UNDERSTANDING FLUCTUATIONS IN THE INS AND OUTS OF THE LABOR FORCE

ETIENNE LALÉ[†]

June 2013

[Preliminary, Do Not Quote – Comments Welcome]

ABSTRACT. Although the size of the labor force is nearly acyclical, worker flows between employment and unemployment on the one hand and inactivity on the other hand fluctuate significantly over the business cycle. After reviewing these facts, this paper lays out a job-search model to uncover the determinants of worker flows between employment, unemployment and inactivity. However rudimentary, the proposed model reproduces the value of these flows across cycles remarkably well. The model further shows that, if aggregate conditions are not a first-order determinant of workers' labor force decisions, then a substantial part of the observed fluctuations in the ins and outs of the labor force can be attributed to two composition effects. First, workers who join the labor force irrespective of the business cycle are more likely to find themselves immediately unemployed rather than employed when the aggregate job-finding rate is below trend. Hence the countercylicality of the probability to move from inactivity into unemployment. Second, high-productivity workers who are less likely to drop from the labor force are more numerous to be drawn into unemployment in times of high job destruction. They thereby contribute to the fall in the probability to leave the labor force from unemployment during recessions. Both mechanisms are supported by microdata from the Current Population Survey.

Keywords: Worker Flows, Labor Force Participation, Business Cycle, Job-Search
JEL codes: E32, J21, J64

I am grateful to Fabien Postel-Vinay and Etienne Wasmer for helpful comments and suggestions on this paper. This work was supported by Region Ile-de-France. All remaining errors are my own.

[†] Sciences-Po Department of Economics, Office H405 – 28 rue des Saints-Pères, 75007 Paris, France
E-mail: etienne.lale@sciences-po.org

1. Introduction

A flurry of recent research indicates that worker flows not only between employment and unemployment but also into and out of the labor force are crucial for labor market dynamics, from various perspectives. Over the business cycle, transitions into and out of the labor force display significant levels of volatility and they play a major role in determining the size of the unemployment pool, particularly during downturns. This holds true especially for the recent recession in the United States, as documented by Hotchkiss et al. (2012), Elsby et al. (2013a) and Shimer (2013). At a lower frequency, flows into and out of the labor force – which are an order of magnitude larger than the flows from employment to unemployment – are key for explaining long-run differences in employment rates across countries (Pries and Rogerson, 2009). Finally, movements from inactivity to unemployment are a first-order contributor to the distribution of participation and unemployment spells over the life-cycle of workers (Villena-Roldan et al., 2011). Altogether, these findings call for a better understanding of movements into (the ins) and out (the outs) of the labor force.

This paper seeks to contribute to this recent line of research on the theoretical and empirical fronts. First, it develops a job-search model that can account for worker flows between employment, unemployment and inactivity. Second, it stresses some of its implications for short-run fluctuations in the ins and outs of the labor force. Third, it provides empirical tests of the mechanisms uncovered by the model against microdata from the Current Population Survey (CPS).

To begin with, the first contribution of the paper is in proposing a job-search model to organize thinking about worker flows between employment, unemployment and inactivity. While this is not the first paper to develop such a model[1], its originality lies in the fact that the model remains relatively simple and in the meantime connects to the data. In the proposed setting, workers experience stochastic shocks to their own idiosyncratic productivity, they forego home production when in the labor market and exercise the option to drop from the labor force when their productivity is too low. Workers can search for jobs when unemployed or when inactive (with different search intensities), implying that a steady-state of the model features worker flows between all three labor market states. Idiosyncratic productivity levels map into wages, which allows to infer the underlying productivity process from data on wage earnings. After feeding this productivity process, the small set of remaining parameters are readily linked to worker flows as observed in the data and the model performs remarkably well in reproducing the value of these flows across cycles. This qualifies the model as a useful tool, potentially relevant to study the short-run, medium-long run and life-cycle issues mentioned before. The rest of the paper focuses on the former: it draws on the model to uncover mechanisms behind the observed short-term fluctuations in the ins and outs of the labor force.[2]

[1]Studies of the labor market where workers are classified into three distinct states – employment, unemployment and inactivity – but that do not address explicitly the issue of worker flows between these states include Andolfatto and Gomme (1996), Alvarez and Veracierto (1999), Veracierto (2008) and Shimer (2013). Papers that focus on these flows are Garibaldi and Wasmer (2005), Pries and Rogerson (2009), Krusell et al. (2011, 2012) and Mankart and Oikonomou (2011). As explained hereafter, the model analyzed by Garibaldi and Wasmer (2005) is the closest in spirit to the one developed here but has fewer connections with the data; on the other hand, the study by Krusell et al. (2012) shares some of the objectives of the present paper but proceeds within a much different model and along different lines.

[2]None of the papers that develop models of worker flows between employment, unemployment and inactivity use it as a point of departure for an empirical analysis of the corresponding transitions observed in microdata.

The second step of the analysis is to study a typical recession through the lens of the job-search model. Because this class of models is notoriously unable to reproduce the observed volatility of key labor market variables (Shimer, 2005), the analysis does not undertake the task of feeding the model with aggregate shocks to, say, productivity. Instead, it studies how the ins and outs of the labor force depart from their steady-state values in response to shocks that affect movements between unemployment and employment by the magnitude of a typical recession. This different route allows to single out two implications of the model for fluctuations in the ins and outs of the labor force, both of which can be viewed as composition effects. First, if workers do not (fully) delay labor market entry when the job-finding rate is below trend, then they are more likely to be found unemployed rather than employed following labor market entry. This can rationalize part of the observed fluctuations in the probabilities to join employment and unemployment from inactivity. Second, when job destruction rates are high (typically: early on in recessions), an unusually large number of high-productivity workers are drawn into the unemployment pool. Since these workers have a lower than average tendency to drop from the labor force, part of the observed fall in the probability to leave the labor force from unemployment during downturns can be attributed to this compositional shift.

Informed by the dynamic implications of the model, the paper then turns to microdata from the Current Population Survey to gauge their relevance. The empirical analysis consists primarily in estimating and decomposing elasticities of the probability of movements into and out of the labor force with respect to the unemployment rate (a cyclical indicator).

A first set of results supports the prediction that fluctuations in the ins of the labor force stem mostly from changes in the probability to move into unemployment conditional on labor market entry, not from changes in the decision to enter the labor force. First, the elasticity of the probability to join the labor force with respect to the unemployment rate is found to be rather low: this is consistent with aggregate conditions not being a first-order determinant of workers' labor force decisions. Second, the elasticity to move into unemployment rather than employment following labor market entry is about four times larger and is similar across workers with different characteristics, including their degree of labor force attachment. This is in line with the view that fluctuations in regaining employment are driven predominantly by the aggregate job-finding rate (see Elsby et al., 2010; Shimer, 2012).

A second set of results validates the view that shifts in the heterogeneous characteristics of the unemployed are one primary driving force behind the observed fluctuations in the outs of the labor force. That is, flows from unemployment into either employment or inactivity exhibit similar cyclical behaviors in the aggregate. However, when measured at the individual level, the elasticity of the probability to drop from unemployment into inactivity is up to four times lower than that of the probability to move into employment. Moreover, it becomes essentially zero after controlling for workers' skills (proxied by wages in their previous job) and the latter confirm that high-productivity individuals are less likely to drop from the labor force. This is noteworthy because a number of studies (Baker (1992) and Shimer (2012), among others) have cast doubt on the potential of composition effects to be a major source of fluctuations in worker flows.

This paper is related to various strands of the literature on worker flows. First, it connects to the set of studies that build dynamic, frictional models of the labor market to account for worker flows between employment, unemployment and inactivity (see footnote 1). An analogous model was first developed by Garibaldi and Wasmer (2005). They used it to address a range of questions such as the definition of marginally-attached workers, the feedback of labor supply decisions onto labor demand and the effects of various labor market policies. Relative to them, the present paper puts more emphasis on the quantitative performance of the model and studies its implications for fluctuations in the ins and outs of the labor force. Instead of shocks to the valuation of leisure, it posits that changes in labor force status reflect shocks to productivity that can be inferred from wages and are governed by an autoregressive stochastic process.[3] Finally it does not rule out the possibility of genuine transitions from inactivity to employment. Krusell et al. (2011, 2012) on the other hand also consider a more quantitative model with remarkable empirical performances. Their approach is more complex than the one adopted here since workers in their model engage into precautionary savings: changes in labor market status thus stem from shocks to idiosyncratic productivity as well as income and substitution effects. The model analyzed here achieves similar performances with only one dimension of heterogeneity across workers.

The fact that this paper connects to the literature on job-search models and addresses business cycle issues suggests a potential link with the numerous studies devoted to the cyclical performance of the canonical Mortensen-Pissarides model. A cautionary note is in order here: as noted before, the model is not tailored to make worker flows respond endogenously to aggregate productivity shocks. The poor cyclical performances of the Mortensen-Pissarides model are one reason for this choice. Another reason is that Tripier (2004), Veracierto (2008), Mankart and Oikonomou (2011) and Shimer (2013) all establish that models where workers flow between employment, unemployment and inactivity have a tendency to generate pro-cyclical unemployment rates.[4] Undertaking a conventional business cycle analysis within the job-search model would thus also involve dealing with this counterfactual result. Presumably, this would not add to the main results of the paper.

This paper is also related to empirical studies of worker flows, at least on two different levels. Firstly, the results point to composition effects playing a significant role in fluctuations along certain margins of the labor market. In the language of the model, the unemployment pool shifts towards high-productivity workers during recessions. This is in line with Mueller (2012) who finds that this pattern holds true for the United States over the past three decades. Regarding the empirical analysis, this translates into a significant attenuation effect of individual controls on the elasticity of the probability to leave the labor force from unemployment. This corroborates Elsby et al. (2013b) who address a similar empirical question with a different methodology. Their and the present paper can be viewed as reviving the "heterogeneity hypothesis" of Darby et al. (1986) to stress its relevance when looking at the ins and outs of the labor force.

Finally, the paper contributes to the empirical literature on worker flows on another level. By drawing on microdata from the Current Population Survey, it highlights the behavior of worker flows in and out of the labor force differently from studies that draw on aggregate stock-flow equations (e.g. Petrongolo and Pissarides, 2008; Elsby et al., 2013a,b; Shimer, 2012). To do so, it tackles the question of measurement error in reported labor force transitions at the individual level. Another paper that

[3] Garibaldi and Wasmer (2005) assume a stochastic process with less persistence to maintain tractability.

[4] This is because inactive workers enter the labor market during expansions and thereby inflate the unemployment pool.

addresses this issue is Elsby et al. (2013b). Their methodology or the one developed here can be used to recode transitions at the individual level and are fruitful because analyzing worker flows from microdata allows to study hypotheses such as heterogeneity and duration dependence that are more difficult to test with the aggregate stock-flow approach.

The paper is organized as follows. Section 2 reviews some empirical facts characterizing the ins and outs of the labor force. Section 3 introduces the job search model used to decode these flows. The model is calibrated in section 4 and its dynamic implications are exhibited in section 5. Section 6 contains the empirical results and relates them to the model predictions. Section 7 concludes.

2. U.S. Labor Market Facts

This section summarizes the set of facts that the rest of the paper seeks to analyze. Establishing these facts is the goal of a large literature which is still active notably through controversies regarding various measurement issues.[5] This section does not aim at taking a stand on these controversies and, since the discussion of measurement issues in data from the CPS is deferred to section 6, it does not add to this literature. Instead, it draws on the time series of worker flows computed by Shimer (2012) as they have become authoritative and organizes a series of facts for the purposes of the next sections.

2.1. **The ins and outs of the labor force over time.** Figure 1 shows the time series of the ins and outs of the labor force computed by Shimer (2012) for the years 1967 to 2007.[6] The upper graphs (a) plots movements into the labor force (the ins) which may occur either into employment or unemployment; the lower graphs (b) plots movements out of the labor force (the outs) which may occur either from employment or unemployment. In addition the dashed lines for the years 1994 onwards show the corresponding worker flows computed from the data used in section 6 and which cover the most recent recession. Each time series represents the (quarterly average of the) monthly probability to move between the corresponding labor market states.

Figure 1 reveals stark contrasts in the behavior of the different series. Beginning with the ins, worker flows from inactivity into employment are larger than those from inactivity into unemployment: they average 4.33 percent and 3.43 percent, respectively. These series exhibit similar levels of volatility but move in opposite directions: during recessions, workers become less likely to move from inactivity into employment and more likely to move from inactivity into unemployment. This was particularly pronounced during the 2007-2009 recession. Finally the two time series were also trented in opposite directions before the 1980s; this pattern vanished away from that period onwards.

Turning to the outs, the first thing to note from figure 1 is that flows from employment into inactivity are non-negligible (2.99 percent on average), although they have been trending downwards throughout almost the entire period. In particular, these flows are larger than the corresponding movements

[5]The much debated measurement issues include: (i) the role of time aggregation in the measurement of monthly flows, (ii) spurious labor market transition when comparing individual labor market statuses in two consecutive periods, (iii) margin errors, i.e. the discrepancy between measured labor market stocks and those implied by measured labor market flows and (iv) sample attrition and the resulting potential bias in transition rates.

[6]This data was constructed by Robert Shimer. For additional details, please see Shimer (2012) and his webpage http://sites.google.com/site/robertshimer/research/flows.

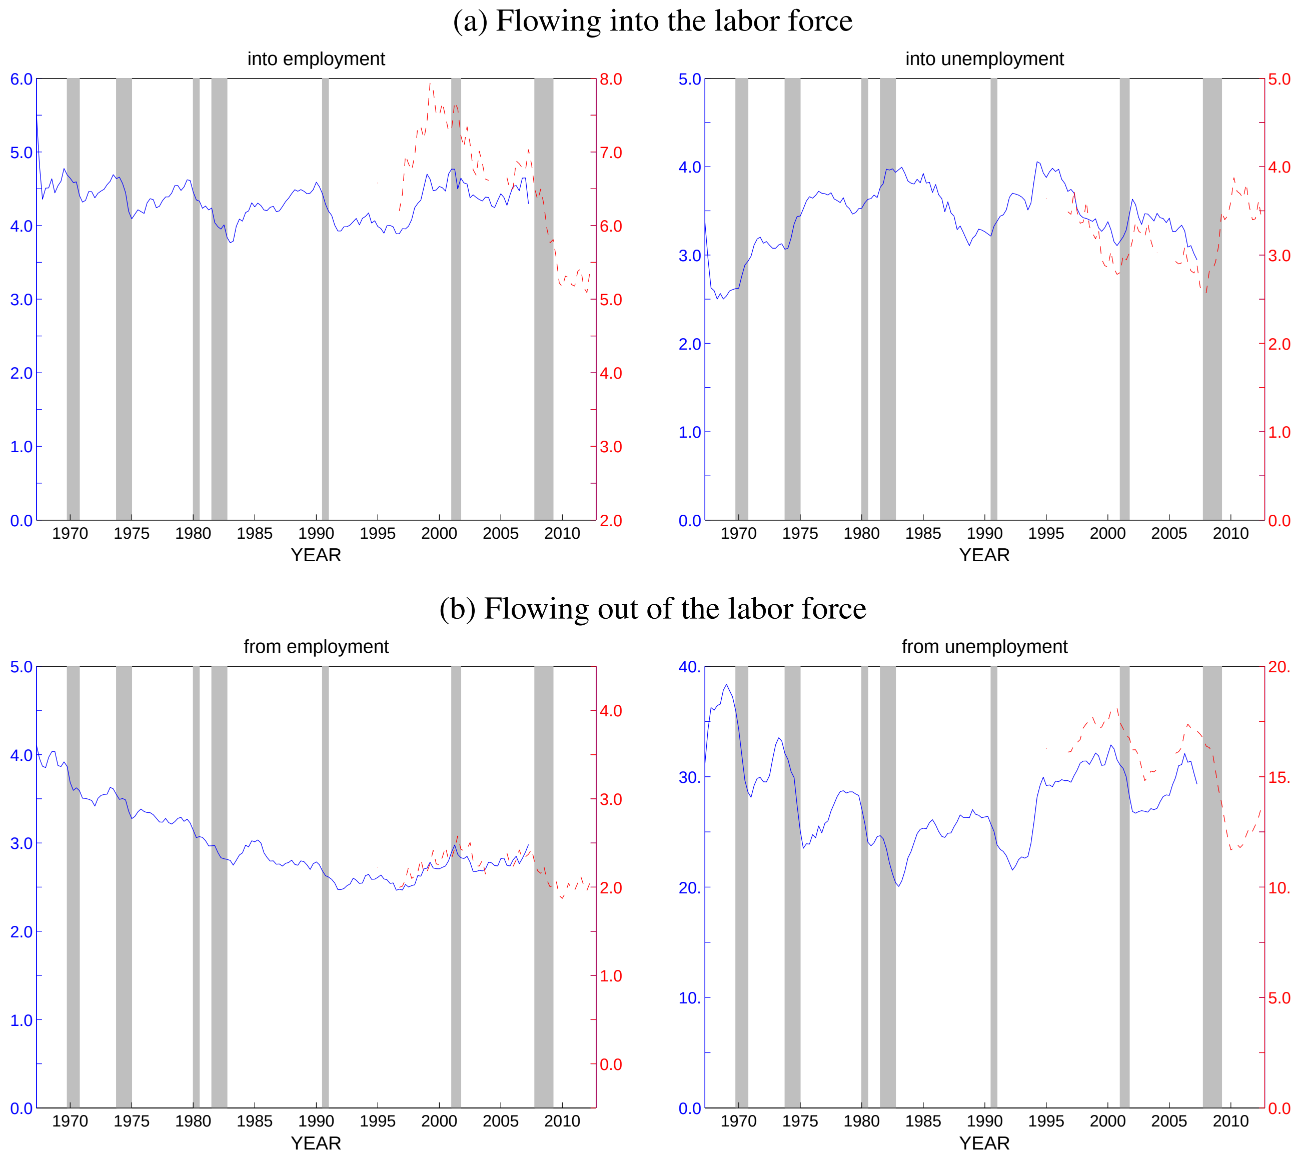

Figure 1. The Ins and Outs of the Labor Force: Monthly Transition Probabilities
The solid lines show the monthly transition probabilities computed by Shimer (2012); they are plotted against the left axis. The dashed lines show the monthly transition probabilities computed from data from the CPS used in section 6, adjusted for sample attrition and classification error (see appendix B); they are plotted against the right axis. Each time series is MA smoothed. Gray bands indicate NBER recession periods.

from employment into unemployment (1.98 percent, see table 1). Second, flows from employment into inactivity look almost acyclical. Table 1 in the next subsection confirms this visual impression by showing that this is the least cyclical of the six time series of worker flows between employment, unemployment and inactivity. Finally, the flows from unemployment into inactivity average 27.98 percent over the years 1967 to 2007. They exhibit trends that appear similar to those of the employment-to-inactivity time series. More strikingly, these flows are strongly countercyclical: during recessions, unemployed workers become less likely to drop from the labor force. This feature seemed exacerbated by the recent recession too, as documented in more details by Rothstein (2011), Hotchkiss et al. (2012) and Elsby et al. (2013a).

2.2. **The ins and outs and the cycle.** To provide a more systematic picture of fluctuations in the ins and outs of the labor force, this section computes a set of statistics that would inform conventional business cycle analysis of the labor market. To do so, it uses a series for productivity, namely the seasonally adjusted quarterly series of real output per person in the non-farm business sector (PRS85006163) of the Bureau of Labor Statistics (BLS). This and the other time series in this section are then taken in log as deviations from their Hodrick-Prescott trend estimated with a smoothing parameter of 10^5. The statistics characterizing the behavior of the ins and outs of the labor force over the business cycle are displayed in table 1. In addition, the table reports the corresponding statistics for the other labor market flows to compare them to the ins and outs of the labor force.

Looking first at the standard deviation of the different series (relative to the standard deviation of productivity), table 1 allows to rank the components of the ins and outs of the labor force with respect to their volatility. The outs from employment (E-I) reveal the least volatile, the outs from unemployment (U-I) are the most volatile; the ins fall in between in terms of volatility and turn out to have close standard deviations. When compared to the two other worker flows, what is striking in table 1 is that movements out of unemployment into inactivity (U-I) are almost as volatile as those in the opposite direction, i.e. towards employment (U-E). This pattern does not hold for transitions out of employment since the corresponding flows into inactivity (E-I) are almost twice less volatile than flows into unemployment (E-U).

Correlations with labor productivity confirms the impression conveyed by the gray bands indicating recession periods in figure 1. That is, movements from employment into inactivity (E-I) are not very responsive to the business cycle whereas all other components of the ins and outs of the labor force co-move substantially with productivity. The probability to leave the labor force from unemployment (U-I) exhibits the largest co-movements: it drops during recessions with a magnitude similar to that of the probability to regain employment from unemployment (U-E). The two components of the ins (I-E and I-U) also have large correlations with productivity and they co-move in opposite directions. Expansions are thus accompanied by a burst in movements from inactivity into employment and a decline of close magnitude in movements from inactivity into unemployment. The different time series are not very persistent, with the exception of the flows out of unemployment (U-I and U-E): for these two series the auto-correlation is as large as that of labor productivity.

Finally, table 1 reports the behavior of labor market stocks (the unemployment pool and its sum with the employment pool) to contrast them with the underlying flows. These stocks correspond to the published series of the BLS, not to the stocks that are consistent with the labor market flows reported in table 1.[7] The columns for labor market stocks reproduces the well-known fact that the size of the unemployment pool is highly volatile, persistent and counter-cyclical. They also show that labor force participation is procyclical, but only mildly: the correlation with labor productivity is below one fourth. This is striking given what the rest of the table documents: large fluctuations in the ins and outs of the labor force maintain the size of the labor force almost constant over the cycle.

[7]These time series are obtained via the seasonally adjusted monthly series for labor force participation rate (LNS11300000) and of the unemployment rate (LNS14000000). These time series are aggregated to quarterly frequency by averaging over the three monthly values. I use the published series instead of the labor market stocks implied by the set of stock-flow equations because the latter exhibit too many erratic jumps that affect their measured auto-correlations.

Table 1. Labor Market Stocks and Flows: Summary Statistics

	Flows						Stocks	
	E-U	E-I	U-E	U-I	I-E	I-U	U	E+U
Average across cycle	1.98	2.99	32.11	27.98	4.33	3.43	3.87	64.50
Standard deviation relative to that of P	4.771	2.495	6.247	5.271	3.318	4.090	8.710	0.285
Correlation with P	-0.826	0.336	0.782	0.786	0.642	-0.701	-0.873	0.240
Auto-correlation	0.596	0.356	0.871	0.835	0.458	0.659	0.949	0.877

NOTE: E: employment; U: unemployment; I: inactivity. Worker flows are the time series computed by Robert Shimer. Worker stocks are the time series from the BLS. Averages are reported in percent (first row). The other statistics (standard deviation, auto-correlation and correlation with labor productivity) are computed after taking each time series in log as deviation from an HP trend with smoothing parameter 10^5. P is the HP-filtered seasonally adjusted quarterly series of real output per person in the non-farm business sector from the BLS. The standard deviation of P over the corresponding period is 0.018 and the auto-correlation is 0.885.

2.3. **Towards the model.** Is a job-search model likely to rationalize the facts summarized in the previous subsections? The similar volatilities and co-movements of opposite signs in the ins of the labor force suggests that the answer is "yes": they point to an underlying job-finding rate that allocates new labor market entrants to employment and with complementary probability to unemployment based on aggregate conditions. However, the procyclicality of the unemployment-to-inactivity probability seems difficult to reconcile with central tenets of job-search theory: indeed, why would workers choose to stick more to the labor market when the job-finding rate is below trend?

The next sections will give a job-search model a chance to meet this challenge. Because it is unusual for this class of models to allow for movements between three states (employment, unemployment and inactivity), the framework requires some premises that are discussed in this subsection.

One fact that proves difficult to rationalize is the positive flow from inactivity to employment: if a worker is nonemployed this period and not seeking a job, how is it that he/she can move to employment directly in the next period? A popular explanation for this is time-aggregation, i.e. the idea that inactivity-to-employment transitions involve an intervening spell of unemployment which is not recorded in the data due to its brevity. However, the fact that I-E flows dominate I-U flows requires time-aggregation to be a severe measurement problem if it is to account for all transitions from inactivity to employment. The job-search model developed hereafter will instead proceed under the milder assumption that time-aggregation is only one piece of the story.

The other piece of explanation that will be provided is that inactive workers search for jobs too, but less actively/efficiently than their unemployed counterparts. The preferred interpretation for this is that they are not classified as unemployed because their search efforts are sufficiently close to zero for not being detected by the statistical agency – a possibility discussed by Garibaldi and Wasmer (2005). This distinction between inactivity and unemployment does not contradict the finding that those are "behaviorally distinct labor force states" (Flinn and Heckman, 1983) and is not inconsistent with the official definition of unemployment (i.e. actively searching for a job). Furthermore, as noted by Jones and Riddell (1999), some job-searchers are appropriately classified as inactive because they resort only to passive search methods. There is yet another possibility for the observed transitions from inactivity to employment: it is that "jobs bump into people" in the words of Garibaldi and Wasmer (2005). The model will also allow for this interpretation. That is, job-search theory contends that work opportunities result from search efforts, which implies that search effort *on average* across inactive workers should be positive but does not rule out that some of them receive a job offer without exerting *any* search effort.

3. A Simple Job-Search Model

This section presents the job-search model used to decode worker flows throughout the rest of the paper. The model can be viewed as the Mortensen and Pissarides (1994)'s framework modified so as to account for transitions not only between employment and unemployment but also in and out of the labor force. Unlike them, however, the model does not seek to endogenize labor demand.[8] Instead, the emphasis is put on the supply side that makes the model a useful companion to analyze the ins and outs of the labor force.

3.1. **Economic Environment.**

Individuals. The economy is populated by a continuum of infinitely-lived workers of total mass equal to one. Workers are risk-neutral and they maximize

$$\mathbb{E}_0 \sum_{t=0}^{+\infty} \beta^t \left[c_t - (\pi_E n_{E,t} + \pi_U n_{U,t})\right] \tag{1}$$

where c_t denotes consumption and $n_{i,t}$ with $i \in \{E,U\}$ is labor supply along the extensive margin. Workers can be in three mutually exclusive states: employed (E), unemployed (U) or inactive (I), and $n_{E,t}$ (resp. $n_{U,t}$) takes the value of one if the worker is employed (resp. unemployed) and zero otherwise. π_E (resp. π_U) measures the opportunity cost of participating in the labor market as an employed (resp. unemployed) person. One can think of these costs as foregone home production, the utility of which is normalized to zero.

Job-search. While nonemployed, workers search for jobs. At the aggregate level, the probability to receive a work opportunity is given by a constant $\lambda > 0$. At the individual level, this translates into per-period probabilities to get a job offer as a function of search effort:

$$\lambda_i = \lambda s_i, i \in \{U,I\} \tag{2}$$

The probabilities in (2) aim at capturing the fact that inactive workers search for jobs (as discussed in subsection 2.3), but with a different search intensity than their unemployed counterparts. Intuitively, given workers' preferences as defined by (1), setting $s_U > s_I$ and $\pi_U > 0$ implies than an unemployed

[8]Accommodating the model to allow for endogenous labor demand is easily done by picking a job creation cost and a matching function that deliver the value of the job-finding rate consistent with the equilibrium of this economy. With the usual Cobb-Douglas matching function, these modifications would imply calibrating three additional parameters, but this would not add any value to the results for the steady-state version of the model.

worker trades a higher probability to regain employment in the future against foregone home production in the current period. Proposition 2 below establishes that this trade-off is what effectively determines workers' decision to participate in the labor market.

Productivity. Each individual worker is endowed with an idiosyncratic productivity level denoted by x. Productivity x evolves over time according to a first-order autoregressive process

$$x_{t+1} = (1-\rho)\bar{x} + \rho x_t + \varepsilon_{t+1} \tag{3}$$

where ρ belongs to the $(0,1)$ interval and $\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$. Hereafter, $F(.|x)$ denotes the transition function for x, i.e. $F(x'|x) = \Pr\{x_{t+1} < x'|x_t = x\}$. Changes in productivity levels are supposed to reflect shocks that alter workers' return to market activity. Disability shocks for instance would be one extreme form of these shocks. ρ and σ_ε are thought as parameters that can be estimated from panel data on wages.

Wages. In order to connect the model to the data, one needs to specify how the stochastic process for x can be inferred from an econometrician's point of view. This is done by assuming that work opportunities in the model reflect actual jobs. A job yields a wage $\omega(x)$ to a worker whose current productivity is x and is terminated exogenously with probability δ per period. Following much of the literature, the following assumption is made in order to determine wages:

> **Assumption A1:** Wages are determined by a Nash bargain between workers and fictitious employers operating under a free-entry condition.

In assumption (A1), the word "fictitious" is meant to recall that the model does not seek to endogenize labor demand. Meanwhile, wages are equivalent to what would be obtained in general equilibrium with employers who would receive $x - \omega(x)$ in the current period and have an outside option of zero when bargaining with workers. The Nash bargaining game is fully described by assigning a bargaining power of $\gamma > 0$ to workers.

3.2. **Recursive Formulation.** Workers' decision problems can be formulated in recursive form. The asset values of being employed, unemployed and inactive are denoted by $E(.)$, $U(.)$ and $I(.)$, respectively. For expositional purposes, it is also useful to define the asset value of being nonemployed denoted by $N(.)$. It is given by:

$$N(x) = \max\{U(x), I(x)\} \tag{4}$$

The list $(I(.), U(.), E(.))$ is then determined by a system of Bellman equations:

$$I(x) = \beta \int \left(\lambda_I \max\left\{E(x'), N(x')\right\} + (1-\lambda_I) N(x')\right) dF(x'|x) \tag{5}$$

$$U(x) = -\pi_U + \beta \int \left(\lambda_U \max\left\{E(x'), N(x')\right\} + (1-\lambda_U) N(x')\right) dF(x'|x) \tag{6}$$

$$E(x) = \omega(x) - \pi_E + \beta \int \left((1-\delta) \max\left\{E(x'), N(x')\right\} + \delta N(x')\right) dF(x'|x) \tag{7}$$

This recursive representation of workers' problem is convenient because it highlights the two economic decisions of interest. Equation (4) encloses the choice of being unemployed vs. being inactive in the current period. In equations (5)–(7), the maximum operator corresponds to the decision of working vs. not working.

3.3. **Equilibrium.** The model is in partial equilibrium. In what follows, the focus will be primarily on the distribution of workers across employment, unemployment and inactivity. This distribution stems from exogenous labor market frictions but also from endogenous decisions. Two results map these decisions into the partial equilibrium of the economy under the following assumption:

Assumption A2: Labor market frictions are moderate in the sense that $\delta + \gamma\lambda_i < 1$, $i \in \{U, I\}$

Assumption (A2) is essentially a restriction on the range of parameter values. It squares with the evidence on worker flows reported in section 2 and is sufficient to make workers' decisions well-behaved in the model. Specifically, one can show that:

Proposition 1. *Under assumptions (A1) and (A2), the list* $(I(.), U(.), E(.))$ *exists and is unique.*

Proof. See appendix A. □

The proof of proposition 1 is instrumental in computing the partial equilibrium of the model because it ensures that standard solution algorithms (e.g. value function iteration) solve the system of Bellman equations (5)–(7). One is interested in this solution because of proposition 2:

Proposition 2. *There exists a unique threshold productivity* x_P^* *(resp.* x_W^**) above which workers decide to participate in the labor market (resp. to work). These threshold levels are the solution to:*

$$U(x_P^*) = I(x_P^*) \tag{8}$$

$$E(x_W^*) = N(x_W^*) \tag{9}$$

Proof. See appendix A. □

The interpretation of equations (8) and (9) is straightforward: they define reservation policies for labor force participation and work decisions, respectively. By plugging equations (5) and (6) into (8), one can also see that x_P^* solves: $\pi_U = \beta(\lambda_U - \lambda_I)\int \max\{E(x') - N(x'), 0\}\, dF(x'|x)$. That is, x_P^* equates the opportunity cost of unemployment with the net returns to being unemployed in the form of future employment prospects.

It shall be emphasized that proposition 2 states that x_P^* and x_W^* are well-defined but does not rank them. In principle, both $x_P^* < x_W^*$ and $x_P^* > x_W^*$ are possible. If $x_P^* < x_W^*$, then employed workers whose productivity switches to $x \in [x_P^*, x_W^*]$ separate from their job into unemployment. On the other hand if $x_P^* > x_W^*$, then worker who endogenously terminate their jobs flow directly into inactivity. These simple insights will help calibrate the model, which favors the case $x_P^* > x_W^*$.

Making use of the thresholds in proposition 2, it is straightforward to write the law of motion for the distribution of workers across employment, unemployment and inactivity as well as across productivity levels. Further assuming that productivity levels are distributed over a finite, discrete support, the model economy is Markovian and the corresponding law of motion has an ergodic distribution. In what follows, I refer to this time-invariant distribution as the equilibrium of the model.

4. CALIBRATION

In order to proceed further with the job-search model, this section lays out a calibration strategy. Then it characterizes the resulting steady-state equilibrium.

Calibration strategy. The calibration targets are measured at (at least) a monthly frequency, which implies that the model needs to operate at a frequency higher than a month to allow for time-aggregation. A period is thus set to be one week (more precisely: one forty-eighth of a year) and the model-generated data are then aggregated appropriately when compared to the calibration targets. The discount factor β is set to 0.9992, consistent with an annualized interest rate of 4%.

The model allows for a few normalizations. Firstly, implicit in workers' preferences defined by (1) is that the utility of home production is zero: this makes $\pi_E n_{E,t} + \pi_U n_{U,t}$ measure the opportunity cost of participating in the labor market. Second, λ_U and λ_I are both proportional to λ through (2); λ can thus be normalized to one. Finally because ρ and σ_ε are the only parameters of interest for the productivity process (3), one can set $\bar{x}$ equal to one.

The calibration strategy is then in two steps: (i) parametrize the productivity process and (ii) calibrate the remaining parameters to match selected moments from the data.

As for the first step, it draws on Chang and Kim (2007)'s estimates based on data from the Panel Study of Income Dynamics. The authors estimate a wage-earning process that accounts for selection in employment and show how estimates of the underlying productivity process follow from first differencing this wage process. They report an annual persistence component of 0.809 and a standard deviation of 0.348 for innovation after correcting for sample selection bias. Adjusted to the model period[9], this yields $\rho = 0.9956$ and $\sigma_\varepsilon = 0.0925$. Given $\bar{x}$, ρ and σ_ε the transition function $F(.)$ can then be computed using standard techniques for discretization.[10]

Before moving to the second step, a caveat is in order regarding an auxiliary parameter that relates to wages: γ, the bargaining power of workers. This does not connect readily to observable data, and a transparent strategy is thus to set it exogenously to some popular value in the literature. Hence the choice of $\gamma = 0.70$: this is the elasticity of the job-filling rate with respect to labor market tightness that is most frequently used in calibrated version of the job-search model, and it is also the value of γ under the Hosios-Pissarides condition.

The remaining parameters to be calibrated are: δ, s_U, s_I, π_E, π_U. As for the calibration targets, there is a total of six moments that one can use from the data, namely the average across periods of the six time series of worker flows. Because these averages imply steady-state values for employment, unemployment and inactivity, one may also choose to replace a calibration target for flows by a calibration target for stocks. Hereafter the strategy draws on information for stocks and flows: it targets the E-U, U-E, I-E transition probabilities and the E and U stocks. The rest of the subsection explains how each of the remaining parameters is linked to the calibration targets.

---

[9] The weekly values ρ and σ_ε are related to the annual parameters ρ_a and σ_a estimated by Chang and Kim (2007) through: $\rho = \rho_a^{1/48}$ and $\sigma_a = \sigma_\varepsilon \frac{1}{48} \sqrt{\sum_{j=1}^{48} \left(\sum_{k=0}^{48-j} \rho^k \right)^2}$.

[10] In the computations, Tauchen (1986)'s method is used to approximate the autoregressive process. Productivity is discretized by means of a grid with 500 evenly-spaced points.

To account for worker flows, it is useful to return to proposition 2. Given that the parametrized productivity process renders x highly persistent and that all jobs are alike, one does not expect workers to quit their current job in order to search actively for a new one. Therefore, separations that are followed by a transition into unemployment should result from exogenous job destruction whereas endogenous separations should be followed by a transition to inactivity: this is the case $x_P^* > x_W^*$ in proposition 2. A direct implication is that δ can be disciplined by the transition probability from employment into unemployment. Another consequence is that the ratio of the U-E transition probability to the I-E transition probability is equal to the ratio of s_U to s_I. s_I can thus be calibrated to match the I-E flow in level and s_U be adjusted to reproduce the U-E flow.

Finally, as for the opportunity cost of market activity, the following observations pin down π_E and π_U. First, x_W^* depends on the surplus value $E-N$, which in turn depends directly on π_E. Second, the participation margin x_P^* is given by π_U once the relative search intensity $s_U - s_I$ has been fixed (equation (8)). π_E and π_U can thus be disciplined to match the employment and unemployment stocks. The table below reports the parameter values obtained via this calibration strategy:

β	π_E	π_U	δ	s_U	s_I	ρ_x	σ_x	γ
0.9992	0.70	0.45	0.007	0.156	0.043	0.996	0.093	0.70

Steady-state. Table 2 shows worker flows in the model at a weekly frequency and also aggregated at a monthly frequency and compares the latter with the data. The fit to the data is remarkably good: flows out of employment and out of inactivity are almost identical to their empirical counterparts. The flow from unemployment to employment (U-E) is also close to the empirical value by virtue of the calibration strategy. Finally the quantitative performance is somewhat lower for the U-I flow (which is a free moment from the data) but remains satisfactory. For instance Krusell et al. (2011, 2012) also find it more difficult to account for this flow, but their model-generated moment is farther away from the empirical value of the U-I flow than it is in table 2.

Comparing the weekly flows and the monthly flows is also instructive. At a weekly frequency, the I-E flow is lower than the I-U flow; the opposite is true after aggregation to a monthly frequency. This is because the weekly U-E probability is relatively large: it increases the likelihood for a worker who joins the labor force into unemployment to transit into employment, which may eventually be recorded as a I-E transition at a monthly frequency. The fact that the weekly I-E and I-U flows are close also implies that, according to the model, time-aggregation is only one piece of explanation behind I-E transitions. This is in line with the discussion in subsection 2.3.

Some final remarks are in order regarding the effects of idiosyncratic productivity shocks. Intuitively, the magnitude of these shocks drives exits towards inactivity. As for the calibration strategy, this suggests that σ_ε could have been included as one of the calibrated parameters to target either the E-I or the U-I flow. But given the model fit, the calibrated σ_ε would result close to the value used here. It should also be noticed that these shocks explain why the weekly employment-to-unemployment probability is lower than the job destruction probability δ. This is because among workers who are exogenously displaced from their jobs, a fraction also experience a negative productivity shock and decide to exit the labor force instead of searching actively for a job.

Table 2. Worker Flows in the Model and in the Data

	Model							Data (monthly)		
	Weekly			Monthly						
	E	U	I	E	U	I		E	U	I
E	98.82	0.51	0.67	94.64	1.95	3.41	E	95.03	1.98	2.99
U	8.39	86.21	5.39	31.45	48.36	20.19	U	32.11	39.90	27.98
I	0.74	1.02	98.23	4.23	3.78	92.00	I	4.33	3.43	92.24

NOTE: E: employment; U: unemployment; I: inactivity. Worker flows in the data are computed as the average over the period 1967-2007 of the time series computed by Robert Shimer and displayed in the first row of table 1.

5. Dynamic Implications

The good quantitative performance of the job-search model qualifies it as a useful construct to study worker flows. This section therefore expands on its implications for fluctuations in the ins and outs of the labor force. To do so, it proceeds in two steps. First, it analyzes a typical recession through the lens of the job-search model. A recession here is viewed as a fall in the job-finding rate and a burst in job destructions. These two changes are analyzed in isolation so as to single out the implications of the model. Second, since shocks to the job-finding and to the job destruction rates are not disconnected in the data, the model is then fed with time series that reproduces their empirical comovements (as observed in data for the U.S. labor market over the years 1967 to 2007). This results in model-generated time series for the ins and outs of the labor force than can be compared to their empirical counterparts.

5.1. **Recessions as transient shocks.** In this first series of numerical experiments, the model economy is hit by a multiple shock to either the aggregate job-finding rate λ or the job destruction rate δ. These parameters are shifted once and for all by 10 percent and, after the shift, they instantaneously return to their normal position. The response of worker flows are then traced for one year following the shock. Because such shocks have a probability of zero, they cannot be anticipated: workers' decision rules are thus unchanged throughout the experiments.

A decrease in the job-finding rate. Figure 2 plots the response of the inactivity-to-unemployment probability (upper graph) to a one-time decrease in the aggregate job-finding rate λ. It also shows the resulting change in the mean productivity level of the pool of unemployed workers (lower graph). Both time series are expressed in percentage point relative their steady-state value. Hereafter the qualitative and quantitative implications of the model are discussed in turn.

The upper graph in figure 2 illustrates that when the job-finding rate falls below trend, the model predicts that more inactive workers will be observed joining the labor force directly into unemployment. As indicated above, workers' decision rules remain fixed to their steady-state values in these experiments: the model therefore does not claim that the higher probability to move from inactivity to unemployment reflects the behavioral response of workers to bad economic times. Rather, the model suggests that if workers cannot fully delay their decision to return to the labor market, then the I-U flow will increase mechanically due to the fact that new entrants move less rapidly into employment relative to normal economic times – a form of composition effect.

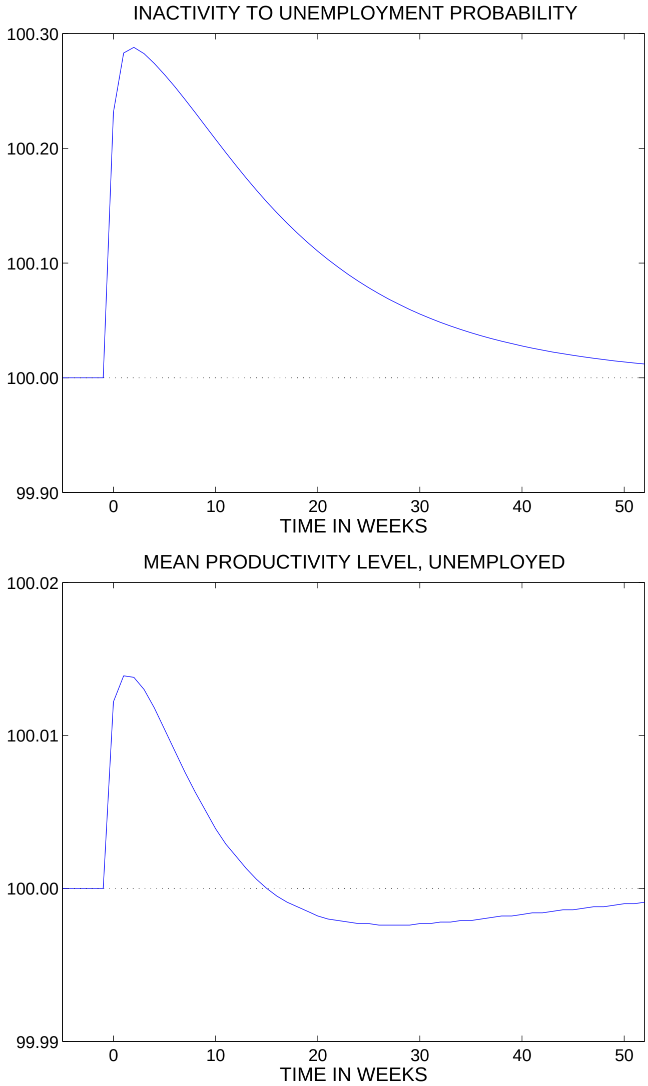

Figure 2. A one-time decrease in the job-finding rate: implications for the I-U flow
Each time series is in percentage point relative to the value in steady-state equilibrium

Another implication of the model is that, following a drop in the job-finding rate, the skill composition of the unemployment pool initially shifts towards high-productivity workers and then shifts in the opposite direction before returning to its normal position. This is the result of two effects. On the one hand, workers who are displaced from their jobs exogenously are more productive than the average unemployed: when the job-finding rate is unusually low, they grow more numerous in the unemployment pool and contribute to the upward shift. On the other hand, workers who enter the labor force are less productive than the average unemployed: as they flow into the unemployment pool, they shift the average productivity level downwards. Hence the pattern observed in the lower graph of figure 2.

Quantitatively, the effects described here are extremely modest. Although λ is shifted downwards by 10 percentage points in week 0, the I-U probability jumps by no more than 0.30 percentage point. Changes to the skill composition of the unemployment pool are also limited but the figures are more difficult to interpret since they are the sum of two countervailing forces. Meanwhile, to summarize the

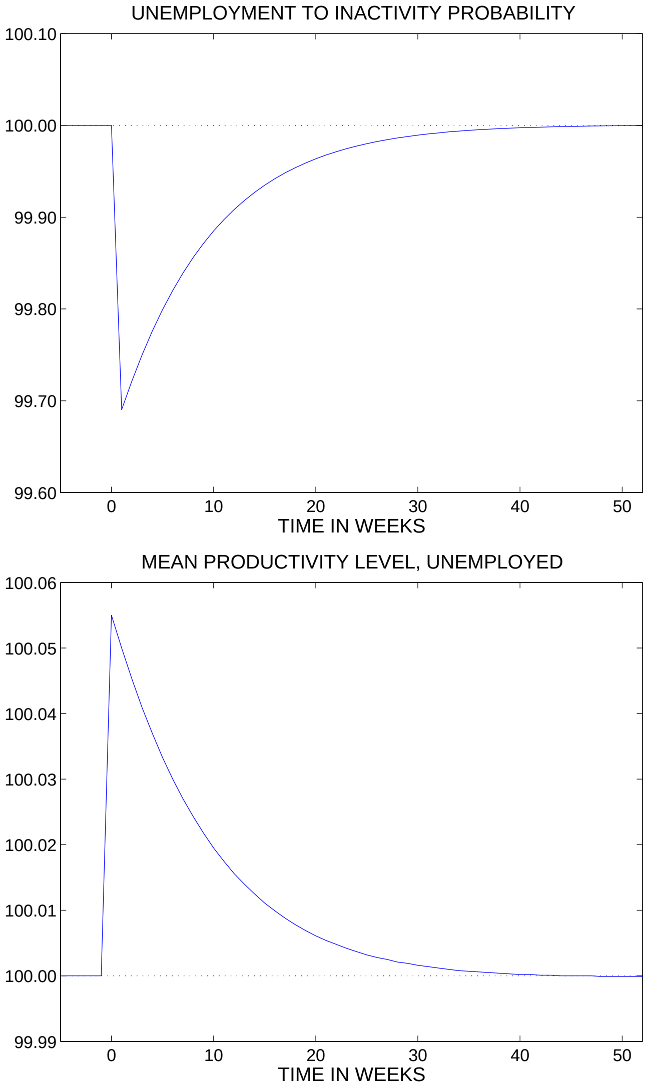

Figure 3. A one-time increase in the job destruction rate: implications for the U-I flow
Each time series is in percentage point relative to the value in steady-state equilibrium

model implications, it predicts mostly an increase in the skill level of the unemployment pool since the shift below the steady-state value after week 15 is tiny. This prediction is noteworthy because it may reinforce the effect of an increase in the job destruction rate analyzed in the next paragraphs.

An increase in the job destruction rate. Figure 3 is the analogon of figure 2: it depicts the model response to a one-time increase in the job-destruction rate δ. The upper graph plots the response of the unemployment-to-inactivity probability and the lower graph plots the change in the mean productivity level of the pool of unemployed. However, unlike the previous numerical experiments, shifts to the skill composition of the unemployment pool will prove to be a driver of changes in the observed worker flows rather than a consequence of them.

When δ becomes unusually high, the probability for the unemployed to leave the labor force falls below its value in normal economic times. It shall be recalled that in steady-state δ drives exits from employment into unemployment, not movements from unemployment into inactivity. Moreover, since workers decision rules remain unchanged throughout the experiment, the observed fall in the

U-I probability depicted in figure 3 does not represent a change in the probability to drop from the labor force at the individual level. Instead, the observed movement is entirely due to a composition effect: high job destruction rates draw an unusually large number of high-productivity workers into the unemployment pool and these workers happen to be less likely to drop from the labor force.

As in the previous numerical experiments, the effects turn out to be quantitatively limited: the U-I probability falls by roughly 0.30 percentage point as a result of the composition effect caused by a one-time jump of δ by 10 percentage points. Meanwhile, what is worth emphasizing here is that both experiments point to the skill composition of the unemployment pool becoming skewed towards high-productivity workers during recessions. As for the empirical analysis in section 6, this indicates that one expects individual heterogeneity to matter quantitatively for the flows from unemployment into inactivity.[11]

5.2. **Recessions as in the data.** The first series of numerical experiments allow the model to tell only part of its implications regarding fluctuations in the ins and outs of the labor force. First of all, the behaviors of the job-finding and job destruction rates differ markedly in turbulent economic times: early on in recessions, job destruction rates increase tremendously but they then return to their pre-recession levels rapidly. The job-finding rate on the other hand reacts more sluggishly initially but remains below trend for prolonged periods of time. Second, in the first numerical experiments changes to the job-finding and job destruction rates are considered in isolation whereas in the data these changes are inter-related. The previous numerical experiments thus do not allow the effects of recessions on the different margins of the model to reinforce each other. This motivates the second series of experiments which do more full justice to recessions as observed in the data.

This second series of experiments is performed by means of the data presented in section 2. The empirical time series of the unemployment-to-employment and employment-to-unemployment flows yield two series λ_t and δ_t covering the period 1967-2007. The latter can thus be used to feed the model and obtain model-generated time series for the ins and outs of the labor force.[12] This is done, again, under the assumption that workers' decisions remain unchanged relative to their steady-state values: this gives a good approximation of the model predictions regarding the ins and outs of the labor force over the past four decades.

[11]On a different note, the fact that both experiments predict a shift in the skill composition of the unemployment pool towards high-productivity workers suggests that endogenizing labor demand to conduct conventional real business cycle analysis would only reiterate the finding that the Mortensen-Pissarides cannot replicate movements in key labor market variables such as the vacancy to unemployment ratio. For instance, under the hypothesis that high-productivity workers yield larger profits to employers in discounted present value terms, a recession would be accompanied by a burst in vacancy posting, thereby increasing the job-finding rate and generating a positively-slopped Beveridge curve. A similar point was made recently by Mueller (2012).

[12]There are three intermediary steps to perform this experiment. First, the series λ_t and δ_t obtained from the data are quarterly series. Two weekly series are thus created by interpolating the quarterly series linearly between quarters. Second, the resulting weekly series correspond to monthly probabilities. They are converted to weekly values using the correspondence in table 2. Third, the start date $t = 0$ of the series may well correspond to an upturn or to a downturn of the U.S. business cycle rather than an intermediary period. This makes the steady-state distribution of workers across the different states of the model economy a poor distribution to initialize the experiment. To alleviate this problem, one can plug the starting value λ_0 and δ_0 into the model economy and iterate on the law of motion of the distribution until the economy comes close to the start date $t = 0$ with respect to the U-E and E-U flows.

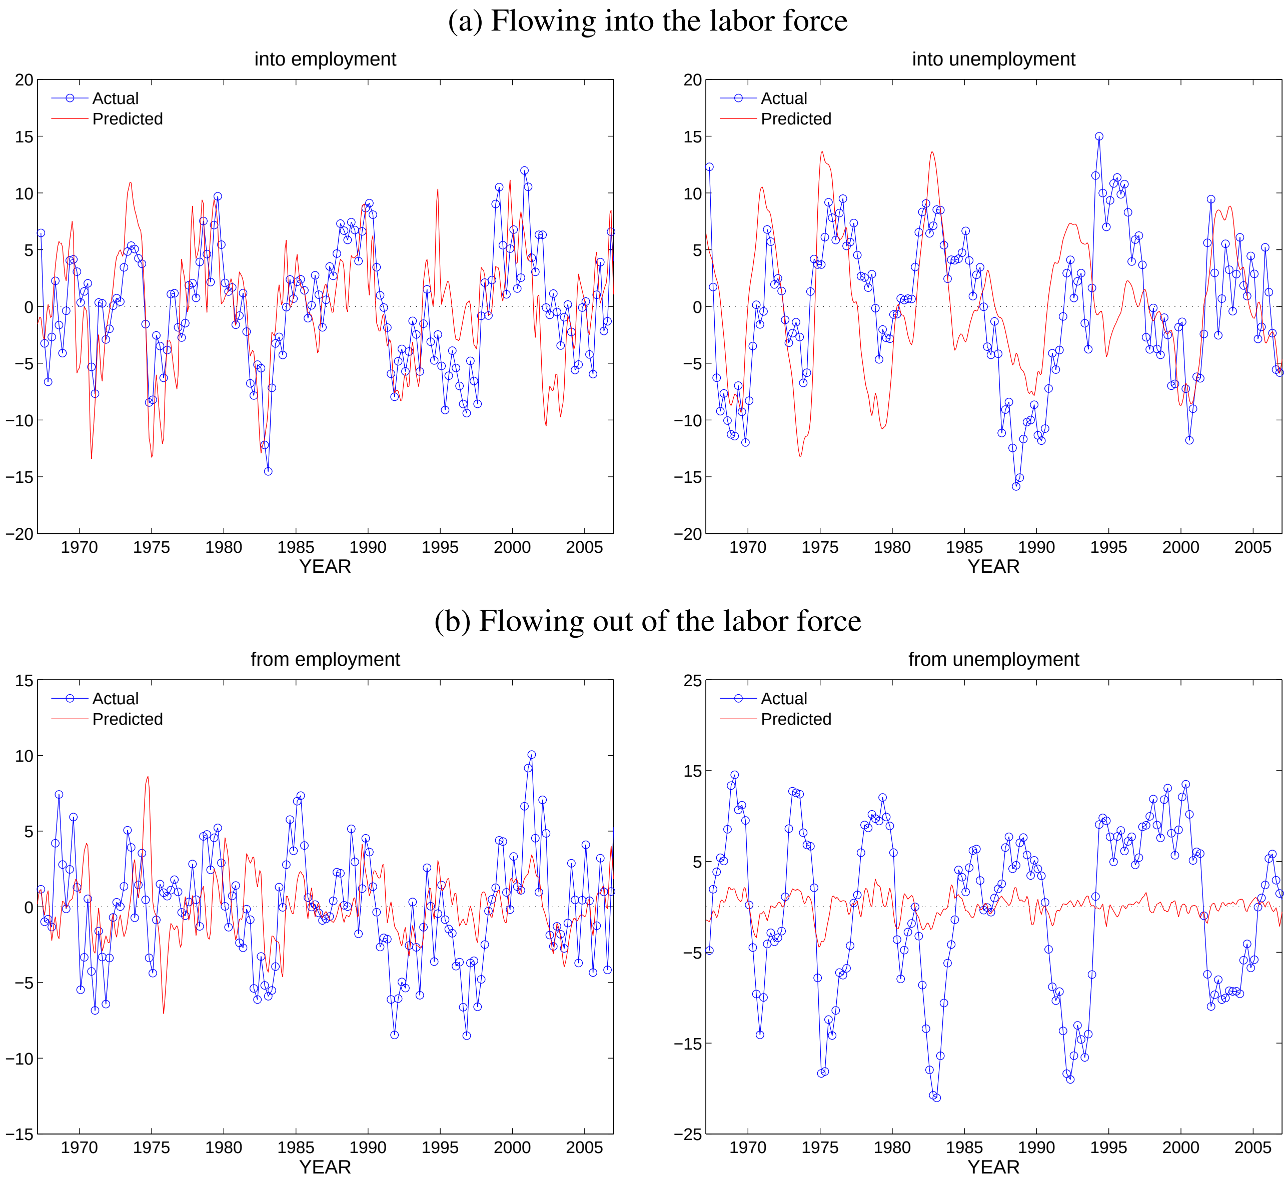

Figure 4. Fluctuations in the ins and outs of the labor force: model vs. data
The dots are from the monthly transition probabilities displayed in figure 1. The solid lines are from the model-generated flows aggregated at a monthly frequency. Each time series is in log as deviation from an HP trend with smoothing parameter 10^5. Each time series is then multiplied by 100. The time series for the data are MA smoothed.

Figure 4 plots the model-generated series of the ins and outs of the labor force and compares them to their empirical counterparts from figure 1. To emphasize fluctuations, these series are reported in log as deviations from their HP trend. The model-generated series are aggregated at a monthly frequency since the empirical time series correspond to monthly values (measured only every quarter).

Focusing first on fluctuations in the ins of the labor force (the upper graphs (a) in figure 4), the fit to the data looks remarkably good. As for the inactivity-to-employment flow, this is not surprising: since fluctuations for this time series are caused directly by $\lambda_t s_I$, they are only marginally a product of the numerical experiment. Meanwhile, it shall be recalled that the series λ_t is backed up from data on the U-E flow, not the I-E flow. This shows that cyclical changes in the U-E flow predict almost perfectly those in the I-E flow, which dovetails with the view that fluctuations in regaining employment are driven mostly by the aggregate job-finding rate.[13]

What is less expected in the upper graphs (a) is the good fit of the inactivity-to-unemployment flow. Like the I-E probability, these changes accrue from movements in λ_t (through the complementary probability $1 - \lambda_t s_I$). However, the model-generated series for this flow would not line up so closely to its empirical counterparts if workers' decision to participate in the labor force were to fluctuate over the business cycle too. In other words, assuming as in the numerical experiment that workers do not time their labor force decisions based on aggregate conditions holds as a good approximation for understanding the inactivity-to-unemployment flow. Subsection 6.2 in the empirical portion of the paper will substantiate this finding.

With respect to the outs of the labor force (the lower graphs (b) in figure 4), the fit is less satisfactory. Yet, the employment-to-inactivity flow lines up closer to its empirical counterpart than expected. The underlying reason is more difficult to fathom than for the other flows but seems to reflect a composition effect. That is, during good times the aggregate job finding rate is above trend and draws many inactive workers into employment. Since they are less productive than average, they drive the average productivity level in the employment pool downwards and thereby increase the probability to move from employment into inactivity. Hence the observed correlation.

Finally, as for the unemployment-to-inactivity flow, the numerical experiment underpredicts the volatility for this series substantially: it is about four times larger in the data. The timing of movements between the model-generated and empirical series is nevertheless good, at least for the first four recessions where one observes a decline in the U-I probability early on in the downturn. As for the underlying mechanism, the second numerical experiment from subsection 5.1 informs us that these movements result from a compositional shift among the unemployed. It also indicates that this effect cannot be large when compared to movements in the job destruction rate and this finding is corroborated here. In sum, the model suggests that there is room for the composition effect described before to account for the decline in the unemployment-to-inactivity flow during recessions. Subsection 6.3 will provide supporting empirical evidence for this mechanism.

6. EMPIRICAL RESULTS

The goal of this final section is to substantiate two implications of the job-search model regarding fluctuations in the ins and outs of the labor force. First, the model contends that fluctuations in the I-U flow reflect mostly changes in the probability to move into unemployment conditional on labor market entry, not changes in the decision to enter the labor force. Since movements in the I-E probability are the converse of those in I-U probability, focusing on the latter is sufficient to uncover the drivers of fluctuations in the ins of the labor force. Second, it predicts that worker heterogeneity (especially heterogeneity in skill levels) plays a non-negligible role in explaining fluctuations in the U-I probability. Since the model does not deliver strong predictions regarding movements in the E-I

[13]This, of course, would be a direct implication of time-aggregation if all movements from inactivity to employment were to involve an intervening spell of unemployment. This line of reasoning applies only partly here (see section 4).

probability, the empirical analysis of the outs of the labor force will be restricted to changes in the unemployment-to-inactivity flow.

6.1. **Data and sample.** The empirical analysis draws on microdata from the Current Population Survey. The CPS has informed the majority of studies on worker flows in the United States, which makes it a firsthand source of information for the purpose of this section. The data, sample disposition and methodology to measure transitions in and out of the labor force are only briefly described here since a detailed presentation is provided in appendix B.

This section uses the basic monthly files of the CPS for the period 1994-2012. Although the monthly files are available over a longer period of time, the methodology to recode worker transition relies on features of the CPS that were introduced only upon the 1994 re-design of the survey. The methodology further requires to restrict the sample to individuals who can be matched across four consecutive months of the survey. This is potentially an important restriction but individual survey weights can be adjusted to take sample attrition into account. Finally, the procedure to match individuals across surveys requires basic demographic information such as age, sex and race. A handful of individuals whose basic demographic information is missing are thus dropped from the sample. The analysis then focuses on individuals of working age (16-64 years old) with valid labor force status (employed, unemployed or inactive) in all four months when they are interviewed.

Workers' transitions as measured hereafter refer to changes in labor force status that may occur between their second and third month of interview. To minimize measurement error in recorded transitions, a rich set of information from the first to fourth month of interview is exploited. First and foremost, the methodology reconciles the labor market status of nonemployed workers with the information provided with respect to their number of weeks spent out of employment. The latter is considered of greater reliability because it is obtained via a dependent-interview procedure in the CPS. The correction procedure also draws on information about the main reason for being currently unemployed or the primary activity when out of the labor force provided by these individuals. Finally, it discriminates between self-reported labor force statuses vs. labor force statuses obtained via proxy respondents upon examining sequences of three consecutive labor force status. For example the sequence "I-U-I" (the individual is recorded as being inactive, then unemployed and then inactive again) is recoded as "I-I-I" when information in the first and third months were self-reported but was obtained via proxy respondents in the second month. Appendix B explain the difference between this correction strategy and the one proposed by Elsby et al. (2013a).

The objective of the correction strategy is to obtain a dataset where one can study the determinants of workers' transitions at the individual level. The resulting dataset can also be used to build time series of transition probabilities in and out of the labor force. This is illustrated in figure 1 which plots the time series computed after implementing the correction strategy (dashed lines) along with the corresponding series by Shimer (2012) covering a longer period of time (solid lines). As discussed in appendix B, the main differences induced by the correction strategy are in the levels of flows from inactivity to employment (I-E) and from unemployment to inactivity (U-I). The cyclical behavior of all series however does not look different in the series of Robert Shimer and in the data used hereafter.

6.2. **Fluctuations in the ins of the labor force.** To analyze fluctuations in the ins of the labor force, the probability to observe one such transition into unemployment is decomposed as

$$\Pr\{\sigma_{t+1}=U|\sigma_t=I\}=\Pr\{\sigma_{t+1}=U|\sigma_{t+1}\neq I,\sigma_t=I\}\times\Pr\{\sigma_{t+1}\neq I|\sigma_t=I\} \tag{10}$$

where σ_t denotes labor force status at time t. $\Pr\{\sigma_{t+1}=U|\sigma_{t+1}\neq I,\sigma_t=I\}$ is the probability to move into unemployment rather than employment conditional on labor market entry. According to the job-search model, this probability reacts strongly to aggregate conditions. $\Pr\{\sigma_{t+1}\neq I|\sigma_t=I\}$ is the probability to join the labor force in the current period, which is predicted to be much less sensitive to aggregate labor market conditions. These predictions will be tested against the data by estimating the elasticity of each probability with respect to the unemployment rate.

The decomposition in equation (10) is sensible for at least two reasons. First, it is dictated by the high persistence of the inactivity state (e.g. table 1). Second, it squares with the way participation is measured in labor force surveys: individuals are asked whether they are available for work and would be willing to take a job. The job-search model also makes sense of these probabilities: $\Pr\{\sigma_{t+1}\neq I|\sigma_t=I\}$ corresponds to an economic choice at the individual level (the participation threshold x_P^*) whereas $\Pr\{\sigma_{t+1}=U|\sigma_{t+1}\neq I,\sigma_t=I\}$ is governed by labor market frictions (the probability $1-\lambda_I$) taken as given by the worker.

In order to bring the probabilities in (10) to the data, the following model is postulated:

$$\Pr\left\{\sigma_{i,s,t+1}\neq I|\sigma_{i,s,t}=I;X_{i,t},u_{s,t}\right\}=\Phi\left(X_{i,t}\vartheta+\log\left(u_{s,t}\right)\eta\right) \tag{11}$$

In equation (11), $\Phi(.)$ is the c.d.f. of the Normal distribution. $X_{i,t}$ is a vector of dummy variables for sex, age (with individuals grouped into five age brackets of equal length), race (white; black; hispanic or other), marital status (married; divorced, separated or widowed; never married) and educational level (high school dropouts; high school graduates; some college; college and higher education). In addition to demographic controls, a set of dummies is also included for the main reason for being out of the labor force (disabled or ill; in school; household care; retired; other). Finally $\log(u_{s,t})$ is the log of the seasonally adjusted monthly series of the unemployment rate in State s at time t[14]; it is thought as a proxy for aggregate labor market conditions. A similar model is postulated for $\Pr\left\{\sigma_{i,s,t+1}=U|\sigma_{i,s,t+1}\neq I,\sigma_{i,s,t}=I;X_{i,t},u_{s,t}\right\}$.[15]

Equation (11) and its counterpart for the probability to move into unemployment rather than employment following labor market entry are estimated via maximum likelihood[16] where the (log of the) State-level unemployment rate is instrumented with the (log of the) national unemployment rate. The elasticities of interest are then evaluated at a given vector $\left(\overline{X}_{i,t},\overline{u}_{s,t}\right)$. As a benchmark, this vector corresponds to the mean of all variables across sample periods. In order to illustrate the effect of the business cycle, the elasticities are also computed using averages over the six-month period preceding

[14]These time series are obtained from the BLS webpage for local area unemployment statistics and correspond to the series labeled LASSTxx000003, where xx is the numeric code for each State. See http://www.bls.gov/schedule/archives/laus_nr.htm.

[15]Linear models for equation (11) estimated via 2-stage least squares regression yielded similar results.

[16]The estimates reported in this section are obtained via the Broyden-Fletcher-Goldfarb-Shannon method. Other quasi-Newton method for maximizing the log-likelihood gave similar results.

Table 3. Maximum Likelihood Estimates of Elasticities with Respect to the Unemployment Rate for the Ins of the Labor Force

	Elasticity of $\Pr\{\sigma_{t+1} \neq I \vert \sigma_t = I\}$				
	(I)	(II)	(III)	(IV)	(V)
All periods	-0.158	-0.163	-0.168	-0.203	-0.194
	(0.014)	(0.014)	(0.014)	(0.015)	(0.015)
Pre-recession	-0.157	-0.161	-0.166	-0.200	-0.192
	(0.014)	(0.014)	(0.014)	(0.014)	(0.015)
Post-recession	-0.160	-0.164	-0.170	-0.205	-0.197
	(0.014)	(0.015)	(0.015)	(0.015)	(0.016)
Observations	839,112	839,112	839,112	839,112	826,171
Log-likelihood	-159,139	-148,670	-141,740	-140,542	-125,710

	Elasticity of $\Pr\{\sigma_{t+1} = U \vert \sigma_{t+1} \neq I, \sigma_t = I\}$				
	(I)	(II)	(III)	(IV)	(V)
All periods	0.517	0.545	0.529	0.563	0.574
	(0.021)	(0.022)	(0.022)	(0.021)	(0.022)
Pre-recession	0.563	0.593	0.576	0.618	0.629
	(0.025)	(0.025)	(0.026)	(0.026)	(0.026)
Post-recession	0.475	0.499	0.485	0.514	0.525
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Observations	77,427	77,427	77,427	77,427	75,620
Log-likelihood	-40,669	-40,326	-38,735	-38,384	-37,554

NOTE: Top panel: elasticity of the probability to join the labor force w.r.t. unemployment. Bottom panel: elasticity of the probability to move into unemployment following labor force entry w.r.t. unemployment. "All periods" refers to the 1994-2012 period; "Pre-recession" refers to the 6 months preceding the 2001 and 2007-2009 recessions; "Post-recession" refers to the 6 months following the 2001 and 2007-2009 recessions. (I) does not control for demographics; (II) controls for sex and age; (III) controls for sex, age, race and marital status; (IV) controls for sex, age, race, marital status and educational level; (V) controls for sex, age, race, marital status, educational level, and main reason for being out of the labor force. Standard errors in parentheses are clustered at the individual level in the Probit models and are then computed via the delta method.

the recessions covered by the sample (i.e. the 2001 and 2007-2009 downturns) and the six-month periods that followed. Table 3 reports the elasticities estimated for the two components in the right-hand side of equation (10).[17]

Focusing first of the top panel of table 3, the elasticity of the probability to join the labor force with respect to aggregate labor market conditions turns out to be rather low (around 0.15-0.20 in absolute values). This lines up well with the assumption in section 5 that aggregate conditions play only a secondary role in workers' choice to join the labor force. These elasticities are negative, which

[17]Standard errors in the Probit models are clustered at the individual level. This is because individuals in the CPS may complete up to two series of four consecutive interviews. When this occurs, they appear twice in the dataset if they also meet the sample criteria. See appendix B for details about the rotating structure of the CPS.

also suggests that workers are more reluctant to join the labor market during downturns. Meanwhile, moving along the rows of the table, one finds no significant differences between elasticities estimated at different points of the business cycle. Differences are more pronounced when moving along the columns of the table (i.e. after introducing controls for individual characteristics), but there again the estimated elasticities remain small in magnitude.

Conversely, the bottom panel of table 3 shows positive elasticities for the probability that new entrants find themselves unemployed rather than employed and those are almost four times larger than the elasticities in the upper panel. Differences within a given column point to the business cycle playing a more significant role in allocating new entrants to the unemployment pool: quantitatively, this becomes as important as controlling for individual characteristics. These findings thus support the view from the numerical experiments of section 5 that the aggregate job-finding probability drives the bulk of fluctuations in the probability to regain employment. This dovetails with much of the literature that finds that business cycle fluctuations in entries into employment are governed predominantly by an aggregate job-finding rate (e.g. Elsby et al., 2010; Shimer, 2012).

As a robustness check for the findings presented in table 3, the elasticities for the ins of the labor force are also computed for workers with different degree of labor force attachment. Specifically, the sample is broken down into two subsamples: one that comprises individuals who were in the labor force in the previous year and the other that comprises those who were already out of the labor force.[18] The estimation procedure described before is then repeated with these two subsamples. The results are displayed in table 4.

Table 4 first confirms that the probability to join the labor force is not very sensitive to aggregate labor market conditions. Indeed, for individuals with low and high labor force attachment, the elasticities are similar when evaluated at different points of the business cycle. Another finding that lines up well with the experiments of section 5 is that aggregate conditions matter less to understand the labor force decisions of individuals who resemble more the workers in the job-search model, i.e. those who were previously in the labor force. For these individuals, the elasticities are never higher than 0.16 in absolute values. Finally, individual characteristics play a more significant role for those who were previously out of the labor force (the elasticities almost double when moving along the columns of table 4), although the estimates are less precise.

Turning to the probability to move into unemployment conditional on labor market entry, the elasticities in the bottom panel of table 4 reveal remarkably similar for individuals with low and high labor force attachment. In fact, there are no significant difference across samples. Furthermore, controlling for individual characteristics has a negligible impact on the estimates, unlike in the top panel. For instance among individuals with low labor force attachment, the elasticities do not increase by more than 0.03 points when moving from column I to column V whereas they almost double in the upper panel of the table. Altogether, the estimates in tables 3 and 4 provide strong support for the first composition effect uncovered by the job-search model.

[18] As explained in footnote 17, respondents in the CPS may complete up to two series of four consecutive interviews. Those are the observations used for the computations in table 4. Labor force attachment is measured by a dummy variable that takes the value of one if the individual was in the labor force in his/her second interview and zero otherwise.

Table 4. Maximum Likelihood Estimates of Elasticities with Respect to the Unemployment Rate for the Ins of the Labor Force: Workers with Different Degree of Labor Force Attachment

	Elasticity of $\Pr\{\sigma_{t+1} \neq I \mid \sigma_t = I\}$									
	Previously in the labor force					Previously out of the labor force				
	(I)	(II)	(III)	(IV)	(V)	(I)	(II)	(III)	(IV)	(V)
All periods	-0.113	-0.112	-0.113	-0.145	-0.158	-0.108	-0.180	-0.185	-0.210	-0.187
	(0.026)	(0.026)	(0.026)	(0.027)	(0.028)	(0.037)	(0.038)	(0.039)	(0.039)	(0.040)
Pre-recession	-0.112	-0.111	-0.112	-0.143	-0.156	-0.107	-0.178	-0.183	-0.208	-0.185
	(0.025)	(0.026)	(0.026)	(0.026)	(0.027)	(0.036)	(0.038)	(0.038)	(0.038)	(0.040)
Post-recession	-0.114	-0.112	-0.114	-0.146	-0.160	-0.108	-0.181	-0.186	-0.212	-0.188
	(0.026)	(0.027)	(0.027)	(0.027)	(0.029)	(0.037)	(0.039)	(0.039)	(0.039)	(0.041)
Observations	206,727	206,727	206,727	206,727	202,626	199,459	199,459	199,459	199,459	197,677
Log-likelihood	-58,671	-57,215	-55,245	-54,713	-50,319	-12,258	-9,376	-7,996	-7,883	-6,151

	Elasticity of $\Pr\{\sigma_{t+1} = U \mid \sigma_{t+1} \neq I, \sigma_t = I\}$									
	Previously in the labor force					Previously out of the labor force				
	(I)	(II)	(III)	(IV)	(V)	(I)	(II)	(III)	(IV)	(V)
All periods	0.526	0.564	0.548	0.583	0.597	0.566	0.572	0.553	0.577	0.590
	(0.037)	(0.039)	(0.039)	(0.038)	(0.040)	(0.052)	(0.053)	(0.053)	(0.053)	(0.054)
Pre-recession	0.572	0.614	0.596	0.640	0.655	0.625	0.630	0.609	0.640	0.655
	(0.044)	(0.046)	(0.046)	(0.046)	(0.047)	(0.063)	(0.064)	(0.064)	(0.064)	(0.065)
Post-recession	0.485	0.517	0.503	0.532	0.547	0.512	0.518	0.500	0.520	0.532
	(0.031)	(0.032)	(0.033)	(0.032)	(0.033)	(0.042)	(0.043)	(0.043)	(0.043)	(0.043)
Observations	27,751	27,751	27,751	27,751	27,145	9,206	9,206	9,206	9,206	9,002
Log-likelihood	-14,386	-14,243	-13,731	-13,595	-13,272	-5,015	-4,982	-4,728	-4,695	-4,603

NOTE: Top panel: elasticity of the probability to join the labor force w.r.t. unemployment. Bottom panel: elasticity of the probability to move into unemployment following labor force entry w.r.t. unemployment. "All periods" refers to the 1994-2012 period; "Pre-recession" refers to the 6 months preceding the 2001 and 2007-2009 recessions; "Post-recession" refers to the 6 months following the 2001 and 2007-2009 recessions. (I) does not control for demographics; (II) controls for sex and age; (III) controls for sex, age, race and marital status; (IV) controls for sex, age, race, marital status and educational level; (V) controls for sex, age, race, marital status, educational level, and main reason for being out of the labor force. Robust standard errors in parentheses are computed via the delta method.

6.3. **Fluctuations in the outs of the labor force.**

Empirical approach. As indicated in the opening paragraph of this section, the study of the outs of the labor force can be limited to fluctuations in the U-I probability. As for the term $\Pr\{\sigma_{t+1} = I | \sigma_t = U\}$, breaking it into two components in the fashion of equation (10) does not appear as a sensible choice. First, this would imply estimating the probability of moving into inactivity rather than employment conditional on leaving the unemployment pool: according to the job-search model, almost no worker finds himself/herself in the position of making such a choice. Second, table 2 indicates that the probability to move into either one of the different labor market states is strongly balanced for those currently unemployed. This must be accommodated by the statistical model.

The previous remarks calls for allowing σ_{t+1} to take any value of the set $\{I, U, E\}$ when $\sigma_t = U$. The following polytomic model is thus postulated:

$$\{\sigma_{t+1} | \sigma_t = U\} = \arg \max_{\sigma \in \{I,U,E\}} \{\Upsilon^* (\sigma) | \sigma_t = U\} \tag{12}$$

with $\Upsilon^* (\sigma)$ the propensity of workers to be observed in state σ (see e.g. Cameron and Trivedi, 2005). In the spirit of equation (11), this latent variable is brought to the data by assuming that:

$$\Upsilon^*_{i,t+1} (\sigma) = X_{i,t} \theta_\sigma + \log (u_t) \phi_\sigma + \nu_{\sigma,i,t} \tag{13}$$

In equation (13), both θ and ϕ are allowed to be choice-specific. The error terms $\nu_{\sigma,i,t}$ are assumed to be distributed according to the multivariate Normal distribution with free covariance parameters.[19] The vector $X_{i,t}$ has the same set of demographic controls as in the previous subsection. In addition, it also includes dummy variables for the main reason for being currently unemployed (job loss; end of temporary job; job leaver; new or re-entrant into the labor force). Finally, $\log (u_t)$ is the log of the seasonally adjusted monthly series of the national unemployment rate at time t. This more aggregated series is used instead of the unemployment rate at the State-level because instrumental variable methods are not available for this class of models. Besides, the endogeneity of unemployment at the local level may be more of a serious concern when analyzing transitions out of the unemployment pool.[20]

The model specified by equations (12)–(13) allows to estimate the elasticity of the unemployment-to-employment probability (U-E) and that of the unemployment-to-inactivity probability (U-I) with respect to aggregate conditions. These two elasticities are of interest to test the relevance of the second composition effect uncovered by the job-search model. Indeed, as emphasized in section 2, the U-E and U-I probabilities exhibit cyclical behaviors that are strikingly similar in the aggregate. Informed by the dynamic implications of the model, one instead expect to find a significantly lower elasticity for the U-I probability than for the U-E probability when estimated via equations (12)–(13). This shall be true especially after controlling for individual characteristics, particularly if those proxy for workers' idiosyncratic productivity levels.

[19] A multinomial Probit model is preferred over multinomial Logit because of its higher flexibility. The small number of alternatives for σ also makes it attractive from an estimation standpoint.

[20] To be sure, the specification in equation (13) alleviates concerns regarding endogeneity issues, but I shall also mention that those do not seem to matter in quantitative terms. Replacing the series of the national unemployment rate with that of unemployment at the State level delivered virtually the same coefficients (and more precise estimates).

Table 5. Maximum Likelihood Estimates of Elasticities with Respect to the Unemployment Rate for the Outs of the Labor Force

	Elasticity of $\Pr\{\sigma_{t+1}=E\|\sigma_t=U\}$				
	(I)	(II)	(III)	(IV)	(V)
All periods	-0.792	-0.775	-0.781	-0.804	-0.815
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Pre-recession	-0.690	-0.676	-0.682	-0.704	-0.701
	(0.014)	(0.014)	(0.015)	(0.015)	(0.014)
Post-recession	-0.869	-0.849	-0.854	-0.877	-0.897
	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)
	Elasticity of $\Pr\{\sigma_{t+1}=I\|\sigma_t=U\}$				
	(I)	(II)	(III)	(IV)	(V)
All periods	-0.436	-0.377	-0.394	-0.347	-0.260
	(0.021)	(0.021)	(0.021)	(0.021)	(0.022)
Pre-recession	-0.356	-0.301	-0.317	-0.272	-0.176
	(0.019)	(0.020)	(0.020)	(0.020)	(0.021)
Post-recession	-0.493	-0.433	-0.449	-0.401	-0.323
	(0.022)	(0.023)	(0.023)	(0.023)	(0.024)
Observations	174,351	174,351	174,351	174,351	147,829
Log-likelihood	-150,663	-149,029	-148,119	-147,648	-132,131

NOTE: Top panel: elasticity of the probability to move from unemployment to employment w.r.t. unemployment. Bottom panel: elasticity of the probability to move from unemployment to inactivity w.r.t. unemployment. "All periods" refers to the 1994-2012 period; "Pre-recession" refers to the 6 months preceding the 2001 and 2007-2009 recessions; "Post-recession" refers to the 6 months following the 2001 and 2007-2009 recessions. (I) does not control for demographics; (II) controls for sex and age; (III) controls for sex, age, race and marital status; (IV) controls for sex, age, race, marital status and educational level; (V) controls for sex, age, race, marital status, educational level, and main reason for being currently unemployed. Standard errors in parentheses are clustered at the individual level in the multinomial Probit models and are then computed via the delta method.

Results. Table 5 reports the estimates of the elasticity for the U-E and U-I probabilities obtained with various set of controls and at different points in the business cycle. The results are readily summarized: at any point in the cycle, the elasticity of the U-I probability is substantially lower than that of the U-E probability. For instance it is up to four times lower when controlling for individual characteristics (column V of the table). Another finding in line with the numerical experiments of section 5 is that the attenuation effect of demographic controls is significant for the U-I probabilities (the estimates are cut by almost 50 percent when moving along the columns of the table), but not for the U-E probability. For the latter, the idea that all workers face the same aggregate job-finding rate seems, again, to hold as a good description of labor market dynamics. In light of the discussion of the elasticities uncovered by equations (12)–(13), table 5 validates that the similar cyclical behaviors of the U-I and U-E probabilities mask substantial composition effects. Those are key for explaining the observed fluctuations in the flow from unemployment to inactivity.

Doe the attenuation effect depicted in table 5 reflect shifts in the skill composition of the unemployment pool? To address this question and bring the empirical analysis of the outs of the labor force further closer to the job-search model, one would need to control for workers' idiosyncratic productivity levels. A natural proxy for this is workers' wage in their previous job. In the CPS, this information is available in the fourth month of interview; thus, it can be used to analyze the outs of the labor force among respondents who completed exactly two series of four consecutive interviews.[21] Specifically, what the rest of this section analyzes are the elasticities computed on this subsample of the CPS after introducing wages in equation (13). Two versions of wages are studied in isolation: the log of usual weekly earnings and the log-residual obtained from Mincerian regressions.[22] The results of these computations are presented in table 6.[23]

Looking first at column (I) in table 6, one notices that restricting the analysis to individuals with non-missing wage data induces some sample selection: the elasticities for both the U-E and U-I probabilities reveal significantly lower than in table 5.[24] With this difference in levels, the results provided in this table then confirm that controlling for idiosyncratic productivity levels dampens the responsiveness to aggregate conditions for the U-I probability more than for the U-E probability. Indeed, moving along the columns of the table (with either the log-wage or the log-residual wage), the elasticity for the probability to leave the labor force from unemployment becomes essentially zero whereas it remains significantly larger in absolute values and negative for transitions to employment. And although the estimates in the middle panel are not significant at conventional significance levels, they move in the expected direction as further controls for workers' skill levels are included in the model. In sum, the table confirms that shifts in the heterogeneous characteristics of the unemployed matter for fluctuations in the outs of the labor force. This concurs with Elsby et al. (2013a) who used a different methodology to measure the impact of various observable characteristics on the cyclicality of the unemployment-to-inactivity probability.

Finally, one further result from table 6 is that high-productivity individuals have higher labor force attachment as measured by their probability to drop from unemployment into inactivity. This is noteworthy because Mueller (2012) documents a compositional shift of the unemployment pool towards these workers during downturns. Altogether, these findings provide an account for the cyclicality of the outs of the labor force (from unemployment) that resorts only to the composition effect argument.

---

[21] The approach adopted here is hence to control for the wage observed in the fourth month of interview to analyze transitions occurring during the second series of four consecutive interviews. It shall be noted that information on wages is also available for those in their eighth month of interview (in the language of the CPS, the wage variables are available for outgoing rotation groups). I choose to use lagged wages instead to avoid potential selection effects.

[22] The wage variables are constructed as follows. Firstly, usual weekly earnings are trimmed by removing the bottom and top 1 percent of workers as well as those who work less than 10 hours or more than 80 hours per week on a regular basis. The log of the trimmed variable is the wage used in table 6. Second, based on this variable, OLS regressions are run within each cross section of the CPS. The covariates in these regressions include a quartic polynomial of age, dummy variables for sex, race, marital status and educational level, eight occupation dummies and finally dummy variables for those working part-time and for government workers and self-employed workers. The residual predicted by these OLS regressions is the residual wage used in table 6.

[23] To save on space, the elasticities are evaluated only at the average of the covariates across all periods but not at different points in the cycle. Patterns similar to those in table 5 emerge when computing the elasticities before and after recessions.

[24] In addition to the restriction to respondents in their second series of four consecutive interviews, the analysis is necessarily limited to workers who were in employment in their fourth month of interview.

Table 6. Maximum Likelihood Estimates of Elasticities with Respect to the Unemployment Rate for the Outs of the Labor Force: Workers with Different Skill Levels

	Elasticity of $\Pr\{\sigma_{t+1}=E\|\sigma_t=U\}$					
	(I)	(II)	(III)	(IV)	(V)	(VI)
Unemployment	-0.191	-0.191	-0.182	-0.182	-0.186	-0.165
	(0.024)	(0.024)	(0.024)	(0.024)	(0.025)	(0.023)
Wage		0.052	0.076	0.072	0.068	0.053
		(0.008)	(0.009)	(0.010)	(0.010)	(0.010)
Unemployment	-0.195	-0.195	-0.191	-0.191	-0.199	-0.175
	(0.024)	(0.024)	(0.024)	(0.024)	(0.025)	(0.023)
Residual wage		0.045	0.042	0.048	0.054	0.039
		(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
	Elasticity of $\Pr\{\sigma_{t+1}=I\|\sigma_t=U\}$					
	(I)	(II)	(III)	(IV)	(V)	(VI)
Unemployment	-0.122	-0.123	-0.083	-0.097	-0.060	0.014
	(0.091)	(0.093)	(0.094)	(0.093)	(0.094)	(0.094)
Wage		-0.437	-0.328	-0.311	-0.271	-0.222
		(0.034)	(0.037)	(0.038)	(0.039)	(0.040)
Unemployment	-0.117	-0.119	-0.058	-0.076	-0.022	0.048
	(0.091)	(0.092)	(0.094)	(0.093)	(0.093)	(0.094)
Residual wage		-0.220	-0.274	-0.291	-0.313	-0.271
		(0.042)	(0.043)	(0.043)	(0.044)	(0.045)
Model: Equation (13) + *wage*						
Observations	21,532	21,532	21,532	21,532	21,532	19,407
Log-likelihood	-19,126	-19,032	-18,968	-18,914	-18,901	-16,650
Model: Equation (13) + *residual wage*						
Observations	21,467	21,467	21,467	21,467	21,467	19,344
Log-likelihood	-19,071	-19,053	-18,948	-18,883	-18,853	-16,601

NOTE: Top panel: elasticity of the probability to move from unemployment to employment. Middle panel: elasticity of the probability to move from unemployment to inactivity. Bottom panel: Estimation statistics. "Unemployment" refers to the elasticity with respect to unemployment; "Wage" (resp. "Residual wage") refers to the elasticity with respect to usual weekly earnings in the previous job (resp. the residual of usual weekly earnings). (I) does not control for individual characteristics but is estimated on the sample with non-missing wage data. (II)–(VI) control for log-wages. In addition: (II) does not control for demographics; (III) controls for sex and age; (IV) controls for sex, age, race and marital status; (V) controls for sex, age, race, marital status and educational level; (VI) controls for sex, age, race, marital status, educational level, and main reason for being currently unemployed. Robust standard errors in parentheses are computed via the delta method.

7. Conclusion

This paper advances the literature on worker flows between employment, unemployment and inactivity. First, it shows that a simple job-search model with heterogeneity in skill levels that connects to the data goes a long way in accounting for idiosyncratic worker trajectories between the three labor market states. Second, it points to two mechanisms that help explain the observed movements in the ins and outs of the labor force over the business cycle: (i) fluctuations in the economy-wide job-finding rate along with aggregate conditions playing only a secondary role in workers' labor force decisions and (ii) compositional shifts in the unemployment pool towards high-productivity workers during downturns. Third, it provides empirical evidence based on microdata from the Current Population Survey in support of these two composition effects.

This conclusion suggests a number of avenues for future work. First, the idea that workers do not time their labor force decisions based on aggregate conditions may be restrictive for certain subgroups of the population. For instance on both ends of the spectrum, school-to-work transitions and early retirement decisions are likely to be responsive to the business cycle. Analyzing those in details would provide a more complete picture of fluctuations in the ins of the labor force. Second, the composition effects emphasized in the paper account for only part of the fluctuations in the outs of the labor force. Other determinants thus need to be uncovered. One such likely contributor is the incidence of marginally attached and discouraged workers. This was not investigated in this paper and future work should dwell on its quantitative implications, notably in relation with the recent recession (see the BLS periodic publications about measures of labor underutilization, and also Elsby et al. (2010) and Hotchkiss et al. (2012) for suggestive empirical evidence). Third, Elsby et al. (2013a) and the present paper both draw on microdata from the Current Population Survey to uncover mechanisms behind individual movements into and out of the labor force. Another avenue for future research is hence to take advantage of richer datasets (such as the Survey of Income and Program Participation) to substantiate the findings regarding the determinants of these transitions.

From a more policy-oriented point of view, the model developed in this paper would be a relevant point of departure to investigate the effect of policy tools that may affect several margins of the labor market at once. One such example are unemployment benefits and their implications over the business cycle. For instance Fujita (2011) and Rothstein (2011) suggest that part of the persistence of the high unemployment rates during the recent recession can be attributed to automatic extensions of UI benefits triggered at the State level that retained workers in the labor force. On the other hand, recent research that analyzes whether and how these benefits should be conditioned on aggregate conditions (e.g. Landais et al. 2010; Mitman and Rabinovich 2011) consistently ignores the participation margin. A natural extension of these analyzes would thus draw on the model developed in this paper.

Finally, while this paper has focused attention on short-term fluctuations in the ins and outs of the labor force, a number of patterns worth analyzing emerge when looking at these fluctuations over the medium-long run, over the life-cycle or across countries. By virtue of its simplicity and good quantitative performance, the job-search model developed here may prove a useful companion to investigate these interesting topics.

REFERENCES

Alvarez, F. and Veracierto, M. (1999). Labor-market policies in an equilibrium search model. *NBER Macroeconomics Annual*, pages 265–304.

Andolfatto, D. and Gomme, P. (1996). Unemployment insurance and labor-market activity in Canada. *Carnegie-Rochester Conference Series on Public Policy*, 44(1):47–82.

Baker, M. (1992). Unemployment duration: Compositional effects and cyclical variability. *American Economic Review*, 82(1):313–21.

Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge university press.

Chang, Y. and Kim, S. (2007). Heterogeneity and aggregation: Implications for labor-market fluctuations. *American Economic Review*, 97(5):1939–1956.

Darby, M., Haltiwanger, J., and Plant, M. (1986). The ins and outs of unemployment: the ins win. *NBER Working Paper No. 1997*.

Elsby, M. W., Hobijn, B., and Sahin, A. (2010). The labor market in the Great Recession. *Brookings Papers on Economic Activity*.

Elsby, M. W., Hobijn, B., and Sahin, A. (2013a). On the importance of the participation margin for labor market fluctuations. *Federal Reserve Bank of San Francisco Working Paper*.

Elsby, M. W., Hobijn, B., and Sahin, A. (2013b). Unemployment dynamics in the OECD. *Review of Economics and Statistics*, 95(2).

Flinn, C. and Heckman, J. (1983). Are unemployment and out of the labor force behaviorally distinct labor force states? *Journal of Labor Economics*, 1(1):28–42.

Fujita, S. (2011). Effects of extended unemployment insurance benefits: evidence from the monthly cps. *Federal Reserve Bank of Philadelphia Working Paper*.

Garibaldi, P. and Wasmer, E. (2005). Equilibrium search unemployment, endogenous participation, and labor market flows. *Journal of the European Economic Association*, 3(4):851–882.

Hotchkiss, J., Pitts, M., and Rios-Avila, F. (2012). A closer look at nonparticipants during and after the Great Recession. *Federal Reserve Bank of Atlanta Working Paper*.

Jones, S. and Riddell, W. (1999). The measurement of unemployment: an empirical approach. *Econometrica*, 67(1):147–162.

Krusell, P., Mukoyama, T., Rogerson, R., and Şahin, A. (2011). A three state model of worker flows in general equilibrium. *Journal of Economic Theory*, 146(3):1107–1133.

Krusell, P., Mukoyama, T., Rogerson, R., and Şahin, A. (2012). Is labor supply important for business cycles? *NBER Working Paper*.

Landais, C., Michaillat, P., and Saez, E. (2010). Optimal unemployment insurance over the business cycle. *NBER Working Paper*.

Madrian, B. C. and Lefgren, L. J. (2000). An approach to longitudinally matching Current Population Survey (CPS) respondents. *Journal of Economic and Social Measurement*, 26(1):31–62.

Mankart, J. and Oikonomou, R. (2011). Household search and the aggregate labor market. *London School of Economics Working Paper*.

Mitman, K. and Rabinovich, S. (2011). Pro-cyclical unemployment benefits? optimal policy in an equilibrium business cycle model. *PIER Working Paper*.

Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *Review of Economic Studies*, 61(3):397–415.

Mueller, A. (2012). Separations, sorting and cyclical unemployment. *IZA Discussion Paper No. 6849*.

Petrongolo, B. and Pissarides, C. A. (2008). The ins and outs of european unemployment. *American Economic Review*, 98(2):256–262.

Poterba, J. and Summers, L. H. (1986). Reporting errors and labor market dynamics. *Econometrica*, 54(6):1319–38.

Pries, M. and Rogerson, R. (2009). Search frictions and labor market participation. *European Economic Review*, 53(5):568–587.

Rothstein, J. (2011). Unemployment insurance and job search in the Great Recession. *Brookings Papers on Economic Activity*.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2):127–148.

Shimer, R. (forthcoming, 2013). Job search, labor force participation, and wage rigidities. *Advances in Economics and Econometrics, Tenth World Congress (edited by D. Acemoglu, M. Arellano and Eddie Dekel)*.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics letters*, 20(2):177–181.

Tripier, F. (2004). Can the labor market search model explain the fluctuations of allocations of time? *Economic Modelling*, 21(1):131–146.

Veracierto, M. (2008). On the cyclical behavior of employment, unemployment and labor force participation. *Journal of Monetary Economics*, 55(6):1143–1157.

Villena-Roldan, B., Janiak, A., and Choi, S. (2011). Unemployment, participation and worker flows over the life cycle. *Society for Economic Dynamics 2011 Meeting Papers*.

APPENDIX A. MODEL APPENDIX

In order to take advantage of assumption (A1), it is convenient to formulate the problem of a fictitious employer matched to a worker with productivity x and with an outside option of zero. The asset value of employment for this employer is

$$J(x) = x - \omega(x) + \beta \int (1-\delta) \max\left\{J\left(x'\right), 0\right\} dF\left(x'|x\right) \tag{14}$$

Denoting by $S(.)$ the total surplus of a job to a worker-employer pair, i.e. $S(.) = J(.) + E(.) - N(.)$, the Nash bargaining assumption implies that

$$E(x) - N(x) = \gamma S(x) \text{ and } J(x) = (1-\gamma) S(x) \tag{15}$$

Combining this surplus sharing rule with equations (7) and (14) yields

$$S(x) = x - \pi_E - N(x) + \beta \int \left(N\left(x'\right) + (1-\delta) \max\left\{S\left(x'\right), 0\right\}\right) dF\left(x'|x\right) \tag{16}$$

where

$$\begin{aligned} N(x) \;=\; \max\Big\{ & -\pi_U + \beta \int \left(N\left(x'\right) + \lambda_U \max\left\{\gamma S\left(x'\right), 0\right\}\right) dF\left(x'|x\right), \\ & \beta \int \left(N\left(x'\right) + \lambda_I \max\left\{\gamma S\left(x'\right), 0\right\}\right) dF\left(x'|x\right) \Big\} \end{aligned} \tag{17}$$

after combining equations (4), (5) and (6). Plugging this last equation into (16) makes $S(.)$ the fixed point of the mapping Γ defined by:

$$\begin{aligned} \Gamma(S)(x) \;=\; x - \pi_E + \min\Big\{ & \pi_U + \beta(1-\delta-\lambda_U\gamma) \int \max\left\{S\left(x'\right), 0\right\} dF\left(x'|x\right), \\ & \beta(1-\delta-\lambda_I\gamma) \int \max\left\{S\left(x'\right), 0\right\} dF\left(x'|x\right) \Big\} \end{aligned} \tag{18}$$

The next step is thus to demonstrate that Γ has a unique fixed point by proving that $\Gamma(.)$ is a contraction mapping. First, it is fairly straightforward to show that equation (18) defines a monotone map. First-order stochastic dominance for F owing to equation (3) implies that, for any $S_1(.)$ and $S_2(.)$ such that $S_2(x) \geq S_1(x)$ for all x, we have

$$\int \max\left\{S_2\left(x'\right), 0\right\} dF\left(x'|x\right) \geq \int \max\left\{S_1\left(x'\right), 0\right\} dF\left(x'|x\right) \tag{19}$$

Under assumption (A2), i.e. $1-\delta-\lambda_U\gamma \geq 0$ and $1-\delta-\lambda_I\gamma \geq 0$, it follows that

$$\Gamma(S_2)(x) \geq \Gamma(S_1)(x) \tag{20}$$

Second, one can show that $\Gamma(.)$ discounts and has modulus strictly lower than one. Indeed, for any positive constant k,

$$\begin{aligned} \Gamma(S+k)(x) \;=\; x - \pi_E + \min\Big\{ & \pi_U + \beta(1-\delta-\lambda_U\gamma) \int \max\left\{S\left(x'\right)+k, 0\right\} dF\left(x'|x\right), \\ & \beta(1-\delta-\lambda_I\gamma) \int \max\left\{S\left(x'\right)+k, 0\right\} dF\left(x'|x\right) \Big\} \end{aligned} \tag{21}$$

Because

$$\int \max\left\{S\left(x'\right)+k,0\right\} dF\left(x'|x\right) \leq \int \max\left\{S\left(x'\right)+k,k\right\} dF\left(x'|x\right) = \int \max\left\{S\left(x'\right),0\right\} dF\left(x'|x\right)+k \tag{22}$$

the following inequality holds

$$\begin{aligned} \min\Big\{ & \pi_U+\beta\left(1-\delta-\lambda_U\gamma\right)\int \max\left\{S\left(x'\right)+k,0\right\} dF\left(x'|x\right), \\ & \beta\left(1-\delta-\lambda_I\gamma\right)\int \max\left\{S\left(x'\right)+k,0\right\} dF\left(x'|x\right)\Big\} \\ \leq \min\Big\{ & \pi_U+\beta\left(1-\delta-\lambda_U\gamma\right)\int \max\left\{S\left(x'\right),0\right\} dF\left(x'|x\right), \\ & \beta\left(1-\delta-\lambda_I\gamma\right)\int \max\left\{S\left(x'\right),0\right\} dF\left(x'|x\right)\Big\} \\ & +\beta\max\left\{1-\delta-\lambda_U\gamma,1-\delta-\lambda_I\gamma\right\}k \end{aligned} \tag{23}$$

In turns, this gives

$$\Gamma\left(S+k\right)\left(x\right) \leq \Gamma\left(S\right)\left(x\right)+\beta\max\left\{1-\delta-\lambda_U\gamma,1-\delta-\lambda_I\gamma\right\}k \tag{24}$$

Thus, $\Gamma(.)$ discounts with modulus $\beta\max\{1-\delta-\lambda_U\gamma,1-\delta-\lambda_I\gamma\}<1$. This completes the proof that $\Gamma(.)$ is a contraction mapping. Therefore, the surplus function $S(.)$ exists and is unique, and so does $N(.)$, the fixed point to the contraction mapping in (17). Given that $E(.)=\gamma S(.)+N(.)$, existence and unicity carries over for $E(.)$ and then for $I(.)$ and $U(.)$ through equations (5) and (6), respectively. This demonstrates proposition 1.

A few additional remarks establish proposition 2. First, since F satisfies the Feller property, $\Gamma(.)$ maps the set of bounded, continuous and increasing functions onto itself. $S(.)$ is thus bounded, continuous and monotonously increasing with its argument. Provided there exist some $\underline{x}$ and $\overline{x}$ such that $S(\underline{x})<0$ and $S(\overline{x})>0$, respectively, the surplus sharing rule $E(x)-N(x)=\gamma S(x)$ ensures that a unique x_W^* solves $E(x_W^*)-N(x_W^*)=0$ (equation (9), proposition 2).

Finally for the participation decision, it is defined through equation (17) by

$$P\left(x\right)=\beta\left(\lambda_U-\lambda_I\right)\int \max\left\{E\left(x'\right)-N\left(x'\right),0\right\} dF\left(x'|x\right)-\pi_U \tag{25}$$

where workers decide to join the labor force when $P(x)\geq 0$. The Feller property for F, the properties of $S(.)$ aforementioned (and thereby of $E(.)-N(.)$) and the assumption that $\lambda_U>\lambda_I$ (which holds true in the calibration) then imply that $P(.)$ is bounded, continuous and monotonously increasing. Provided there exist some $\underline{x}$ and $\overline{x}$ such that $P(\underline{x})<0$ and $P(\overline{x})>0$, respectively, a unique x_P^* solves $P(x_P^*)=0$, which is equivalent to equation (8) in proposition 2.

To clarify the calibration of the model, it is useful to write the law of motion for the measures of workers in employment, unemployment and inactivity denotes by $\mu_E(.)$, $\mu_U(.)$ and $\mu_I(.)$, respectively. As discussed in the text, the calibration implies that $x_E^*<x_P^*$. Letting $X_E=\{x\in X/x<x_E^*\}$ and $X_P=\{x\in X/x<x_P^*\}$ and denoting by $\overline{X}_i$ their complementary sets, i.e. $\overline{X}_i\cup X_i=X$ and $\overline{X}_i\cap X_i=\varnothing$

with $i \in \{E,P\}$, the law of motion writes

$$\mu_E\left(\overline{X}_E\right)' = \int_{\overline{X}_E}\int_X \left((1-\delta)\,\mu_E(x) + \lambda_U\mu_U(x) + \lambda_I\mu_I(x)\right) dF\left(x'|x\right) dx' \tag{26}$$

$$\mu_U\left(\overline{X}_P\right)' = \int_{\overline{X}_P}\int_X \left(\delta\mu_E(x) + (1-\lambda_U)\,\mu_U(x) + (1-\lambda_I)\,\mu_I(x)\right) dF\left(x'|x\right) dx' \tag{27}$$

$$\begin{aligned}\mu_I(X_P)' = &\int_{X_P\cap\overline{X}_E}\int_X \left(\delta\mu_E(x) + (1-\lambda_U)\,\mu_U(x) + (1-\lambda_I)\,\mu_I(x)\right) dF\left(x'|x\right) dx' \\ &+ \int_{X_E}\int_X \left(\mu_E(x) + \mu_U(x) + \mu_I(x)\right) dF\left(x'|x\right) dx'\end{aligned} \tag{28}$$

and $\mu_E(X_E)' = 0$, $\mu_U(X_P)' = 0$, $\mu_I\left(\overline{X}_P\right)' = 0$. In equations (26)–(28), the symbol $'$ indicates one period forward. The time-invariant distribution of workers is obtained by iterating on the discretized versions of these equations until pointwise convergence occurs.

Appendix B. Data Appendix

B.1. **Data.** The data used in section 6 come from the basic monthly files of the Current Population Survey made available at the NBER webpage (http://www.nber.org/data/cps_basic.html). The CPS is particularly well-suited to study workers' transitions in the labor market over the business cycle because it consists of short panels of very large size and is administered at a monthly frequency. Over the past decades, it has become the dominant source of information to document worker flows in the United States.

The CPS has a rotating panel structure that is exploited upon building the dataset for section 6. Individuals in the survey are interviewed for four consecutive months, remain out of the sample for eight months and are then interviewed again for four consecutive months. Respondents in rotation groups 1 to 3 and 5 to 7 (i.e. in their own first to third and then fifth to seventh month of interviews) can thus be matched one month forward so as to study potential changes in their labor market status in two consecutive months. However, not all individuals remain in the sample for all months of planned interviews. Sample attrition thus needs to be taken into account.

In 1994, the CPS underwent a significant re-design. Among the changes introduced, a dependent-interview procedure was introduce to measure the number of weeks that nonemployed individuals spent out of employment. It is well known that information obtained via such procedures is more reliable for the study of labor market trajectories at the individual level. For this reason, in section 6 the analysis is restricted to the 1994 onwards period so as to exploit the enhanced features of the CPS.

B.2. **Sample.** To measure labor force transitions in two consecutive months, individuals in the CPS are matched across surveys using household and person identifiers along with the age-sex-race filter prescribed by Madrian and Lefgren (2000). The resulting matching rates in two consecutive months are plotted in figure B1 separately for rotation groups 1 to 3 and 5 to 7. In both groups, the matching rates reveal fairly high relative to matching rates that are typically achieved with the CPS. As expected, they are slightly higher for rotation groups 5 to 7 since individuals in these groups are likely to be less mobile than individuals in their first four months of interviews. Finally, figure B1 suggests that

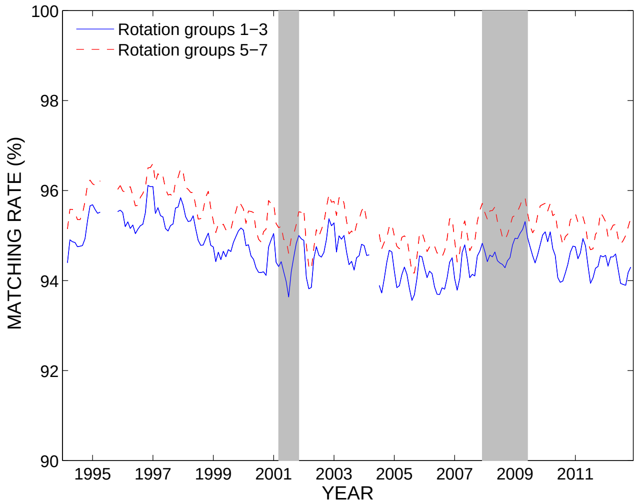

Figure B1. Matching rate in two consecutive months of the month files of the CPS The solid (resp. dashed) line shows the matching rate in rotation groups 1 to 3 (resp. 5 to 7). Gray bands indicate NBER recession periods.

matching rates are mildly countercyclical: one plausible explanation for this pattern is that lower residential mobility during downturns reduces sample attrition in the survey.

Matching individuals across two consecutive months is a first step in the sample disposition of the final dataset used in section 6. Indeed, to minimize measurement error in recorded transitions, the correction strategy will exploit the full sequence of labor force statuses observed in four consecutive months. The matching procedure thus needs to be repeated and the final dataset retains only individuals who were successfully matched in four consecutive months. Observe that individuals who completed eight months of interviews may thus appear twice in the final dataset.

B.3. **Adjustments.**

Sample attrition. To take sample attrition into account, individual observations in the final dataset are re-weighted using the following procedure. The cross-section comprising all individuals from the final dataset surveyed at date t[25] is merged with the corresponding monthly file of the CPS[26]. Defining an indicator that takes the value of one if the observation is in the final dataset and 0 otherwise, a Logistic regression is run against a quartic polynomial of age interacted with educational and race dummies. The model is estimated separately for male and female and for rotation groups 1 to 3 and 5 to 7. It is then used to predict the probability p^* that an individual is in the final dataset. The adjusted weight ω^* is finally obtained as $\omega^* = \frac{1-p^*}{p^*}\omega$ where ω is the original longitudinal weight from the CPS. The procedure is repeated for each cross-section t of the final dataset. These adjusted weights are used for establishing the series plotted in figure 1 (dashed lines) and also to compute averages for the covariates across all sample periods or before and after recessions upon evaluating the different elasticities in section 6.

[25]Date t corresponds to the date of the second interview since transitions will be measured from date t to date $t+1$.
[26]Individuals from the basic monthly file who do not meet the basic sample requirements (non-missing basic demographic information and aged 16 to 64 years-old) are dropped upon merging the files.

Classification error. Individual transitions are measured between the second and third month of interview but, eventually, they are established on the basis of the full sequence of labor force statuses from the first to fourth month of interview of the individual. This is done in an effort to minimize measurement error in recorded labor market transitions.

The first step of the correction strategy reconciles workers' labor force status with the (more reliable) information provided with respect to their number of weeks spent out of employment:

- For those unemployed in month t and who have been continuously unemployed for more than 5 weeks at this point, their labor force status in month $t-1$ is corrected to "unemployed" (U) when it is different. This filter is applied backwards, i.e. it recodes the labor force status in the third month, then in the second and finally in the first month.
- For those inactive in their fourth month of interview and who last worked at a job more than one year ago, their labor force statuses in the previous months are recoded as "inactive" (I) if they are disabled, ill or retired in the fourth month of interview or if they left their previous job for disability reasons or to return to school.[27] Otherwise, it is recoded as "unemployed".

The second step of the correction strategy exploits the fact that, in each month of interview, the CPS also indicates whether the labor force status of the individual was self-reported or obtained via proxy respondents.[28] Labor force status recorded via proxy respondents may generate a substantial amount of noise in workers' transitions. Hence the following filters:

- Whenever: (i) the labor force status σ in months $t-1$ and $t+1$ is self-reported and is the same in these two months, (ii) the labor force status σ' in month t is obtained via proxy respondents and (iii) the full sequence of labor force statuses is not of the form "$\sigma\sigma'\sigma\sigma'$" (for instance "IEIE"), then: σ is considered as more reliable. Thus σ' in month t is recoded as σ.
- Sequences of the form "$\sigma\sigma'\sigma\sigma'$" are more problematic. For those, the labor force status in month t of interview is recoded as the one of the adjacent months only when the interview in month t is the only one for which the labor force status of the individual was not self-reported.

The correction strategy implemented here differs from the one proposed by Elsby et al. (2013a). Theirs also draws on repeated (more than two) observations of the labor force status of CPS respondents but recodes systematically any sequence that is considered as "suspicious" (for instance "IIUI", "EIUI", etc.) without basing this choice on auxiliary information. As for the first step of the correction strategy, it extends the methodology implemented by Fujita (2011). This author analyzes transitions out of unemployment and thus only reconciles labor force status with information of the number of continuous weeks spent in unemployment.

[27] This second filter draws on two variables that are available only for outgoing rotation groups. The first one indicates the duration elapsed since the last spell of employment, but only indicates whether this is below or above one year. The other one lists various reasons for leaving the most recent job. The categories are somewhat redundant with the main reason for being currently out of the labor force. This allows to check for consistency between answers.

[28] In any given month of the CPS, only between 45 percent and 50 percent of individuals self-report information about their current labor force status. Moreover, when looking at four consecutive months of interviews, less than one-third of respondents provide information for themselves in all months, while another third of respondents never self-report information. This reporting behavior and its potential effects are an under-documented aspect of the CPS.

Table B1. Labor Force Status in the CPS: Original vs. Recoded

		Recoded		
		E	U	I
	E	99.56	0.19	0.25
Original	U	0.00	97.83	2.17
	I	0.19	1.50	98.31

NOTE: E: employment; U: unemployment; I: inactivity. The data are from the pooled cross-sections of the CPS for the year 2000. The table reads as follow: among those classified as employed in the original CPS file, 99.56% remain classified as employed, 0.19% are re-classified as unemployed and 0.25% are re-classified as inactive after implementing the correction strategy.

Table B1 shows the consequences of the correction strategy for labor force stocks in a given cross section of the Current Population Survey. The table is organized in the spirit of the seminal study by Poterba and Summers (1986) to emphasize that the correction strategy does *not* induce a large revision of the measured stocks. However, as shown in figure 1, the correction has a significant impact on some labor market flows, namely the flows from inactivity to employment (I-E) and from unemployment to inactivity (U-I). The differences in levels between the dashed and the solid lines are not caused only by the correction strategy: for instance the dashed lines are not adjusted for time-aggregation and margin error. However, the differences line up well with those obtained by Elsby et al. (2013a) with their own correction strategy. This suggests that both their methodology and the one developed here can be fruitfully used to recode labor market transitions in the CPS.