The Alpha Beta Gamma of the Labor Market*

Victoria Gregory
FRB Saint Louis

Guido Menzio
New York University and NBER

David Wiczer
Stony Brook University

April 1^{st} 2021

Abstract

Based on patterns of employment transitions, we identify three different types of workers in the US labor market: α's β's and γ's. Workers of type α make up over half of all workers, are most likely to remain on the same job for more than 2 years and, when they become unemployed, typically find a new job within 1 quarter. Workers of type γ comprise less than one-fifth of workers, have a low probability of staying on the same job for more than 2 years and, when they become unemployed, face a high probability of remaining jobless for more than 1 year. Workers of type β are in between αs and γ's. The earnings losses caused by displacement are relatively small and transitory for α-workers, while they are large and persistent for γ-workers. During the Great Recession, excess unemployment for α-workers rose by little and was reabsorbed quickly; unemployment for γ-workers rose by 20 percentage points and was not reabsorbed 4 years after its peak. We use a search-theoretic model of the labor market to rationalize the different patterns of employment transitions across types. The model naturally explains both the variation in the consequences of job displacement across types, and the variation in the dynamics of unemployment during the Great Recession. Our view is that several puzzling micro and macro phenomena about the labor market are driven by the behavior of the small group of γ-workers.

JEL Codes: E24, O40, R11.

Keywords: Search frictions, Unemployment, Business Cycles.

*We are grateful to Jean Flemming, Simon Mongey and Iacopo Morchio for their comments. The views expressed herein are solely those of the authors and do not necessarily reflect those of the Federal Reserve Bank of St. Louis or the Federal Reserve System. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1819. (CBDRB-FY21-P1819-R8907,CBDRB-FY20-P1819-R8772, CBDRB-FY20-P1819-R8293) This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.

1 Introduction

Using a large panel dataset of US workers, we document the existence of systematic differences in the pattern of labor market transitions of different workers. Applying standard machine-learning techniques, we find that workers can be classified into three typologies—which we dub α-, β- and γ-workers—featuring different lengths of unemployment spells, different distributions of job tenures, and different frequencies of job-to-job transitions. Motivated by these findings, we develop a heterogeneous-worker version of the search-theoretic framework of Diamond (1982), Mortensen (1970) and Pissarides (1985) that is capable of rationalizing the differences in behavior across different types of workers. We then use our model to shed light on a number of labor market phenomena, both at the micro and the macro level.

We use the Longitudinal Employer-Household Dynamics (LEHD) dataset between 1997 and 2014 to create a record of the transition across employment states (unemployment, employment, and different employers) for individual workers. We then use a k-means algorithm to classify workers into types that share a similar pattern of transitions, in terms of duration and frequency of unemployment spells and distribution of job tenures. The algorithm identifies three distinct types of workers—α, β and γ—and assigns every individual worker to one of these types. On average, α workers have the shortest unemployment spells, the highest likelihood of reaching 2 years of tenure on the same job, and the highest labor income while employed. In contrast, γ workers have the longest unemployment spells, a very high likelihood of leaving a job within the first quarter from its inception, a very low likelihood of reaching two years of tenure on a job, and the lowest income while employed. Workers of type β are in between α's and γ's. Mirroring the well-established fact that observables explain very little of wage dispersion, we find that age, sex and education have low power in predicting a worker's type.

Motivated by these empirical findings, we develop a search-theoretic model of the labor market with different types of workers. The basic structure of the model is as in Menzio and Shi (2011). Specifically, firms pay a cost to maintain vacancies and advertise the terms of trade offered to workers filling each of their vacancies. Workers choose where to direct their search, in the sense that they choose to seek vacancies offering particular terms of trade. Vacancies and workers offering and demanding the same terms of trade come together through a bilateral matching process with constant returns to scale. Upon matching, a firm and a worker begin production. Productivity depends on the quality of the firm-worker match, which is initially unknown and learned over time. The model generates, as equilibrium outcomes, the probability that an unemployed worker becomes employed (UE rate), the probability that an employed worker becomes unemployed (EU rate), and the probability that a worker moves from one employer to another (EE rate). Similarly, the model generates, as an equilibrium outcome, the tenure-length distribution of jobs.

On top of the basic structure of the model, we introduce three different types of workers. A worker's type is permanent and affects the baseline level of productivity (so as to capture differences in labor earnings across types), the distribution of match qualities and the speed at which match quality is discovered (so as to capture differences in the distribution of tenure lengths), and different probabilities of searching the market (so as to capture different unemployment rates). Even though the model features multiple layers of heterogeneity—namely, the distribution of workers of different types across unemployment and employment in matches of different quality—the assumption of directed search makes the equilibrium block-recursive and, hence, allows us to easily solve the model in and out of steady-state.

We calibrate the model to match the prevalence of workers of type α, β and γ in the data, as well as the distinguishing features of different types: average earnings, average duration of unemployment, tenure distribution of jobs, and fraction of jobs ending into unemployment or into another job at different tenure lengths. The calibrated model matches quite well all of the data targets and, in doing so, it provides a coherent interpretation for the observed differences between different types. For instance, the calibrated model reveals that α-workers sample from a distribution of match qualities that is approximately normal and with low variance. Since their sampling distribution has little downside and little upside, α-workers are likely to remain in the same job for a long period of time. In contrast, γ-workers sample from a distribution of match qualities that is approximately exponential. Because such a sampling distribution has a large density at the left and a thick right tail, γ-workers have a higher incentive to search. As a result, they keep moving from job to job (with or without intervening spells of unemployment) until they find a match whose quality is in the right tail. These match quality distributions impact the distribution of earnings within type and affect how long different worker type stay in certain jobs.

The existence of different types of workers can help us better understand micro-level phenomena in the labor market. Using the classification of individual workers into types generated by the k-means algorithm, we show that the earnings consequences of displacement of long-tenured workers varies a great deal across types. For α-workers, the earning losses dissipate fairly quickly—they are equal to 29% of pre-displacement earnings after 1 year and 10% after 5 years. For γ-workers, the earning losses are much more persistent—they are equal to 69% of pre-displacement earnings after 1 year and 50% after 5 years. Using the calibrated model, we show that the theory reproduces well the type-specific magnitude and pattern of earning losses. According to the theory, the earning losses of displaced workers of type γ are more persistent because γ-workers are slower at finding jobs and because γ-workers are more likely to have to sample several jobs before finding one that is above their reservation quality.

We also use our model to revisit the well-documented fact that the UE rate falls dramatically as a function of the duration of unemployment. We show that the existence of different types of workers with a very different UE rate (30% per month for α-workers, 15% for β-workers, and 10% for γ-workers) generates by itself a sharp decline in the UE rate. This is because γ-workers make up a larger and larger fraction of the unemployment pool as duration increases. Specifically, we find that worker heterogeneity alone implies that the UE rate falls by approximately 40% after 24 months of unemployment.

Lastly, we examine the macroeconomic implications of worker heterogeneity. We start by considering the effect of a transitory, negative shock to the aggregate component of productivity.

For α-workers, the shock causes a relatively small decline in the UE rate and a relatively small increase in the EU rate and, hence, a relatively small increase in the unemployment rate. For γ-workers, in contrast, the shock causes a relatively large decline in the UE rate and a relatively large and persistent increase in the EU rate and, hence, a large and persistent increase in their unemployment rate. The response of the UE, EU and unemployment rates for β-workers lies between the response of α and γ-workers. Intuitively, the UE rate of γ-workers is most responsive to the shock because the net value of employment of γ-workers is smallest. The EU rate of γ-workers is most responsive and most persistent to the shock because there are more γ-workers who are employed in matches of marginal quality and, once they become unemployed, they need to sample the most jobs before finding one that is acceptable.

At the aggregate level, the productivity shock generates a large and persistent increase in the unemployment rate. The response of the aggregate unemployment rate is large because of the large increase in the number of γ-workers who, on impact, lose their job. The response of the aggregate unemployment rate is persistent because it takes a long time for γ-workers who lost their job to find another acceptable job. In contrast, the productivity shock generates a decline in average labor productivity that is smaller and shorter-lived as compared to the shock. The response of average labor productivity is muted because the workers who keep their job are disproportionately α-workers, who are workers with a higher baseline productivity. The response of average labor productivity is more transitory than the shock because the α-workers who do lose their job find another acceptable job very quickly.

We then proceed to compare the above predictions of the model with the behavior of the US labor market during the Great Recession of 2007-2009 and the subsequent recovery. The Great Recession featured a large increase in the unemployment rate (about 6 percentage point) in the face of a modest decline in average labor productivity (about 4% decline from trend). Moreover, the aftermath of the Great Recession featured a slow recovery of unemployment (2 percentage points above the pre-recession level in 2013) in the face of a quick recovery of average labor productivity (back on trend in 2011 and then above trend). Indeed, "jobless recoveries" have been a feature of the US labor market over the last 30 years, and present one of the most challenging puzzles in the macro-labor literature.

Using the LEHD and the classification of individual workers produced by the k-means algorithm, we measure the increase of the unemployment rate of different types of workers during the Great Recession of 2009-2010, as well as the speed of the recovery of the unemployment rate of different types during the recovery of 2011-2014. We find that the increase in the unemployment rate was small (3 percentage points) and short-lived (vanished 1 year into the recovery) for α-workers. The increase in the unemployment rate was intermediate in size (7 percentage points) and of average duration (all but vanished 3 years into the recovery) for β-workers. And the increase in the unemployment rate was large (20 percentage points) and very persistent (still 10 percentage points 4 years into the recovery) for γ-workers.

These empirical findings are broadly consistent with the predictions of the model in response to a negative shock to the aggregate component of productivity. We show that—by assuming that the aggregate shock hit more severely γ and β-workers than α-workers—the predictions of the model align very well with the behavior of the US economy. Specifically, the model correctly predicts the large and persistent increase in the unemployment rate, which as in the data, is driven by the size and by the persistence of the increase in the unemployment rate of γ-workers. Further, the model correctly predicts a modest decline in average labor productivity followed by a very rapid recovery. Indeed, as in the data, the model predicts that average labor productivity bottoms out at 4% below trend and returns to trend within a year.

Related Literature. Our approach to studying individual labor market transitions and aggregate labor market dynamics has a parallel in the consumption and earning dynamics literature. The consumption and earning dynamics literature has traditionally focused on ex-post heterogeneity across households, heterogeneity produced by different realizations of a common stochastic income process. Analogously, the search-theoretic literature on labor market transitions has typically focused on ex-post heterogeneity across workers, heterogeneity produced by different realizations of a common job search process (see, e.g., Pissarides 1985, Mortensen and Pissarides 1994, or Menzio and Shi 2011). More recently, the consumption and earning dynamics literature has recognized and incorporated the view that ex-ante heterogeneity caused by differences in income processes is just as important as ex-post heterogeneity (see, e.g., Guvenen 2007, Guvenen and Smith 2014, or Guvenen, Karahan, Ozkan and Song 2015). Analogously, our paper documents the existence of systematic differences across workers in their pattern of labor market transitions and examines the implications of this type of heterogeneity for understanding aggregate labor market dynamics.

The idea that workers differ with respect to their search process and that this might matter for understanding labor market phenomena has been floating around for a while. For instance, statistical studies on the duration dependence of the UE rate have long concerned themselves with the possibility that workers may be heterogeneous with respect to their UE rate. Only recently, though, has evidence for ex-ante heterogeneity in search processes been sought more broadly. Morchio (2020) documents that most of unemployment spells can be attributed to the same small group of workers, thus hinting at systematic differences in the EU rate. Ahn and Hamilton (2020) document changes in the shape of the UE rate as a function of unemployment duration during the Great Recession. They conjecture that these changes are caused by changes in the composition of worker types into the unemployment pool. In this paper, we confirm their conjecture. Methodologically, the paper closest to ours is Hall and Kudlyak (2020), who use short panel data to classify workers into types based on their transitions between employment and unemployment. There are several important differences between our paper and Hall and Kudlyak (2020). First, using almost 20 years of panel data, we are able to classify individual workers into types using much more evidence on their behavior. Second, having information about employers, we are able to classify individual workers based not only on transitions between employment and unemployment, but also across employers. Third, having data that covers the Great Recession and its aftermath, we are able to document the behavior of different types of workers over the business cycle. Moreover, we do not stop at documenting differences, but we root them into differences in fundamentals in the context of an equilibrium model of search.

The paper does not only make a contribution to the measurement of heterogeneity in the search process of different workers. The paper also contributes substantively to our understanding of some phenomena in labor economics. The paper contributes to the literature on the effect of displacement on labor earnings. We document that the earning losses from displacement are very different for different types of workers, thus highlighting the importance of heterogeneity to understand the "average" findings of Jacobson, LaLonde and Sullivan (1993), Davis and von Wachter (2011), and Flaaen, Shapiro and Sorkin (2019). Moreover, we show that the magnitude and persistence of earnings losses is exactly what one would expect from a model calibrated to match the pattern of transitions of different types of workers. The paper also contributes to the literature trying to explain why the UE rate declines so dramatically with unemployment duration. We show that heterogeneity in the UE rate of different types of workers can by itself account for almost all of the decline in the UE rate. This finding, emphasizing dynamic composition, is consistent with the results reported in Mueller, Spinnewjin and Topa (2019).

The paper contributes to the macro literature trying to make sense of the behavior of the labor market during the Great Recession and, more generally, the emergence of jobless recoveries and the vanishing of the comovement between unemployment and productivity. The lack of comovement between labor productivity and unemployment motivated research exploring the possibility that labor market fluctuations are driven not by productivity shocks but by self-fulfilling expectation shocks (Kaplan and Menzio 2016), discount factor shocks (Hall 2017, Kehoe, Midrigan and Pastorino 2019), or by correlated equilibria (Golosov and Menzio 2020). In this paper, we document that—during the Great Recession—the persistence of high unemployment was caused by one group of workers who need a remarkably long time to find suitable employment after losing their job. As the workers in this group have relatively low earnings, this explains why average labor productivity recovered much more quickly than unemployment. That is, the correlation between the search process and the earnings of different types of workers leads to a natural explanation of jobless recoveries and weak comovement between productivity and unemployment.

The rest of this paper proceeds as follows. In Section 2, we describe how we categorize workers in the LEHD into one of three types via a k-means clustering method. Section 3 presents the model we build that incorporates features of these three worker types and Section 4 outlines our calibration strategy. We show how the model rationalizes several microeconomic and macroeconomic phenomena of the labor market in Sections 5 and 6, respectively.

2 Documenting Heterogeneity

In this section, we show that there exist different typologies of workers in the US labor market. In Section 2.1, we describe the administrative data that we use for the empirical analysis. In Section 2.2, we describe the method that we use to classify workers into different groups based on their pattern of employment transitions. In Section 2.3, we describe the distinctive features of different types of workers.

2.1 Data

We use data from the Longitudinal Employer-Household Dynamics (LEHD) program. The LEHD contains quarterly information about individual employment histories, which includes an identifier for the individual, an identifier for the individual's employer (a state-level SEIN), and the quarterly earnings of the individual from each employer (as measured by pre-tax labor earnings as in a W2 tax form). The LEHD does not report employment in the military or in the federal government, self-employment, contracting work or other forms of employment not covered by Unemployment Insurance.[1]

Our extract from the LEHD covers 17 states whose available data covers the entirety of our reference period between 1997 and 2014.[2] The 17 states include some large economies, such as California, Illinois and Texas. Overall, the labor force from the 17 states represents about 40% of the total labor force in the US. Our extract from the LEHD is a 2% random sample of individuals. Since we want to focus on workers with a strong attachment to the labor force, we purge our sample from all individuals who have an earning gap of more than 2 years between two consecutive employment episodes. The purged sample contains about $692,000$ unique individuals, or about 0.5% of the US labor force and 0.65% of private sector employment. Given this requirement for being included into our sample, we classify all spells without earnings as unemployment rather than non-employment. Obviously, our definition of unemployment, which is based on outcomes, does not coincide with the definition of unemployment from the Bureau of Labor Statistics (BLS), which is based on intent.

In order to deal with individuals entering and exiting the labor force and with censored spells, we create a 2-year window at the beginning of the reference period 1997-2014 (i.e. 1997-1998) and a 2-year window at the end of the reference period (i.e. 2013-2014). If an individual is employed in the first quarter of 1999, we start his record with that job. If the individual was employed in the first quarter of 1999, we know whether his job has lasted more than 2 years (which is the highest duration bin we use for classifying job spells) or when it did start (because of the 2-year window at the beginning of the reference period). In either case, we start the record of the individual at the beginning of his job. If an individual was unemployed in the first quarter of 1999, we know whether his unemployment spell lasted less or more than 2 years. If it lasted less than 2 years, we start the record of the individual with the beginning of that unemployment spell. If it lasted more than 2 years, we have no record of prior employment for the individual and we start his record from his first job in the reference period. That is, we assume that the individual was out of the labor force prior to his first job.

Symmetrically, if an individual was employed in the last quarter of 2012, we know whether his job lasted more than 2 years (since we track the worker until the end of 2014) or when the job ended. In either case, we end the record of the individual with that job. If an individual was unemployed in the last quarter of 2012, we know whether his unemployment spell would

[1] See Abowd et al. (2009) for more details about the LEHD.

[2] The 17 states are California, Colorado, Hawaii, Idaho, Illinois, Indiana, Kansas, Maine, Maryland, Missouri, Montana, Nevada, North Dakota, Tennessee, Texas, Virginia and Washington.

last more than 2 years, in which case we stop the record of the individual with the end of the last recorded job, or less than 2 years, in which case we stop the record of the individual at the end of the current unemployment spell.

Having determined the start and end of the record of each individual in the sample, we proceed as follows. We measure the duration of each job as the number of quarters during which the individual reports earnings from a particular employer. We measure the duration of each unemployment spell as the number of quarters during which the individual does not report any earnings. We impute unemployment spells of less than a quarter for two types of employer transitions: those that happen without a period in which there were earnings from both employers or if earnings in the quarter in which both employers paid the worker were lower than the minimum of earnings in the two adjacent quarters. In these situations, we assign an unemployment spell of duration 0.5 quarters. Otherwise, we characterize it as a direct job-to-job transition. We classify every job into 4 bins: jobs lasting no more than 1 quarter; jobs lasting more than 1 quarter but no more than 4 quarters; jobs lasting more than 4 quarters but no more than 8 quarters; and jobs lasting more than 8 quarters. Similarly, we classify every unemployment spell into 3 bins: spells lasting no more than 1 quarter (which includes imputed spells); spells lasting more than 1 quarter but no more than 4 quarters; spells lasting more than 4 quarters but less than 8 quarters.

For each individual, we then compute the following statistics: (i) the fraction of jobs lasting less than 1 quarter, between 1 and 4 quarters, between 4 and 8 quarters, and more than 8 quarters; (ii) the fraction of unemployment spells lasting less than 1 quarter, between 1 and 4 quarters, and between 4 and 8 quarters; (iii) the total quarters of unemployment as a fraction of the total number of quarters on record; (iv) the total number of different jobs as a fraction of the total number of quarters on record. These statistics paint a picture of the pattern of employment transitions of a particular individual. Statistic (iii) tells us how much time the individual spends in unemployment. Statistic (ii) tells us the distribution of unemployment durations for an individual. Statistic (i) tells us the distribution of job durations for an individual. And statistic (iv) together with (iii) tells us about direct job-to-job transitions.

2.2 Classifying workers

In order to classify individual workers into groups, we use a k-means algorithm, a standard tool in machine learning (see, e.g., Friedman, Hastie and Tibshirani 2017) that is becoming more and more commonplace in economics (e.g., Bonhomme, Lamadon, and Manresa 2020). The k-means algorithm allows us to group workers into clusters based on their similarity with respect to the pattern of their employment transitions and to uniquely assign each worker to a group. The optimal number of clusters is chosen using the cross-validation method proposed by Wang (2010).

Let i denote an individual in our sample. Let $s_{1,i}$, $s_{2,i}$, $s_{3,i}$ and $s_{4,i}$ denote the distribution of job durations for individual i. Let $s_{5,i}$, $s_{6,i}$ and $s_{7,i}$ denote the distribution of unemployment durations for individual i. Let $s_{8,i}$ denote the fraction of time spent by individual i in unemployment. Let $s_{9,i}$ denote the number of jobs of individual i per unit of time. All statistics are expressed as ratios with respect to their population-wide standard deviation. The four statistics describing the distribution of job durations $s_{1,i}$, $s_{2,i}$, $s_{3,i}$ and $s_{4,i}$ are multiplied by 1/4 and the three statistics describing the distribution of unemployment durations $s_{5,i}$, $s_{6,i}$ and $s_{7,i}$ are multiplied by 1/3.

For a given the number J of clusters, the assignment of individuals to clusters is a mapping $j(i)$ from an individual $i \in \{1, 2, ...N\}$ to a cluster $j \in \{1, 2, ..J\}$ that solves the following optimization problem

$$\min_{j(i)} \sum_{j=1}^{J} \sum_{i=1}^{N} \sum_{k=1}^{9} \mathbf{1}[j = j(i)](s_{k,i} - s^*_{k,j})^2, \tag{2.1}$$

where $s^*_{k,j}$ denotes the average of statistic $s_{k,i}$ across all i's assigned to cluster j. In words, (2.1) states that the assignment of individuals to clusters is such that the squared distance between the statistics of an individual and the average statistics in that individual's cluster are minimized.

We solve the minimization problem in (2.1) using a simple iterative process that combines well-known k-means clustering algorithms with a heuristic for trying starting clusters. To initialize the iteration, we select one of the dimensions describing individual histories. We rank individuals along the selected dimension and divide them into J clusters of equal size. That is, individual i is assigned to cluster 1 if he ranks in the lowest $1/J$ percent of the population along the selected dimension. Individual i' is assigned to cluster 2 if he ranks in the second lowest $1/J$ percent of the population along the selected dimension, etc... Having created an assignment $j^0(i)$, we then compute the average $s^0_{k,j}$ of statistic s_k for all individuals i assigned to cluster j. In the n-th step of the iteration, with $n = 1, 2, \ldots.$, we solve (2.1) using $s^{n-1}_{k,j}$ instead of $s^*_{k,j}$ in the objective function. The solution of (2.1) is an updated assignment $j^n(i)$ of individuals to clusters. Using the assignment $j^n(i)$, we compute an updated average $s^n_{k,j}$ of statistic s_k for all individuals i assigned to cluster j. We continue the iteration process until we reach a fixed point. We loop over each dimension used to initialize the assignment to heuristically check that the fixed point is unique and, hence, the solution to (2.1).

In order to choose J, the number of clusters, we follow the cross-validation approach proposed by Wang (2010). We randomly divide our sample of individuals into three subsamples, S_0, S_1 and S_2. The subsamples S_1 and S_2 are for training, and each of them accounts for 25% of individuals. The subsample S_0 is for validation, and it accounts for the remaining 50% of individuals. For any $J \geq 2$, we solve the optimization problem (2.1) on the subsample S_1 and obtain the cluster-specific averages $s^1_{k,j}$. Analogously, we solve the optimization problem (2.1) on the subsample S_2 and obtain the cluster-specific averages $s^2_{k,j}$. Then, we solve the assignment problem (2.1) on the validation sample S_0 using $s^1_{k,j}$ instead of $s^*_{k,j}$ in the objective function. This gives us an assignment $j^1(i)$ of individuals in S_0 to clusters. We do the same using $s^2_{k,j}$ and obtain a different assignment $j^2(i)$ of individuals in S_0 to clusters. We choose J so as to minimize the number of individuals in S_0 who are assigned to different clusters based on $s^1_{k,j}$

	α-workers	β-workers	γ-workers
Population share	0.57	0.26	0.17
Job duration			
<1Q	0.139	0.201	0.360
1Q-4Q	0.187	0.233	0.321
5Q-8Q	0.236	0.238	0.191
>8Q	0.439	0.329	0.119
Unemployment duration			
<1Q	0.794	0.390	0.553
1Q-4Q	0.157	0.558	0.315
5Q-8Q	0.049	0.052	0.133
Fraction of time unemployed	0.036	0.096	0.292
Demographics			
Birth year	1963	1964	1967
Fraction female	0.479	0.486	0.463
Fraction some college	0.620	0.559	0.496

Table 1: Descriptive statistics for each worker type

and $s^2_{k,j}$, i.e.

$$\min_{J\geq 2} \sum_{i=1}^{N_0} 1[j^1(i) \neq j^2(i)]. \tag{2.2}$$

On the one hand, the criterion (2.2) penalizes a J that is so low that some workers are grouped with very different individuals in S_1 and S_2. On the other hand, the criterion (2.2) penalizes a J that is so high that the features of each cluster are sensitive to the different subsamples S_0 and S_1. The result of this exercise was that having 3 clusters is optimal in this sense.

2.3 Classification outcomes

Table 1 reports the outcomes of the classification process described above. We identify three different types of workers in our sample, which we dub α, β and γ. Workers of type α represent the majority of individuals in our sample (57%), while workers of type β are 26%, and workers of type γ are 17%.

Different types of workers have very different patterns of labor market transitions. First, consider the distribution of job durations for different types. For α-workers, the fraction of job spells lasting less than 1 quarter is 14% and the fraction of job spells lasting more than 2 years is 44%. For β-workers, the fraction of job spells lasting less than 1 quarter is 20% and the fraction of job spells lasting more than 2 years is 33%. For γ-workers, the fraction of job spells lasting less than 1 quarter is 36%, and the fraction of job spells lasting more than 2 years is 12%. That is, α-workers are 50% more likely to remain in the same job for more than 2 years than β-workers and 400% more likely than γ-workers. Conversely, γ-workers are 100% more likely to leave a job within 1 quarter than β-workers, and 300% more likely than α-workers.

Second, consider the distribution of unemployment durations for different types. For α-

workers, the fraction of unemployment spells lasting less than 1 quarter is 79% and the fraction of spells lasting more than 1 year is only 5%. For β-workers, the fraction of unemployment spells lasting less than 1 quarter is 39% and the fraction of spells lasting more than 1 year is 5%. For γ-workers, the fraction of unemployment spells lasting less than 1 quarter is 55% but the fraction of spells lasting more than 1 year is 13%. That is, α-workers typically have short unemployment spells and very rarely have long ones; β-workers are less likely to have short unemployment spells but they also have very few long ones; γ-workers are much more likely to experience long unemployment spells compared to αs and βs. The average time spent in unemployment is 3.5% for an α-worker, 9.6% for a β-worker, and 29.2% for a γ-worker.

Overall, our classification of workers into types paints a clear picture. Workers of type α are likely to remain on a job for a long period of time and, when they do move into unemployment, they find a new job very quickly. Workers of type β are less likely to remain on a job for a long period of time and, when they move into unemployment, it takes them longer to find a new job. However, workers of type β are unlikely to be stuck into unemployment for more than a year. Workers of type γ are most likely to leave a job within two years. When they do become unemployed, they face a significant probability of becoming long-term unemployed.

Using the output of the k-means algorithm, we can assign each individual worker in the data to a particular type. We can then compute the average of some observable characteristics for different types of workers. We find that the average birth year is 1963 for α-workers, 1964 for β-workers and 1967 for γ-workers. These small age differences seem to rule out the possibility that the algorithm simply picks out differences in transition patters that are related to the life cycle of a worker. We find that the fraction of women is 48% for α-workers, 49% for β-workers, and 46% for γ-workers. Again, we take this evidence that the algorithm is not splitting workers based on gender. Lastly, the fraction of workers with some college education is 62% for α-workers, 56% for β-workers, and 50% for γ-workers. Along this dimension, there are starker differences, and yet having some college education hardly predicts a worker's type.

3 Modeling Heterogeneity

In this section, we propose a theory to make sense of the heterogeneity among workers observed in the data.

3.1 Environment

The labor market is populated by a positive measure of workers and firms. Workers are ex-ante heterogeneous with respect to their type $i = 1, 2, \ldots I$, which affects their productivity, unemployment income, and their search and learning processes. A worker of type i maximizes the present value of income (measured in units of output), discounted by the factor $\rho \in (0, 1)$. A worker of type i earns some income b_i when he is unemployed, and some income w_i when he is employed. The unemployment income b_i is a combination of unemployment benefits, transfers, and the value of leisure. The employment income w_i is determined by the worker's employment contract. The measure of workers of type i is $\mu_i \in [0, 1]$ and the total measure of workers is 1.

Firms are ex-ante homogeneous. A firm maximizes the present value of profits, discounted by the factor ρ. A firm operates a constant returns to scale technology which turns the labor supply of a worker of type i into xy_iz units of output, where $x \in X \subset \mathbb{R}_+$ is a component of productivity that is common to all firm-worker pairs, $y_i \in Y \subset \mathbb{R}_+$ is a component that is common to all pairs of firms and workers of type i, and $z \in Z \subset R_+$ is a component that is specific to a particular firm-worker pair. The first component of productivity is stochastic, and it is a source of cyclical fluctuations. The second component of productivity is permanent, and it is the source of persistent differences in the productivity of different types of workers. The last component of productivity is permanent, and it is the source of worker's mobility from job to job. We refer to this component of productivity as the quality of a firm-worker match. We assume that the quality of a firm-worker match is initially unknown but, eventually, is observed.

The labor market is organized in a continuum of submarkets indexed by the vector $s = \{v, i\}$, where $v \in R$ denotes the lifetime income promised by firms to workers hired in submarket s, and $i \in \{1, 2, ...I\}$ denotes the type of workers hired by firms in submarket s. Associated with each submarket $s = \{v, i\}$, there is an endogenous vacancy-to-applicant ratio $\theta_i(v) \in \mathbb{R}_+$. If a worker searches in submarket s, he finds a vacancy with probability $p(\theta_i(v))$, where p is a strictly increasing, strictly concave function with $p(0) = 0$ and $p(\infty) = 1$. Similarly, a vacancy in submarket s finds an applicant with probability $q(\theta_i(v))$, where q is a strictly decreasing function with $q(\theta) = p(\theta)/\theta$, $q(0) = 1$ and $q(\infty) = 0$.

Time is discrete. At the beginning of a period, the state of the economy is described by the exogenous state, which is the aggregate component of productivity, and by the endogenous distribution of workers across employment states. Formally, the state of the economy is given by $\psi \equiv \{x, u_i, n_i, g_i\}$, where $x \in X$ is the aggregate component of productivity, $u_i \in [0, 1]$ is the measure of workers of type i who are unemployed, $n_i \in [0, 1]$ is the measure of workers of type i who are employed in a match of unknown quality, and $g_i : Z \to \mathbb{R}_+$ is a function such that $g_i(z)$ denotes the measure of workers of type i who are employed in a match of known quality z.

Every period comprises four stages: learning, separation, search and production. In the first stage, a worker and a firm who are in a match of unknown quality learn the idiosyncratic component of productivity z with probability $\phi_i \in [0, 1]$, where ϕ_i is allowed to depend on the worker's type i. The idiosyncratic component of productivity z is a random draw from a probability distribution function $f_i : Z \to \mathbb{R}_+$ with a mean normalized to 1. The sampling function f_i also depends on the worker's type i.

In the second stage, a match between a worker and a firm breaks up with probability $d \in [\delta, 1]$ The probability d is specified by the employment contract regulating the relationship between the worker and the firm. The lower bound δ represents the probability that the worker has to leave the match for exogenous reasons (e.g., firm closure or worker relocation).

In the third stage, a worker gets the opportunity to search the labor market with a probability that depends on his employment status and his recent employment history. If a worker is unemployed, he gets to search the labor market with probability $\lambda_u^i \in (0, 1]$. If the worker is employed, he gets to search the market with probability $\lambda_e^i \in (0, 1]$. If the worker just became unemployed during the previous separation stage, we assume that he cannot search. The probabilities λ_u^i and λ_e^i are allowed to depend on the worker's type i. Whenever the worker gets to search, he chooses which submarket s to visit. In the same stage, firms choose how many vacancies to open in submarket $s = \{v, i\}$ at the unit cost $k_i > 0$.

Workers and firms searching in submarket $s = \{v, i\}$ meet bilaterally. When a firm and a worker of type i meet in submarket s, the firm offers to the worker an employment contract that is worth v in lifetime income. If the worker accepts the offer, he becomes employed by the firm under the rules of the contract. If the worker rejects the offer, which is an off-equilibrium event, he returns to his previous employment status. When a firm and a worker of a type different from i meet in submarket s, the firm refuses to hire the worker.[3]

In the last stage, an unemployed worker enjoys an income equal to b_i units of output, where b_i is allowed to depend on a worker's type. A worker employed in a match of unknown quality produces, in expectation, xy_i units of output. A worker employed in a match of known quality z produces xy_iz units of output. The worker's consumption is w_i, where w_i is determined by the employment contract controlling the relationship between the worker and the firm. After production and consumption take place, next period's aggregate component of productivity, $\hat{x}$, is drawn from the probability density function $h : X \times X \to \mathbb{R}_+$ with $h(\hat{x}, x)$ denoting the probability density of $\hat{x}$ conditional on x.

We assume that the contracts offered by firms to workers are bilaterally efficient, in the sense that they maximize the joint income of the firm-worker match. As discussed in Menzio and Shi (2011), this assumption is consistent with several contractual environments. Consider two cases. In the first case, a contract can specify the worker's wage, the worker's search strategy on the job (i.e. in which submarket to search) and the worker's quitting strategy (i.e. when to move into unemployment) contingent on the history of the match and the economy. In this case, the contract space is rich enough to independently control the allocative decisions of the match and the distribution of the value of the match between the firm and the worker. Given this contractual environment, the firm finds it optimal to offer a contract such that the allocative decisions maximize the joint income of the match, and such that the wages provide the worker with the lifetime income v. In the second case, a contract can specify a sign-on transfer and then a wage contingent on the history of the match and the economy. The worker is then free to follow his preferred search and quitting strategy. In this case, the firm finds it optimal to offer a contract such that the worker is the residual claimant of output (and, hence, makes allocative decisions to maximize the joint income of the match) and a (possibly negative) transfer such

[3]We assume that a worker knows his own type and so does the market. The second part of the assumption may appear unrealistic to some readers, but it does greatly simplify the model. In particular, the assumption allows us to abstract from issues of signaling—the worker distorting his behavior so as to convince the market that his type is better than what it actually is—as well as from issues of inference—the firms trying to assess the probability distribution of a worker's type by examining his employment history and performance on the job.

that the worker's lifetime income is v.

It is useful to pause to motivate our approach to modeling worker's heterogeneity. We assume that the worker's type affects his productivity, his search process, and the speed at which he learns the quality of his match. Intuitively, we assume that the type affects the worker's productivity, y_i, in order to reproduce the differences in the average earnings of different types. We assume that the type affects the worker's probability of searching the market, λ_u^i and λ_e^i, so as to reproduce the differences in the average UE and EE rates of different types. We assume that the type affects the distribution of match qualities from which the worker samples, f_i, and the speed at which the worker learns the quality of a particular match, ϕ_i, so as to reproduce differences in the job-duration distribution of different types. We also assume that the type affects the worker's unemployment income, b_i, because unemployment benefits depend on labor income.

3.2 Equilibrium

To formally define an equilibrium, we need to introduce a few additional pieces of notation. Let $U_i(\psi)$ denote the lifetime income for a worker of type i who is unemployed at the beginning of the production stage. Let $\tilde{V}_i(\psi)$ denote the sum of the lifetime income for a firm and a worker of type i who, at the beginning of the production stage, are in a match of unknown quality. Let $V_i(z,\psi)$ denote the sum of the lifetime income for a firm and a worker of type i who, at the beginning of the production stage, are in a match of known quality z. Lastly, let $\theta_i(v,\psi)$ denote the equilibrium tightness of submarket $s = \{v, i\}$.

The value $U_i(\psi)$ of unemployment for a worker of type i is given by

$$U_i(\psi) = b_i + \rho \mathbb{E}_{\hat{\psi}} \left[U_i(\hat{\psi}) + \lambda_u^i \max_v \left\{ p(\theta_i(v,\hat{\psi})) \left(v - U_i(\hat{\psi}) \right) \right\} \right]. \tag{3.1}$$

In the current period, the worker's income is b_i. In the next period, the worker finds a job with probability $\lambda_u^i p(\theta_i(v,\hat{\psi}))$. In this case, the worker's continuation value is v. The worker does not find a job with probability $1 - \lambda_u^i p(\theta_i(v,\hat{\psi}))$. In this case, the worker's continuation value is $U_i(\hat{\psi})$. Note that, since search is directed, the worker chooses v so as to maximize his lifetime income.

The joint value $V_i(z,\psi)$ of a match of quality z between a firm and a worker of type i is given by

$$\begin{aligned} V_i(z,\psi) &= x y_i z \\ &+ \rho \mathbb{E}_{\hat{\psi}} \left[\max_{d \in [\delta,1]} d U_i(\hat{\psi}) + (1-d) \left[V_i(z,\hat{\psi}) + \lambda_e^i \max_v \left\{ p(\theta_i(v,\hat{\psi})) \left(v - V_i(z,\hat{\psi}) \right) \right\} \right] \right] \end{aligned} \tag{3.2}$$

In the current period, the sum of the worker's income and firm's profit is $x y_i z$, the output of the match. In the next period, the worker moves into unemployment with probability d. In this case, the worker's continuation value is $U_i(\hat{\psi})$ and the firm's continuation value is 0. The worker moves from the current job to a new job with probability $(1-d)\lambda_e^i p(\theta_i(v,\hat{\psi}))$. In this case, the worker's continuation value is v and the firm's continuation value is 0. The worker and the firm remain together with probability $(1-d)(1-\lambda_e^i p(\theta_i(v,\hat{\psi})))$. In this case, the firm's and worker's joint continuation value is $V_i(z,\hat{\psi})$. Note that, since employment contracts are bilaterally efficient, d and v are chosen so as to maximize the joint value of the match.

The joint value $\tilde{V}_i(\psi)$ of a match of unknown quality between a firm and a worker of type i is given by

$$\begin{aligned}
&\tilde{V}_i(\psi) = xy_i \\
&+\rho\phi_i \mathbb{E}_{z,\hat{\psi}}\left[\max_{d\in[\delta,1]} dU_i(\hat{\psi}) + (1-d)\left[V_i(z,\hat{\psi}) + \lambda_e^i \max_v \left\{p(\theta_i(v,\hat{\psi}))\left(v - V_i(z,\hat{\psi})\right)\right\}\right]\right] f_i(z) \\
&+\rho(1-\phi_i)\mathbb{E}_{\hat{\psi}}\left[\max_{d\in[\delta,1]} dU_i(\hat{\psi}) + (1-d)\left[\tilde{V}_i(\hat{\psi}) + \lambda_e^i \max_v \left\{p(\theta_i(v,\hat{\psi}))\left(v - \tilde{V}_i(\hat{\psi})\right)\right\}\right]\right]
\end{aligned} \tag{3.3}$$

In the current period, the expected output of the match is xy_i. In the next period, the firm and the worker learn the quality z of their match with probability ϕ_i. The worker leaves the match for unemployment with probability d. In this case, the joint continuation value is $U_i(\hat{\psi})$. The worker searches on-the-job and finds a new job with probability $(1-d)\lambda_e^i p(\theta_i(v,\hat{\psi}))$. In this case, the joint continuation value is v. The worker and the firm remain together with probability $(1-d)(1-\lambda_e^i p(\theta_i(v,\hat{\psi})))$. In this case, the joint continuation value is $V_i(z,\hat{\psi})$. The firm and the worker do not learn the quality of their match with probability $1-\phi_i$. Conditional on not knowing the quality of their match, if the worker and the firm remain together, their joint continuation value is $\tilde{V}_i(\hat{\psi})$. Note that, since employment contracts are bilaterally efficient, the choice of d and v is contingent on whether the quality of the match is learned or not and, if it is, the observed z.

The tightness $\theta_i(v,\psi)$ of submarket $s=\{v,i\}$ is such that

$$k_i \geq q(\theta_i(v,\psi))\left[\tilde{V}_i(\psi) - v\right], \tag{3.4}$$

and $\theta_i(v) \geq 0$, with the two inequalities holding with complementary slackness. The left-hand side of (3.4) is the cost to a firm from opening a vacancy in submarket x. The right-hand side is the benefit to the firm from opening a vacancy in submarket x. The benefit is the probability that the firm fills its vacancy, $q(\theta_i(v))$, times the firm's value from filling a vacancy, $\tilde{V}_i(\psi) - v$, i.e. the joint value of a match between the firm and a worker of type i net of the lifetime utility promised by the firm to the worker. Condition (3.4) then states that the cost and benefit of a vacancy in submarket x must be equal if the vacancy-to-applicant ratio is strictly positive. The cost of a vacancy is greater than or equal to the benefit if the vacancy-to-applicant ratio is equal to zero. In submarkets with some applicants, the condition guarantees that the tightness is consistent with firm's profit maximization. In submarkets without applicants, the condition pins down the agents' expectations about the tightness.

We now turn to characterizing the solution to the search and separation problems in (3.1), (3.2), (3.3). The search problem for a worker of type i who currently is an employment state with value v_0 is given by

$$D_i(v_0, \psi) = \max_v p(\theta_i(v, \psi))(v - v_0). \tag{3.5}$$

For any v such that $\theta_i(v, \psi) > 0$, (3.4) implies that v is equal to $\tilde{V}_i(\psi) - k_i/q(\theta_i(v, \psi))$ and, hence, the objective function in (3.5) is equal to $p(\theta_i(v, \psi))(\tilde{V}_i(\psi) - v_0) - k\theta_i(v, \psi)$. For any v such that $\theta_i(v, \psi) = 0$, $p(\theta_i(v, \psi)) = 0$ and, hence, the objective function in (3.5) is also equal to zero or, equivalently, to $p(\theta_i(v, \psi))(\tilde{V}_i(\psi) - v_0) - k\theta_i(v, \psi)$.

The above observations allow us to rewrite the search problem in (3.5) as

$$D_i(v_0, \psi) = \max_v -k_i\theta_i(v, \psi) + p(\theta_i(v, \psi))(\tilde{V}_i(\psi) - v_0). \tag{3.6}$$

Notice that, for all $\theta \geq 0$, there exists a v such that $\theta_i(v, \psi) = \theta$. Thus, by changing the choice variable from v to θ in (3.6), we do not enlarge the choice set. Conversely, for all v, there exists a $\theta \geq 0$ such that $\theta = \theta_i(v, \psi)$. Thus, by changing the choice variable from v to θ in (3.6), we do not shrink the choice set. Since the choice set is the same whether the worker chooses v or θ, we can rewrite (3.6) as

$$D_i(v_0, \psi) = \max_{\theta \geq 0} -k_i\theta + p(\theta)(\tilde{V}_i(\psi) - v_0). \tag{3.7}$$

In words, the worker chooses the tightness θ of the submarket in which to search so as to maximize the probability of meeting a firm, $p(\theta)$, times the value of entering a new match, $\tilde{V}_i(\psi) - v_0$, net of the cost of opening tightness θ vacancies, $k_i\theta$.

The solution to the worker's search problem in (3.7) satisfies the following necessary and sufficient condition for optimality

$$k_i \geq p'(\theta)(\tilde{V}_i(\psi) - v_0), \tag{3.8}$$

and $\theta \geq 0$, where the two inequalities hold with complementary slackness. In words, (3.8) states that, if the worker searches in a submarket with a strictly positive tightness, the cost, k_i, of searching in a submarket with a marginally higher tightness must be equal to the benefit, $p'(\theta)(\tilde{V}_i(\psi) - v_0)$, of searching in a submarket with a marginally higher tightness. If the worker searches in a submarket with zero tightness, the marginal cost must be greater of equal to the marginal benefit. We denote as $\theta^*_{i,u}(\psi)$ the optimal search strategy for a worker of type i who is unemployed. That is, $\theta^*_{i,u}(\psi)$ denotes the solution to (3.8) for $v_0 = U_i(\psi)$. Similarly, we denote as $\theta^*_{i,e}(z, \psi)$ the optimal search strategy for a worker of type i who is employed in a match of known quality z. That is, $\theta^*_{i,e}(z, \psi)$ is the solution to (3.8) for $v_0 = V_i(z, \psi)$. By evaluating (3.8) at $v_0 = \tilde{V}_i(\psi)$, it is clear that a worker of type i who is employed in a match of unknown quality finds it optimal to search in a submarket with zero tightness – hence, matches of unknown quality have no incentive to search.

Next, we turn to the characterization of the separation problems in (3.2) and (3.3). The optimal separation probability for a firm and a worker of type i who are in a match with some joint value v_0 is determined by the sign of the following inequality

$$U_i(\psi) \lesseqgtr v_0 + \lambda_e^i D_i(v_0, \psi). \tag{3.9}$$

The left-hand side is the firm's and worker's joint value of breaking up at the separation stage. The right-hand side is the firm's and worker's joint value of remaining together at the separation stage. If the left-hand side is greater than the right-hand side, then the joint value of the match is strictly increasing in d and the optimal separation probability is 1. Otherwise, the optimal separation probability is δ. We denote as $d_i^*(z)$ the optimal separation probability for a firm and a worker in a match of known quality z. That is, $d_i^*(z)$ denotes the optimal separation probability for $v_0 = V_i(z, \psi)$. Since $V_i(z, \psi)$ is strictly increasing in z, it follows that there exists a reservation quality $R_i(\psi)$ such that $d_i^*(z) = 1$ for all $z < R_i(\psi)$ and $d_i^*(z) = \delta$ for all $z \geq R_i(\psi)$. Similarly, we denote as $\tilde{d}^*$ the optimal separation probability for a firm and a worker in a match of unknown quality.

Lastly, we formulate the laws of motion for the distribution of workers across employment states. The law of motion for the measure u_i of workers of type i who are unemployed is given by

$$\begin{aligned} \hat{u}_i \quad &= u_i(1 - \lambda_u^i(\theta_{i,u}^*(\psi))) \\ &+ \textstyle\sum_z \left[(g_i(z) + n_i \phi_i f_i(z)) d_i^*(z, \psi)\right] + n_i(1 - \phi_i)\tilde{d}_i^*(\psi). \end{aligned} \tag{3.10}$$

The left-hand side of (3.10) is the measure of unemployed workers at the beginning of next period. The first term on the right-hand side of (3.10) is the measure of workers who are unemployed at the beginning of the current period and do not find a job. The second term sums the measure of workers who are employed in a match of known quality at the beginning of the current period and become unemployed at the separation stage with the measure of workers who are employed in a match of unknown quality at the beginning of the current period, discover the quality of their match at the learning stage, and become unemployed at the separation stage. The last term on the right-hand side of (3.10) is the measure of workers who are employed in a match of unknown quality at the beginning of the current period, do not discover the quality of their match at the learning stage, and become unemployed at the separation stage.

The law of motion for the measure n_i of workers of type i who are employed in a match of unknown quality is given by

$$\begin{aligned} \hat{n}_i \quad &= n_i(1 - \phi_i)(1 - \tilde{d}_i^*(\psi)) + u_i \lambda_u^i(\theta_{i,u}^*(\psi)) \\ &+ \textstyle\sum_z (g_i(z) + n_i \phi_i f_i(z))(1 - d_i^*(z, \psi)) \lambda_e^i(\theta_{i,e}^*(z, \psi)). \end{aligned} \tag{3.11}$$

The left-hand side of (3.11) is the measure of workers employed in a match of unknown quality at the beginning of next period. The first term on the right-hand side of (3.11) is the measure of workers who are employed in a match of unknown quality at the beginning of the current period, do not discover the quality of their match at the learning stage, and remain on their job. The second term is the measure of workers who are unemployed at the beginning of the current period and find a job at the search stage. The last term is the measure of workers who are employed at the beginning of the current period and move to a new job during the search stage.

The law of motion for the measure $g_i(z)$ of workers of type i who are employed in a match of known quality z is given by

$$\begin{aligned} \hat{g}_i(z) \quad &= g_i(z)(1 - d_i^*(z,\psi))(1 - \lambda_e^i p(\theta_{i,e}^*(z,\psi))) \\ &+n_i \phi_i f_i(z)(1 - d_i^*(z,\psi))(1 - \lambda_e^i p(\theta_{i,e}^*(z,\psi))). \end{aligned} \tag{3.12}$$

The left-hand side of (3.12) is the measure of workers employed in a match of quality z at the beginning of next period. The first term on the right-hand side is the measure of workers who are employed in a match of quality z at the beginning of the period and remain on their job. The second term is the measure of workers who are employed in a match of unknown quality at the beginning of the period, discover that the quality of their match is z during the learning stage, and remain on their job.

We are now in the position to define a Recursive Equilibrium.

Definition 1. *A Recursive Equilibrium (RE) is given by value functions $\{U_i, \tilde{V}_i, V_i\}$, policy functions $\{d_i^*, \tilde{d}_i^*, \theta_{i,u}^*, \theta_{i,e}^*\}$, and a transition probability function $\Phi(\hat{\psi}|\psi)$ for the aggregate state of the economy such that: (i) $\{U_i, \tilde{V}_i, V_i\}$ satisfy (3.1), (3.2) and (3.3); (ii) $\{d_i^*, \tilde{d}_i^*, \theta_{i,u}^*, \theta_{i,e}^*\}$ satisfy the optimality conditions (3.8) and (3.9); Φ is consistent with the laws of motion (3.10), (3.11) and (3.12) for $\{\hat{u}_i, \hat{n}_i, \hat{g}_i\}$ and with the probability distribution for $\hat{x}$.*

We can also define a Block Recursive Equilibrium.

Definition 2. *A Block Recursive Equilibrium (BRE) is a RE in which the value and policy functions depend on the aggregate state of the economy $\psi = \{x, u_i, n_i, g_i\}$ only through the exogenous state of productivity x and not through the endogenous distribution of workers across employment states $\{u_i, n_i, g_i\}$.*

As proved in Menzio and Shi (2011), there exists a unique BRE and no other RE. Moreover, the BRE is efficient in the sense that it decentralizes the solution of the problem of a utilitarian social planner. Let us provide some intuition for the existence and uniqueness of a BRE. A BRE exists because search is directed, and the search and production processes feature constant returns to scale. Consider the equilibrium condition (3.4) for the tightness of submarket $s = \{v, i\}$. Since the production process features constant returns to scale, the value to the firm from filling a vacancy in submarket s depends on the aggregate productivity x and on the promised value v and not on the distribution of workers $\{u_i, n_i, g_i\}$. Since the search process is directed, the probability that a worker contacted in submarket s is willing to accept the value offered by the firm is always equal to one. Since the search process features constant returns to scale, the probability of contacting a worker in submarket s only depends on the tightness of the submarket and not on the measure of workers searching there. The above observations imply that the tightness of submarket s depends on x but not on the distribution of workers

$\{u_i, n_i, g_i\}$. In turn, this implies that the value functions—as well as the associated policy functions—are all independent of the endogenous distribution of workers across employment states. The uniqueness of the BRE follows from the fact that the value functions can be combined in a single operator and this operator is a contraction.

4 Calibration

In this section, we calibrate our model with heterogeneous workers using the empirical evidence emerging from the LEHD. In Section 4.1, we describe and explain the calibration strategy. In Section 4.2, we present and discuss the calibration outcomes.

4.1 Calibration strategy

We calibrate the parameters of the model using moments generated by the model at its non-stochastic steady state. The non-stochastic steady-state is defined as the steady-state associated with a version of the model in which the aggregate component of productivity x is kept constant and equal to the unconditional mean of its stochastic process $h(\hat{x}|x)$. We then calibrate the parameters of the model by comparing the moments generated by the model at its non-stochastic steady-state with the same moments computed in the data during the period preceding the Great Recession. For this reason, our calibration targets differ somewhat from Table 1, which uses the whole sample period.

Let us begin by reviewing the parameters that need to be calibrated. The parameters describing the production process are: (i) the unconditional mean x^* of the aggregate component of productivity, which we normalize to 1; (ii) the component of productivity y_i that is specific to a worker of type i; (iii) the probability ϕ_i that a worker of type i and a firm discover the quality of their match; (iv) the distribution f_i of the component of productivity z that is specific to a match between a particular worker of type i and a particular firm. We specialize f_i to be a Weibull distribution with shape parameter ω_i and scale parameter σ_i that is appropriately relocated so as to have a mean of one. The Weibull distribution is flexible and, depending on the shape parameter ω_i, its shape can resemble an exponential, a log-normal, a normal, or a left-skewed distribution.

The parameters describing the search process are: (i) the probability λ_u^i that a worker of type i can search the labor market when unemployed; (ii) the probability λ_e^i that a worker of type i can search the labor market when employed; (iii) the probability $p(\theta)$ that an applicant meets a vacancy as a function of the tightness θ; (iv) the probability δ that a firm-worker match breaks up for exogenous reasons. We specialize $p(\theta)$ to have the form $p(\theta) = \min\{\theta^\gamma, 1\}$, where γ denotes the elasticity of the job-finding probability with respect to tightness and is set to 0.5.

The parameters describing the demographic structure and the preferences are: (i) the measure μ_i of workers of type i; (ii) the worker's and firm's discount factor ρ; (iii) the sum b_i of the unemployment income and the value of leisure for workers of type i. We specialize b_i to be of the form $\zeta + r\mathbb{E}[xy_i z]$, where ζ denotes the value of leisure and r denotes the fraction of the average productivity $\mathbb{E}[xy_i z]$ that is replaced by unemployment benefits.

Let us now turn to the calibration strategy. To be consistent with the empirical analysis presented in Section 2, we calibrate the model to have three types of workers, which we refer to as α, β and γ. The measure of workers of type α, β and γ is set equal to the fraction of workers of each type in the data, i.e. 57% of workers are α, 26% are β and 17% are γ.

We choose the model period to be one month and the discount factor ρ sets the annual real interest rate in the model to 5%. We choose the vacancy cost k_i so that the non-stochastic steady-state of the model reproduces the average unemployment rate of the different worker types in the period preceding the Great Recession. Specifically, the average unemployment rate is 4.2% for α-workers, 12.5% for β-workers, and 28.8% for γ-workers. The choice of the calibration target is natural, as the vacancy cost affects the speed at which unemployed workers find a job and, in turn, the unemployment rate.

The scale σ_i of the match quality distribution is chosen so that the model reproduces the empirical fraction of firm-worker matches that terminates before exceeding 24 months of tenure. The fraction of matches lasting no more than 24 months is very different for different types of workers: it is about 50% for α-workers, 60% for β-workers, and 85% for γ-workers. The choice of the calibration target is natural, as the dispersion of the match quality distribution affects the probability that the quality of the match is such that the worker wants to move into unemployment, as well as the probability that the quality of the match is such that the worker wants to search on the job. In turn, these probabilities determine the probability that the firm-worker match terminates before exceeding a 24-month tenure. We also choose the exogenous termination probabilities δ_i to match the empirical UE rates by type. These rates are 30% per month for α-workers, 15% per month for β-workers, and 10% per month for γ-workers. In the model, δ_i directly effects the number of flows into unemployment, which feeds into the denominator of the UE rate. Matching the UE rate also helps for replicating the empirical distribution of unemployment durations by type, even though these are not explicitly targeted.

The probability λ_u^i with which an unemployed worker gets to search the labor market is normalized to 1 for all worker types. The probability λ_e^i with which an employed worker gets to search the labor market is chosen so that the model reproduces the empirical fraction of firm-worker matches that last no more than 24 months and terminate with the worker moving directly to another employer. For α-workers, this fraction is 21% (about one-half of matches that last no more than 24 months). For β-workers, it is 18% (about one-third of matches that last no more than 24 months). For γ-workers, it is 22% (or about one-fourth of matches that last no more than 24 months). Again, the choice of the calibration target is natural, as λ_e^i affects the speed at which workers can directly move from one employer to another.

The shape ω_i of the match-quality distribution and the probability ϕ_i with which a firm-worker pair discovers the quality of their match are chosen to reproduce the whole shape of the tenure distribution. Specifically, the parameters are chosen so as to minimize the distance between the model and the data with respect to: (i) the fraction of firm-worker matches that terminate before exceeding 3 months of tenure, 12 months of tenure, and 24 months of tenure; (ii) the fraction of firm-worker matches that terminate with the worker moving to another employer before exceeding 3, 12 and 24 months of tenure; (iii) the fraction of firm-worker matches that terminate with the worker moving into unemployment before exceeding 3, 12 and 24 months of tenure. That is, ω_i and ϕ_i are chosen to fit the shape of the tenure distribution (unconditional, and conditional on the type of termination). The shape of the tenure is quite different for different types. For instance, for α-workers, the fraction of matches ending within the first 3 months is lower than the fraction of matches ending between 13 and 24 months. For γ-workers, the fraction of matches ending within the first 3 months is much higher than the fraction of matches ending between 13 and 24 months. Our choice of these calibration targets for ω_i and ϕ_i is natural, as ϕ_i determines how quickly low-quality matches are identified, and ω_i determines the right-tail of the match-quality distribution (and, hence, the incentives to searching for a better match).

We normalize the component of productivity y_α that is specific to α-workers to 1. We choose the component of productivity y_i for $i = \{\beta, \gamma\}$ so that the model-generated ratio between the average productivity among employed workers of type i and the average productivity among employed workers of type α is equal to the empirical ratio between the average earnings of employed workers of type i and the average earnings of employed workers of type α. The attentive reader may have noticed that in the calibration of y_i we compare productivity in the model with earnings in the data. We do so because computing productivity is easier than computing wages and, for the most common specification of the wage process (e.g., the wage is set equal to some fraction of the worker's productivity as in Bagger et al. 2014 or Menzio, Telyukova and Visschers 2016), the difference between productivity and wage turns out to be negligible as search frictions are small.

Lastly, we need to calibrate the parameters associated with the unemployment income. Shimer (2005) argues that unemployment income should be set to 40% of average productivity, as this is the typical replacement rate in the US. Hagedorn and Manovskii (2008) point out that unemployment income should also include the value of leisure. Hall and Milgrom (2008) argue that, on average, the ratio between unemployment income (unemployment benefits plus value of leisure) is about 65% of employment income. Based on these observations, we choose the replacement ratio r of unemployment benefits for workers of type i to be equal to 40% of the average productivity among employed workers of type i. We then choose the value of leisure ζ so that the ratio between unemployment income and labor productivity is, on average, equal to 65%.

4.2 Calibration outcomes

Table 2 reports the calibrated values of the parameters. Some comments about the calibration outcomes are worthwhile. First, note that different types of workers face a different match-quality distribution, as illustrated in Figure 1. For α-workers, the match-quality distribution is a Weibull with shape 4.5 and scale 0.25. The distribution is approximately normal, with a

Parameter	Value	Description
β	0.996	discount factor
b_i	(0.638, 0.452, 0.316)	flow unemployment income
y_i	(1, 0.623, 0.459)	type-specific productivity
α_i	(4.515, 3.941, 0.640)	shape of f_i
σ_i	(0.058, 0.143, 0.082)	standard deviation of f_i
ϕ_i	(0.307, 0.233, 0.229)	probability match quality is discovered
λ_e^i	(0.151, 0.493, 0.641)	probability an employed worker searches
λ_u^i	1	probability an unemployed worker searches
δ_i	(0.006, 0.009, 0.005)	exogenous separation probabilities
k_i	(3.536, 8.332, 4.564)	vacancy posting cost
γ	0.5	elasticity of job-finding rate wrt tightness
χ	0.004	exogenous labor market exit probability

Table 2: Model parameters

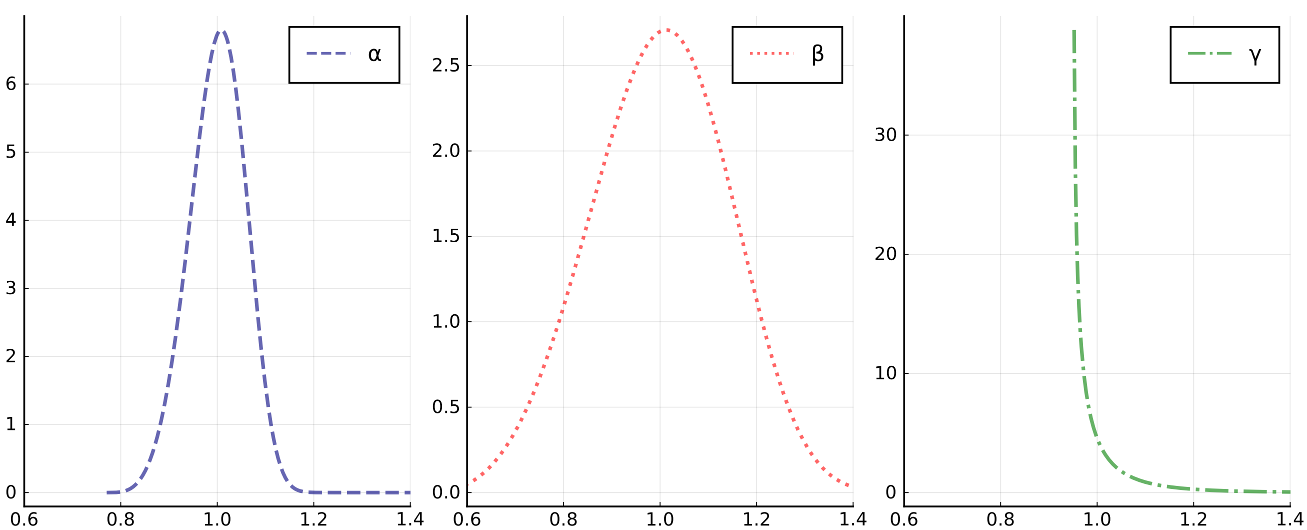

Figure 1: Distributions from which z is drawn for different types of workers

	α-workers					
	Data			**Model**		
	Total	Ends in E	Ends in U	Total	Ends in E	Ends in U
<1Q	0.122	0.077	0.045	0.099	0.088	0.010
1Q-4Q	0.162	0.087	0.075	0.193	0.111	0.082
5Q-8Q	0.207	0.101	0.105	0.138	0.050	0.087
>8Q	0.510	0.258	0.252	0.571	0.394	0.177
	β-workers					
	Data			**Model**		
	Total	Ends in E	Ends in U	Total	Ends in E	Ends in U
<1Q	0.180	0.141	0.040	0.149	0.138	0.011
1Q-4Q	0.209	0.143	0.066	0.302	0.206	0.096
5Q-8Q	0.219	0.147	0.072	0.154	0.065	0.089
>8Q	0.392	0.289	0.103	0.395	0.298	0.098
	γ-workers					
	Data			**Model**		
	Total	Ends in E	Ends in U	Total	Ends in E	Ends in U
<1Q	0.354	0.279	0.075	0.308	0.305	0.003
1Q-4Q	0.309	0.219	0.090	0.435	0.404	0.031
5Q-8Q	0.192	0.132	0.060	0.081	0.051	0.031
>8Q	0.144	0.103	0.041	0.176	0.137	0.039

Table 3: Employment duration moments

mean of 1, a standard deviation of 0.06, a skewness of -0.17, and a 90-50 percentile ratio equal to 90% of the 50-10 percentile ratio. For β-workers, the match-quality distribution is a Weibull with shape 3.9 and scale 0.55. The distribution is approximately normal, with a mean of 1, a standard deviation of 0.14, a skewness of -0.07, and a 90-50 percentile ratio equal to 93% of the 50-10 percentile ratio. For γ-workers, the match-quality distribution is a Weibull with shape 0.64 and scale 0.04. The distribution is approximately exponential, with a mean of 1, a standard derivation of 0.08, a skewness of 4.12, and a 90-50 percentile ratio that is 6 times larger than the 50-10 percentile ratio. In words, the match-quality distributions for α and β workers are approximately normal with different variances. The match-quality distribution for γ-workers is left-skewed with a compressed left-tail and a long right-tail.

The differences between the match-quality distributions are needed to rationalize the differences in the observed behavior of different types of workers. Workers of type α have a 50% probability of remaining on their job for more than 2 years, and a 10% probability of leaving their job within the first 3 months. In order to rationalize these facts, the match-quality distribution has a relatively thin left-tail—so that the fraction of matches that are below the reservation quality is low—and a relatively thin right-tail—so that workers in average-quality matches have no incentive to search on the job. Workers of type γ, in contrast, have a 15% probability of remaining on their job for more than 2 years, and a 35% probability of leaving their job within the first three months. In order to rationalize these facts, the match-quality

	Additional Targeted Moments					
	α-workers		**β-workers**		**γ-workers**	
	Data	**Model**	**Data**	**Model**	**Data**	**Model**
Unemployment rate	0.042	0.042	0.125	0.124	0.288	0.297
UE rate	0.300	0.301	0.150	0.151	0.100	0.101
Replacement rate	0.650	0.643	0.750	0.747	0.900	0.890
	Unemployment Duration (Not Targeted)					
	α-workers		**β-workers**		**γ-workers**	
	Data	**Model**	**Data**	**Model**	**Data**	**Model**
<1Q	0.697	0.514	0.315	0.280	0.485	0.193
1Q-4Q	0.240	0.467	0.628	0.555	0.380	0.497
>4Q	0.070	0.019	0.058	0.166	0.135	0.310
Average	1.082	1.104	1.888	2.213	2.791	3.301

Table 4: Additional moments, targeted and non-targeted

distribution has a relatively thick left-tail—so that the fraction of matches that are below the reservation quality is high—and a relatively thick right-tail—so that workers have an incentive to search on the job until they find a match in the right-tail. We compare the complete set of moments for job durations in the model and data in Table 3, and the fit is overall very good for all types of workers and durations.

Second, let us comment on the calibrated unemployment income. For α-workers, the calibrated unemployment income is 0.64, which is approximately equal to 64% of their average labor productivity. For β-workers, the calibrated unemployment income is 0.45, which is approximately equal to 74% of their average labor productivity. For γ-workers, the calibrated unemployment income is 0.32, which is approximately equal to 89% of their average labor productivity.

Lastly, we wish to comment on the calibrated search process. For α-workers, the cost of maintaining a vacancy is 3.5 units of output per month, the probability of searching for a vacancy is 1 off the job and 15% on the job. These parameters (together with the others) imply an unemployment rate of 4.2%, an UE rate of 30% per month, an EE rate of 0.6% per month, and an EU rate of 0.9% per month. For β-workers, the cost of maintaining a vacancy is 8.3 units of output per month, the probability of searching for a vacancy is 1 off the job and 49% on the job. These parameters imply an unemployment rate of 12.4%, a UE rate of 15% per month, an EE rate of 0.8% per month, and an EU rate of 1.7% per month. For γ-workers, the cost of maintaining a vacancy is 4.6 units of output per month, the probability of searching for a vacancy is 1 off the job, and 64% on the job. These parameters imply an unemployment rate of 29.7%, a UE rate of 10% per month, an EE rate of 0.4% per month, and an EU rate of 3.8% per month. Note that we assumed that the job-finding probability $p(\theta)$ is the same for all worker types and, hence, any differences in the unemployment rate that are not accounted by differences in the value of a new match are captured by differences in the vacancy cost. Alternatively, we could have assumed that the vacancy cost was the same for all workers, and let the job-finding probability $p(\theta)$ have a different efficiency for different workers.

The search process parameters also impact the distribution of unemployment duration. Although not targeted, the model fit for this is shown in the bottom panel of Table 4. The model generally does well on this dimension, although it generates slightly longer durations than in the data.

5 Micro Implications of Heterogeneity

In this section, we use our calibrated model to illustrate how worker heterogeneity matters for understanding two well-documented phenomena in labor economics. In Section 5.1, we show that the earnings losses for displaced workers are, on average, large and persistent, but average earnings losses hide dramatic differences across worker types. We then show that our calibrated model can easily explain these findings. In Section 5.2, we show that our calibrated model can explain much of the decline of the UE rate as a function of unemployment duration—bringing new data to support the role of dynamic composition in this phenomenon.

5.1 Earning losses of displaced workers

In this subsection, we use the LEHD, our classifications of workers into types, and our calibrated model to analyze the effect of displacement on earnings. We find that, when a long-tenured worker is displaced, the decline in earnings is large and persistent—consistent with what has been documented in Jacobson, LaLonde and Sullivan (1993). We find that the decline in earnings is larger during a recession—consistent with what has been documented in Davis and von Wachter (2011). We find that average earnings losses hide substantial heterogeneity, as the effect of displacement on earnings varies dramatically across types. We also find that our model successfully reproduces these empirical findings and offers a simple explanation for the magnitude and persistence of earning losses for displaced workers.

Using the LEHD, we identify those who, after reaching 3 years of tenure with the same firm, move into unemployment. We also identify workers who, after reaching 3 years of tenure with the same firm, either remain with the firm or directly move to another firm. We pool these workers and, for each of them, we identify their type i. We then run the following regression in the spirit of Jacobson, LaLonde and Sullivan (1993)

$$\log y_{j,t} = a_j + b \cdot age_{j,t} + \sum_{k=0,1,..} c_{k,i(j)} \cdot \mathbf{1}_t[k] + \epsilon_{i,t},$$

where $y_{j,t}$ are the earnings of worker j in quarter t, a_j is a worker-specific fixed effect, $age_{j,t}$ is the age of the worker at date t, $\mathbf{1}_t[k]$ is an indicator function that takes the value 1 if and worker separated from the firm in quarter $t-k$ and 0 otherwise, and $\epsilon_{i,t}$ is a mean-zero error. The coefficients of interest are $c_{k,i(j)}$, which can be interpreted as the percentage change between the wage of a worker of type i displaced k quarters ago and the wage that the same worker

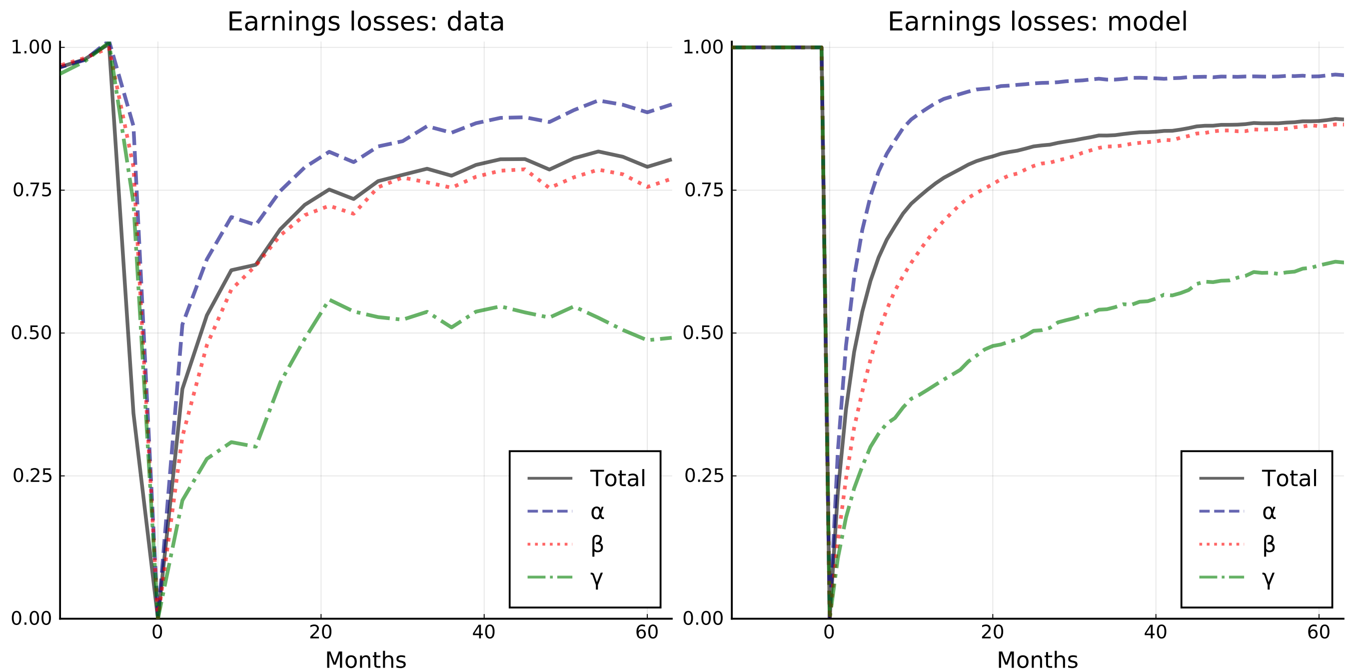

Figure 2: Earnings losses from job separation in steady state.

would have earned had he not been displaced, where this counterfactual wage is based on the worker's fixed effect adjusted for the change in age.[4]

The left panel of Figure 2 illustrates our findings.[5] The black line is the average path of earnings for a displaced worker, i.e. the path of estimated c_k in a regression that does not allow these coefficients to differ across types. Consistent with Jacobson, LaLonde and Sullivan (1993), we find that the earnings losses following displacement are large and persistent. Specifically, the average earnings are about 61% of pre-displacement earnings after 1 year and reach 81% of pre-displacement earnings after 5 years. The blue line ($c_{k,\alpha}$), the red line ($c_{k,\beta}$) and the green line ($c_{k,\gamma}$) are the average paths of earnings for displaced workers of different types and, as one can immediately see, are very different. For α-workers, earnings losses are about 29% after 1 year, and 10% after 5 years. For β-workers, earnings losses are about 42% after 1 year, and 22% after 5 years. For γ-workers, earnings losses are about 69% after 1 year, and still about 50% after 5 years. Note that, since workers with long tenure are self-selected into high-quality matches, their earnings should not be expected to fully recover.

Figure 3 is the analogue of Figure 2 using a sample of workers who are displaced during the Great Recession. Figure 3 confirms the findings of Davis and von Wachter (2013), as it shows that the average effect of displacement on earnings is larger and more persistent during the recession than before the recession. Moreover, the figure shows that the same is true for every type of worker. The largest change, in both model and data comes early, noting especially the

[4]Because the precise timing of the loss in the data is somewhat ambiguous for workers who do not experience a full quarter of non-employment, given time aggregation, we do not include any earnings from these workers during the quarter in which we infer a separation.

[5]For comparison with the model in the right panel, we interpolate this quarterly data to a monthly frequency.

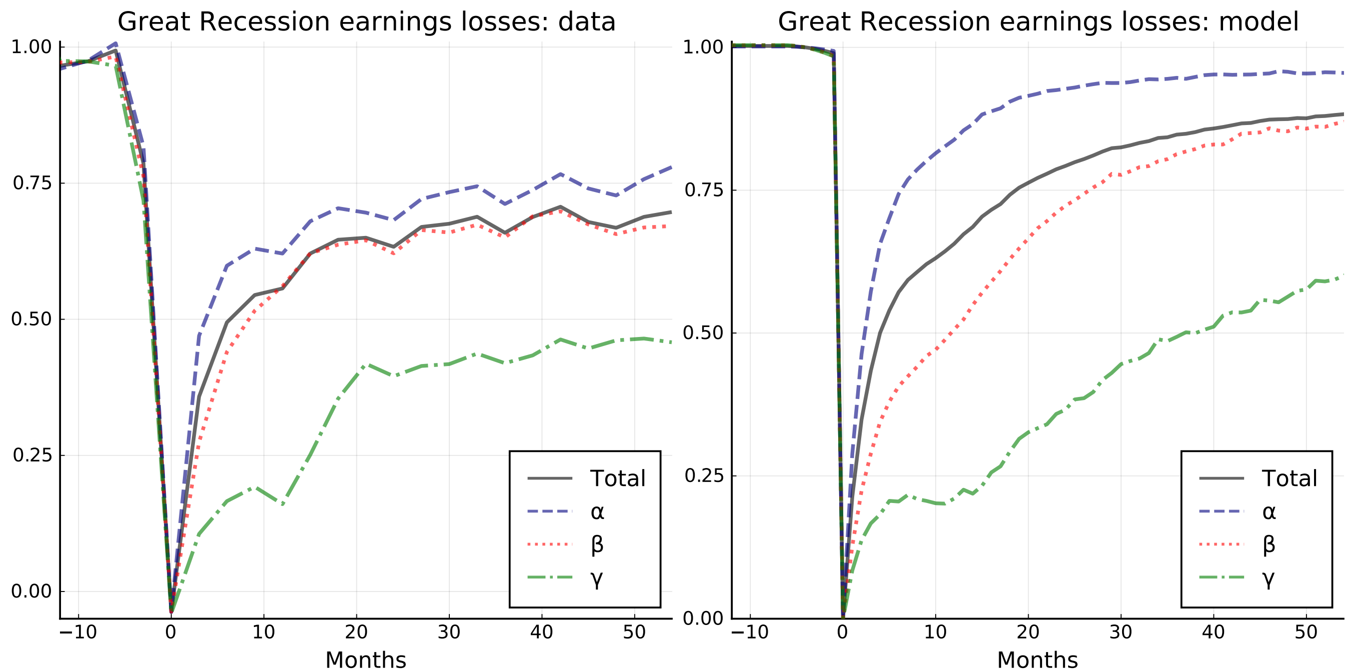

Figure 3: Earnings losses from job separation during the Great Recession

γ-types in the first year after separation.

The magnitude and persistence of earnings losses following displacement has long been seen as a puzzle in the macro-labor literature (see, e.g., Davis and von Wachter 2013 or Jarosch 2016). Indeed, it seems hard to rationalize the fact that it takes workers so many years to recoup their earnings losses following a displacement. In order to address this "puzzle," we turn to our calibrated model. The right panel of Figure 2 reports the model-generated path of labor productivity for workers who, after reaching 3 years of tenure with the same firm, move into unemployment. As one can see, the model reproduces quite well the empirical path of earnings following a displacement. Specifically, in the model as in the data, α-workers display the fastest recovery from displacement (their productivity is 32% lower than its pre-displacement level after 4 quarters, and 5% lower after 5 years), β-workers display a slower recovery (productivity is 60% lower after 4 quarters, and 14% lower after 5 years), and γ-workers display the slowest recovery (productivity is 73% lower after 4 quarters, and still 38% lower after 5 years). The right panel of Figure 3 is the model-generated path of labor productivity for displaced workers during a recession, modeled–as we shall detail in the next section—as an aggregate decline in productivity. Again, the predictions of the model line up with the empirical evidence, as the decline in productivity is larger during a recession for all worker types. Figure 4 puts the model's steady state and recession earnings losses on the same axes for easier comparison.

The model offers a simple explanation for the magnitude and persistence of average earnings losses and for why they differ across the types of workers and the cycle. Workers of type α have a relatively high UE rate and they face a relatively compressed match-quality distribution. If an α-worker is displaced from a firm, it will take him relatively little time to find a new firm (since his UE rate is high) and the quality of his match with this new firm is likely to be of

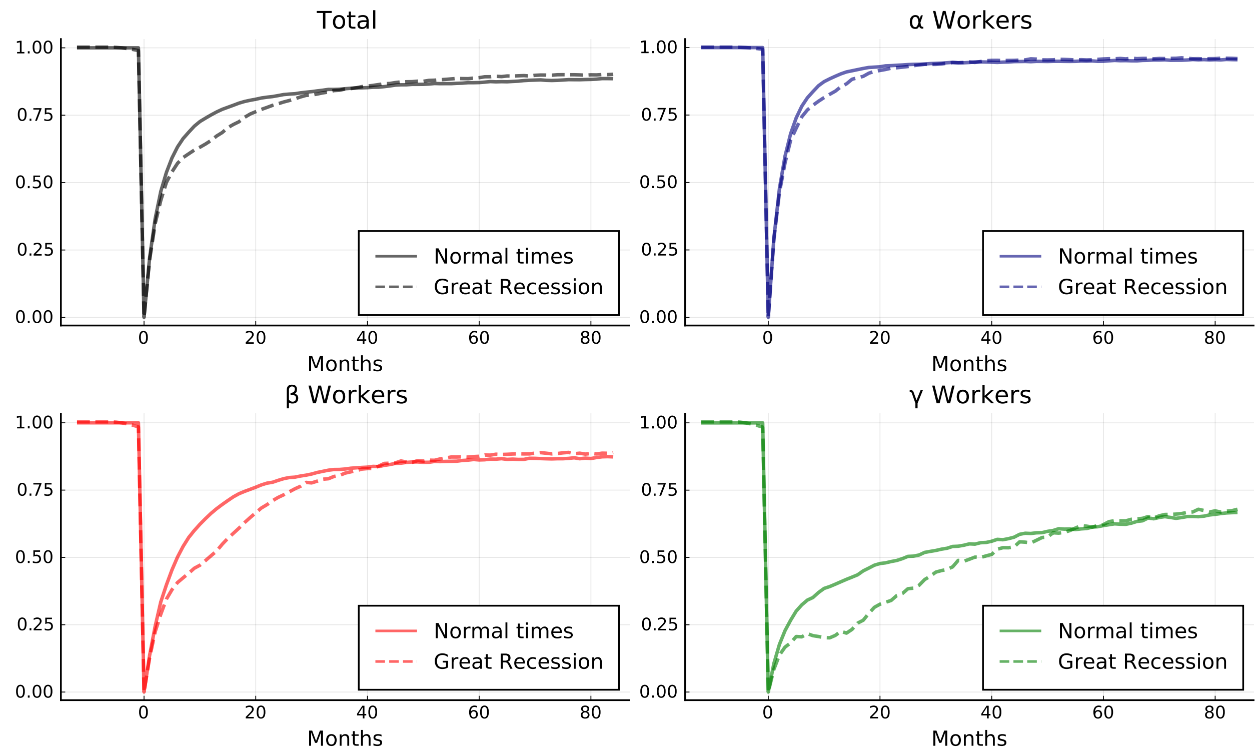

Figure 4: Model-implied earnings losses from unemployment: steady state vs. the Great Recession

similar quality as the match with the old firm (since the match-quality distribution is relatively compressed). In contrast, workers of type γ have a relatively low UE rate, face a match-quality distribution with a relatively thick right-tail, and their probability of keeping a particular job for more than 2 years is very low. If a γ-worker has reached a long tenure with a particular firm, he is likely to be in a match that is at the very top of the quality distribution. If the worker is displaced from that firm, it will take him a relatively long time to locate a new firm (since his UE rate is low) and the quality of his match with this new firm is unlikely to be as good as the match he held before being displaced (since the quality of his match with the old firm was at the very top of the distribution). Indeed, only after experimenting with several firms and after experiencing multiple spells of unemployment, the worker will return to the same level of productivity as before the displacement.

Then, to go from Figure 2 to 3, the recession exacerbates these forces especially for γ-type workers in two ways. The productivity threshold in z for an acceptable match moves up in the distribution f_γ, making it less likely that they'll keep the job once the productivity has been learned. This makes them more exposed to the lowered job-finding rates in recession than if they were only facing one job search. Indeed, in the data the average γ-type job looser in the Great Recession had to find three more jobs by 2013 compared to one more job for the average α-type. This feature amplifies the earnings-effect that comes from their UE rate being more cyclically sensitive, as we will explore in the next section.

Let us briefly comment on the implications of our findings. First, we showed that displace-

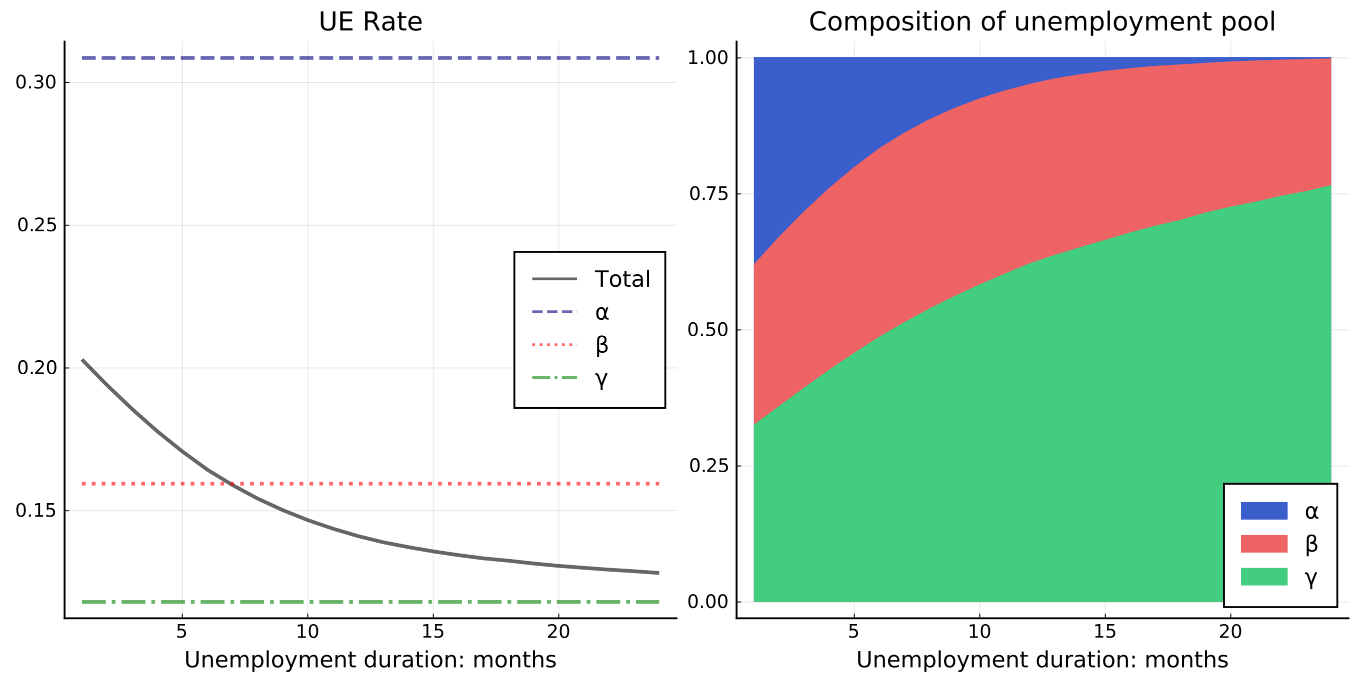

Figure 5: Left panel: type-specific and aggregate UE rates by duration; right panel: composition of unemployment pool by duration

ment from a job has relatively transitory effects on earnings for α-workers, more persistent effects for β-workers, and dramatically large and persistent consequences for γ-workers. Thus, when evaluating the impact of displacement during either aggregate or industry-specific crises, one should identify the type of workers that are being affected. Second, we showed that our model can reproduce the effect of displacement for different types of workers. In our model, the effect of displacement on earnings is caused by the loss of search capital, which is different for different types of workers, and by the rate of accumulation of search capital, which is different for different types of workers. Thus, in order to explain the magnitude and persistence of earnings losses caused by displacement, one need not resort to human capital depreciation, stigmatization, or other exotic sources of scarring.

5.2 Duration dependence of UE

In this section, we use our model to study the relationship between the UE rate and the duration of an unemployment spell. While it is well-known that the UE rate declines with the duration of unemployment, there is still disagreement on whether the decline is caused by a decline in the UE rate of individual workers (true duration dependence) or by the fact that workers with the lowest UE rate become a larger and larger fraction of the pool of unemployed workers (dynamic selection). Using our calibrated model, we show that worker heterogeneity alone accounts for a large decline in the UE rate.

The left panel of Figure 5 plots the UE rate for workers of type α (blue line), β (red line), and γ (green line) as a function of unemployment duration. Figure 5 also plots the overall UE rate (black line) as a function of unemployment duration. By design of the model, the

UE rate for a particular type of worker is independent of unemployment duration. The UE rate is constant at 30% for α-workers, constant at 15% for β-workers, and constant at 10% for γ-workers. Because of the heterogeneity in the UE rate of different types of workers, the overall UE rate declines with unemployment duration, and sharply so. In the first month of unemployment, the overall UE rate is 19%. After 12 months of unemployment, the overall UE rate is only 13%.

The logic behind our findings is standard dynamic selection. The right panel of Figure 5 shows how the composition of the unemployment pool changes with duration. In the first month of unemployment, the fraction of α-workers is 38%, the fraction of β-workers is 30%, and the fraction of γ-workers is 31%. Since the probability of exiting unemployment is highest for α-workers, intermediate for β-workers, and lowest for γ-workers, the composition of the unemployment pool changes with duration. After 12 months of unemployment, the fraction of α-workers is 5%, the fraction of β-workers is 32%, and the fraction of γ-workers is 63%. Since the composition of the unemployment includes fewer α's and more γ's, the overall UE rate falls.

Figure 5 features a decline in the overall UE rate that is similar to the decline observed in the data. For example, Mueller, Spinnewjin and Topa (2019) find that the UE rate is about 30% in the first quarter of unemployment and falls to about 13% in the third and fourth quarters of unemployment. Therefore, our model suggests that much of the decline in the UE rate is a natural consequence of worker heterogeneity, rather than being caused by stigmatization or losses of human capital. This is the same conclusion reached by Mueller, Spinnewjin and Topa (2019), albeit through very different arguments. We use a panel dataset to classify workers into different types and show that they have very different UE rates. Mueller, Spinnewjin and Topa (2019) use data on expected UE rates and their correlation with actual transitions from unemployment to employment to argue that differences in the UE rate across workers can account for much of the decline in the UE rate over unemployment duration. Our conclusion is also in line with the findings in Jarosch and Pilossoph (2019), who argue that the UE rate declines with unemployment duration because the quality of the unemployment pool becomes progressively worse.

6 Macro Implications of Heterogeneity

In this section, we examine the implications of our calibrated model of worker heterogeneity for business cycle analysis. In Section 6.1, we study the effect of a negative aggregate productivity shock to different types of workers and to the aggregate labor market. We find that the response of the UE, EU and unemployment rates is very different for different types of workers. We show that their responses are shaped by the the large and long-lasting effect of a negative productivity shock on γ-workers. Indeed, we show that the increase in aggregate unemployment is large compared to the shock and that the unemployment rate recovers more slowly than the underlying productivity shock. We also show that changes in the composition of employment by worker type and by match quality lead to a "double cleansing" effect of recession and, in turn, to a smaller and less persistent decline in average labor productivity as compared to the

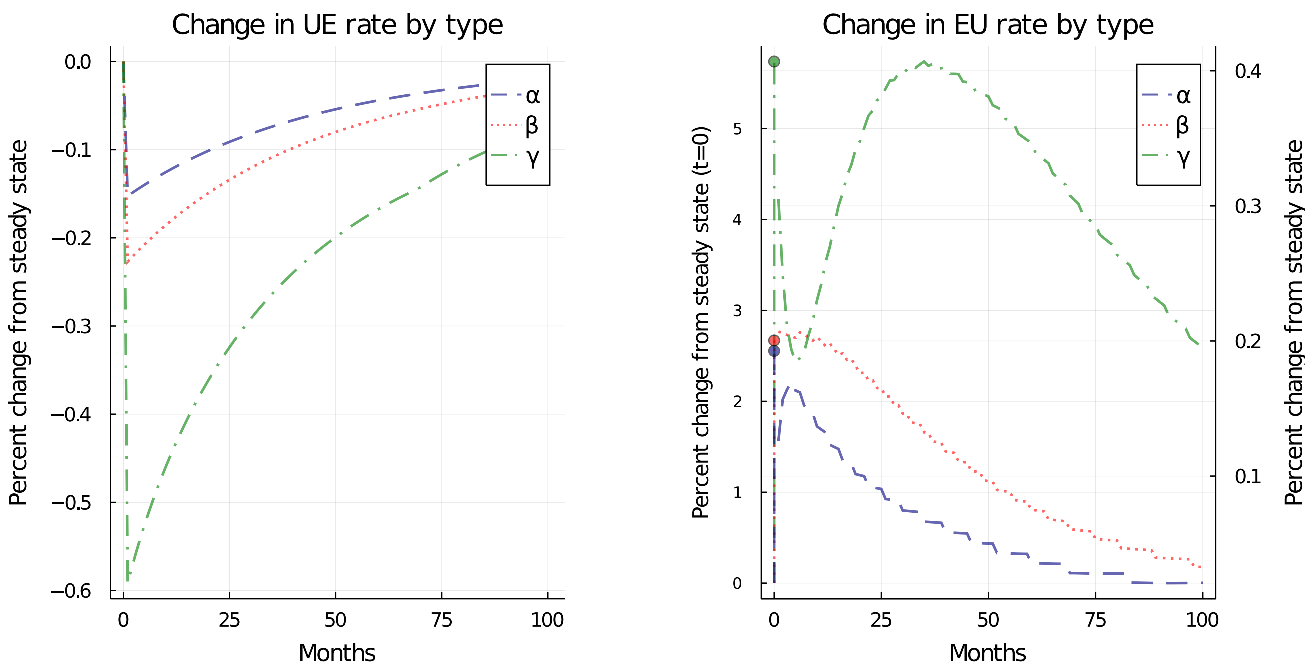

Figure 6: Left panel: model-implied UE rates by type; right panel: model-implied EU rates by type (left axis refers to the response on impact, right axis refers to everything thereafter)

underlying shock. In Section 6.2, we show that the predictions of the model are qualitatively consistent with the behavior of the US labor market during and after the Great Recession. Moreover, we show that there is a particular set of type-specific productivity shocks such that our model matches not only qualitatively, but also quantitatively, the behavior of the US labor market. In particular, the model can generate a large increase in unemployment followed by a slow recovery and a relatively small and highly transitory decline in average labor productivity.

6.1 Aggregate productivity shock

As a starting point, we examine the effect of a 10% decline in the aggregate component of productivity y. The left panel in Figure 6 plots the response of the UE rate by type. The UE rate declines by 15.2% (4.5 percentage points) for α workers, by 22.7% (3.4 percentage points) for β workers, and by 59% (6 percentage points) for γ workers. There is a simple explanation behind the finding that the response of the UE rate is smallest for α-workers and largest for γ-workers. The UE rate depends on the tightness of the submarket visited by unemployed workers and, in turn, on the value of moving from unemployment to employment. The value of moving from unemployment to employment depends on the difference between employment output and unemployment income. This difference is largest for α workers and, hence, a decline in the aggregate component of productivity has the smallest percentage effect on the value of moving from unemployment to employment. In contrast, the difference between employment output and unemployment income is smallest for γ workers and, hence, the same decline in the aggregate component of productivity has the largest percentage effect on the value of moving from unemployment to employment. This is the same logic behind the well-known feature of

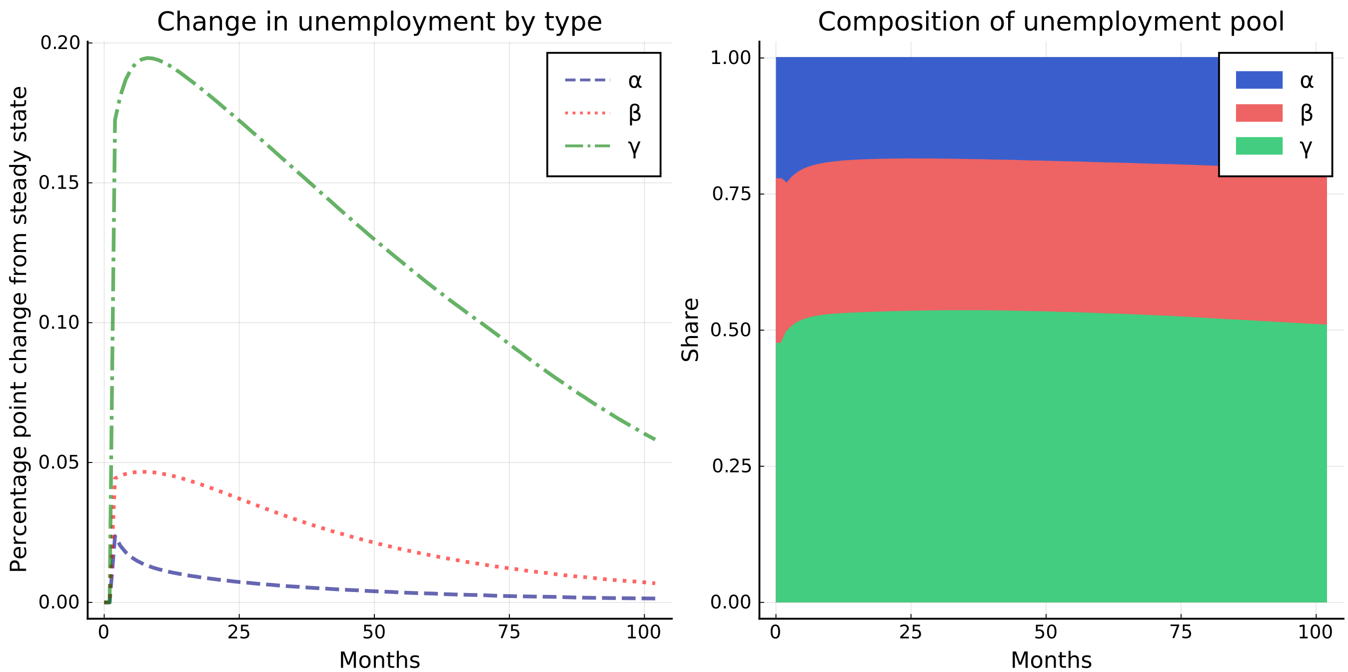

Figure 7: Model-implied unemployment rates and composition in response to a 10% decline in aggregate productivity

Pissarides (1985) that the elasticity of the UE rate to productivity shocks is increasing in the difference between productivity and unemployment income (see, e.g., Hagedorn and Manovskii 2008 or Ljundquivst and Sargent 2017).

The right panel in Figure 6 plots the response of the EU rate by type. The decline in the aggregate component of productivity increases the reservation quality. On impact, the increase in the reservation quality causes a number of marginal matches to be destroyed. Over time, the increase in the reservation quality causes an increase in the fraction of new matches to be destroyed once their quality is observed. Since the density of the cross-sectional distribution of match-qualities is highest for γs and lowest for αs, the initial increase in the EU rate is highest for γs and lowest for αs. The flow of workers into unemployment is reabsorbed over time. However, the process is different for different types. For γ-workers, the process typically involves sampling and abandoning several jobs before finding a match above the reservation quality. For α-workers, the process typically involves finding just one job. For this reason, the EU rate reverts back to normal more slowly for γ-workers than for α-workers. This is the same logic behind the different effect of displacement on earnings for different types of workers that we documented and interpreted in Section 5.

The left panel in Figure 7 illustrates the response of the unemployment rate by type. As one can see, the increase in the unemployment rate to the aggregate productivity shock is very different for different types of workers. The increase in the unemployment rate of α-workers is about 2 percentage points (from 4.2% to 6.5%), the increase in the unemployment rate of β-workers is about 5 percentage points (from 12.4% to 17%), and the increase in the unemployment rate of γ-workers is about 19 percentage points (from 30% to 49%). As explained above, α-workers have the least responsive UE rate and the least responsive EU rate and, hence, the smallest increase in the unemployment rate. In contrast, γ-workers have the most responsive UE rate and the most responsive EU rate and, hence, the largest increase in the unemployment rate. Also, notice that increase in unemployment fades most rapidly for α-workers and most slowly for γ-workers. As explained above, this is because it takes much less for displaced α -workers to find a match above the reservation quality than it does for γ-workers.

The right panel in Figure 7 illustrates the evolution of the composition of the unemployment pool. First, notice that the non-stochastic steady-state composition of the unemployment pool is different from the composition of the population. Specifically, while α-workers are 55% of the population, they represent only 25% of the unemployment pool. Conversely, γ-workers are only 15% of the population, but they represent nearly 50% of the unemployment pool. When the productivity shock hits the economy, a large flow of workers flows into the unemployment pool. Since the inflow is disproportionately made of γ-workers, the composition of the unemployment pool tilts even further towards them. Throughout the recovery, displaced workers flow back out of unemployment. Since the outflow is disproportionately made of α and β workers, the composition of the unemployment pool tilts even more towards γ workers. Eventually, displaced γ workers find suitable matches and exit the unemployment pool, but the process is slow and, even after several years, there is an excess of γ-workers.

The changing composition of the unemployment pool is important to understand the aggregate response of the labor market to the productivity shock. Consider the aggregate UE rate. The aggregate UE rate falls by 30%. While most of the decline is due to the fact that the UE rate falls for each type of worker, some of the decline is due to the fact that the composition of the unemployment pool tilts towards workers of type γ. Moreover, since the unemployment pool takes a long time to purge itself of the excess γs, the UE rate recovers more slowly than any one of the type-specific UE rates. Now, consider the aggregate EU rate. The initial flow of workers into unemployment contains a disproportionate fraction of γ workers. Since the process by which the displaced γ-workers involves trying and quitting several jobs, their EU rate is the slowest to return back to normal and the aggregate EU rate inherits this slowness. Overall, our calibrated model is consistent with the view of the Great Recession proposed by Ahn and Hamilton (2020): the increase in the EU rate at the beginning of a recession contains a large fraction of low-UE workers and, hence, amplifies the decline in the UE rate and slows down its recovery.

The left panel of Figure 8 illustrates the response of the aggregate unemployment rate, and decomposes it into the contribution of different types of workers. The aggregate unemployment rate increases by 5.4 percentage points (from 10.6 to 16.1%). On impact, about 54% of the increase in the aggregate unemployment rate is due to the increase in the unemployment rate of γ-workers, 21% to the increase in the unemployment rate of β-workers, and 25% to the increase in the unemployment rate of α-workers. The increase in the unemployment rate of α workers fades first. The increase in the unemployment rate of β workers fades second. And, eventually, the aggregate unemployment rate is almost entirely due to the increase in the unemployment rate of γ workers, which fades very slowly. The unemployment rate of γ-workers is thus a key

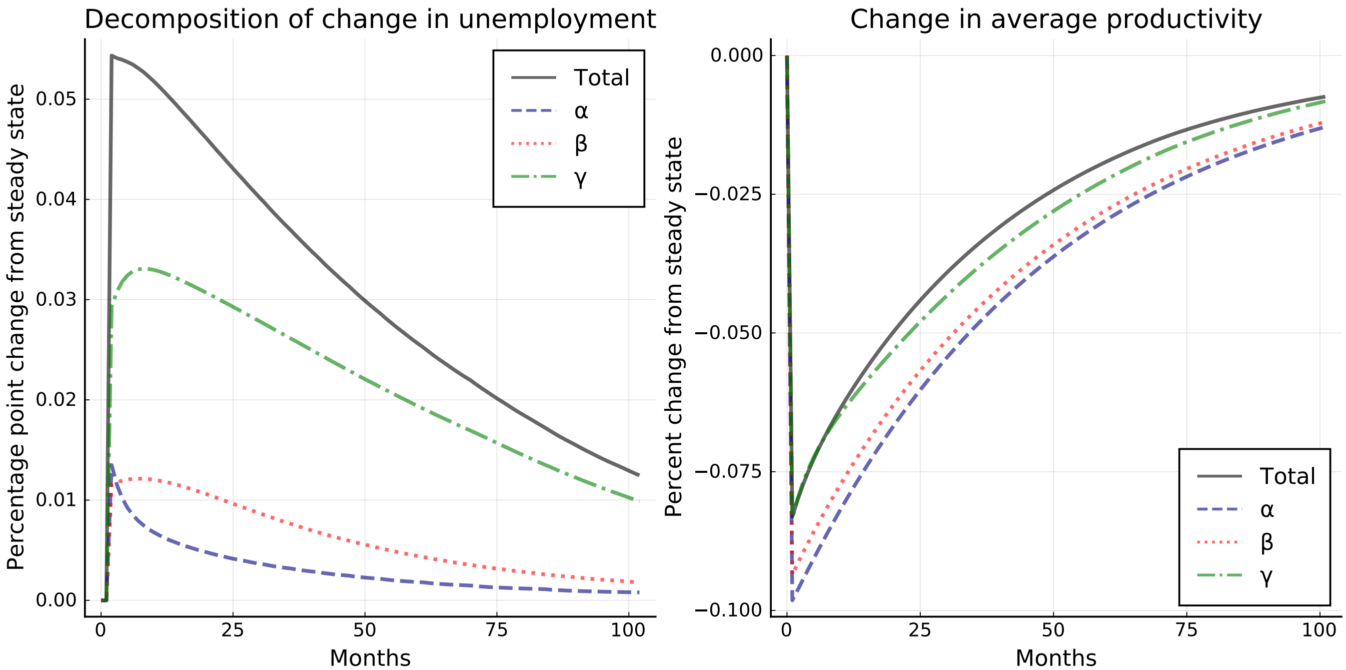

Figure 8: Left panel: decomposition of the response of the aggregate unemployment rate; right panel: aggregate and type-specific productivity of labor

determinant of the response of aggregate unemployment, in the sense that it drives much of its increase and causes it to recover slowly.

The right panel of Figure 8 illustrates the average productivity of different types of workers, and the aggregate average productivity of labor. The decline in the average productivity of each type of worker is smaller than the decline in the aggregate component of productivity y. This is because it is the lowest quality matches that are destroyed during the recession. Moreover, the decline in the aggregate average productivity is smaller and less persistent than the decline in the aggregate component of productivity. Again the intuition is simple. The workers who are displaced during the recession are predominantly γ-workers, who have the lowest productivity among all types. As a result, the composition of non-displaced workers becomes more tilted towards α and β-workers, who have higher productivity. Throughout the recovery, the displaced α and β-workers regain employment more quickly and, hence, the composition of employed worker becomes even more tilted towards α and β-workers, thus speeding up the recovery of average labor productivity relative to the underlying shock. Overall, the response average labor productivity is dampened and its recovery quicker because of a double cleansing effect: Worker types with higher productivity are more likely to weather the storm of the recession, and—for a given worker type—those in higher quality matches are more likely to do so.

6.2 The Great Recession

The left panel of Figure 9 plots the evolution of the unemployment rate by type in the LEHD during and after the Great Recession. The rise and decline in the unemployment rate is very different for different types. During the recession, the increase in the unemployment rate is

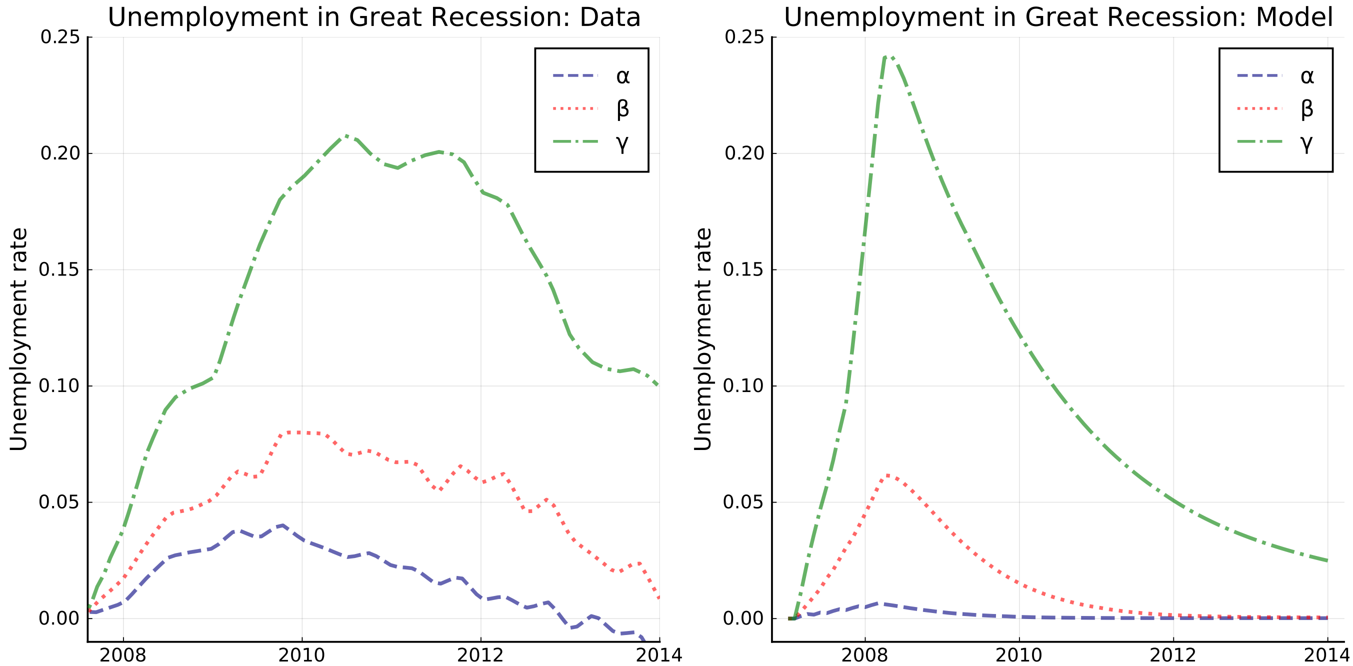

Figure 9: Increase in unemployment rates for each worker type during the Great Recession in the model and data

about 3 percentage points for α-workers, about 7 percentage points for β-workers, and about 21 percentage points for γ-workers. During the recovery, the decline in the unemployment rate is fastest for α-workers (their unemployment rate is basically back to its pre-recession level already 2012), the decline is less rapid for β-workers (their unemployment rate is essentially back to its pre-recession level in 2014), and slowest for γ-workers (their unemployment rate is still 10 percentage points above its pre-recession level in 2014).

Comparing Figure 8 with Figure 9, it is immediate to see the striking resemblance between the dynamics of the unemployment rate during and after the Great Recession and the dynamics of the unemployment rate in response to an aggregate productivity shock. In particular, the model successfully captures the fact that the increase in the unemployment rate is highest for γ-workers and lowest for α-workers. Moreover, the model successfully captures the fact that the recovery of the unemployment rate for γ-workers is slowest, and the recovery is fastest for α-workers.

We now want to demonstrate that our model can rationalize not only qualitatively, but also quantitatively the behavior of the economy during and after the Great Recession. We assume that, at the onset of the recession, the economy was at the steady-state associated with an aggregate component of productivity y equal to y^*. We then progressively lower the aggregate component of productivity for 5 quarters. We assume that the decline in productivity is different for different types. Specifically, the decline in productivity reaches a maximum of 5% for α-workers and a maximum of 12% for β and for γ-workers. After reaching its peak, the decline in productivity reverts back towards y^* at the rate of 10% per month. In words, we model the cause of the Great Recession as a series of bad shocks followed by a quick recovery,

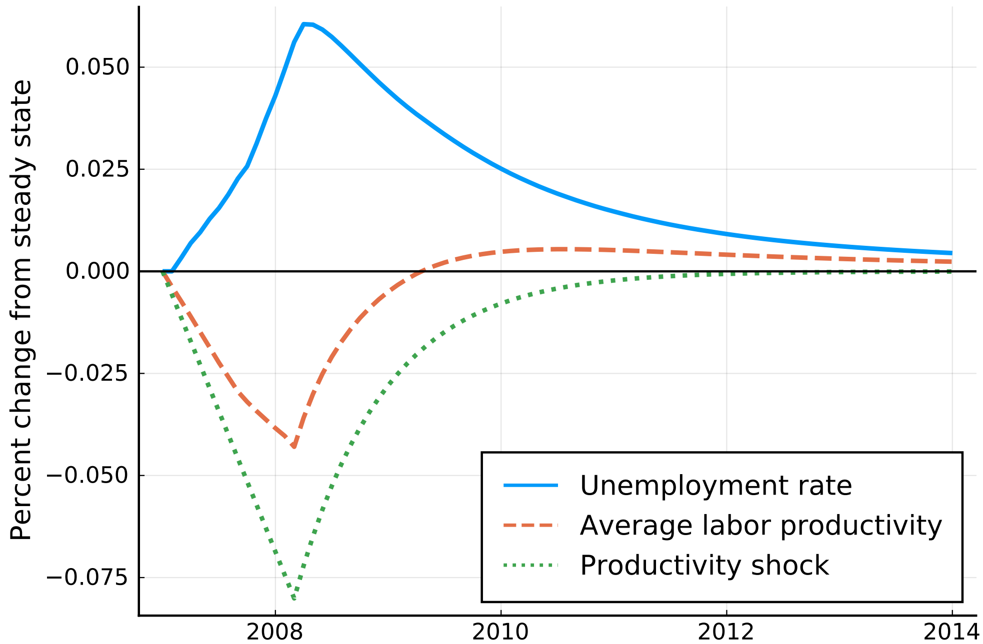

Figure 10: Model-implied aggregate unemployment rates, average labor productivity, and type-weighted productivity shock over the Great Recession

where the impact of the shock on the productivity of β and γ-workers is stronger than for α-workers. This may be because different types of workers are employed in different sectors of the economy, or because different types of workers are employed in different tasks with different exposure to the shock.

The right panel of Figure 9 shows the response of the unemployment rate by type in the model. The unemployment rate of α-workers increases by 0.7 percentage points. The unemployment rate of β-workers increases by 6 percentage points. The unemployment rate of γ-workers increases by 24 percentage points. The order of magnitude of these increases lines up with what we measure in the LEHD data. In the recovery, the unemployment rate of α-workers is back to its pre-recession level after 1 year, the unemployment rate of β-workers is essentially back to its pre-recession level after 3 years, and even after 4 years the unemployment rate of γ-workers is still around 5 percentage points above its pre-recession level. These responses are quantitatively similar to the dynamics the unemployment rate measured in the LEHD.

Figure 10 shows the response of the average labor productivity. The decline in the average labor productivity peaks at about 4%. Then it recovers quickly, surpassing its pre-recession level within a year of reaching its peak. This is precisely the path of average labor productivity that we observed in the data. As documented in Kaplan and Menzio (2016), average labor productivity declined (relative to trend) for 4 consecutive quarters reaching a minimum of 4% below trend. Then, the average labor productivity started recovering (relative to trend), it was back on trend 3 quarters after reaching its minimum, and it then overshoot the trend for a couple of years before returning to trend again.

Figure 10 also shows the response of the aggregate unemployment rate and the productivity shock (weighted by the measure of different worker types). The aggregate unemployment rate increases by 6 percentage points and then recovers very slowly. The average productivity of labor falls by 4% and recovers very quickly. When the average productivity of labor is back on trend, the aggregate unemployment rate is 4 percentage points above its pre-recession level. When the average productivity of labor has been steadily above trend for about 2 years, the aggregate unemployment rate is still 1.5 percentage points above its pre-recession level. As explained in the previous subsection, the sluggishness of the unemployment recovery is due to the fact that displaced γ-workers take a long time and many attempts before finding a suitable match. The difference in the speed of recovery between unemployment and labor productivity generates a phase during which both labor productivity and unemployment are above trend.

Compared with the shock, the decline in the average productivity of labor is half as large and it recovers much more quickly. This is because of the cleansing effects described in the previous subsection, which here are magnified by the fact that low-productivity types (βs and γs) are more exposed to the shock than high-productivity types (αs). Specifically, since the shock leads to a much larger decline in employment for β and γ-workers than for α-workers, the average productivity of labor falls by as little as half as the aggregate component of productivity. Moreover, since displaced α-workers are much faster at regaining employment, the average productivity of labor recovers quickly and, for a period of time, overshoots its pre-recession level.

The exercise presented in this subsection shows that our model with heterogeneous workers can explain several seemingly puzzling features of the US business cycle, which were evident during the Great Recession and its aftermath. Our model naturally produces large unemployment fluctuations out of relatively small productivity shocks because the employment of γ-workers is very sensitive to changes in the economic environment. Thus, it addresses the so-called Shimer puzzle (Shimer 2005) without resorting to sticky wages (see, e.g., Hall 2005, Kennan 2010, Menzio and Moen 2010). Our model produces naturally slow unemployment recoveries because it takes a long time for displaced γ-workers to find another suitable match. Thus, it addresses the so-called job-less recovery puzzle (see, e.g., Bachmann 2007). Our model naturally dampens average labor productivity relative to productivity shocks because of changes in the composition of the employment pool, which contributes to our explanation of the Shimer puzzle. Lastly, our model can produce phases in which unemployment and labor productivity are both above trend. Hence, our model provides a novel potential explanation for the weak contemporaneous correlation between unemployment and labor productivity without resorting to non-fundamental shocks (e.g., Kaplan and Menzio 2016, Golosov and Menzio 2020, Fernandez-Villaverde et al. 2020) or to shocks to the discount rate (e.g., Hall 2017, Kehoe, Midrigan and Pastorino 2019, or Martellini, Menzio and Visschers 2020).

It is important to clarify that the exercise presented is not a measurement exercise in the spirit of Kydland and Prescott (1983). In fact, we do not have independent measures of the productivity shocks for each type of worker which we then input into the model. Rather, the exercise in this subsection is a fitting exercise and it is only meant to show that there is a set of type-specific productivity shocks that generates a response of unemployment and average productivity that is consistent with the data. Clearly, this is an achievement of a lower-level, but an achievement nonetheless, as it shows that many of the puzzling features of the Great Recession can be explained with a rather parsimonious model and a set of shocks that is consistent with the data when aggregated.

The only somewhat-major discrepancy between the behavior of the model and the data is the increase in the unemployment rate for α-workers. In the model, we need a small productivity shock for α-workers to make the decline in the average labor productivity as modest as it is in the data. This small productivity shock generates an increase in the unemployment rate of α-workers of only 1 percentage point compared to the measured increase of 3 percentage points. Put it in the perspective of the large differences in the response of unemployment across types, this deviation is small. Moreover, our measure of unemployment is not direct (as it requires comparing the behavior of two cohorts of worker at different times). For this reason, it may very well be that the difference between 3 and 1 percentage points is mostly a matter of measurement error.

7 Conclusions

In this paper, we presented evidence that workers systematically and persistently differ in their labor market transition rates and we incorporated this observation into a search-theoretic framework with business cycle fluctuations. We showed that a relatively small number of the most unstable workers are responsible for the majority of the Great Recession's rise in unemployment, and its persistence. Within the model, this is because recessions displace the workers who are least able to find another stable employer match. To be consistent with their steady state patterns of short job spells and relatively frequent and long unemployment spells, the model postulates that they draw match productivity from a heavy right-tailed distribution, but do not immediately learn about its realization. This prompts them to cycle through employers chasing a high-enough match quality and makes recessions, when UE rates are lower and required match productivities are higher, particularly harmful to their employment rate.

In the data, we identified these workers using a k-means clustering approach in the LEHD, exploiting its long and high frequency panels of employment history. Because we often observe several spells of employment and unemployment for a single worker, we can essentially calculate a worker-specific hazard function for transitions out of a job and out of unemployment. The clustering approach finds a pattern linking high early exit rates from jobs to frequent and long unemployment (identified as γ-type workers), low early exit rates from employment with sparse unemployment experience (identified as α-type workers) and an important, middle set who mix these characteristics (identified as β-type workers).

The model generated these patterns within a directed search framework similar to Menzio and Shi (2011) but in which we introduced ex ante heterogeneity in the form of discrete types. These types differ in several dimensions, but most notably the distribution from which they draw match-specific productivity. Because β-type and γ-type workers drew from heavy right-tailed distributions, and only gradually learn about their draw, they experience short tenured job matches as they repeatedly sample the distribution looking for a high draw. When the recession hit, these workers were more likely to be in a low productivity match and then, upon separating, had more difficulty finding a high-enough productivity match. This generates the persistently high unemployment we saw in the Great Recession, as it is especially difficult for these workers to repeatedly sample in search of an acceptable match. Tying to micro-level findings, this model with ex ante heterogeneity also rationalized the larger earnings losses upon separation among the γ-type workers and also the downward sloping average UE hazard profiles.

References

[1] Abowd, J. M., B. E. Stephens, L. Vilhuber, F. Andersson, K. L. McKinney, M. Roemer, and S. Woodcock. 2009. “The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators,” in *Producer Dynamics: New Evidence from Micro Data*, National Bureau of Economic Research, Inc, NBER Chapters, 149–230.

[2] Ahn, H., and J. Hamilton. 2019. “Heterogeneity and Unemployment Dynamics” *Journal of Business & Economic Statistics.*

[3] Bagger, J., F. Fontaine, F. Postel-Vinay, and J. Robin. 2014. “Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics.” *American Economic Review*, 104: 1551-1596.

[4] Barlevy, G. 2002. “The Sullying Effect of Recessions.” *Review of Economic Studies*, 69: 65-96.

[5] Bonhomme, S., T. Lamadon, and E. Manresa. 2020. “Discretizing Unobserved Heterogeneity,” *Econometrica.*

[6] Cahuc, P. F. Postel-Vinay and J. Robin. 2006. “Wage Bargaining with On-the-Job Search: Theory and Evidence.” *Econometrica*, 74: 323-364.

[7] Davis, S. and J. Haltiwanger. 2014. “Labor Market Fluidity and Economic Performance.” NBER Working Paper 20479.

[8] Davis, S., and T. von Wachter. 2011. “Recessions and the Costs of Job Loss.” *Brookings Papers on Economic Activity.*

[9] Diamond, P. 1982. “Wage Determination and Efficiency in Search Equilibrium,” *Review of Economic Studies*, 49: 217-227.

[10] Flaeen, A., M. Shapiro, and I. Sorkin. 2019. “Reconsidering the Consequences of Worker Displacements: Firm versus Worker Perspective.” *AEJ: Macroeconomics*, 11 (2): 193-227.

[11] Gali, J., and T. van Rens. 2019. “The Vanishing Procyclicality of Labor Productivity,” Manuscript, University of Warwick.

[12] Gavazza, A., S. Mongey, and G. Violante. 2018. “Aggregate Recruiting Intensity.” *American Economic Review*, Forthcoming.

[13] Gertler, M., and A. Trigari. 2009. “Unemployment Fluctuations with Staggered Nash Wage Bargaining.” *Journal of Political Economy*, 117: 38-85.

[14] Golosov, M., and G. Menzio. 2020. “Agency Business Cycles.” *Theoretical Economics*, 15:123-158.

[15] Gregory, V., Menzio, G., and D. Wiczer. 2020. "Pandemic Recession: L- or V-Shaped?" *Quarterly Review*, 40: 1–28.

[16] Guvenen, F. 2007. "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" *American Economic Review*, 97 (3): 687-712.

[17] Guvenen, F., F. Karahan, S. Ozkan, and K. Song. 2015. "What do data on millions of US workers reveal about life-cycle earnings risk?" NBER Working Paper 20913.

[18] Guvenen, F., and A. Smith. 2014. "Inferring Labor Income Risk and Partial Insurance From Economic Choices." *Econometrica*, 82 (6): 2085-2129.

[19] Hagedorn, M., and I. Manovskii. 2008. "The Cyclical Behavior of Unemployment and Vacancies Revisited." *American Economic Review,* 98: 1692–706.

[20] Hall, R. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95: 50-65.

[21] Hall, R. 2017. "High Discounts and High Unemployment," *American Economic Review*, 107: 305-330.

[22] Hall, R., and M. Kudlyak. 2019. "Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics Paper." FRB San Francisco WP No. 19-05.

[23] Hall, R., and P. Milgrom. 2008. "The Limited Influence of Unemployment of the Wage Bargain." *American Economic Review,* 98: 1653–74.

[24] Jacobson, L., R. LaLonde, and D. Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review*, 83: 685-709.

[25] Jarosch, G., and L. Pillosoph. 2019. "Statistical Discrimination and Duration Dependence in the Job Finding Rate." *Review of Economic Studies*, 86: 1631–1665.

[26] Kaplan, G., and G. Menzio. 2016. "Shopping Externalities and Self Fulfilling Unemployment Fluctuations." *Journal of Political Economy*, 124: 771-825.

[27] Kehoe, P., V. Midrigan and E. Pastorino. 2019. "Debt Constraints and Employment." *Journal of Political Economy*, 127: 1926-1991.

[28] Kennan, J. 2010. "Private Information, Wage Bargaining and Unemployment Fluctuations." *Review of Economic Studies*, 77: 633-664.

[29] Kydland, F., and E. Prescott. 1982. "Time to Build and Aggegate Fluctuations." *Econometrica*, 50: 1345-1370.

[30] Ljungqvist, L., and T. Sargent. 2017. "The Fundamental Surplus." *American Economic Review*, 107: 2630-65.

[31] Martellini, P., G. Menzio, and L. Visschers. 2020. "Revisiting the Hypothesis of High Discounts and High Unemployment." *Economic Journal*.

[32] Menzio, G., and S. Shi. 2010. "Block Recursive Equilibria for Stochastic Models of Search on the Job." *Journal of Economic Theory*, 145: 1453-1494.

[33] Menzio, G., and S. Shi. 2011. "Efficient Search on the Job and the Business Cycle." *Journal of Political Economy*, 119: 468-510.

[34] Menzio, G., I. Telyukova, and L. Visschers. 2016. "Directed Search over the Life Cycle." *Review of Economic Dynamics*, 19: 38-62.

[35] Moen, E. 1997. "Competitive Search Equilibrium." *Journal of Political Economy*, 105: 694-723.

[36] Morchio, I. 2020. "Work Histories and Lifetime Unemployment." *International Economic Review* 61.

[37] Mortensen, D. 1970. "Job Search, the Duration of Unemployment, and the Phillips Curve." *American Economic Review*, 60(5): 847–62.

[38] Mortensen, D. and C. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 61: 397-415.

[39] Mueller, A., J. Spinnewijn, and G. Topa. 2019. "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence and Bias." *American Economic Review*.

[40] Pissarides, C. 1984. "Efficient Job Rejection." *Economic Journal*, 94: 97-108.

[41] Pissarides, C. 1985. "Short-run equilibrium dynamics of unemployment, vacancies, and real wages." *American Economic Review*, 75: 676-690.

[42] Pissarides, C. 2000. *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA.

[43] Postel Vinay, F., and J. Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica*, 70: 2295-2350.

[44] Shimer, R. 2005 "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review*, 95: 25-49.

[45] Wang, J. 2010. "Consistent selection of the number of clusters via crossvalidation," *Biometrika*, 97: 893–904.