Worker Flows under Mismatch

William B. Hawkins*

Yeshiva University

This version: February 1, 2013

Preliminary and incomplete

Abstract

I study a dynamic model of matching between heterogeneous workers and heterogeneous jobs. Workers and jobs are attached to different labor markets, and within each labor market, workers and jobs are assigned to different locations, representing heterogeneity in the skills of workers and the skill requirements of jobs. A match between a worker and a job in a given labor market produces flow output that depends on the quality of their match, which is a decreasing function of the distance between them. The worker and the job split this output in a fixed ratio. Matching is frictionless within each labor market. Each period, the set of matches between workers and firms is such that no worker-job pair strictly prefer to be matched with each other than with their current partners. The model is consistent with evidence from the Current Population Survey on the cross-sectional and cyclical behavior of job-to-job transitions. Notably, a worker who recently experienced a job-to-job transition is more likely to undergo another.

JEL Codes: E24, J41, J63, J64.

Key words: mismatch, unemployment, vacancies, matching, job-to-job transitions

*Addresses: Department of Economics, Yeshiva University, 245 Lexington Ave (LX711), New York, NY 10016; `whawkins@yu.edu`. I thank Jongsuk Han for excellent research assistance, and Carlos Carillo-Tudela for very helpful comments. All errors and omissions are my own. This paper is still preliminary; please do not circulate or cite it without permission.

1 Introduction

Why do unemployed workers and vacant jobs coexist in the labor market? The answer given by models of stock-flow matching (Taylor, 1995; Coles and Muthoo, 1998; Coles and Smith, 1998; Ebrahimy and Shimer, 2010) or mismatch (Shimer, 2007) is that a given worker and a given job may not be suitable to match productively, so that agents may be left unpaired once all productive matches are exhausted. The notion of a productive match in these models, however, is very stylized: conditional on being able to match, all worker-firm matches are equally good. In real labor markets, on the other hand, while many workers may be able to produce output when employed in a specific job, some workers can produce more than others. This means that a worker may be unemployed not because he is unemployable or unlucky, but rather simply because every job in the market can find another worker who is more productive in that job than the first worker is. Similarly, even though there may exist a worker who would be highly productive if matched with some particular job, the job may remain vacant or filled by a worker who is a bad match for it if the worker has an even more productive match available. In this paper I investigate the implications of such heterogeneity for worker flows between employment and unemployment and from job to job, both over time for each individual worker and in the aggregate in response to changes in aggregate productivity.

To do this, I build on an existing model of mismatch (Shimer, 2007) by allowing for rich microeconomic heterogeneity in the productivity of worker-job pairs. The most important contribution of the model relative to existing stock-flow and mismatch models is that it allows for complex patterns of flows between unemployment and employment and from job-to-job. Job-to-job transitions can be associated both with wage gains and wage cuts, the first as workers move to better matches that have become available either due to the creation of a new job or to a change in matching patterns of other workers, the second when an existing match ends, either exogenously or endogenously, so that the worker is compelled to accept a less productive alternative. Evidence from the Current Population Survey shows that a worker who recently started a new job (moving either from unemployment or from another job) is more likely to undergo a further job-to-job transition and more likely to separate to unemployment; the model is consistent with these patterns. Additionally, the model is also consistent with the procyclicality of job-to-job transitions first noted by Fallick and Fleischman (2004).

In the model, there are a large number of local labor markets, each representing the combination of a particular occupation and geographic location. A worker can match only with a job in the same local labor market; moreover, her productivity is not constant across all matches in her local labor market, but rather varies with the identity of the job. Better matches produce more output, and are more likely to be observed in equilibrium than worse matches. Matching within a local labor market is frictionless: all firms and all workers are fully aware of each others' characteristics, and the process of matching is costless and instantaneous. In Section 2.3 I characterize matching conditional on the sets of workers and firms in a specific local labor market, and give a sufficient condition for the uniqueness of the resulting matching.

Workers and jobs are hit by shocks which cause them randomly to enter and exit labor markets. The resulting changes in the numbers of workers and jobs in each local labor market, together with the associated changes in the sets of potential matches, lead to changes in the set of matched worker-job pairs. This generates flows between employment and unemployment, as well as job-to-job transitions. In Section 2.4 I characterize the patterns of worker flows in an aggregate steady-state economy. I also study the comparative statics of job-to-job flows in response to changes in aggregate productivity, and show that higher aggregate productivity is associated with higher job-to-job flows.

A feature of the benchmark model is that it generates too many job-to-job transitions, relative to transitions between employment and unemployment, relative to U.S. data. This is because the arrival or departure of a worker or the creation or destruction of a job sometimes causes a cascade of other employment flows in the relevant local labor market, and especially of job-to-job flows. For example, the departure of one worker may free up a job that is a better match for a second worker than her current job. This causes a job-to-job transition that then frees up another job that may in turn be a very good match for yet a third worker, and so on a local labor market. These vacancy chains arise because of the absence of labor market frictions in the model, so that workers may change jobs in response even to very tiny gains in match quality. In Section 5 I show that allowing for frictions within local labor markets reduces the ratio of the frequency of job-to-job transitions to that of transitions in and out of unemployment. This allows me to propose this ratio as a measure of the extent of frictions in the labor market.[1]

[1]Because the unemployment rate is determined by the EU and UE transition rates, conditional on the unem-

The concept of matching and of job-to-job transitions proposed in this paper can be contrasted with alternative ways of modeling frictional matching. In search theoretic models, both with random search (Diamond, 1982; Pissarides, 1985; Mortensen and Pissarides, 1994; Pissarides, 2000) and directed search (Moen, 1997; Shimer, 1996; Burdett et al., 2001), trading frictions impede the consummation of potentially profitable trades. To put it in overly simplistic terms, search theory posits that a worker is unemployed either because she cannot find a vacancy (even though there are vacant jobs), or because each time she does locate a vacancy, so too do other unemployed workers who turn out to be luckier than she is. Such an explanation, however, ignores the heterogeneity of workers and jobs—and in particular, the match-specific heterogeneity—that casual empiricism suggests is at least as important in many labor markets. A model of mismatch between heterogeneous workers and jobs, as in this paper, suggests that a worker is unemployed because even though she knows about all the potential employers in the economy, other workers are better suited to those jobs than she is. At least in some labor markets (including those for academic labor), the second account of unemployment seems at least as reasonable as the first.

It is also worth comparing the view of job-to-job flows in search models with that in this paper. The earliest search models did not allow for job-to-job transitions, but rather assumed that all worker flows are between employment and unemployment. Allowing for search on the job as well as by unemployed workers naturally generates job-to-job transitions as workers seek to rise on a job ladder, a story that can be told using either random search (Burdett, 1978; Burdett and Mortensen, 1998; Postel-Vinay and Robin, 2002) or directed search (Delacroix and Shi, 2006; Menzio and Shi, 2011). Because it becomes progressively more and more difficult to find better matches the further up the ladder a worker has risen, job ladder models naturally predict that workers who have recently transitioned from unemployment to employment experience job-to-job transitions more frequently than workers who have been continuously employed for a long time. However, the mismatch model in the current paper can also generate this feature: workers who have been continuously employed for a long time are likely in a high-quality match that was not easily disrupted by the random

ployment rate, this measure is a simple transformation of the ratio of the EE transition rate to the UE transition rate, also recently proposed as a measure of labor market frictions in a job ladder model by Gautier and Teulings (2011). More generally, the idea that reduced 'churn' of workers between jobs can be associated with an increase in mismatch in recessions is a familiar one (for a recent example, see Lazear and Spletzer, 2012), but the idea that the ratio of job-to-job transitions to job-to-unemployment transitions is itself a way to measure the degree of frictions in the labor market is of more recent vintage.

arrivals and departures of other workers and jobs from the local labor market, and such a match is therefore relatively unlikely to end, whether by a transition to unemployment or to another job.

The structure of the remainder of the paper is as follows. In Section 2 I introduce the model. Section 2.3 describes static equilibrium within a single local labor market, and Section 2.4 studies worker flows over time in an aggregate steady state. I calibrate the model in Section 3, and then compare the results of the benchmark model to evidence on worker flows from the Current Population Survey in Section 4. Section 5 introduces labor market frictions as described above, Section 6 discusses efficiency, and Section 7 studies the robustness of the results to the way in which heterogeneity is modeled. Finally, Section 8 concludes briefly.

2 Benchmark Model

2.1 Environment

Time is continuous. All agents in the economy are risk neutral and discount the future at rate r.

The economy contains $\hat{L}$ local labor markets. There is a constant number $\hat{M}$ of workers in the economy. At any time t, each worker is attached to a particular local labor market. Which local labor market a specific worker is attached to at any time is random, with all markets *ex ante* equally likely. A worker remains attached to her local labor market until she is affected by an exogenous quit shock, which forces her to move to another, randomly chosen, local labor market. A quit shock arrives at Poisson rate $q > 0$.

There are a large number of risk-neutral, profit-maximizing, perfectly competitive firms in the economy, which own a stock of existing jobs and can create new jobs at any time at a cost of $k > 0$. At any time t, there are $\hat{N}(t)$ jobs in the economy. When a job is created, it is allocated to a particular local labor market; which local labor market is determined randomly at the time the job is created. The job remains attached to the same local labor market as long as it continues to exist. A job terminates upon the arrival of an exogenous job destruction shock affecting that specific job. A job destruction shock arrives at Poisson rate $l > 0$.

Write $M = \hat{M}/\hat{L}$ and $N(t) = \hat{N}(t)/\hat{L}$. I consider the limiting case of a large economy with many local labor markets, while thinking of $M > 0$ as a parameter and $N(t) > 0$ as an endogenous variable. M and $N(t)$ can be interpreted respectively as the mean numbers of workers and jobs per local labor market. Because workers are allocated randomly across labor markets, as are jobs, a straightforward application of the law of large numbers ensures that the distribution of local labor markets by number of workers is Poisson. That is, the fraction of local labor markets with precisely $i \in \mathbb{N} \equiv \{0, 1, 2, \ldots\}$ workers is $\tilde{\pi}(i; M) = e^{-M} M^i / i!$. Similarly, the distribution of local labor markets according to the number of jobs each contains is also Poisson, so that the fraction of markets with precisely $j \in \mathbb{N}$ jobs is $\tilde{\pi}(j; N(t) = e^{-N(t)} N(t)^j / j!$. Workers and jobs are independently allotted to different local labor markets, so that the fraction of markets with precisely i workers and j jobs, denoted $\pi(i, j; M, N(t))$, is simply the product of $\tilde{\pi}(i; M)$ and $\tilde{\pi}(j; N(t))$, that is,

$$\pi(i, j; M, N(t)) = \frac{e^{-(M+N(t))} M^i N(t)^j}{i! j!}.$$

The structure of the economy just described, including the process for allocating workers and jobs to local labor markets, the structure of quit and job destruction shocks, and the nature of job creation are deliberately borrowed from the model of Shimer (2007); this will allow my results easily to be compared with his. It follows from the argument given in Section II.C of that paper that the independent Poisson distributions of workers and jobs across labor markets are preserved by each of these processes.

2.2 Local Labor Markets

An individual local labor market consists of a circle of unit circumference. Locations on the circle are indexed by real numbers $x \in [0, 1)$. Each worker and each job located in the local labor market are attached to a specific location on the circle. The circle is endowed with a distance function $d : [0, 1)^2 \to [0, 1)$. A match between a worker located at x and a firm located at y produces flow output $f(x, y) = [1 - d(x, y)]\, z$. Here z is the level of aggregate labor productivity, which is constant; I will consider comparative statics with respect to z in Section 3.3. I assume that $d(\cdot)$ is the Euclidean metric restricted to the circle, that is, $d(x, y)$ is the length of the shorter of the two arcs connecting x and y:[2]

$$d(x,y) = \min\{|x-y|, 1-|x-y|\}. \tag{1}$$

Unmatched workers and jobs produce zero flow output.

I assume that a worker's location x is determined at the time he first arrives in the local labor market (which occurs after he receives a quit shock in his previous labor market and randomly enters a new labor market). His location within his labor market does not change while he attached to it. Similarly, a job's location y is determined when it is first created and allocated to a particular local labor market, and does not change before the job is destroyed. The random draw of the location x for a worker moving into a new local labor market has cumulative distribution function $F_x : [0,1) \to [0,1]$. Analogously, new jobs draw their location y from a cumulative distribution function $F_y : [0,1) \to [0,1]$. In the benchmark model, I assume that $F_x(\cdot)$ and $F_y(\cdot)$ are atomless.[3]

The output $f(x,y)$ produced by a match between a worker at location x and a firm at location y is split in a fixed ratio between the two parties; the worker and the firm each receive half of this output.

2.3 Matching

At any time t, the i workers and j jobs located within an individual labor market are matched. Although some matches are more productive than others, all matches produce positive net output, and so $\min\{i,j\}$ matches will be formed. The matching process is instantaneous and frictionless. I require that the matching process in each local labor market satisfy the following condition: for every pair consisting of a worker i and a job j which are not matched to each other, either i is matched to some job j' such that the match (i,j') produces at least as much output as the match (i,j), or j is matched to some worker i' such that the match (i',j) produces strictly higher output

[2] The assumption that $d(\cdot)$ is the standard metric on the circle is without loss of generality since, as will become evident below, it is the ordinal, rather than the cardinal, properties of $d(\cdot)$ which are relevant for the way in which workers and firms match within a fixed labor market. The cardinal properties of $d(\cdot)$ *are* relevant for the profitability of each match, and therefore also for the expected flow profit of each job conditional on the aggregate number of jobs $N(t)$. They therefore are important for determining the steady-state value of $N(t)$, as well as for how $N(t)$ responds to aggregate productivity shocks. However, because I will study the effects of changes in aggregate conditions by varying $N(t)$ directly to match changes in the aggregate unemployment rate, rather than by studying the effects of aggregate productivity shocks, this effect will not be relevant in my analysis.

[3] This ensures that in all but a measure zero of local labor markets, all agents (workers and jobs) are found in distinct locations.

than the match (i, j). I call this call a matching which satisfies this condition *stable*, following the use of this term by Gale and Shapley (1962). It is straightforward to establish that in almost every local labor market, this condition alone determines uniquely the resulting match.

Proposition 1. *In almost every local labor market, there is a unique stable matching.*

The proof of Proposition 1 can be found in the Appendix. The matching is straightforward to construct according to the following 'greedy' algorithm, which specializes the deferred acceptance algorithm of Gale and Shapley (1962) to the environment studied here. (The full deferred acceptance algorithm is not required because the preferences of workers over firms and firms over workers are aligned, in the sense that they are derived from the same underlying distance function $d(\cdot)$.)[4]

1. Let $\mathcal{U}$ be the set of locations of workers not yet matched, and $\mathcal{V}$ the set of job locations. Start with $\mathcal{U}$ equal to the set of all worker locations and $\mathcal{V}$ equal to the set of all job locations.

2. While both $\mathcal{U}$ and $\mathcal{V}$ are nonempty, do:

   (a) Calculate $\arg\min\{d(x, y) \mid x \in \mathcal{U}, y \in \mathcal{V}\}$.

   (b) Match the worker at x with the job at y. Replace $\mathcal{U}$ with $\mathcal{U} \setminus \{x\}$ and $\mathcal{V} \setminus \{y\}$.

3. Any remaining unmatched workers are unemployed; any remaining unmatched jobs are vacant.

2.4 Aggregate Unemployment and Vacancies

Because all possible matches produce strictly positive flow output, the short side of the market determines the number of matches. That is, if there are i workers and j jobs in some local labor market, then there are $\min\{i, j\}$ matches formed. If $i > j$, there are $i - j$ unemployed workers and no vacant jobs, while if $j > i$, then there are $j - i$ vacant jobs and no unemployed workers.

[4] The proof of Proposition 1 given in the Appendix is direct. However, note that a second proof can be given using results from the literature on stable matching. Suppose that in a local labor market all workers have strict preference rankings over jobs, and vice versa; this occurs almost surely because there are finitely many workers and jobs. Then the set of stable matchings forms a lattice partially ordered by the preferences of either workers or firms. The Gale-Shapley deferred acceptance algorithm with workers proposing generates the worker-optimal stable matching, while that with jobs proposing generates the job-optimal stable matching. Because worker preferences over jobs and job preferences over working are derived from the same underlying distance function $d(\cdot)$, however, it is immediate that the deferred acceptance algorithm is equivalent to the algorithm given in the text, independent of whether workers or jobs propose matches. It follows immediately that there is a unique stable matching.

Aggregate employment $E(N(t))$, unemployment $U(N(t))$ and vacancies $V(N(t))$ can be calculated by summing unemployment and vacancies across labor markets:

$$E(M, N(t)) = \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \min\{i, j\} \pi(i, j; M, N(t));$$
$$U(M, N(t)) = \sum_{i=0}^{\infty} \sum_{j=0}^{i} (i - j) \pi(i, j; M, N(t));$$
$$V(M, N(t)) = \sum_{i=0}^{\infty} \sum_{j=i}^{\infty} (j - i) \pi(i, j; M, N(t)).$$

Note that $E(t)+U(t) = \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} i\pi(i, j; M, N(t)) = \sum_{i=0}^{\infty} i\tilde{\pi}(i; M) = M$, consistent with the notion that the average number of workers per labor market is M. Similarly $E(t)+V(t) = N(t)$. Also note that the functional forms of $E(\cdot)$, $U(\cdot)$, and $V(\cdot)$ are the same as those which obtain in Shimer (2007). It follows that, as in that paper, the unemployment rate $u = U(M, N(t))/M$ is strictly increasing in M and strictly decreasing in $N(t)$, while the vacancy rate $v = V(M, N(t))/N(t)$ is strictly decreasing in M and strictly increasing in $N(t)$. (Note that this implies that the steady-state Beveridge curve is downward-sloping.) For any $u \in (0, 1)$ and $v \in (0, 1)$, one can find a unique pair $(M, N(t))$ which is consistent with an unemployment rate of u and a vacancy rate of v.[5] More precisely, the mapping between $(M, N(t))$ and (u, v) is identical to that which obtains in Shimer's paper, so that the model described here generates the same Beveridge curve as in that paper. Setting $M = 244.2$ and $N = 236.3$ generates an unemployment rate of 5.4 percent and a vacancy rate of 2.3 percent, consistent with data from the Job Openings and Labor Turnover Survey for the period December 2000 through April 2006. Increasing N steadily from 233 to 243 reproduces Shimer's Figure 1; the unemployment rate falls from 6.3 percent to 3.8 percent and the vacancy rate rises from 1.8 percent to 3.4 percent.

It should be emphasized that the fact that the results of the model for the aggregate unemployment and vacancy rates are identical with those of Shimer (2007) is by design. The contribution of the current model is its ability to account for flows of workers between employment and unemployment and from employer to employer, and I next consider this.

[5] These results are proved as Propositions 3 and 4 in Shimer (2007).

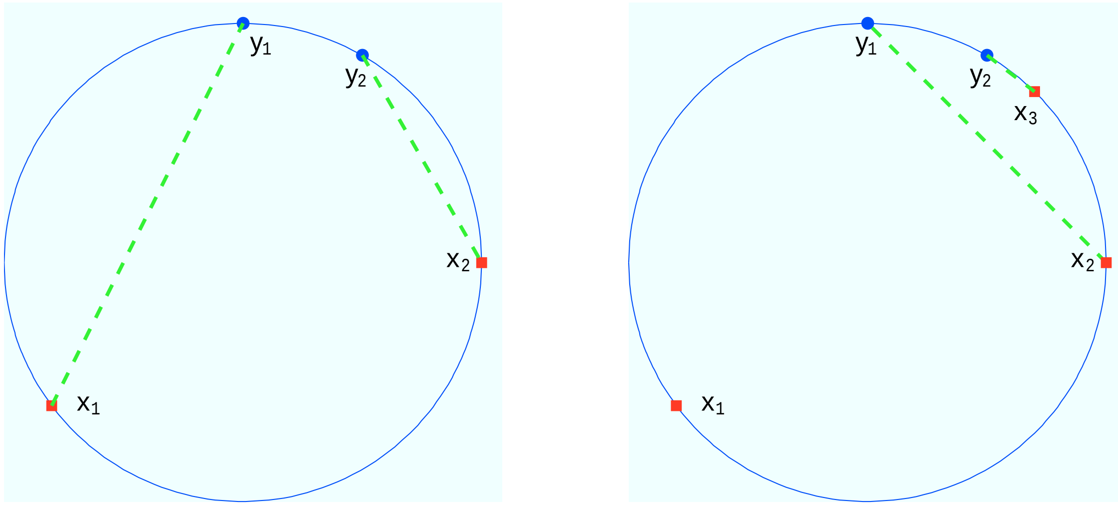

Figure 1: Left panel: Initial matching: worker 1 matched to job 1 and worker 2 matched to job 2. Right panel: worker 2 matched to job 1 and worker 3 matched to job 2.

2.5 Worker Flows

According to Proposition 1, the matching between workers and jobs in a particular local labor market depends only on the sets of locations of workers and jobs and is almost surely unique. Therefore the matching between firms and workers will not change while the set of agents in the local labor market remains unchanged. It follows that worker flows from employment to unemployment (EU transitions), from unemployment to employment (UE transitions), and from job to job (EE transitions) will occur only when either a worker arrives in or leaves the local labor market due to a quit shock, a new job is created and enters the market, or an existing job is destroyed.

Of course, when a shock occurs, some worker flows must occur by construction. When a filled job is destroyed, the worker who was previously its employee necessarily experiences either an EE or an EU transition, depending on whether he immediately finds a new job or not. When an employed worker experiences a quit shock, the affected worker again undergoes an EE or EU transition; when a quit shock affects an unemployed worker, then it is possible that the worker remains unemployed in the new local labor market, but also possible that a UE transition occurs. However, the employment status of more than the worker most directly affected by the shock can also change. This is because the addition or subtraction of an agent from the model can change the best matches available for other agents.

An example is shown graphically in Figure 1. The left panel shows a local labor market with

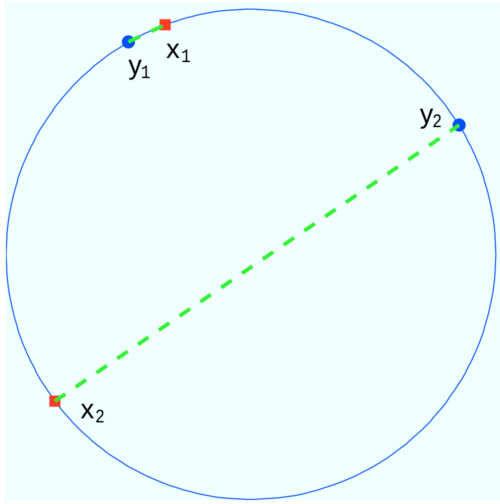

Figure 2: A strong match and a fragile match.

two workers (shown as red squares and labeled x_1 and x_2) and two jobs (shown with blue circles and labeled y_1 and y_2). The stable matching requires that the best worker-job match, namely that between worker 2 and job 2, form; this leaves worker 1 to match with job 1, even though the match quality is low. The resulting matches are shown by green dashed links. In the right panel is shown the situation after a new worker arrives in the market due to quitting another market. Worker 3 is a much better match for job 2, so the new stable matching pairs these two agents. Job 1 is now able to match with worker 2, who is a better match than worker 1 but was not available previously. Worker 1 is left unemployed. Thus, worker 1 experiences an EU transition, worker 2 an EE transition, and worker 3 experiences either a UE or an EE transition depending on her employment status in her previous labor market.

It is intuitive that workers who are badly matched are more likely to experience a transition, either to unemployment or to another job, in response to the addition and subtraction of workers and jobs from the local labor market. Figure 2 shows why this is.

In Figure 2, worker 1 and job 1, located respectively at x_1 and y_1, are well-matched. The only types of shocks that can break their employment relationship are a quit shock affecting worker 1, a job destruction shock affecting job 1, the arrival of a new worker who is closer to job 1 than worker 1 is, or the creation of a new job that is closer to worker 1 than job 1 is. Because worker 1 and job 1 are very near each other, there is only a small Poisson arrival rate of a worker who would be closer to job 1 than worker 1 is. Similarly, there is a low arrival rate of a better job match for worker 1 than job 1. Therefore, worker 1 is unlikely to experience any kind of transition (either to unemployment or job-to-job), and likewise job 1 is unlikely to experience any kind of flow (either becoming vacant or experiencing a churning flow in which the identity of its employee changes): this match is strong. In fact, even if a better worker match for job 1 arrives, worker 1 will still not undergo an EU transition, since if job 1 is no longer available to her, worker 1 can then match with job 2, with which she is a better match than is worker 2. For the arrival of other workers to cause an EU transition for worker 1, at least two workers must arrive, one who is a better match for job 1 than worker 1 is, and one who is a better match for job 2 than worker 1 is. For the arrival of other workers to cause an EE transition for worker 1 requires only that a worker arrive who is a better match for job 1 than worker 1; but as already observed, this is a low-probability event. Finally, a quit shock involving worker 2 or a job destruction shock involving job 2 will also not cause the worker 1-job 1 match to break. In summary, worker 1 and job 1 are unlikely to undergo employment flows.

Worker 2, on the other hand, is much worse-matched than is worker 1. In fact, this match is nearly the worst possible (the pair are in diametrically opposite locations), so the arrival of a new worker in the local labor market in essentially *any* location will cause an EU transition for worker 2. (The new worker might replace worker 2 in job 2 directly, or he might replace worker 1 in job 1, but in that case, as discussed in the previous paragraph, worker 1 will then transition to occupy job 2, displacing worker 2.) Likewise, the arrival of a new job will very likely cause an EE transition for worker 2, since almost any new job would be a better match for worker 2 than job 2 is.

As a result of the forces discussed above, the model has two clear predictions for the pattern of worker transitions. First, match quality in the model is closely related to the wage. A match between a worker and a nearby job generates a high flow output, and because this output is split in a fixed ratio, it therefore also generates a high flow wage for the worker (as well as a high flow profit for the firm which owns the job). Because high match-quality jobs are less vulnerable to shocks that break the match, this implies that a transition to unemployment, or to a new job, is less likely for a worker who earns higher wages. (Of course, the configuration of other workers and jobs in the labor market is also important for the probability a given job ends, in addition to the flow output of the match itself; thus wages and transition probabilities are not perfectly correlated.)

Prediction 1. *Workers who earn higher wages are less likely to experience both EU and EE transitions than workers who earn low wages.*

Another prediction for worker flows can be obtained without using wage data by using lagged worker flows as a proxy for match quality. While it is possible that a worker who has recently transitioned from unemployment to employment is in a high quality match, it is unlikely; a worker who has been in an employment spell (possibly including transitions from employer to employer), however, is much more likely to be in a high-quality match. To gain intuition for this, again consider Figure 2. In that example local labor market, worker 1, who is well matched, is unlikely to experience either an EU or EE transition relative to worker 2, who is badly matched. Thus, any statement that is more likely to be true of worker 2 than of worker 1 is predictive of future EU and EE transitions. How could the situation in Figure 2 have arisen?

One possibility is that one of the four agents is a new arrival. The recent arrival of one of the two workers does not lead to differential worker flows, but if it was either of the two jobs that recently arrived in the local labor market, then since worker 1 is a better match for both job 1 and job 2, the previous configuration must have been that worker 1 was matched to the one job and worker 2 was unemployed. Thus worker 1 may have remained employed with job 1 (if that was the incumbent job) or undergone an EE transition (if job 2 was the incumbent job); worker 2 experienced a UE transition in either case.

A second possibility is that there was previously a third job in the labor market; call it job 3. Since the worker 1-job 1 match is high quality and the worker 2-job 2 match is low quality, it is likely that job 3 was not as good a match for worker 1 as job 1, but a better match for worker 2 than job 2. In this case, worker 1 experienced no transition and worker 2 an EE transition. Finally, a third possibility is that there was previously a third worker in the market; call him worker 3. In this case, similar reasoning suggests the most likely previous matching had worker 1 matched to job 1 and worker 3 matched to job 2. This implies that worker 1 experienced no transition and worker 2 a UE transition.

In summary, in the situation of Figure 2, knowing that a worker has recently experienced a UE transition clearly predicts that that worker is worker 2, not worker 1, and so is in a low-quality match and likely to undergo both future EU and EE transitions. Knowing that a worker has recently experienced an EE transition is less diagnostic: either worker 1 or worker 2 could recently have experienced an EE transition.

Of course, the preceding discussion is about a particular local labor market only, and for the sake of intuition, I considered a labor market with only a very small number of workers and jobs. In the calibrated model, the average labor market has more than 200 workers and jobs. The only significant feature present in labor markets with so many more agents on each side and not in Figure 2 is that match quality is a much stronger selection force. In large local labor markets, jobs can intuitively be divided into good matches, which are seldom affected by EE flows, and fragile matches with lower output that are much more likely to be affected by EE transitions.

To generate a prediction about whether past worker flows predict future worker flows for the full model, requires first aggregating over different local markets conditional on i, the number of workers in the market, and j, the number of jobs; one then needs to aggregate over different (i, j) pairs.[6] Doing this analytically is intractable; however, to the extent that Figure 2 is representative, the following prediction is naturally suggested.

Prediction 2. *A worker who recently began a new job, either from unemployment or from a new job, is more likely to experience both EU and EE transitions than workers who has not recently experienced a labor market transition.*

These two predictions need to be investigated numerically in a calibrated version of the model. After a brief discussion of the structure of the model, I turn to this in Section 3 below.

2.6 Discussion

Some of the assumptions of the model require some further discussion. The model is stylized, and for the sake of comparability, builds closely on the work of Shimer (2007). (Shimer's model is very similar to the case in which all potential worker-firm matches in a particular local labor market produce the same output, that is, $d(x, y) \equiv 0$, or equivalently in which the distributions F_x and F_y are each atomic. However, allowing for a worker to view multiple possible jobs as equally good

[6] Aggregation is not straightforward: having experienced a UE transition requires having been unemployed, which requires that the worker be in a local labor market with at least as many workers as jobs; this suggests that future transitions to unemployment are also more likely, *ceteris paribus*. However, being in such a labor market and currently being employed, despite the availability of other unemployed workers, suggests the worker may be in a good match rather than a bad one, making both EU and EE transitions less likely, rather than more. The net effect of these competing effects is not obvious.

matches, or vice versa, requires tie-breaking assumptions to determine which matches will form. I turn to this in Section 5.)

The heterogeneous locations of workers and jobs within a local labor market can be interpreted as representing differences either in worker skills and job skill requirements, or alternatively in non-pecuniary job characteristics (for example, schedule flexibility, geographic location within a local labor market, or workplace culture) which are valued differently by different workers. To give a concrete example, think of the market for secretarial services in a single town as represented by a single local labor market in the model. Even within this narrowly-disaggregated local labor market, not all jobs are identical. A person with skills as a medical secretary might be able to work as a legal secretary if necessary, but his productivity would be higher if he worked in a position that made better use of his specific skills. A trained legal secretary who lives on the north side of town can commute across town to a job on the south side if no jobs are available on the north side, but the output of the match (net of the worker's cost of commuting) of the match would be higher if a position were available closer to her home.

I do not allow for workers to change their locations in response to local labor market conditions, and this restriction applies both to movement from one local labor market to another and to movement from one location to another within an individual local labor market. The presence of relatively mild mobility costs make this assumption palatable. In the numerical work in Section 3 below, no worker could improve their expected lifetime income by more than 5 percent of the average lifetime income by moving to another randomly-chosen local labor market, and by no more than an additional 0.5 percent of the average lifetime income by moving to the best possible location within the best possible labor market. Thus, if mobility costs exceed 5.5 percent of average lifetime income, then even if worker search were perfectly directed, no mobility would be observed.

The assumption that wages are determined by splitting the flow output is made for the sake of tractability. Since the goal of the paper is to understand the pattern of worker flows conditional on the aggregate unemployment rate (and therefore conditional on the average number of jobs per local labor market, $N(t)$), the results are robust to any wage determination process that does not affect each worker's preference ordering over jobs or each job's preference ordering over workers.[7] What

[7] Similarly to the effect of changing $d(\cdot)$ in footnote 2 above, changing the way in which output is split between worker and firm changes the profitability of each specific match, therefore the expected profitability of a job, therefore the return to entry conditional on $N(t)$, and therefore the way in which $N(t)$ responds to a change in aggregate labor

would affect the results, however, would be any wage determination mechanism which changed the ordinal rankings of workers over jobs or of jobs over workers. It would be interesting to investigate alternative assumptions on wage setting of this type. For example, one possibility would be to allow firms to make take-it-or-leave-it offers to workers of their choice. Intuitively, this would lead to lower wages for workers who have close substitutes and higher wages for workers who are harder to replace. For example, consider a local labor market containing precisely one worker (located at $x = 0.3$), and precisely two jobs, located at $y_1 = 0.7$ and $y_2 = 0.71$. In the benchmark model, the resulting matching is between the worker and job 1; job 2 remains vacant. The flow output of $1 - |0.7 - 0.3| = 0.6$ divided equally between the worker and job 1. Allowing for Bertrand competition between the two jobs would lead to job 2 being willing to offer the entire output it would produce if it matched with the worker, namely 0.59, as a wage. The resulting matching would then still have the worker being matched to job 1, but the wage would be much higher. Understanding the effect of such wage determination mechanisms is left for future research.

3 Calibrated Model

In this section I describe how I calibrate and solve the model. I then discuss the results for worker flows.

The model has very few parameters to calibrate. As observed in Section 2.4 above, there is a bijective relation between the average numbers of workers and jobs per local labor market $(M, N(t))$ and the unemployment and vacancy rates $(u(t), v(t))$. I set $M = 244.2$ and $N = 236.3$ to be consistent with a steady-state unemployment rate of 5.4 percent and a vacancy rate of 2.3 percent, following Shimer (2007) who chooses these targets to match information from the CPS and from the Job Opening and Labor Turnover Survey (JOLTS) of the Bureau of Labor Statistics respectively, for the period December 2000 through April 2006. The only remaining two parameters are the arrival rates of quit and layoff shocks, q and l. As in Shimer (2007), the division of total separation shocks into quit and layoff shocks appears quantitatively irrelevant for the moments considered here (conditional on the sum $q + l$). I choose $q + l$ to be 0.28 at monthly frequency, and somewhat arbitrarily divide this total separation rate as $q = l = 0.14$. This ensures that the

productivity. However, because I consider the effect of changes in $N(t)$ directly, this does not matter for the results.

monthly separation probability in the model (that is, the probability that a worker who is employed in month t is unemployed in month $t+1$) is 0.0255, consistent with a 10 percent instantaneous EU transition rate when measured at quarter. (As observed by Shimer (2005), because the job-finding rate is large, time aggregation means that the separation probability differs substantially from the separation rate in a continuous-time model.)

To see the implication of the model for flows, I simulate a model-generated time series for the labor market experience of a single worker; for concreteness I will refer to this worker in this paragraph as worker A. To begin, I draw an initial local labor market from the cross-sectional distribution of islands, and add worker A to it. Let i be the number of workers and j the number of jobs in the market. I draw locations on the circle for the $i+j$ agents on the island, calculate the initial matching, and record whether worker A is employed or unemployed, and, if employed, the identity of and distance from his job. I then draw the time of arrival of the first shock to affect the local labor market; this shock arrives at rate $qi + lj + qM + lN$ since each worker (including worker A) independently gets hit by a quit shock at rate q and each job by a layoff shock at rate l, while new workers arrive from other markets at rate qM and new jobs enter at rate lN. Because I am solving for a steady state, I do not alter the aggregate number of jobs $N(t) \equiv N$ over time; this is consistent with free entry since the aggregate rates of job destruction and job destruction must be equal to keep the free-entry condition satisfied in steady state. In all cases except that of a quit shock affecting worker A, upon the arrival of the shock, I perform the relevant action (for example, if a layoff shock is indicated, I randomly select a job and remove it). If worker A is affected by a quit shock, I discard the existing labor market and redraw a new one from the unconditional distribution, again randomly assigning locations to the agents, including worker A. I then repeat this process until one month has elapsed. I then again record the labor market status of worker A and record it. I continue this period for $T = 100,000$ months[8]

Note that because I only record the state of the labor market each month, time aggregation affects the interpretation of the results in two main ways. The first is standard and concerns the relationship between the EU transition rate and the EU transition probability. A worker who loses her job at the beginning of a month and finds a new job within the month will be counted as undergoing an EE transition, not an EU transition followed by a UE transition, even if she did

[8]This small number will be increased in a later version of this paper.

Type of flow	EU	UE	EE	EE with wage rise	EE with wage cut
Monthly probability	0.0255	0.560	0.464	0.229	0.235

Table 1: Worker transition probabilities

have a short unemployment spell. The second interpretation issue is that I described the model under the assumption that (re-)matching occurs only once per month: the implicit assumption was that workers and jobs newly arrived to a local labor market remain unemployed until the beginning of the month following their arrival. An alternative assumption would be that every change in the set of workers and firms in an individual labor market caused the labor market to reallocate the set of worker-job matches completely. This would cause a large increase in the instantaneous transition rates (since the expected number of shocks to the average labor market, $2qM + 2lN$ per month, is large). However, if UE, EU, and EE transitions were measured by observing a worker's labor force status and employer identity once per month, as done in practice in the CPS, this modification has no effect on the reported transition rates. I prefer the first assumption since the second implies implausibly high instantaneous transition rates; however, it should be noted that the lack of intra-month reallocation is implicitly a form of labor market friction that reduces churning flows.

3.1 Worker flows

I use the resulting labor market history for worker A to measure the frequency of EU, UE, and EE flows, together with their characteristics. The results are summarized in Table 1.

The probability of an EU transition between months $t-1$ and t (conditional on being employed in month $t-1$) in the resulting series is 0.0239 at monthly frequency. This is close to the target of 0.0256; however, the results are preliminary due to simulation variability (only a small time series has been calculated). The probability of a UE transition conditional on unemployment in month $t-1$ is 0.472. Because the model matches the US unemployment rate, flow balance requires that if it matches the separation probability of an employed agent, it must also match the job finding probability of an unemployed agent. The model generates a large number of job-to-job transitions; the probability of a job-to-job transition each month for an employed worker exceeds 0.4; job-to-job transitions are more than 14 times as frequent as job-to-unemployment transitions.

Among workers experiencing an EE transition, wage cuts are common, nearly as common as wage gains: 48.6 percent of job-to-job transitions are associated with wage cuts. Both wage cuts and wage rises tend to be small. The probability of a wage rise exceeding 5 percent is 0.0198, and that of a wage rise above 10 percent is 0.0110. The corresponding wage cut probabilities are 0.0196 and 0.0108 respectively. Accordingly, the median log wage rise is 0.0026, as is the median log wage cut.

3.2 Cross-sectional Predictions

I now investigate whether the key cross-sectional predictions of the model, Prediction 1 and Prediction 2, are supported by model-generated data. To do this, I use the model-generated simulation of the life history of an individual worker already constructed. To generate a regression sample, I mimic the structure of the CPS, which allows me to observe a four-month panel of the employment status and employer identity of a worker. I construct a model analog to the CPS using the following procedure. I draw a month number t_1 uniformly from $\{4, 5, \ldots, T\}$, and record the employment status of the worker, the identity of his employer, and his wage for months $t_1 - 3$ through t_1. I then repeat this procedure with a new random month number t_2, again drawn uniformly from $\{4, 5, \ldots, T\}$. I continue until I have drawn a sample the same size as the CPS sample used in Section 4 below, that is, around 1.8 million workers. I then use this data to investigate Prediction 1 and Prediction 2.

Sample	Employed at $t-1$	
	(1) EE trans.	(2) EU trans.
$\log w_{i,t-1}$	-2.12**	-0.514**
Constant	0.381**	0.0188**
R^2	0.034	0.022

Bootstrap standard errors will be reported in a future draft.
* $p < 0.05$, ** $p < 0.01$.

Table 2: Patterns of worker flows in the benchmark model: using wage data

Table 2 gives evidence for Prediction 1, showing that being in a better match significantly reduces a worker's likelihood of experiencing an employment transition. In Column (1), I use OLS to regress a dummy variable which takes the value 1 if worker i undergoes an EE transition between period $t-1$ and period t on a constant and on worker i's log wage in period $t-1$:[9]

$$\mathbf{1}(\text{EE trans.}_{it}) = c + a \log w_{i,t-1} + \varepsilon_{i,t} \tag{2}$$

High wage workers are less likely to undergo EE transitions, although the effect is economically reasonably small. Moving from the 10th percentile of the employed wage distribution to the median reduces the probability of an EE transition from 0.412 to 0.384; moving to the 90th percentile reduces the probability further to 0.381. Column (2) of Table 2 reports the results of a similar regression with EU transitions as the dependent variable. Again, high-wage workers are significantly less likely to transition from employment to unemployment.

Sample	Employed at $t-1$		Employed at $t-1$		Employed at $t-3$ to $t-1$	
	(1) EE trans.	(2) EU trans.	(3) EE trans.	(4) EU trans.	(5) EE trans.	(6) EU trans.
New job	0.192**	0.019**			0.189**	0.0136**
New spell			0.114**	0.0770**		
Constant	0.285**	0.0126**	0.398**	0.0207**	0.285**	0.0126**
R^2	0.036	0.004	0.002	0.011	0.036	0.002

Bootstrap standard errors will be reported in a future draft. * $p < 0.05$, ** $p < 0.01$.

Table 3: Patterns of worker flows in the benchmark model

Table 3 investigates Prediction 2, that is, whether future transitions are less likely for workers who have recently experienced either a UE or an EE transition. Recall that the intuition for this result is that past transitions proxy for low match quality. The regressions run here mimic those that can be run in the CPS data, to be investigated in Section 4 below. I treat the CPS data as providing four-month panels of worker employment history, and, motivated by this, I study whether employment transitions between months $t-3$ and $t-2$ or between months $t-2$ and $t-1$ predict employment transitions between months $t-1$ and month t in model-generated data. Columns (1)-(6) report slightly different specifications and slightly different data selection criteria.

In Column (1), I report the results of a regression of a dummy variable for an EE transition between periods $t-1$ and t on a dummy variable which takes the value 1 if the job the worker holds

[9]I will switch from this linear probability model specification to a probit in a future version of the paper.

in period $t-1$ commenced no earlier than period $t-2$. (That is, the worker was unemployed in period $t-3$ or $t-2$, or had a different employer in period $t-3$ or period $t-2$ than in period $t-1$.) A new job significantly increases the probability of an EE transition, by 19.2 percentage points. Column (2) reports a similar regression with the dependent variable being a dummy variable for an EU transition rather than an EE transition. A worker in a new job has more than double the probability of transitioning to unemployment than one who has held his job longer.

Columns (1) and (2) report the combined effect of any kind of past worker transition, either EE or UE, on future transitions. It is also interesting to investigate the effects of recent unemployment and recent EE transitions without unemployment separately. Columns (3) and (4) investigate the ffect of a recent UE transition. These columns are analogous to Columns (1) and (2), except that the independent variable is now a dummy variable which takes the value 1 only if the worker recently experienced a UE transition (that is, was unemployed in month $t-3$ and/or month $t-2$). This variable also is associated with a significant increase in the probability of both EE and especially EU transitions. Finally, Columns(5) and (6) investigate the effect of EE transitions for workers who have not recently experienced unemployment. These columns replicate columns (1) and (2) but restrict the sample to consider only those workers who were never unemployed in months $t-3$ through $t-1$. The results are similar, although comparing columns (2), (4), and (6) shows that a recent UE flow is much more predictive than a recent EE flow is for a future EU flow. This is consistent with the idea that workers who recently found a job having been unemployed are likely to be badly matched on average, and so more likely to experience a transition back into unemployment.

3.3 Aggregate Productivity and Worker Flows

I now turn to investigating the business cycle properties of worker flows in the model.

In the mismatch structure, the average number of jobs per local labor market, $N(t)$, is a sufficient statistic for the contemporaneous values both of the stocks of employed and unemployed workers and for the instantaneous UE, EU, and EE transition rates. Therefore, given a stochastic process for the unemployment rate, the structure of the model completely determines the stochastic processes for these variables also. The most parsimonious way to understand the implication of the model for the cyclical pattern of worker flows, therefore, is to choose the process of $N(t)$ to match the time series behavior of the unemployment rate and use the model to trace out the resulting pattern of worker flows.

The exercise just described is, of course, not an equilibrium exercise since the number of jobs is taken as exogenous. This can be endogenized using free entry of jobs. Firms can create a job at a cost of $k > 0$, and this job is allocated to a random location. Any shock which, *ceteris paribus*, makes jobs more profitable will increase the number of jobs in the economy, and therefore reduce the unemployment rate. Conditional on an exogenous stochastic process for labor productivity z, and conditional on a fully-specified wage determination rule that determines the profitability of each job, one can determine the endogenous response of job creation to a given productivity shock history, and use this in turn to generate cyclical variation in the stocks of employed and unemployed workers and the EU, UE, and EE transition probabilities. Just as in the familiar Mortensen-Pissarides environment, the wage determination process is key for the volatility of the model's responses to productivity shocks.[10] Performing the simpler exercise described in the previous paragraph allows me to sidestep these issues while still fully illustrating the model's implications for worker flows.

In a future version of the paper I plan to show the implications of the model for worker flows using the full post-war history of the US economy, but for now, I show only the comparative static effect of reducing the unemployment rate by one percentage point from its benchmark calibration target, from 5.6 percent to 4.6 percent.[11] This requires increasing the average number of jobs per local labor market from 236.3 to 240.4. Table 4 shows the effect of this change on the rates of worker flows in such a boom. For convenience, the first row of the table repeats the benchmark results from **??**; the second row shows the effect of the increase in the number of jobs. The EU transition rate is countercyclical, and the UE transition rate procyclical; interestingly, the fall in the EU transition probability is much larger than that in the UE probability. In this model, booms entail lower unemployment mostly through a reduction in the EU separation rate. Although the EE transition rate is weakly procyclical, the net amount of worker separations (that is, the sum of the EU and EE transition probabilities) falls somewhat. Workers transitioning from job to job are slightly more likely to experience a wage increase in a boom than in normal times, but the effect is

[10] The importance of wage determination for volatility in the Mortensen-Pissarides model is familiar from the work of Shimer (2005), Costain and Reiter (2008), and Hagedorn and Manovskii (2008). Hawkins (2012) shows that the same is true in the mismatch environment of Shimer (2007).

[11] Note, however, that because the unemployment rate is a sufficient statistic for contemporaneous flows, the usual issue that the response to a permanent and a transitory shock might differ does not arise in this model.

Type of flow	EU	UE	EE	EE with wage rise	EE with wage cut
Monthly probability: $u = 5.6\%$	0.0255	0.560	0.464	0.229	0.235
Monthly probability: $u = 4.6\%$	0.0220	0.572	0.466	0.231	0.235

Table 4: Worker transition probabilities after productivity shock

very small.

In summary, the model does a fair job of matching the cyclicality of EU, UE, and EE transition rates, consistent with evidence from Shimer and from Fallick and Fleischman (2004). Of course, the job-to-job transition rate in the benchmark model is much higher than its empirical counterpart, so I will repeat this robustness exercise in the frictional model in Section 5 below.

4 Cross-Sectional Evidence from the Current Population Survey

I now examine whether the patterns of worker flows predicted by the model match their empirical counterparts.

To do this, I use data from the Current Population Survey (CPS) of the Bureau of Labor Statistics. The CPS interviews a random sample of households using a panel structure that allows studying the pattern of worker flows over time for an individual worker. A household is interviewed for four successive months, then not interviewed for eight months, and then finally again interviewed for each of the following four months. I use data from 1994 through 2008. I do not use the information that a worker is seen for two separate four month periods, but instead treat each four-month rotation as an independent observation. I apply only minimal sample selection criteria by requiring that workers be aged between 16 and 65, and either employed or unemployed in each month of the relevant four-month period. From the information in the CPS, I use standard demographic information on the household (age, race, and education). Note that I do not currently use wage information.

After pooling all the available data, I run OLS regressions as in (2). I control for demographic variables by including a cubic in age, dummies for gender and for educational attainment (the omitted categories are male and high school graduates). I also include a full set of year dummies (not reported). The results, the exact empirical analog to Table 3, are reported in Table 5. As before, columns (1), (3), and (5) have as dependent variable a dummy for an EE transition in month 4 of the panel; columns (2), (4), and (6) study EU transitions. In columns (1)-(4), the sample selection criterion is simply that the worker was employed in month 3 (and so could possibly undergo an EE or EU transition from month 3 to month 4); in columns (5) and (6) I restrict the sample only to workers who were always employed in months 1, 2, and 3. Finally, columns (1), (2), (5), and (6) use as an independent variable a dummy for whether the worker has a new job in month 3 (that is, was not employed in at least one of months 1 and 2 with the month-3 employer, but could have been either employed at a different employer or unemployed); columns (3) and (4) consider instead the effect of having been unemployed in either month 1 or 2.

The results are remarkably encouraging, with the signs of all coefficients on the variables of interest (specifically whether the worker recently started a new job or a new employment spell) all positive, just as in the model-generated data in Table 3. In addition, the coefficients on the demographic controls have the expected signs: older workers are less likely to make employment transitions, and less educated workers are more likely to transition from employment to unemployment and less likely to move from job to job. The results are robust to controlling to various slight modifications, including controlling for potential experience instead of age.

This evidence is in support of the model, and contrary at least to some extent to several other familiar models of worker flows. First, the first generation of search models (Pissarides, 1985; Shimer, 2005) assumed exogenous separations and do not predict that workers who have recently started jobs should be more likely to transition to unemployment. The standard version of the model with endogenous separations and no on-the-job search (Mortensen and Pissarides, 1994) predicts that the transition rate to unemployment should be *lower*, not higher, for workers who recently experienced. In addition, these models lack a notion of job-to-job transitions. Second, the job ladder model of Burdett (1978) and Burdett and Mortensen (1998) naturally predicts that EE transition rates should be higher for workers who are better matched, but it is not obvious whether workers who have recently experienced an EE transition should be expected to be better matched (after all, they recently moved up the ladder, and it's easier for a worker to move up the ladder the further she still has to move). Likewise, the job ladder model does not in its simplest form predict that UE or EE transitions should predict EU transitions. In the

Sample	Employed at $t-1$		Employed at $t-1$		Employed at $t-3$ to $t-1$	
	(1) EE trans.	(2) EU trans.	(3) EE trans.	(4) EU trans.	(5) EE trans.	(6) EU trans.
Age	-0.0063**	-0.0026**	-0.0068**	-0.0025**	-0.0065**	-0.0026**
	(0.0002)	(0.0001)	(0.0002)	(0.0001)	(0.0002)	(0.0001)
Age^2	0.0001**	0.0000**	0.0001**	0.0000**	0.0001**	0.0000**
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Age^3	-0.0000**	-0.0000**	-0.0000**	-0.0000**	-0.0000**	-0.0000**
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
Female	-0.0007**	-0.0022**	-0.0007**	-0.0021**	-0.0005*	-0.0016**
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0001)
Black	-0.0011**	0.0045**	-0.0010**	0.0040**	-0.0008*	0.0030**
	(0.0004)	(0.0003)	(0.0004)	(0.0003)	(0.0004)	(0.0003)
$<$ HS	0.0004	0.0066**	0.0004	0.0060**	0.0007	0.0056**
	(0.0004)	(0.0003)	(0.0004)	(0.0003)	(0.0004)	(0.0003)
Some coll.	0.0019**	-0.0030**	0.0020**	-0.0026**	0.0019**	-0.0023**
	(0.0003)	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0002)
Coll. grad.	0.0001	-0.0055**	0.0001	-0.0050**	-0.0000	-0.0046**
	(0.0003)	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0002)
New job	0.0457**	0.0353**			0.0470**	0.0121**
	(0.0004)	(0.0003)			(0.0005)	(0.0003)
New spell			0.0404**	0.0820**		
			(0.0007)	(0.0005)		
Constant	0.1247**	0.0552**	0.1360**	0.0547**	0.1277**	0.0546**
	(0.0023)	(0.0017)	(0.0023)	(0.0016)	(0.0023)	(0.0015)
Observations	1803782	1803782	1803782	1803782	1765162	1765162
R^2	0.011	0.012	0.007	0.019	0.010	0.005

Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$. All regressions include a full set of year dummies. Omitted educational category is high school graduate.

Table 5: Patterns of worker flows in the CPS

5 Labor Market Frictions

[To be written. The key idea is that the benchmark model generates too high a ratio of EE transitions to EU transitions. In this section I model labor market frictions by assuming that the distributions F_x and F_y no longer have continuous support, but instead are atomic. More concretely, I'm going to assume that instead of being uniform on the full perimeter of a circle, F_x and F_y are instead supported on a regular n-gon with vertices at n equally-spaced points on the circle. I think of this as modeling a friction because I think of the true underlying distributions as being unchanged, but I assume that the discreteness of the model distributions models the idea that it is not worth switching a match that is not too bad. How bad is 'not too bad' depends on n.

I make the tie-breaking assumption that a match between a worker at x and a job at y only breaks up if either worker x is hit by a quit shock, job y by a layoff shock, or if due to a change in the sets of workers and jobs in the economy, fewer matches between workers at x and jobs at y are generated by the stable matching at t than were generated in the period $t-1$ matching.

For $n = 1$, the model now reduces to Shimer (2007) precisely, and generates too few EE transitions relative to EU transitions. As $n \to \infty$, the model reduces to the benchmark model, and generates too many. Preliminary numerical work suggests that $n = 3$ matches the data best. Conclusion to this section: the ratio of EE transitions to EU transitions measures the extent of labor market frictions in this type of model.]

6 Efficiency

[To be written. The key observation is that the stable matching calculated above does not solve the planner's problem of maximizing output. However, it is straightforward to solve the planner's problem also. The results for flows are similar.]

7 Ranking

[To be written. This is a new idea. The idea is that instead of heterogeneity on the circle, the workers and jobs in a local labor market are ranked. Worker 1 is best, worker i is worst. Similarly, job 1 is best, job j is worst. Assume supermodularity in production, so that positive assortative matching is optimal (Becker, 1973). Thus, this is the case where the heterogeneity is not so much match-specific (like in the benchmark model) but worker- and firm-specific.

Also assume that when a new worker arrives, he or she is equally likely to slot into any of the $i+1$ possible ranks (better than worker 1, better than worker 2 but worse than worker 1, better than worker 3 but worse than worker 2, ..., worse than worker i). Similarly for jobs. This assumption is very convenient because it means the number of flows can be calculated in closed form. A worker who is worse than worker i_0 causes all workers from $i_0 + 1$ through i to shift one rung down in the matching ladder. Thus the expected number of EE transitions associated with the arrival of a single worker is $i/2$ (since every number from 0 through i is equally likely). Of course, it's not this simple in discrete time.

Of course, this assumption on the productivity distribution of new arrivals is not consistent with the obvious way of ranking agents on a line, where the probability a new arrival is better than worker 1 (say) depends on worker 1's position in the CDF F_x. I can also investigate this case.

Intuition suggests the main results carry over, with perhaps even more EE transitions relative to EU transitions than before. We can also introduce frictions in the same way.]

8 Conclusion

The mismatch model of Shimer (2007) makes an important contribution to our understanding of the microfoundations of the aggregate relationship between unemployment and vacancies, by showing how the empirical relationships between unemployment and vacancy rates, as well as UE and EU transition rates, can be consistent with a simple microfoundation. However, the model is very stylized, particularly in its treatment of heterogeneity. In this paper I investigate the effects of enriching the model to allow for a much richer notion of worker and firm heterogeneity. The resulting model delivers a rich theory of worker flows, both between employment and unemployment and from job to job. This theory delivers an interesting alternative to job ladder models in the style of

A Omitted Proofs

Proof of Proposition 1. Let the number of workers in the local labor market be i and the number of jobs be j. Write the set of worker locations as $\{x_k\}_{k=1}^{i}$ and the set of job locations as $\{y_l\}_{l=1}^{j}$. Because the distributions F_x and F_y are atomless, with probability one all the x_k are distinct, as are all the y_l. Write $I_i = \{1, 2, \ldots, i\}$ and $J_j = \{1, 2, \ldots, j\}$.

Let (k_1, l_1) be the index pair that minimizes $d(x_k, y_l)$ over all pairs $(k, l) \in I_i \times J_j$. Since worker k_1's most preferred employer is job l_1 and job l_1's most preferred employee is worker k_1, stability requires that this worker-job pair be matched. If either $i = 1$ or $j = 1$, this establishes already that there is a unique stable matching; if there are any remaining workers, they will be unemployed, and any remaining jobs will be vacant.

If $i, j \geq 2$, define $I_{i-1} = I_i \setminus \{k_1\}$ and $J_{j-1} = J_j \setminus \{l_1\}$. Let (k_2, k_2) be the index pair that minimizes $d(x_k, y_l)$ over all pairs $(k, l) \in I_{i-1} \times J_{j-1}$. Stability requires that worker k_2 be matched to job l_2, since each prefers to be matched to the other than to any other potential partner except for worker k_1 or job l_1, who are already matched. This pair satisfies the stability condition: it cannot be that worker k_2 prefers some job other than l_1 to l_2, by construction of (k_2, l_2), and although it is possible that worker k_2 prefers job l_1 to job l_2, still, job l_1 prefers its matched employee k_1 to worker k_2 by construction of (k_1, l_1). A very similar argument shows that job l_2 cannot prefer any worker other than k_1 to k_2, and that this does not violate stability since worker k_1 prefers job l_1 to job l_2.

If $i, j \geq 3$, define $I_{i-2} = I_{i-1} \setminus \{k_2\}$ and $J_{j-2} = J_{j-1} \setminus \{l_2\}$. Let (k_3, l_3) be the index pair that minimizes $d(x_k, y_l)$ over all pairs $(k, l) \in I_{i-2} \times J_{j-2}$. Stability requires that worker k_3 be matched to job l_3. As before, it cannot be that worker k_3 prefers any jobs except possibly for l_1 or l_2 to job l_3, but this does not violate stability since by construction job l_1 prefers worker k_1 to worker k_3 and job l_2 prefers worker k_2 to worker k_3. Analogously, the preferences of worker l_3 do not violate stability either.

Continue inductively until all workers are matched (if there are at least as many jobs as workers) or all jobs are matched (if there are at least as many workers as jobs). The resulting matching is uniquely determined by this process. To see that it is stable, consider an arbitrary pair (k, l) which are not matched. By construction either k is matched to some job l', or l is matched to some worker

k' (since all agents on the short side of the market are matched). Suppose that k was matched earlier than l in the algorithm just described (including the possibility that l is unmatched). Then worker k preferred job l' to job l, which was vacant when the match between k and l' was formed. Alternatively, suppose that l was matched earlier than k. Then l preferred worker k' to worker k. In neither case is stability violated. This completes the proof that this matching is stable; uniqueness follows from the construction. $\square$

References

Becker, G. S., "A Theory of Marriage: Part I," *Journal of Political Economy* 81 (July 1973), 813–846.

Burdett, K., "A Theory of Employee Job Search and Quit Rates," *American Economic Review* 68 (March 1978), 212–220.

Burdett, K. and D. T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review* 39 (May 1998), 257–73.

Burdett, K., S. Shi and R. Wright, "Pricing and Matching with Frictions," *Journal of Political Economy* 109 (October 2001), 1060–1085.

Coles, M. G. and A. Muthoo, "Strategic Bargaining and Competitive Bidding in a Dynamic Market Equilibrium," *Review of Economic Studies* 65 (April 1998), 235–260.

Coles, M. G. and E. Smith, "Marketplaces and Matching," *International Economic Review* 39 (February 1998), 239–254.

Costain, J. S. and M. Reiter, "Business Cycles, Unemployment Insurance, and the Calibration of Matching Models," *Journal of Economic Dynamics and Control* 32 (2008), 1120–1155.

Delacroix, A. and S. Shi, "Directed Search on the Job and the Wage Ladder," *International Economic Review* 47 (May 2006), 651–699.

Diamond, P. A., "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90 (October 1982), 881–894.

Ebrahimy, E. and R. Shimer, "Stock-Flow Matching," *Journal of Economic Theory* 145 (July 2010), 1325–1353.

Fallick, B. and C. A. Fleischman, "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," Finance and Economics Discussion Series 2004-34, Federal Reserve Board, 2004.

Gale, D. and L. Shapley, "College Admissions and the Stability of Marriage," *American Mathematical Monthly* 69 (January 1962), 9–15.

Gautier, P. and C. Teulings, "Sorting and the output loss due to search frictions," unpublished, 2011.

Hagedorn, M. and I. Manovskii, "The Cyclical Behavior of Unemployment and Vacancies Revisited," *American Economic Review* 98 (September 2008), 1692–1706.

Hawkins, W. B., "Wage Determination and Labor Market Volatility under Mismatch," unpublished, May 2012.

Lazear, E. P. and J. R. Spletzer, "Hiring, Churn and the Business Cycle," Working Paper 17910, National Bureau of Economic Research, Cambridge, MA, March 2012.

Menzio, G. and S. Shi, "Efficient Search on the Job and the Business Cycle," *Journal of Political Economy* 119 (June 2011), 468–510.

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

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

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

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

Postel-Vinay, F. and J.-M. Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica* 70 (November 2002), 2295–2350.

Shimer, R., *Contracts in a Frictional Labor Market*, Ph.D. thesis, MIT (1996), chapter 1.

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

———, "Mismatch," *American Economic Review* 97 (September 2007), 1074–1101.

Taylor, C. R., "The Long Side of the Market and the Short End of the Stick: Bargaining Power and Price Formation in Buyers', Sellers', and Balanced Markets," *Quarterly Journal of Economics* 110 (August 1995), 837–855.