look more like the high-productivity workers in the model, would look more like low-productivity workers after losing their jobs. While the model presented here does not allow for this possibility, it is clear that it would serve to reinforce the model's amplification of shocks by increasing the proportion of low-productivity workers in the unemployment pool following a downturn in aggregate productivity.

While the available data may not provide a straightforward way to identify the correct calibration of the model, a rough check on the reasonableness of the the level and variation of the unemployment rate of low-productivity workers in the model can be obtained by comparing with data on different demographic groups. For the three targeted values of average μ — 0.35, 0.45, and 0.55 — the average unemployment rates, respectively, for the low-productivity workers were approximately 12%, 16%, and 20% (these numbers were relatively invariant to the value of $\sigma(\mu)$). For comparison, among male and female African-Americans, ages 20 to 24, the average unemployment rate is 20.8% [3]. Among males of all ethnicities, ages 16 to 24, the average unemployment rate is 13.1%

In terms of variation of the unemployment rate, the standard deviation of the (log) unemployment rate among low-productivity workers for the three values of average μ in the model is 0.09, 0.11, and 0.13 for $\sigma(\mu) = 0.04$ and 0.14, 0.17, and 0.20% for $\sigma(\mu) = 0.08$. For comparison, for African-Americans, ages 20 to 24, the standard deviation of the (log) unemployment rate is 0.212. For males of all ethnicities, 16-24, the standard deviation is 0.164. Taken together, these numbers suggest that the results for the calibration that produced the most amplification, with average μ equal to 0.55 and $\sigma(\mu) = 0.08$, may be reasonable.

6 Conclusion

Fluctuations in the unemployment rate are large and persistent. It is crucial for policymakers to understand both why this is so and to what extent it is efficient. There is a rather broad consensus that fluctuations in unemployment are the outcome of an economic environment with pervasive frictions that is constantly buffeted by shocks. When shocks hit the economy that require that workers be reallocated across firms, industries, and regions, the

[3]Calculated from the CPS, for January, 1972, to May, 2006

frictions prevent the adjustment from occurring immediately. There is less consensus regarding the efficiency of this process. To the extent that a Mortensen and Pissarides (1994) style matching model can capture the observed magnitude and persistence of unemployment rate fluctuations, the fluctuations need not be regarded as inefficient. In particular, if the Hosios (1990) condition is satisfied, then the equilibrium is efficient (in a constrained-efficient sense, since surely efficiency gains could be achieved by reducing frictions).

In contrast, if the magnitude and persistence of unemployment rate fluctuations cannot be reconciled with the Mortensen-Pissarides model, then the efficiency of the fluctuations is in greater doubt. In particular, if the unemployment rate fluctuations can only be understood as the result of some sort of wage rigidity, as suggested by Shimer (2005) and Hall (2005), then the inefficiency of unemployment fluctuations may be substantial. Accordingly, policies to counteract the fluctuations could have a significant impact on welfare.

This paper has shown that a standard Mortensen-Pissarides model modified to allow for worker and match heterogeneity can generate significantly greater (relative to the standard model without heterogeneity) cyclical movements of the job-finding rate and of the unemployment rate in response to the small cyclical movements in productivity that drive fluctuations in the model. The key to the model is that the unemployed are not a representative sample of the population. Instead, the unemployed consist disproportionately of workers who are regularly at the margin between employment and unemployment. In bad times, such workers are less likely to form profitable matches. Firms, recognizing this, open fewer vacancies than in a model without worker heterogeneity. The combination of fewer vacancies and diminished chances of a viable worker-firm match result in a large decline in job-finding rates—and a significant increase in unemployment—during downturns.

Appendix A: Solution Algorithm

This appendix provides details of the algorithm used to compute the 'approximate equilibrium' discussed in section 4.

The match-specific component of productivity, ϵ, is discretized into a grid $\{\epsilon_j\}_{j=1}^{700}$. The grid-points are centered on zero, with minimum and maximum values differing from zero by five times the standard deviation $\sigma(\epsilon)$. After forming the vector of probabilities representing the pdf of ϵ, the vector is rescaled to account for the tiny probability mass that lies beyond the minimum and maximum values of ϵ.

As discussed in the text, the aggregate component of productivity, y^a, is a 20-state markov process. The other aggregate state variable, μ, is assumed to take on a uniform grid of 25 values. The minimum and maximum values of the grid vary with each calibration. Because both the contemporaneous and lagged values of y^a are state variables in the approximate equilibrium, the aggregate component of the state space has dimension 20x20x25.

To select starting values for the coefficients in the law of motion for μ, we take advantage of the fact that the adjustment of μ to changes in y^a is quite rapid, so that the true values of μ from the stochastic equilibrium closely track the sequence of values given by the steady-state equilibria associated with different values of y^a. Accordingly, the steady-state version of the model is solved for each value of the aggregate shock, y^a. Then, for a sequence of 10,000 values of the markov chain y^a, a sequence of corresponding steady-state values of μ is determined from the steady-state solutions. These sequences of 'data' are then used to estimate the coefficients in the law of motion given in equation 7.

All values of the (20x20x25) matrix of starting values for $\theta(\chi)$ are set to one. Given these starting values, the two surplus functions $S_l(\chi)$ and $S_h(\chi)$ (stored as two 700x20x20x25 matrixes) are jointly solved by value function iteration. Given solutions for those value functions, the values of $\theta(\chi)$ that solve the free-entry condition (for each point in the aggregate state space) are computed. The matrix of values for $\theta(\chi)$ is updated by taking a convex combination of the values that solve the free-entry condition and the previous iterate's values. The surplus functions are re-solved and the process is repeated until the values in the $\theta(\chi)$ matrix converge (so that the free-entry condition is satisfied).

Next, the model is simulated in order to assess how the observed law of motion for μ, from the simulated data, relates to the assumed law of motion. To carry out this simulation, an initial distribution of workers (both low- and high-productivity workers) over idiosyncratic employment situations (workers are either unemployed or employed at a particular value of ϵ) is assumed. A markov chain with 10,500 values of y^a is generated. Given this markov chain, a sequence of 10,500 distributions of low- and high-productivity workers over values of ϵ is generated, utilizing from the previous steps the optimal values of $\theta(\chi)$, $\bar{\epsilon}_l$, and $\bar{\epsilon}_h$ (which determine the matching rates and separation rates). From those distributions, a sequence for μ is calculated. The first 500 values are discarded and the remaining 10,000 'data' points for μ and y^a are utilized to estimate the law of motion in equation (7).

The above steps are repeated until a fixed point in the law of motion is achieved (i.e. the law of motion assumed when solving the value functions is sufficiently close to the law of motion estimated

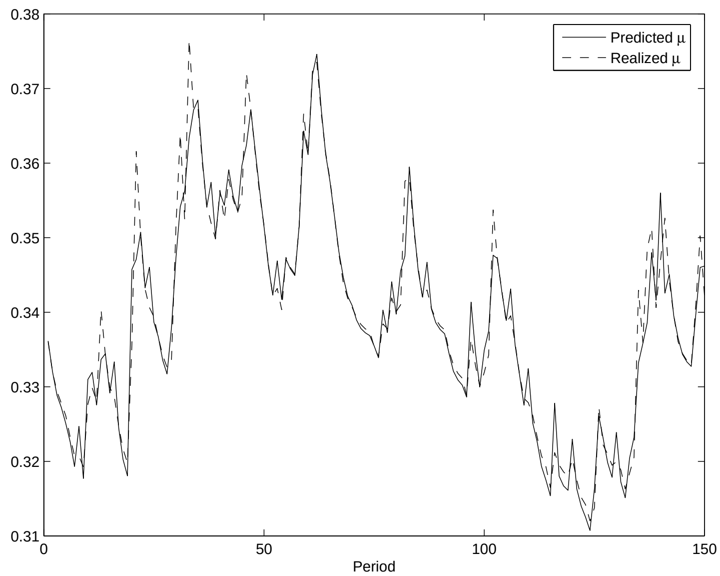

Figure 1: Comparison of the realized value of μ with the value of μ predicted by the law of motion in equation (7), for a simulated sample of 150 observations.

from the simulations). Given a fixed point, the 'goodness of fit' of the law of motion is assessed. If it is deficient, a different form for the law of motion is considered. In practice, the law of motion in equation (7) converged very quickly (which reinforces the arguments for the chosen method of selecting starting values for the law of motion parameters, discussed above). For all of the parameter sets, the standard deviation of the regression error was less than 0.006 and the R^2s were greater than 0.97. Figure 1 shows graphically the quality of the fit. It plots both the predicted and the realized values of μ for a sample with 150 observations (for the calibration with $\bar{\mu} = 0.35$ and $\sigma(\mu) = 0.04$).

Appendix B: Detailed parameter values

Parameter	$\sigma(\mu) = 0.04$			$\sigma(\mu) = 0.08$		
	$\bar{\mu} = 0.35$	$\bar{\mu} = 0.45$	$\bar{\mu} = 0.55$	$\bar{\mu} = 0.35$	$\bar{\mu} = 0.45$	$\bar{\mu} = 0.55$
γ	0.513	0.591	0.699	0.493	0.548	0.623
c	0.193	0.171	0.149	0.184	0.160	0.136
y_l	0.652	0.566	0.488	0.550	0.508	0.472
y_h	1.080	1.080	1.080	1.100	1.100	1.100
s	0.027	0.023	0.019	0.029	0.027	0.025
$\sigma(\epsilon)$	0.100	0.100	0.100	0.058	0.059	0.058
$\sigma(y^a)$	0.034	0.035	0.036	0.035	0.035	0.037
$\rho(y^a)$	0.974	0.974	0.974	0.974	0.974	0.974

Table 4: Details of the parameter values used for the six different calibrations that target two values of the standard deviation of μ and three values of the mean of μ. The table does not show the parameters whose values were held fixed across calibrations (those are shown in table 1 in the main text). In practice, several parameters here — $\sigma(y^a)$, $\rho(y^a)$, y_h — varied minimally across calibrations.

References

Barsky, Robert; Parker, Jonathan A. and Solon, Gary. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *Quarterly Journal of Economics*, *109*(1), February 1994, pp. 1–25.

Blanchard, Olivier Jean and Diamond, Peter. "The Beveridge Curve." *Brookings Papers on Economic Activity*, (1), 1989, pp. 1–76.

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

Cole, Harold L. and Rogerson, Richard. "Can the Mortensen-Pissarides Matching Model Match the Business-Cycle Facts?" *International Economic Review*, *40*(4), November 1999, pp. 933–59.

Costain, James S. and Reiter, Michael. "Business Cycles, Unemployment Insurance, and the Calibration of Matching Models.", 2005. Mimeo, Universidad Carlos III.

Davis, Steven J.; Haltiwanger, John C. and Schuh, Scott. *Job Creation and Job Destruction*. Cambridge, MA: MIT Press, 1996.

Éva Nagypál. "Labor-Market Fluctuations, On-the-Job Search, and the Acceptance Curve.", 2005. Mimeo, Northwestern University.

Gertler, Mark and Trigari, Antonella. "Unemployment Dynamics with Staggered Nash Wage Bargaining.", 2005. Mimeo, New York University.

Hagedorn, Marcus and Manovskii, Iourii. "The Cylical Behavior of Equilibrium Unemployment and Vacancies Revisited.", 2005. Mimeo, University of Pennsylvania.

Hall, Robert E. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review*, *95*(1), 2005, pp. 50–65.

Hosios, Arthur J. "On the Efficiency of Matching and Related Models of Search and Unemployment." *Review of Economic Studies*, *57*, April 1990, pp. 279–98.

Jacobson, Louis S.; LaLonde, Robert J. and Sullivan, Daniel G. "Earnings Losses of Displaced Workers." *The American Economic Review*, *83*(4), Sept 1993, pp. 685–709.

Krusell, Per and Smith, Anthony A. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, *106*(5), October 1998, pp. 867–896.

Mortensen, Dale T. and Éva Nagypál. "More on Unemployment and Vacancy Fluctuations.", 2006. Mimeo, Northwestern University.

Mortensen, Dale T. and Pissarides, Christopher A. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, *61*(3), July 1994, pp. 397–415.

Pissarides, Christopher A. "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages." *American Economic Review*, *75*(4), September 1985, pp. 676–90.

Pries, Michael J. "Persistence of Employment Fluctuations: A Model of Recurring Job Loss." *Review of Economic Studies*, *71*(1), January 2004, pp. 193–215.

Shimer, Robert. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*, *95*(1), March 2005, pp. 25–49.