

Did the Job Ladder Fail After the Great Recession?*

Giuseppe Moscarini[†]
Yale University and NBER

Fabien Postel-Vinay[‡]
UCL, IFS and CfM

August 2013

Abstract

We study employment reallocation across heterogeneous employers through the lens of a dynamic job-ladder model, where more productive employers spend more hiring effort and are more likely to succeed in hiring because they offer more. As a consequence, an employer's size is a relevant proxy for productivity. We exploit newly available U.S. data from JOLTS on employment flows by size of the establishment. Our parsimonious job ladder model fits the facts quite well, and implies 'true' vacancy postings by size that are more in line with gross flows and intuition than JOLTS' actual measures of job openings, previously criticized by other authors. Focusing on the U.S. experience in and around the Great Recession, our main finding is that the job ladder stopped working in the GR and has not yet fully resumed.

1 Introduction

The persistence of high unemployment in the US and many other countries after the 2007-2009 Great Recession (henceforth GR) is currently the central issue for macroeconomic policy around the world. In previous work (Moscarini and Postel-Vinay 2009, 2012, 2013, resp. MPV09, MPV12 and MPV13) we documented empirically and formulated a hypothesis to explain the pattern of employment decline and recovery during and after a typical recession.

*Prepared for the NBER Project on the Labor Market in the Aftermath of the Great Recession. We thank our discussant, Rasmus Lentz, and participants to the Project's Conference at NBER (Cambridge, Mass., May 16-17, 2013), the 2013 Shanghai Macroeconomics Workshop and the 2013 meeting of the Society of Economic Dynamics in Seoul for comments. Special thanks to Charlotte Oslund at the BLS for extensive assistance with JOLTS data. Moscarini also thanks the NSF for support to his research under grant SES 1123021. The usual disclaimer applies.

[†]Address: Department of Economics, Yale University, PO Box 208268, New Haven CT 06520-8268. Tel: +1-203-432-3596. E-mail giuseppe.moscarini@yale.edu. Web <http://www.econ.yale.edu/faculty1/moscarini.htm>

[‡]Address: Department of Economics, University College London, Drayton House, 30 Gordon Street, London WC1H 0AX, UK. Tel: +44(0)207 679 5856. E-mail f.postel-vinay@ucl.ac.uk. Web <https://sites.google.com/site/fabienpostelvinay/>. Postel-Vinay is also affiliated with Sciences Po (Paris), CEPR (London) and IZA (Bonn).

In a nutshell, in a tight labor market high-paying, large employers overcome the scarcity of unemployed job applicants by poaching employees from smaller, less productive and lower-paying competitors, whose employment share then shrinks in relative terms. When the expansion ends, large employers, that were less constrained, have more employment to shed than small ones. In addition, the resulting high unemployment relaxes hiring constraints on all employers, particularly the small ones that are less capable of poaching from other firms. As a result, small employers downsize less in the recession and grow faster (still in relative terms) in the early recovery. According to this hypothesis, in a prolonged phase of high unemployment, as we witnessed since 2009, small firms should be leading the charge in job creation, followed years later by upgrading to larger, better-paying employers.

In this paper, we confront this “dynamic job ladder” hypothesis with more demanding empirical tests. We still adopt employer size as an empirical measure of the job ladder ‘rung’ based on the observed wage/size relationship. We go beyond net worker flows and consider also the model’s implications for gross worker flows (hires, quits, layoffs) and vacancy postings by employer size, measured by newly available semi-disaggregated data by establishment size from the BLS’ Job Openings and Labor Turnover Survey (JOLTS). Specifically, we calibrate the key flow equations of the job ladder model to fit the time series of net and gross employment flows by employer size. We extend our investigation to examine the GR and its aftermath, in comparison with previous cyclical episodes.

We reach the following conclusions. First, the dynamic job ladder model, a parsimonious setup built on some very strong assumptions, such as homogeneous labor and time-invariant heterogeneity in firm-level Total Factor Productivity, does a remarkable job at fitting the dynamics of employment across size classes. The estimated hiring intensity by employer size resembles vacancies by establishment size measured in the JOLTS, but resolves some puzzling aspects of these data, specifically the lack of vacancies at the small employer end.

Second, a comprehensive assessment of the evidence indicates that the job ladder has slowed down considerably since the GR. The drastic decline in labor market turnover affected especially direct movements from smaller, lower-paying to larger, higher-paying employers. Small employers suffered unusual job losses, relative to large employers and a typical recession, mostly through an increase in their layoffs, only partially compensated by resilient vacancy posting and hiring.

We now provide the details. From an aggregate perspective, the GR impacted the labor market as would any (deep) recession: job openings went down across the board, job finding rates plummeted, and layoff rates spiked — albeit temporarily — around the Fall of 2008, as the financial crisis came to a head. As a result, unemployment soared. As we argued and documented in previous work and based on past recessions (MPV09, MPV12, MPV13),

this created conditions that were relatively favorable to small, low-paying, less productive employers. High unemployment meant that there was plenty of cheap labor for them to hire. The collapse in aggregate job market tightness, which largely persists to date, reduced not only the workers' exit rate from unemployment, as is well understood, but also the job-to-job quit rate, meaning that employers at the bottom of the job ladder were losing fewer workers to their larger, more productive, higher-paying competitors.

The available data on gross worker flows and job openings largely support this view. Job-to-job transitions indeed went down markedly during the GR. The 'poaching intensity' (share of new hires that come through a job-to-job transition) declined sharply during and after the GR, especially so for larger employers. Finally, while the share of small establishments in total job openings remained roughly stable throughout the GR (if anything, it went up a little), the vacancy *yield* of small employers sky-rocketed, in sharp contrast to the comparatively modest increase in the vacancy yield of large establishments.

Yet — and this is where the GR differs from past recessions — small employers fared *worse* than large ones in terms of net employment growth. This poor relative job creation performance, which is unusual relative to the previous three decades, was the result of a brutal (temporary) increase in the layoff rate of small employers around the Lehman Brothers episode (September 2008), the peak of the financial crisis. While at that point layoff rates rose sharply at employers of all sizes, small establishments stood out, possibly because they were hit especially hard by the credit crunch. Those among small employers that were still hiring did so relatively easily, and benefited from relatively favorable conditions on the hiring and retention margins, but that was not enough to offset the exceptionally large blow from the credit crunch.

These findings suggest the following, to some extent novel, interpretation of the GR and its aftermath. Small employers, especially existing ones, faced an unusual credit crunch that led to a wave of layoffs, while the class of small employers as a whole, including entrants, kept hiring at a healthy pace. The collapse in hiring by large employers froze job-to-job upgrading, further taming the incentives of small employers to take advantage of abundant unemployment to hire.

In Section 2 we summarize and extend the dynamic job ladder model from MPV13. In Section 3 we present descriptive evidence on labor market flows across employers of different sizes before, during and after the GR. In Section 4 we illustrate our methodology to estimate/calibrate turnover parameters and hiring intensity by firm size in the dynamic job ladder model to replicate the observed net and gross flows of employment by firm size. In Section 5 we discuss our empirical results. Conclusions follow.

2 The dynamic job ladder: a model

2.1 Setup

Time $t = 0, 1, 2, \dots$ is discrete. The labor market is populated by a unit-mass of workers, who can be either employed or unemployed, and by a unit measure of firms. Workers and firms are risk neutral, infinitely lived, and maximize payoffs discounted with common factor $\beta \in (0, 1)$. Firms operate constant-return technologies with labor as the only input and with productivity scale $\omega_t p$, where $\omega_t \in \Omega$ is an aggregate component, evolving according to a stationary first-order Markov process $Q(d\omega_{t+1} | \omega_t)$, and p is a time-invariant, firm-specific component, distributed across firms according to a c.d.f. Γ over some positive interval $[\underline{p}, \bar{p}]$.

The labor market is affected by search frictions in that unemployed workers can only sample job offers sequentially with some probability $\lambda_t \in (0, 1)$ at time t , and while searching enjoy a value of leisure b_t . Employed workers earn a wage, draw each period with probability $s \in (0, 1]$ an i.i.d. opportunity to search on the job, thus face a per-period sampling chance of job offers of $s\lambda_t$. Workers can only send one job application, hence can never receive more than one offer, in each period. All firms of equal productivity p start out with the same labor force. Each employed worker is separated from his employer p and either enters unemployment with probability $\delta_t(p) \in (0, 1)$, that we assume non-increasing in p , or is immediately reallocated to another job, drawn randomly from the available ones, without going through unemployment, with probability $\rho_t \in (0, 1)$. The displacement shock $\delta_t(p)$ encompasses both layoffs and quits to non employment that result in a measurable unemployment spell. The reallocation shock ρ_t captures such events as moves due to spousal relocation, or displacements followed by immediate re-hiring by another employer. We maintain throughout the assumption that the destruction rate and reallocation rate are exogenous and functions of the aggregate productivity state $\delta_t(p) = \delta(p, \omega_t)$, $\rho_t = \rho(\omega_t)$, and the flow value of non production b is constant.

The contact probability λ_t is, in contrast, endogenous and determined as follows. A firm can post $a \geq 0$ job adverts, or spend hiring effort $a \geq 0$, at a cost $c(a)$, with $c(\cdot)$ positive, strictly increasing and convex, continuously differentiable. The firm's hiring effort determines its sampling weight in workers' job search, while total hiring effort determines the rate at which an advert returns contacts with workers. Specifically, let $\mathcal{N}_t : [\underline{p}, \bar{p}] \rightarrow [0, 1]$ denote the cumulated population distribution of employment across firm types. So $\mathcal{N}_0(p)$ is the measure of employment initially at firms of productivity *at most* p , a given initial condition, $\mathcal{N}_t(\bar{p})$ is employment, and $u_t = 1 - \mathcal{N}_t(\bar{p})$ is the unemployment (rate) at time t . Let $a_t(p)$ denote the adverts posted on the equilibrium path by a firm of productivity p , size $L_{t-1}(p)$, in aggregate state $\omega_t, \mathcal{N}_{t-1}$. Defining aggregate firm hiring effort, $A_t = \int_{\underline{p}}^{\bar{p}} a_t(p) d\Gamma(p)$, and

letting η_t denote the probability that an advert receives at least one application, we have:

$$\eta_t A_t = \lambda_t [1 - \mathcal{N}_{t-1}(\bar{p})] + s\lambda_t \mathcal{N}_{t-1}(\bar{p}) + \rho_t \mathcal{N}_{t-1}(\bar{p}) = m(A_t, \mathcal{N}_{t-1}(\bar{p})) \quad (1)$$

where $m(\cdot)$ is a constant-return-to-scale matching function, increasing and concave in both of its arguments. Conditional on a contact with some employer, from unemployment with probability λ_t and from employment with probability $\rho_t + s\lambda_t$, the worker samples the productivity of the job from the cdf

$$\mathcal{F}_t(p) = A_t^{-1} \int_{\underline{p}}^p a_t(q) d\Gamma(q).$$

In each period, the timing is as follows. Given a state ω_{t-1} of aggregate labor productivity and distribution of employed workers \mathcal{N}_{t-1} :

1. production and payments take place at all firms in state ω_{t-1} ; the flow benefit b accrues to unemployed workers;
2. the new state ω_t of aggregate labor productivity is realized;
3. firms post job adverts;
4. each employed worker receives an independent turnover shock, which can be of one of three, mutually exclusive, types: with chance $\delta_t(p)$ exogenous separation to unemployment, with chance ρ_t a random offer from the distribution \mathcal{F}_t of job adverts that the worker is forced to accept, or with chance $s\lambda_t$ a random offer from \mathcal{F}_t that the worker can decline to retain his current job;
5. employed workers can quit to unemployment;
6. each previously unemployed worker receives an offer with probability λ_t .

Finally, we assume that the distribution of firm types, Γ , has continuous and everywhere strictly positive density $\gamma = \Gamma'$ over $[\underline{p}, \bar{p}]$, and that the initial measure of employment across firm types, \mathcal{N}_0 , is continuously differentiable in p . Combining those two assumptions, we obtain that the initial average size of a type- p firm, which is given by $L_0(p) = \frac{d\mathcal{N}_0(p)/dp}{\gamma(p)}$, is a continuous function of p .

2.2 Rank-Preserving Equilibrium

MPV13 study equilibrium in this economy without reallocation shocks ($\rho_t = 0$) under the following assumptions on the strategy space. Each firm chooses and commits to a Markov

employment contract, namely, a state-contingent wage depending on own productivity p , initial own size L_{t-1} and size distribution \mathcal{N}_{t-1} , and aggregate productivity ω_t . The objective of the firm is to maximize the present discounted value of profits, given other firms' contract offers. The firm is further subjected to an *equal treatment constraint*, whereby it must pay the same wage to all its workers. Workers cannot commit not to quit, and always accept an offer, received either while unemployed or already employed, if doing so increases their Bellman value over continuing in the current employment state.

When posting a contract, a firm of productivity p has to forecast how the value $V(p, L_{t-1}, \mathcal{N}_{t-1}, \omega_t)$ provided by the contract to the recipient workers will compare with the values offered by other firms. This calculation is necessary to predict the effect of the contract offer on hiring and retention, thus on the evolution of its size. Each firm must keep track of the entire distribution of contract offers and of contracts currently in the hands of employed workers. The size distribution \mathcal{N}_{t-1} summarizes this information in a Markov contract-posting equilibrium, hence it is a payoff-relevant state variable.

Let $L_t(p)$ denote the size of a firm of type p at time t on the equilibrium path. Labor turnover takes a simple form when equilibrium has the following property:

Definition 1 (Rank-Preserving Equilibrium) *An equilibrium is RANK-PRESERVING (RP) if a more productive firm always pays its workers more: $V_t(p) = V(p, L_{t-1}(p), \mathcal{N}_{t-1}, \omega_t)$ is increasing in p , including the effect of p on firm size, $L_{t-1}(p)$.*

MPV13 prove that the unique stochastic Markov equilibrium of this economy is Rank-Preserving assuming no reallocation shocks, $\rho_t = 0$, type-independent exogenous separations, $\delta_t(p) = \delta_t$, and relatively weak restriction on initial conditions:

Proposition 1 (Ranked initial firm size implies RPE) *If more productive firms are initially weakly larger ($L_0(p)$ is non-decreasing in p), any symmetric Markov contract-posting Equilibrium is necessarily Rank-Preserving, and the ranking of firms' sizes is maintained on the equilibrium path. In addition, more productive firms also spend more hiring effort: $a_t(p) = a(p, L_{t-1}(p), \omega_t, \mathcal{N}_{t-1})$ is increasing in p .*

These results extend to the case of random reallocation shocks ($\rho_t > 0$) that do not change the relative incentives of different employers, and of exogenous separations shocks with type-dependent probability $\delta_t(p)$ that is non-increasing in firm productivity p . Therefore, if all firms start of equal size, e.g. empty, equilibrium turnover in this economy always takes the form of a job-ladder *at all points in time*. Workers being ex-ante homogeneous, an unemployed worker always accepts a job offer. Once employed, he only accepts outside offers from more productive firms than his current employer. On the equilibrium path, the

size $L_t(p)$ of a firm of type p evolves as follows:

$$L_{t+1}(p) = L_t(p) [1 - \delta_{t+1}(p) - \rho_{t+1} - s\lambda_{t+1}\bar{\mathcal{F}}_{t+1}(p)] + \eta_{t+1}a_{t+1}(p) \frac{\lambda_{t+1}[1 - \mathcal{N}_t(\bar{p})] + s\lambda_{t+1}\mathcal{N}_t(p) + \rho_{t+1}\mathcal{N}_t(\bar{p})}{\lambda_{t+1}[1 - \mathcal{N}_t(\bar{p})] + s\lambda_{t+1}\mathcal{N}_t(\bar{p}) + \rho_{t+1}\mathcal{N}_t(\bar{p})} \quad (2)$$

where $\bar{\mathcal{F}}_t(p) = 1 - \mathcal{F}_t(p)$ and cumulated employment is now $\mathcal{N}_t(p) = \int_{\underline{p}}^p L_t(q) d\Gamma(q)$.

In (2), a fraction $\delta_{t+1}(p)$ of the initial $L_t(p)$ workers are displaced exogenously into unemployment, a fraction ρ_t are forcefully reallocated to other employers, and a fraction $s\lambda_{t+1}$ sample a new offer, that they accept only if extended by a more productive/larger employer, in accordance with the Rank-Preserving nature of the unique equilibrium. The rest stay. The firm also hires, by posting $a_{t+1}(p)$ adverts, each contacted with probability η_{t+1} . The chance that this contact turns into a match is the proportion of all job seekers who receive an offer (the denominator of the fraction) who are either unemployed, or currently employed at firms less productive than p , or currently employed and forcefully reallocated (the numerator).

Using (1), we can simplify

$$\eta_t a_t(p) \frac{1}{\lambda_t [1 - \mathcal{N}_{t-1}(\bar{p})] + s\lambda_t \mathcal{N}_{t-1}(\bar{p}) + \rho_t \mathcal{N}_{t-1}(\bar{p})} = \frac{a_t(p)}{A_t} = \frac{\mathcal{F}'_t(p)}{\gamma(p)}$$

Multiplying both sides of (2) by $\gamma(p)$, integrating them between \underline{p} and p and using integration by parts we obtain

$$\mathcal{N}_{t+1}(p) = \mathcal{N}_t(p) [1 - \bar{\delta}_{t+1}(p) - \rho_{t+1} - s\lambda_{t+1}\bar{\mathcal{F}}_{t+1}(p)] + \rho_{t+1}\mathcal{N}_t(\bar{p})\mathcal{F}_{t+1}(p) + \lambda_{t+1}[1 - \mathcal{N}_t(\bar{p})]\mathcal{F}_{t+1}(p)$$

where

$$\bar{\delta}_{t+1}(p) = \frac{1}{\mathcal{N}_t(p)} \cdot \int_{\underline{p}}^p \delta_{t+1}(q) d\mathcal{N}_t(q)$$

is the average separation rate into unemployment from firms of productivity up to p . A fraction $\bar{\delta}_{t+1}(p) + \rho_{t+1} + s\lambda_{t+1}\bar{\mathcal{F}}_{t+1}(p)$ of the employees below productivity p are separated, and a measure $\rho_{t+1}\mathcal{N}_t(\bar{p})$ are reallocated from employment and $\lambda_{t+1}[1 - \mathcal{N}_t(\bar{p})]$ from unemployment, and with chance $\mathcal{F}_{t+1}(p)$ join a firm of productivity at most p .

In MPV13's contractual environment, RP dynamics (2) are implemented by equilibrium offers $V_t(p)$ and advertising rates $a_t(p)$ that solve a system of first-order conditions and Euler equations of the firm's optimal contracting problem.

2.3 The dynamic job ladder

Equilibrium contracts in MPV13 converge, absent aggregate uncertainty, to the stationary wage posting equilibrium. The RP employment dynamics, however, are shared by a larger class of models, in which more productive firms always offer more in equilibrium, under different assumptions about the hiring technology or the contract space. One example is Coles and Mortensen (2011), who relax two constraints imposed by MPV13, full commitment and equal treatment, and solve for one RPE under a specific hiring cost function.

In this class of models, we can express RPE employment dynamics (2) in terms of the employer's rank in the productivity distribution, $x = \Gamma(p)$, rather than of its productivity p itself, which will prove convenient in our empirical application. Let

$$\begin{aligned} N_t(x) &= \mathcal{N}_t(\Gamma^{-1}(x)) \\ F_t(x) &= 1 - \bar{F}_t(x) = \mathcal{F}_t(\Gamma^{-1}(x)) \\ \delta_t(x) &= \delta_t(\Gamma^{-1}(x)) \\ \bar{\delta}_t(x) &= \frac{1}{N_t(x)} \cdot \int_0^x \delta_{t+1}(x') dN_t(x') \end{aligned}$$

denote, respectively, employment at, and fraction of job ads posted by, all employers below the x^{th} quantile in the productivity distribution, which is also (by RPE) the x^{th} quantile in the size distribution, and the separation rate into non-employment from a firm at or below that quantile. We can then rewrite (2) as a *dynamic job ladder equation*

$$\begin{aligned} N_{t+1}(x) &= N_t(x) [1 - \bar{\delta}_{t+1}(x) - \rho_{t+1} - s\lambda_{t+1}\bar{F}_{t+1}(x)] \\ &\quad + \rho_{t+1}N_t(1)F_{t+1}(x) + \lambda_{t+1}[1 - N_t(1)]F_{t+1}(x). \end{aligned} \quad (3)$$

We will estimate turnover rates from this equation using data on employment stocks, net and gross flows by rank of employers in the size distribution. Before doing so, we illustrate the data and descriptive evidence of the empirical relevance of the job ladder model.

3 The dynamic job ladder: descriptive evidence

3.1 Data sources

We examine the cyclical reallocation of employment among firms and establishments, especially around the Great Recession, through the lens of the job ladder model. In equilibrium, the TFP and size of an employer determine its rung on the wage ladder, which in turn predicts its gross flows of hires, separations to non-employment, and quits to other jobs, thus

also its net employment change. Because the size of an employer, unlike its TFP¹, can be accurately measured in the data under relatively weak assumptions, we focus on size, as well as wage, as empirical counterparts of the job ladder rung.

The primary source of information on businesses from the BLS is the Quarterly Census of Employment and Wages (QCEW) program. It publishes a quarterly count of employment and wages reported by employers covering 98 percent of U.S. jobs, both private and public sector, available at the county, MSA, state and national levels by industry. Information is also published at annual frequency (covering the first quarter of the year) by establishment size, in one of ten size classes, with lower bounds 1, 5, 10, 20, 50, 100, 250, 500, and 1,000 employees.² From this we draw the distribution of establishment counts, employment and weekly earnings per worker, all averaged over the first quarter of the year, by establishment size, for each year from 1990 to 2012 included, for the U.S. and all industries combined.³

The Statistics of US Businesses (SUSB) program at the Bureau of the Census publishes annual data on total employment, payroll and (every five years) receipts by firm size, disaggregated in 17-20 size categories, from 1992 to 2010, 2004 excluded (the size classification is coarser in 1992-1993). The Census defines a firm by grouping establishments by legal form and control structure.

The Job Openings and Labor Turnover Survey (JOLTS) comprises about 16,000 establishments, a size-stratified sample from the QCEW frame, surveyed every month according to a rotating panel structure. JOLTS measures job openings, hires, layoffs, quits, and other separations at the establishment level. Recently, the BLS published this information also by size of the establishment, in one of six size classes, with lower bounds equal to 10, 50, 250, 1000, and 5,000 employees. This dataset is central to our exercise. Importantly, JOLTS by size class covers only the private sector, while aggregate JOLTS data cover also the public sector, just like its QCEW frame. This is an important caveat for the GR, where the public sector played a disproportionate role in first buffering employment losses and then dragging on the employment recovery.

In disaggregated JOLTS data, an establishment is assigned to a size class according to the maximum size it attained in the 12 months preceding its inclusion in the sample,

¹In the US, information on sales at the firm level, necessary to compute TFP, is not available for a representative sample of firms from all industries.

²<http://www.bls.gov/cew/cewsize.htm> accessed 2/23/2013.

³For earnings at the national level, all industries, we find two outliers, possibly the result of some coding error in collating the semi-aggregated data, in size class “10-19 employees” in year 1999 and in size class “1,000 employees and up” in 1995. We replace those two values of earnings with the average of the entries in adjacent years for the same size class. Although this averaging introduces measurement error, the year-over-year changes implied by the BLS original entries differ from all the rest of the sample by two orders of magnitude.

independently of how its size changes during the year.⁴ So, within each survey year we know that the identity, hence (in our model) the productivity quantiles of establishments in each size class are fixed. Because this size classification follows an “initial employment” criterion, it is known to be subject to a mean reversion bias, which creates the illusion of a negative size-growth relationship, possibly through a positive correlation between size and quits. The job ladder predicts the opposite, negative correlation between size and quits, as larger employers pay more and retain more workers. This issue is likely to matter more in narrower size classes, at the bottom of the size distribution. We will return to this issue when discussing size misclassification. Finally, in the analysis that follows, we will aggregate the largest two size categories available in the JOLTS sample (1,000-4,999 and over 5,000 employees) into one single category (over 1,000 employees). We do this for two main reasons. First, the largest size cutoff in the CEW sample described above is 1,000 employees. As we get our establishment counts from CEW, we will need to merge information from CEW and JOLTS, which constrains us to use size cutoffs that are available in both data sets. Second, the 5,000+ category in JOLTS is very small (it accounts for less than 2.5% of total employment in the JOLTS sample and probably covers very few establishments), and the data pertaining to this category are somewhat noisy. The loss of information implied by our aggregation of the largest two size classes into one is therefore arguably relatively minor.

Because JOLTS is a survey of employers, it provides a meaningful distinction between layoffs and quits, but not between quits to (or hires from) non-employment and other employers. The key engine of the job ladder are job-to-job quits. We supplement JOLTS with information on gross worker flows from the monthly Current Population Survey (CPS). Specifically, we use the hazard rates of transition between Employment (E), Unemployment (U) and Non-participation (N) estimated by Fallick and Fleischman (2004) from gross flows (using monthly matched files), starting in January 1994 and updated by the authors through February 2013. This series begins with the 1994 re-design of the CPS, which introduced a question on the change of employer that made it possible to measure the EE hazard, and which greatly improved the reliability of employment status and thus reduced margin error.

Finally, also from the worker side, we use the 1996-2008 panels of the Survey of Income and Program Participation (SIPP), which provides information on employment status and employer identity, and coarse information on establishment size class (three classes) at weekly frequency for a representative sample of the U.S. workforce.

⁴JOLTS (re-)sampling dates are December 2000, December 2003, February 2005, March 2006, and every March until 2013. A new JOLTS sample is put in place in the month following each re- sampling.

3.2 The wage-size job ladder


We start by presenting evidence that larger employers pay more, so employer size is a meaningful proxy for its rung on a job ladder. It has been long documented that employer size correlates positively with wage rates, after controlling for observable worker characteristics (Brown and Medoff, 1989). In this paper we focus on employer-level data and take the extreme view that workforce quality is homogeneous across employer sizes, so that any wage differential related to size can be thought of as a wage premium. This is in the spirit of the model we presented earlier. If wage/size differentials reflected entirely different types of workers at employers of different sizes, we would still have to explain why workers sort by the size of their employer.

Establishment size class	I	II	III	IV	V
1 to 4 (omitted)					
5 to 9	-27.62 (0.001)	-17.01 (0.01)	5.59 (0.02)	-24.83 (0.008)	5.93 (0.02)
10 to 19	-13.62 (0.001)	-4.02 (0.01)	27.04 (0.03)	-8.04 (0.01)	34.82 (0.02)
20 to 49	3.80 (0.002)	23.45 (0.01)	50.05 (0.03)	5.70 (0.01)	52.37 (0.03)
50 to 99	28.61 (0.003)	38.38 (0.02)	64.23 (0.05)	29.81 (0.02)	68.96 (0.05)
100 to 249	55.22 (0.004)	48.12 (0.03)	73.24 (0.07)	55.15 (0.02)	81.47 (0.06)
250 to 499	101.94 (0.009)	76.23 (0.06)	88.30 (0.15)	102.19 (0.05)	82.21 (0.13)
500 to 999	158.14 (0.014)	112.73 (0.10)	95.43 (0.27)	156.93 (0.08)	70.93 (0.22)
1000 and up	272.12 (0.02)	226.22 (0.14)	174.40 0.38	263.51 (0.12)	131.98 (0.31)
industry dummies	N	2-digit	2-digit	N	5-digit
state dummies	N	N	Y	Y	N
R^2	0.94	0.87	0.56	0.60	0.78

Source: QCEW and authors' calculations. Dependent variable: average weekly earnings per worker in each establishment (1983\$). Standard errors in parentheses. All regressions include year dummies.

Table 1: The wage-size premium.

Table 1 reports the results of an establishment-level OLS regression of earnings per worker on size and other establishment characteristics. The dependent variable is a measure of real weekly earnings, the ratio between CPI-deflated total quarterly payroll and average employment among all establishments in the “cell”, for each year from 1990 to 2012 included. The cell depends on the specification, and is indicated by which dummies we include among the covariates. So the dependent variable varies across specifications, which are not directly comparable. Size dummies and year dummies are always included. In specification II, each



Source: SUSB.

Fig. 1: Wage differential with respect to firms of size 1-4 employees.

cell includes all establishments in the same size class and 2-digit NAICS industry. In III, each cell includes all establishments in the same size class, 2-digit NAICS industry, and located in the same US state. And so on. Information on geographic location is available only at the 2-digit industry level, due to potential disclosure risk. The regression is weighted by the number of establishments per cell. The results clearly indicate a wage ladder, except at the very bottom when not controlling for industry and location. Because average weekly earnings in the omitted (smallest) size class are about \$330, top-to-bottom pay differentials between largest and smallest establishments are in the order of 80% in specification I, and lower when controlling for industry and location. Employer-level TFP in our model may be in part a result of the industry in which the employer operates, so controlling for industry composition may not be appropriate.

Figure 1 reports evidence from the SUSB on wage/size premia at the firm level. We take total annual payroll per worker for each size class, divide by that of the smallest class 1-4 employees, and subtract one. We omit the size class ‘0’ employees, which includes entrants, because it reports payroll but not employment. The results speak for themselves. The wage differential between the largest and smallest firms is less than 50%, significantly smaller than or establishments, also taking into account that the largest firm size class starts at the higher 1,500 employees cutoff. As we will discuss later in more detail, this may be due to the fact

that large firms that only comprise large establishments may be the highest-paying of all. So, when combining them with equally large multi-establishment firms, their average pay premium declines.

3.3 Employment reallocation in the Great Recession

We now document the evolution of employment in different size classes of employers in recent years, with a special focus on the Great Recession.

Employment and establishment shares by size class. Figure 2 illustrates that employment shares by size of the establishment in QCEW are relatively stable over time, but do exhibit the cyclical pattern documented by MPV12 for firms; namely, the share of larger employers declines in the three recessions in the sample period. The GR is no exception.⁵

Our empirical exercise is based on the assumptions that the distribution of employers by TFP is time-invariant and coincides with their distribution by size. For example, in our theoretical model the median productivity employer is always of median size. One implication of this assumption is that the distribution of establishment counts by size classes should be relatively stable at business cycle frequencies. This is true in JOLTS size data by construction of the dataset, so the identity of the establishments is fixed, at least within each sampling year typically March to February. Across years, Figure 3 illustrates these shares in QCEW, which is the frame from which JOLTS is drawn. Shares are in log scale to make them visible, because the distribution of establishment counts is much more compressed at the low end. We can see a very modest trend and cyclical component. By and large, the distribution of establishment counts is stable, much more so than that of employment.

Gross and net flows by size class. Employment shares provide a limited view of the size of the businesses that were most affected by the Great Recession. As we discussed in MPV12, to avoid the so-called reclassification bias we need to study business dynamics by their initial size. We showed there that the annual growth rate of employment at initially large ($>1,000$ employees) minus small (<50 employees) firms in the US is strongly negatively correlated

⁵We can draw the same information on employment shares by firm size, at a finer degree of size classification, from two BLS datasets, where a firm is identified by a federal tax Employer Identification Number (EIN). First, the Business Employment Dynamics (BED) program collects information on job flows and stocks, from the same QCEW universe, at quarterly frequency starting in 1992, and presents them by size of the parent company. Second, the Current Employment Statistics (CES) program is the well-known monthly “pay-roll survey” of about 145,000 businesses and government agencies from the QCEW frame, representing approximately 557,000 individual work sites. The survey provides timely and detailed industry data on employment, hours, and earnings of workers on nonfarm payrolls. In both datasets, once again, the share of employment at small firms is countercyclical. In the GR, it rose especially in the second, deeper half of the downturn. Results are available upon request.




Fig. 2: Employment shares by establishment size class.




Fig. 3: Shares of establishments by establishment size class.




Fig. 4: Differential employment growth between establishment size classes.

with unemployment in 1979-2010. Here we zoom in on the Great Recession using higher frequency, monthly data updated to cover the post-GR recovery. Figure 4 repeats the exercise using JOLTS data by establishment size (we remark again that these are establishments, not firms). It appears that small establishments were hit especially hard by the credit crunch.

To examine in more detail the nature of these patterns, we can examine gross worker flows. This is a unique advantage of JOLTS and, to the best of our knowledge, we are the first to document the behavior of these flows by employer size around the GR. By definition, net employment growth in JOLTS equals hires minus the sum of layoffs, quits and other separations (such as retirement). The latter category is small and fairly acyclical, thus we focus on hires, layoffs and quits. Figure 5 plots hire rates (new accessions divided by employment) by establishment size. Hire rates begin to decline before the GR. Surprisingly, during the deepest phase of the financial crisis following the Lehmann Brothers episode, hire rates collapse at the larger establishments and not at the smaller ones; they even briefly spike up in the smaller class in late 2008-early 2009.

We will interpret the surprising fact revealed in Figure 5 through the lens of our model. Given that we saw in Figure 4 that the smaller establishments fared worst in terms of net employment growth, especially from the last quarter of 2008 on, it must be the case that their separations rose disproportionately, and more than compensated their brisker hiring




Fig. 5: Hire rates by establishment size class.




Fig. 6: Layoff rates by establishment size class.




Fig. 7: Total employment-to-employment hazard.

pace. We in fact see in Figure 6 that layoff rates rose sharply and temporarily, more so the lower the establishments.

Finally, we turn to quits. One disadvantage of JOLTS is that quits are not broken down by destination, non-employment or other employers, two components that play distinct and crucial roles in our job ladder model. Figure 7 plots the employment-to-employment (EE) transition hazard from the matched files of the monthly Current Population Survey, computed by Fallick and Fleischman (2004) and updated by the authors through February 2013. While the EE rate is clearly procyclical and dropped significantly during the GR, the most striking aspect of the data is the declining trend. Off that trend, the decline in the GR was not especially pronounced, and the recovery afterwards was significant. It is known, however, that EE transitions include involuntary reallocation and other events that reduce worker’s earnings. Therefore, per se they provide only limited information on the extent to which workers climb the job ladder. Our theoretical model explicitly accommodates this possibility through reallocation shocks ρ_t .

The CPS has no information on the size of the worker’s employer. For that, we turn to SIPP, starting with the 1996 panels. In Figure 8 we show the share of all hires that are not from unemployment but directly from other employers, thus entail an EE transition, by size of the hiring “workplace”, the phrasing in the SIPP question that we interpret to




Fig. 8: Share of hires from other employers, by employer size.




Fig. 9: Differential share of hires from other employers, by employer size.

be an establishment. As predicted by the job ladder model, larger employers always hire more from other firms, and especially so late in expansions when the market tightens and competition for workers stiffens. In the GR, this share collapsed for all size groups. Since total hires also declined sharply, this is the strongest evidence that the job ladder came to a grinding halt. Taking the difference between the two series in Figure 9 and detrending carries the point home even more clearly.

Vacancies by size class. We next turn to JOLTS again and describe the behavior of measured job vacancies by size class. Figure 10 reports the time series of total job openings for each JOLTS size class. Figure 11 further shows vacancy *shares* by size class, i.e. vacancies in each size class divided by total aggregate vacancies. If recorded job openings were a good measure of hiring effort (the function $a_t(p)$ in our job ladder model),⁶ then the series plotted on Figure 11 would represent the sampling probabilities of each size class. Next, Figure 12 show vacancy shares divided by the number of establishment in each class.⁷ We refer to those series as the *vacancy weights* by size class. The series on Figure 12 are normalized at one in January 2001 to harmonize scales.

Figure 10 clearly shows that vacancies plummeted across the board during the GR, with vacancy levels seeming to closely track each other across the various size classes. At first glance, Figures 11 and 12 reinforce that impression, as the movements in vacancy shares and weights appear small relative to the absolute decline seen on Figure 10 which, to a first approximation, was uniform. On closer inspection, Figures 11 and 12 further suggest that there is no evidence of a large impact of the financial crisis (post-September 2008) on small establishments: the movements are relatively modest, and the 10-49 employee class shows the largest change, but upwards.

Finally, Figure 13 shows the vacancy yield, namely the ratio between hires and vacancies reported a month before, by establishment size. While it is known that during and after the GR the aggregate yield rose enormously with unemployment duration, so it became as easy for firms to fill vacancies as it was difficult for the unemployed to find work, Figure 13 shows that this phenomenon was more pronounced the smaller the establishment, and the vacancy yield actually *fell* during the acute phase of the GR at the largest establishments. This surprising fact is consistent with the collapse in hires from employment, on which larger establishments rely more.

⁶There are good reasons to believe that they are not, as we discuss below in Section 4.

⁷We take those numbers from CEW.




Fig. 10: Vacancies by establishment size class.




Fig. 11: Vacancy shares by establishment size class.




Fig. 12: Vacancy weights by establishment size class.




Fig. 13: Vacancy yields by establishment size class.

4 The dynamic job ladder: calibration

4.1 Gross and net flows

We now return to our dynamic job ladder model and ask whether it can make sense of this evidence. To calibrate the job ladder model, we use a minimum distance method. Our target empirical moments are gross and net employment flows by size class of the employer observed in JOLTS. Given our strong assumption that employer size is a relevant rung of the job ladder, it is far from obvious that the job ladder dynamic Equation (3) can replicate actual observations on gross and net flows, every month for 12 years. Among many restrictions, our theory predicts that smaller employers should lose a larger proportion of workers to job-to-job quits. Testing all joint restrictions of the job ladder equation is our main goal here.

Beginning with net flows, we saw that total employment growth at employers of productivity rank less than some given threshold x is governed by the RPE by the dynamic Equation (3), repeated here for convenience:

$$N_{t+1}(x) - N_t(x) = - [\bar{\delta}_{t+1}(x) + \rho_{t+1} + s\lambda_{t+1}\bar{F}_{t+1}(x)] N_t(x) + \rho_{t+1}F_{t+1}(x)N_t(1) + \lambda_{t+1}F_{t+1}(x)U_t \quad (3)$$

This net flow equals total hires from non-employment $UE_t(x)$, minus total separations into non-employment $EU_t(x)$, minus total job-to-job quits $QE_t(x)$, plus total job-to-job hires.⁸ Focusing on the first three of those gross flows (the fourth one being redundant, given the net flow), the job ladder model predicts:

$$UE_t(x) = \lambda_{t+1}F_{t+1}(x)U_t \quad (4)$$

$$EU_t(x) = \bar{\delta}_{t+1}N_t(x) \quad (5)$$

$$QE_t(x) = \rho_{t+1}N_t(x) + s\lambda_{t+1} \int_0^x \bar{F}_{t+1}(x') dN_t(x'). \quad (6)$$

Apart from the parameter s (the search intensity of employed relative to unemployed workers), Equations (3) to (6) involve six time series — $\bar{\delta}_t(\cdot)$, λ_t , ρ_t , $F_t(\cdot)$, $N_t(\cdot)$, and U_t —, none of which is directly observed in the data. We now explain how we map monthly empirical observations from JOLTS on net employment growth, total layoffs, and total quits (which combine quits to other employers and quits to non-employment) by establishment size into our time series of interest, based on Equations (3)-(6).

⁸In the notation just laid out, we use the letter U to imply non-employment. The model is silent on any possible distinction between unemployment and non-employment. We will return to this issue momentarily.

4.2 Size and productivity ranks

All gross and net flows described in Equations (3)-(6) are functions of employers' productivity ranks, x . In the absence of a direct measure of employer productivity p , those ranks are not directly observed in the data. Yet a powerful consequence of equilibrium being RP is that, even though an employer's size moves around over time, its rank in the distribution of sizes stays constant and equal to its rank in the productivity distribution.

Assuming for the time being that size is correctly measured, and that establishment size does reflect productivity in RPE (which means assuming that productivity p is an establishment-specific parameter, an assumption that we will address momentarily), establishments in a given JOLTS size class k will be representative of all establishments with productivity ranks between two cutoff values, $x \in [X_{k-1}, X_k]$, with $\{X_k\}_{k=1}^K$ an increasing sequence in $[0, 1]$, which remain fixed so long as the identities of establishments assigned to size class k do not change. In JOLTS, each month except at re-sampling dates, 1/12 of the survey respondents are replaced with ex ante identical establishments, which had the same size and industry at the time of sampling. Under the assumption, which underlies the gradual rotation scheme, that these are statistically equivalent establishments, we can effectively treat the identities and size class membership of the JOLTS establishments as constant between re-sampling times.

Therefore, the JOLTS sample provides observations at (almost) all dates of cumulated employment $N_t(X_k)$, layoffs, and total quits (and, potentially, sampling weights $F_t(X_k)$ — see below), for K productivity quantiles $\{X_k\}_{k=1}^K$ corresponding to as many size classes.⁹ In what follows, we should keep in mind that X_k is the cutoff productivity quantile between size classes k and $k + 1$. With K size classes, this implies that $X_K \equiv 1$. We will also use the convention $X_0 = 0$. We now confront Equations (3)-(6) with the JOLTS sample.

4.3 Separations into non-employment

A survey of employers like JOLTS detects whether a separation is a quit or a layoff from the viewpoint of the surveyed establishment. As workers are neither interviewed nor tracked after separation, measured quits from establishments aggregate quits to unemployment and quits to other jobs, a distinction that is central to the logic of the job-ladder model, where the former are part of total separations into unemployment $\bar{\delta}_{t+1}(x) N_t(x)$, and the latter are upgrades. To estimate $\bar{\delta}_{t+1}(x)$, we thus need some way to break down quits into those to unemployment and those to other employers. To do so, we need worker-side information.

⁹As discussed earlier, the raw JOLTS sample has six establishment size classes: 1 to 9, 10 to 49, 50 to 249, 250 to 999, 1,000 to 4,999, and over 5,000 employees. For reasons discussed earlier, we lump the largest two classes into one.

Focusing first on the aggregate separation rate ($x = 1$), we seek to construct $\bar{\delta}_{t+1}(1)$ based on Equation (5) as the ratio between the total monthly flow from employment to non-employment (EU flow) $EU_t(1)$ and total employment $N_t(1)$. The total EU flow consists of layoffs plus quits into non-employment. We supplement the JOLTS data with the transition rates estimated from CPS by Fallick and Fleischman (2004), updated by the authors through early 2013. For every month t we compute the share σ_t^{FF} of total transitions that are employer-to-employer (EE), as opposed to transitions into non-employment (EU):

$$\sigma_t^{\text{FF}} = \frac{EE_t^{\text{CPS}}}{EE_t^{\text{CPS}} + EU_t^{\text{CPS}}}.$$

All EE transitions are quits in the job ladder model. Assuming that the CPS-based share applies to the workers employed by the JOLTS sample of establishments, we multiply total separations in JOLTS by $1 - \sigma_t^{\text{FF}}$ to obtain an estimate of aggregate separations into non-employment, $EU_t(1)$, that is consistent with the JOLTS data. The corresponding aggregate separation rate is then $\bar{\delta}_{t+1}(1) = EU_t(1) / N_t(1)$.¹⁰

This procedure further gives us the share of aggregate EU separations that are quits. As mentioned earlier, JOLTS has a measure of total layoffs and discharges, which we can subtract from our newly constructed time series $EU_t(1)$ to obtain total quits into non-employment in JOLTS. Introducing the ancillary — yet economically meaningful — parameters $\psi_{t+1}(x)$, defined as the share of quits in total EU separations from employers of productivity rank x , and

$$\bar{\psi}_{t+1}(x) = \frac{1}{\bar{\delta}_{t+1}(x) N_t(x)} \cdot \int_0^x \delta_{t+1}(x') \psi_{t+1}(x') dN_t(x'),$$

the same share of quits in EU separations from employers of productivity rank up to x , we see that, in this notation, $\bar{\psi}_{t+1}(1)$ is the share of quits in aggregate separations into non-employment, $EU_t(1)$, that we obtain from our procedure. The aggregate layoff probability is then $\bar{\delta}_{t+1}(1) \cdot (1 - \bar{\psi}_{t+1}(1))$, and the probability of quitting into non-employment is $\bar{\delta}_{t+1}(1) \bar{\psi}_{t+1}(1)$. Both of those, plus the total aggregate transition rate into non-employment $\bar{\delta}_{t+1}(1)$, are plotted in Figure 14.¹¹ While most of this Figure has the familiar feature of a largely a-cyclical probability of transition into non-employment, the GR stands out as a striking exception, with a sudden (and short-lived) surge in layoffs in the immediate aftermath of the collapse of Lehman Brothers in September 2008.

¹⁰To the best of our knowledge, ours is the first attempt to exploit information from both the employer and the employee sides to draw empirically the distinction between the three main types of separations: layoffs, quits to non-employment, and quits to other employers. Worker surveys such as CPS and SIPP are notoriously plagued by noise in the layoff/quit distinction when the worker loses a job. Administrative datasets do not typically contain information about the reason for separation.

¹¹All the raw JOLTS series are smoothed using a 6-month moving average around each point prior to calibration, to remove the fairly large amount of high-frequency noise in those series.




Fig. 14: Average separation rate into non-employment $\bar{\delta}_{t+1}(1)$ and its components

The Fallick and Fleischman (2004) series are only available at the aggregate level. Therefore, making our quit/layoff distinction operational at lower levels of aggregation ($x < 1$, which we shall need later in the calibration) requires additional assumptions. The identifying assumption that we opt for here is that the probability with which workers quit into non-employment, $\psi_{t+1}(x)\delta_{t+1}(x)$, is independent of their employer's productive type x . That is to say, for all x , $\psi_{t+1}(x)\delta_{t+1}(x) \equiv \bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1)$.¹² Because the theory requires for RPE that the total separation rate into non-employment $\delta_{t+1}(x)$ be non-increasing in productivity rank x , this assumption implies that both total separations and layoff rates are decreasing in productivity (or size) rank x . Both implications hold in the JOLTS data. This additional identifying assumption enables us to construct total separations into non-employment from employers with productivity rank up to X_k , $\bar{\delta}_{t+1}(X_k)N_t(X_k)$, for all cutoff productivity quantiles X_k corresponding to the JOLTS size classes, as the sum of total layoffs from employers in size classes up to k (directly available from the JOLTS data), plus total quits into non-employment from those employers, equal to $\bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1)N_t(X_k)$ by

¹²Any assumption we make at this point is necessarily arbitrary to some degree. An alternative is to assume that the share of EU separations that are quits is independent of productivity, i.e. that $\psi_{t+1}(x) \equiv \bar{\psi}_{t+1}(1)$ for all x . This implies that not only the layoff rate, but also the quit rate into non-employment is decreasing in employer productivity/size, or rank thereof. Results based on this alternative assumption, available upon request, are qualitatively identical, and quantitatively very close, to the ones we present here.

assumption. Given observations on the cumulated employment distribution $N_t(X_k)$, this allows to directly estimate the total probability of transition into non-employment by size class, $\bar{\delta}_{t+1}(X_k)$.

4.4 Job finding probability

Equation (3) applied at the top quantile $x = 1$ gives the law of motion of aggregate employment: $N_{t+1}(1) = [1 - \bar{\delta}_{t+1}(1)] N_t(1) + \lambda_{t+1} U_t$, where $U_t = 1 - N_t(1)$ is non-employment. From this equation, we can back out the baseline job finding rate:

$$\lambda_{t+1} = \frac{N_{t+1}(1) - [1 - \bar{\delta}_{t+1}(1)] N_t(1)}{U_t} = \frac{UE_t(1)}{U_t}.$$

Construction of λ_{t+1} from this equation thus requires knowledge of the stock of non-employed job seekers, U_t . Here again, we call on the Fallick and Fleischman (2004) CPS series, which offers a breakdown of the total non-employment to employment flow ($UE_t(1)$ in our notation) into the flow from unemployment into employment and the flow from inactivity into employment. Taking the (average) ratio of the latter over the former gives us an estimate of the average relative job finding rate of inactive workers, relative to the unemployed. Calling this relative job finding rate s_0 (so that the job finding probability of non-participants is $s_0 \lambda_{t+1}$), we then construct the effective pool of non-employed job seekers as:

$$\frac{U_t}{N_t(1)} = \frac{u_t^{\text{CPS}}}{1 - u_t^{\text{CPS}}} + s_0 \left(\frac{1 - e_t^{\text{CPS}}}{e_t^{\text{CPS}}} - \frac{u_t^{\text{CPS}}}{1 - u_t^{\text{CPS}}} \right),$$

where u_t^{CPS} is the CPS unemployment rate and e_t^{CPS} is the CPS employment-population ratio. The value of s_0 thus calibrated is 0.2, and the resulting job finding rate series is plotted in Figure 15. While it exhibits the familiar cyclicity, including a vertiginous drop during the GR, its level is fairly low because it includes transitions to employment from inactivity, which are a small fraction of the stock of inactive individuals.

4.5 Sampling distribution and job-to-job quits

We now turn to the last gross flow predicted by the job ladder, namely job-to-job quits $QE_t(x)$, given in Equation (6). We show how this equation, combined with the net flow Equation (3) and with the JOLTS data, allows identification of the sampling distribution $F_{t+1}(\cdot)$, the reallocation shock ρ_{t+1} , and the relative intensity of employed search, s .

One easy option to estimate the sampling distribution $F_{t+1}(\cdot)$ would be to set it equal to the observed distribution of job openings by size class, which is readily available from JOLTS. However, the sampling distribution that is consistent with the model will only coincide with




Fig. 15: Job finding rate λ_{t+1}

the empirical distribution of job openings if (a) job openings are measured accurately in JOLTS, and (b) job opening counts are a good measure of actual hiring effort (in particular, all vacancies have equal sampling weights). Both of these are questionable assumptions: for example, Davis, Faberman and Haltiwanger (2010) have recently forcefully argued that neither was true, especially at the low end of the establishment size distribution.

Luckily, the law of motion of employment in RPE offers an alternative solution to estimate $F_{t+1}(\cdot)$. Equation (3) defines the sampling distribution at cutoff productivity quantiles X_k and at all dates as:

$$\begin{aligned} \bar{F}_{t+1}(X_k) = & \\ & \frac{[N_{t+1}(1) - N_{t+1}(X_k)] - (1 - \rho_{t+1})[N_t(1) - N_t(X_k)] + \bar{\delta}_{t+1}(1)N_t(1) - \bar{\delta}_{t+1}(X_k)N_t(X_k)}{\rho_{t+1}N_t(1) + s\lambda_{t+1}N_t(X_k) + \lambda_{t+1}U_t} \end{aligned} \quad (7)$$

that we will use to estimate sampling weights $F_t(\cdot)$, employed search efficiency s , and reallocation shocks ρ_t , using the time series for separation and accession probabilities $\bar{\delta}_{t+1}(X_k)$ and λ_t , and the stock of non-employment from CPS, U_t , all estimated as above, plus the stock of employment $N_t(X_k)$ in size classes up to k from JOLTS. Later, we will gauge how close the estimated sampling distribution from (7) (consistent with RPE employment dynamics by construction) is to the empirical distribution of job openings across size classes.

Knowledge of the sampling distribution $F_t(\cdot)$ allows the construction of total job-to-job quits in any size class k which, following Equation (6), equal

$$\text{QE}_t(X_k) - \text{QE}_t(X_{k-1}) = \rho_{t+1} [N_t(X_k) - N_t(X_{k-1})] + s\lambda_{t+1} \int_{X_{k-1}}^{X_k} \bar{F}_{t+1}(x) dN_t(x), \quad (8)$$

the empirical counterpart of which is total quits in JOLTS size class k , minus quits into non-employment from employers in that size class, which were estimated in subsection 4.3 as $\bar{\psi}_{t+1}(1) \bar{\delta}_{t+1}(1) \cdot [N_t(X_k) - N_t(X_{k-1})]$. Fitting (8) to this JOLTS counterpart at each date t and size class k allows, in principle, identification of both the (constant across dates and classes) search intensity of employed workers s , and the (constant across classes) reallocation shock ρ_t .

This last statement must be qualified as follows. First, in order to limit the computational cost of this calibration, and to attain more precise identification, we further restrict the reallocation probability ρ_t to equal a constant (ρ) times the baseline job finding rate λ_t . While not strictly necessary, this restriction considerably reduces the number of parameters to estimate, from one values of ρ_t for each month in the sample (140 in total) down to a single scalar, ρ . Second, Equation (8) is not directly implementable, as the transformed net flow Equation (7) only gives the sampling distribution at the cutoff quantiles X_k , whereas in principle we would need it over its entire support to calculate the integral in (8). We will approximate $F_t(\cdot)$ for that purpose using shape-preserving spline interpolation between the known nodes $F_t(X_k)$.¹³

4.6 Misclassification

The issue. So far we assumed that an establishment's size, as measured in JOLTS, is the 'relevant' measure of size, in the sense that it reflects the relevant value of productivity p for that establishment. There are at least two reasons to doubt that this is always the case. The first one is random fluctuations in establishment size. While the job ladder model uses a large-number approximation and treats establishment size as evolving deterministically over time, in reality establishment size will fluctuate randomly around the mean value predicted by the job ladder. If, at the time of JOLTS re-sampling, an establishment has an exceptionally high (say) realization of the random component of its size, that establishment may be assigned to the 'wrong' size class, i.e. a size class that reflects its transitory larger size rather than its long-run smaller size. This will be especially true of smaller establishments, both because the large-number approximation is less accurate for small establishment, and also because the small size classes (1-9 and 10-49 employees) are narrower than the larger ones.

¹³Specifically, given that the relevant integration variable in (8) is N_t , we interpolate F_t against N_t at each date.

The second reason to suspect that establishment size does not perfectly reflect the relevant productivity parameter is that many establishments are part of multi-establishment firms. Depending on the degree of decentralization and devolution in the parent firm’s management, the relevant productivity parameter for those establishments may be at the level of the parent firm, in which case the size measure that will reflect productivity is not the size of the establishment, but that of the parent firm, which we do not observe in JOLTS. Indeed, in MPV12 we document from the Census’ Business Dynamics Statistics that the average size of an establishment grows with the size of the parent company, but levels at about 60 employees when the size of the firm reaches 250, and is still about 60 for firms employing over 10,000 workers. So very large firms (national banks and retailers come to mind) own hundreds or even thousands of separate, relatively small establishments, whose workers benefit from the productivity and compensation policy of the parent company.

For both reasons, observed size classes in JOLTS and true ‘productivity classes’ (rungs on the job ladder) may not coincide. We propose to tackle those two issues and reconcile size and productivity classes by modeling misclassification explicitly. To avoid confusion, we now introduce a distinction between *size class* k , defined based on the JOLTS sample as the set of establishments whose observed size falls between two given cutoff values (e.g. 50 to 249 employees), and *productivity class* k , defined as the set of establishments whose unobserved productivity rank falls within the quantile interval $[X_{k-1}, X_k]$.

Modeling misclassification. Consider an establishment with productivity rank x . Suppose that the ‘true’ (or model-predicted) size $\ell_t(x) = dN_t(x)/dx$ is observed with time-invariant probability $1 - \pi(x)$, while with probability $\pi(x)$ what is observed is the true size $\ell_t(x')$ of an establishment with productivity rank x' drawn at random from some distribution M , which, to attain identification, we assume to be independent of x and time. To lighten notation, we now drop the time index, but we should keep in mind that employment measures, observed or reclassified, are time-varying, while reclassification probabilities π and weights M are assumed to be constant over time.

Size classes with misclassification. Next consider size classes. We can define size class k as the set of all establishments whose *observed* size ℓ^o falls within some interval $[\ell(X_{k-1}), \ell(X_k)]$. Establishments assigned to size class k are either establishments with productive ranks $x \in [X_{k-1}, X_k]$ whose true size is observed, or they are from the share of all establishments that are misclassified as establishments with ranks $x \in [X_{k-1}, X_k]$. Observed

employment in size class k is therefore:

$$n_{kt}^o = \int_{X_{k-1}}^{X_k} (1 - \pi(x)) \ell_t(x) dx + m_k \int_0^1 \pi(x) \ell_t(x) dx, \quad (9)$$

where $m_k = M(X_k) - M(X_{k-1})$. Class weights m_k in the misclassification distribution are unknown, and added to the set of parameters to calibrate.

To gain some tractability and amenability to calibration, we further restrict misclassification probabilities to be constant within productivity classes, i.e. we impose $\pi(x) \equiv \pi_k$ for $x \in [X_{k-1}, X_k]$. With this approximation,¹⁴ (9) becomes:

$$n_{kt}^o = (1 - \pi_k) n_{kt} + m_k \sum_{k'=1}^K \pi_{k'} n_{k't},$$

where $n_{kt} = N_t(X_k) - N_t(X_{k-1})$ is true employment in productivity class k . This implies:

$$\begin{pmatrix} n_{1t}^o \\ n_{2t}^o \\ \vdots \\ n_{Kt}^o \end{pmatrix} = \begin{pmatrix} 1 - \pi_1 + \pi_1 m_1 & \pi_2 m_1 & \cdots & \pi_K m_1 \\ \pi_1 m_2 & 1 - \pi_2 + \pi_2 m_2 & \cdots & \pi_K m_2 \\ \vdots & \vdots & \ddots & \vdots \\ \pi_1 m_K & \cdots & \pi_{K-1} m_K & 1 - \pi_K + \pi_K m_K \end{pmatrix} \begin{pmatrix} n_{1t} \\ n_{2t} \\ \vdots \\ n_{Kt} \end{pmatrix} := \mathbf{M} \begin{pmatrix} n_{1t} \\ n_{2t} \\ \vdots \\ n_{Kt} \end{pmatrix}, \quad (10)$$

which in turn implies that ‘true’ employment in productivity class k can be inferred from observed employment in size class k as:

$$\begin{pmatrix} n_{1t} \\ \vdots \\ n_{Kt} \end{pmatrix} = \mathbf{M}^{-1} \begin{pmatrix} n_{1t}^o \\ \vdots \\ n_{Kt}^o \end{pmatrix}.$$

Measurement equations with misclassification. The transition rates λ_t, δ_t are estimated only off aggregate magnitudes and do not depend on misclassification. With our assumption of a productivity-independent probability of quitting into non-employment, neither does said probability ($\bar{\psi}_t(1) \bar{\delta}_t(1)$). Misclassification, however, does affect observed job-to-job quits from establishments in class k . To see how, note that observed total quits from employers in productivity class k are:

$$\begin{aligned} Q_{kt}^o &= \int_{X_{k-1}}^{X_k} [\bar{\psi}_{t+1}(1) \bar{\delta}_{t+1}(1) + \rho_{t+1} + s \lambda_{t+1} \bar{F}_{t+1}(x)] (1 - \pi(x)) dN_t(x) \\ &\quad + m_k \int_0^1 [\bar{\psi}_{t+1}(1) \bar{\delta}_{t+1}(1) + \rho_{t+1} + s \lambda_{t+1} \bar{F}_{t+1}(x)] \pi(x) dN_t(x). \end{aligned}$$

Under the assumption of misclassification probabilities that are constant in each productivity class and time-invariant, $\pi(x) = \pi_k$ for $x \in [X_{k-1}, X_k]$, the expression for total observed quits

¹⁴Note that this is necessarily an approximation, as the boundaries of size classes in terms of productivity, the X_k 's, are likely to change at each JOLTS re-sampling date.

from class k becomes:

$$Q_{kt}^o = (1 - \pi_k) \int_{X_{k-1}}^{X_k} [\bar{\psi}_{t+1}(1) \bar{\delta}_{t+1}(1) + \rho_{t+1} + s\lambda_{t+1} \bar{F}_{t+1}(x)] dN_t(x) \\ + (1 - \pi_k) s\lambda_{t+1} \int_{X_{k-1}}^{X_k} \bar{F}_{t+1}(x) dN_t(x) + s\lambda_{t+1} m_k \sum_{k'=1}^K \pi_{k'} \int_{X_{k'-1}}^{X_{k'}} \bar{F}_{t+1}(x) dN_t(x),$$

so we obtain:

$$s\lambda_{t+1} \begin{pmatrix} \int_{X_0=0}^{X_1} \bar{F}_{t+1}(x) dN_t(x) \\ \vdots \\ \int_{X_{K-1}}^{X_K} \bar{F}_{t+1}(x) dN_t(x) \end{pmatrix} = \mathbf{M}^{-1} \begin{pmatrix} Q_{1t}^o - (\bar{\psi}_{t+1}(1) \bar{\delta}_{t+1}(1) + \rho_{t+1}) n_{1t}^o \\ \vdots \\ Q_{Kt}^o - (\bar{\psi}_{t+1}(1) \bar{\delta}_{t+1}(1) + \rho_{t+1}) n_{Kt}^o \end{pmatrix},$$

where \mathbf{M} is the conversion matrix defined in (10). Denoting the k^{th} element of the vector in the r.h.s. of this last equation by Q_{kt}^* and dividing by employment in productivity class k we thus have that

$$s\lambda_{t+1} \int_{X_{k-1}}^{X_k} \bar{F}_{t+1}(x) \frac{dN_t(x)}{n_{kt}} = \frac{Q_{kt}^*}{n_{kt}}. \quad (11)$$

This equation highlights the importance of introducing misclassification in our JOLTS data. The l.h.s. of (11) is the conditional expectation of $\bar{F}_{t+1}(x)$ within productivity class k . The r.h.s. of (11) is a measure of the rate of job-to-job quits from the size class that are motivated by better offers. The job-ladder model predicts unambiguously that both sides of the equation should be decreasing in size class k : larger employers are more productive, pay more and have an easier time retaining employees. Because $\bar{\psi}_{t+1}(1) \bar{\delta}_{t+1}(1) + \rho_{t+1}$ is constant across size classes k , this requires total quits from size class k to be lower the higher the size class. In the JOLTS data by establishment size, which is split in six size classes, the observed quit rate, Q_{kt}^o/n_{kt}^o actually *increases* between size classes $k = 1$ and $k = 2$ in all months and often also between $k = 2$ and $k = 3$.

4.7 Implementation: summary

For given reallocation shocks ρ_t , search efficiency s and misclassification probabilities and weights \mathbf{M} , using observations on employment stocks and total quits by size class, we can calculate Q_{kt}^* , and the cumulated sampling weights at size cutoffs $\bar{F}_t(X_k)$ from (7) (using $\mathbf{n}_t = \mathbf{M}^{-1} \mathbf{n}_t^o$). We then look for values of ρ_t , s , and \mathbf{M} that minimize the distance between both sides of (11) over the entire sample period.¹⁵

¹⁵In so doing, as mentioned before, we approximate the integral in the l.h.s. of (11) by interpolating F_t between the nodes X_k . Moreover, we add a penalty term to the criterion that we minimize in the estimation (the norm of the difference between the two sides of (11)) to avoid large values of ρ_t that would imply negative corrected net quits Q_{kt}^* at some dates, for the highest productivity class K .

5 Results

We find that no misclassification scheme can easily remedy the basic fact that the total quit rate from the smallest establishment size class in JOLTS, “1-9 employees”, is significantly lower than that from the second-largest class, “10-49 employees”. In the data, it appears that a large group of small establishments have unexpectedly (based on the job-ladder model) low rates of attrition; therefore, their size is not an accurate reflection of their productivity or desirability. The reason may be that small employers are largely of a different nature than larger one, and more likely to ‘break ranks’ and not comply with the job ladder. For example, these small establishments may be young and growing and not have joined yet their long-run size class.¹⁶ At the other end, the largest class of establishments with more than 5,000 employees has a very small sample size in JOLTS and is therefore somewhat noisy.

For both reasons, we aggregate size classes into $K = 4$ classes: 1-49, 50-249, 250-999, and at least 1,000, and calibrate the model on these four. This partition, albeit coarser, still allows for significant heterogeneity, and can be fit quite well by the job-ladder model. While we acknowledge the simple job ladder model’s inability to accurately describe quits at the lower end of the size distribution as an unambiguous failure of the model, we still argue that this model, given its parsimony, does a remarkable job of simultaneously fitting the level and cyclicity of both gross and net unemployment flows by size class.

5.1 Parameter estimates

Estimates of the various rates of separation into non-employment and of the job finding rate were already shown in Sub-sections 4.3 and 4.4, respectively. In this sub-section, we report estimates of the remaining scalar parameters, namely the relative intensity of reallocation shocks $\rho = \rho_t/\lambda_t$ and search by employed workers s , and the elements of the conversion matrix \mathbf{M} , i.e. the misclassification probabilities (π_1, \dots, π_K) and weights (m_1, \dots, m_K) . All those values are gathered in Table 2.

The misclassification probabilities in Table 2 suggest that, when misclassified, a high-productivity employer (from productivity class $K = 4$) has a large (0.679) probability of being mistaken for a smaller establishment. The misclassification weights further show that, when thus misclassified, a high-productivity employer is almost always mistaken for an establishment from size class 1 (1-49 employees). Finally, the probabilities for employers with low or intermediate productivity levels to be misclassified are estimated to equal zero. This finding is consistent with an interpretation of misclassification as arising primarily from the

¹⁶As a manifestation of a similar phenomenon in the Danish matched employer-employee dataset IDA, the wage-size relationship is monotonically increasing except at the very beginning, as very small firms pay higher wages than slightly larger ones. We thank our discussant Rasmus Lentz for pointing out this parallel.

Productivity class k	1	2	3	4
Size range	1-49	50-249	250-999	1,000 plus
π_k	0	0	0	0.679
m_k	0.915	0.086	0	0
ρ	0.015	Sample mean of $\rho\lambda_t = 0.002$		
s	0.213	Sample mean of $s\lambda_t = 0.031$		

Table 2: Parameter estimates

establishment/firm distinction, as some very large — and productive — firms are split into many small establishments, typically no larger than 50 employees.

The relative search intensity of employed workers is estimated at $s = 0.213$, a value which is in the region (although somewhat on the high side) of typical estimates based on micro data. This puts the sample mean monthly probability of receiving an outside offer to 0.031. Finally, the reallocation shock intensity is estimated to equal $\rho = 0.015$. This value may seem small when compared, for instance, to the value of s , however it still implies that the share of EE transitions that are reallocations (as opposed to ‘voluntary’ transitions), is 50% on average. This share is calculated as the sample mean of


$$\frac{\rho\lambda_t N_t(1)}{\rho\lambda_t N_t(1) + s\lambda_t \int_0^1 \bar{F}(x) dN_t(x)}.$$

The relatively large value of this share, given the relatively high odds of receiving an outside offer vs. a reallocation shock ($s:\rho$ is about 14:1), indicates that many offers are rejected by employed workers. This, in turn, is a consequence of the fact that the sampling distribution of productive types $F_t(\cdot)$ is skewed toward the lower end of its support. We now turn to the analysis of that distribution, and the corresponding EE quit patterns.

5.2 Establishment sampling weights and quit patterns

Figure 16 plots the r.h.s. of (11), namely the estimated values of $s\lambda_{t+1}\bar{F}(X_k)$, for $k = 1, \dots, 4$ (solid lines), together with the l.h.s. of (11), Q_{kt}^*/n_{kt} (dashed lines), thus offering a pictorial assessment of the job ladder’s capacity to fit the quit patterns by establishment size observed in the JOLTS sample. Figure 17 further plots the estimated sampling c.d.f. $F(X_k)$ for $k = 1, \dots, 4$ (solid lines), together with $F^{\text{JOLTS}}(X_k)$ (dashed lines), the empirical c.d.f. of job openings, directly taken from the JOLTS data, corrected for misclassification using the probabilities and weights as explained earlier in this section. The vertical dotted lines on Figure 17 indicate JOLTS re-sampling dates.

We can see in Figure 16 that our calibration ensures that the sampling distribution constructed by fitting the RPE dynamic Equation (7) to net employment flow data from



The 'Data' series are corrected for misclassification.
 Shaded areas indicate NBER contractions.
 Source: JOLTS, CPS, and authors' calculations.

Fig. 16: Rate of voluntary EE quit.




Fig. 17: The sampling distribution.




Fig. 18: Calibrated class sampling probabilities.




Fig. 19: Calibrated average class sampling weights, levels.




Fig. 20: Calibrated average class sampling weights, normalized 01/2001 = 1.

JOLTS is by and large consistent with the gross flow data on quits over the decade covered by JOLTS. Although the data exhibit a slight downward trend in the EE quit rates of the highest two productivity classes (3 and 4) which the model fails to fully capture, and the model tends to underestimate the difference between the EE quit rates of classes 3 and 4, we still conclude that the model, including its correction for the misclassification of employers into size classes, offers a remarkably good description of this data, especially considering its parsimony.¹⁷ In particular, EE quit rates, once corrected for misclassification, are indeed neatly ordered by productivity class, as predicted by the job ladder model. We stress that this outcome was not at all guaranteed *ex-ante*.

A further striking lesson from Figure 16 is that job-to-job exit rates from all but the highest productivity class declined sharply during the GR, especially at the lower end, *and* remained low thereafter. Again, our simple job ladder model captures this pattern well, albeit with a slight lag for the lowest productivity class, $k = 1$. This is one of our central findings: the GR was a time when job-to-job quit rates declined sharply, not only in the aggregate as it was already known, but especially from smaller, less productive employers. Because these are always the main source of job-to-job reallocation, we conclude that workers almost stopped climbing the job ladder during the GR, and the recovery was almost absent.

Looking more closely at the calibrated sampling distribution (Figure 17), we first see that the observed distribution of job openings, $F^{\text{JOLTS}}(\cdot)$, vastly underestimates our calibrated $F_t(\cdot)$ for all productivity classes, but more severely so at the lower end of the productivity distribution. This is (qualitatively) consistent with the findings of Davis, Faberman and Haltiwanger (2010), who report that 41.6 percent of all hires occur at establishments with zero posted job opening in the micro data underlying JOLTS, with that proportion ranging from 76.9 percent for the small JOLTS size class down to roughly 7 percent for our largest size class. Second, there is a very slight upward time trend in the sampling distribution at all cutoff points X_k .¹⁸ This is consistent with the empirical observation that the average size of US establishments has declined over recent decades (while that of the average firm has increased, so misclassification in the sense that affects our data has arguably become worse).

Next, Figure 18 shows the sampling probabilities of each productivity class (i.e. $F_{t+1}(X_k) - F_{t+1}(X_{k-1})$ for all k), while Figures 19 and 20 show average sampling weights, i.e. the sampling probabilities divided by the number of employers in each class.¹⁹ Sampling weights

¹⁷One particularly restrictive assumption is that of a constant intensity of employed search, s , across productivity classes. An extended model where workers would choose their search intensity facing a convex cost of search would have it that workers employed at higher productivity employers would search less intensely, as the returns to search are lower for those workers.

¹⁸A linear time trend is found positive and statistically significant for all k in both $F_{t+1}(X_k)$ and $F_{t+1}^{\text{JOLTS}}(X_k)$.

¹⁹Consistently with our procedure to correct for misclassification, we use the number of establishments in

on Figure 20 are normalized to one in January 2001 to harmonize scales. The non-rescaled series, shown on Figure 19, are nicely increasing with productivity class, as predicted by the job ladder model (more productive employers post more vacancies, Prop. 1).

5.3 Discussion

We now take stock of our results. Taken together, Figures 18, 19 and 20 indicate that the sampling weights and probabilities of high-productivity employers increased during both recessions, while those of low-productivity employers declined. This fact in itself is striking in the light of MPV12’s finding that recessions are times when small (or low-productivity) employers are growing relative to large ones. It also suggests that the vacancy *yield* of small employers must have increased by much more than that of large ones during recessions, a hypothesis that finds some support in the raw data (Figure 13). Perhaps even more striking is the sudden *reversal* of this pattern at the end of 2008, immediately after the Lehman Brothers episode: at that point, both the sampling probability and the sampling weight of the high-productivity class collapse, while those of the lowest-productivity class soar, in relative terms. This, combined with a very low baseline job finding rate λ_t (Figure 15) suggests that high-productivity firms essentially stopped hiring at that point, and that what little hiring took place did so at the lower-productivity end of the population of employers. This is indeed what we observe in Figure 21 when plotting JOLTS hire rates by employer size after reclassifications. Even more than in the raw data (Figure 5), hire rates rise sharply and temporarily at the lower end of the size distribution, while upgrading to better jobs slow down considerably, as evidenced by the durably low EE quit rates that ensued (Figure 16). In short, the job ladder failed, especially the upper rungs.

To better understand the sources of this shift in relative hiring effort, in Figure 22 we plot the time series of layoff rates for each of our four size classes after reclassification.

As is well-known, at the aggregate level, layoff rates spiked considerably during the GR, although they were back to normal by late 2009. Therefore, layoffs significantly contributed to the increase in unemployment during the GR, but the persistence of high unemployment in the four years after the end of the GR is entirely accounted for by the failure of job finding rates to recover and the persistent increase in unemployment duration. Less well-known, and in fact novel, is the evidence that we report on layoff rates by size (and, after correction for misclassification, by productivity) class. Even more so than in the raw data (Fig. 6), after reclassifying establishments into productivity classes so as to fit the job ladder model, in Fig. 22 the spike in layoff rates is much sharper among low-productivity employers. The

each size class in CEW, corrected for misclassification using the conversion matrix \mathbf{M} , as our measure of the number of employers in each size class.




Fig. 21: Hire rates by establishment productivity class, corrected for misclassification.




Fig. 22: Layoff rates by establishment productivity class, corrected for misclassification.

contemporaneous shift in vacancy rates towards the bottom of the size distribution that we documented earlier suggests that the employers that were least affected by the GR, especially after September 2008, took advantage of rising unemployment to hire; because in the job ladder model small, low-productivity employers are more dependent on the reservoir of unemployed, they responded more, i.e. cut their vacancies by less. In addition, recall that the job ladder has a hard time fitting the raw data at the very low end of the size distribution, as quit rates from very small establishments are low relative to those in the two subsequent size classes. This observation suggests very significant heterogeneity among small establishments. Some are small because unproductive. Others are temporarily small but very productive and attractive because still growing. Indeed, Fort, Haltiwanger, Jarmin, and Miranda (2012) draw a sharp distinction between the cyclical dynamics of net employment growth at young and old small firms in Census data, that break down net employment flows by age and size, but lacks information on gross workers flows. So it appears that the small class as a whole shed much more employment by actively laying workers off, but also hired more by taking advantage of high unemployment and the dynamism of young employers.

To summarize: during the GR all employers temporarily raised their layoff rates, experienced slower attrition, and reduced their vacancy postings and hire rates; small employers laid off more and simultaneously reduced less their hiring effort, even hired more, but also experienced more of the decline in job-to-job quits, because hires at the top almost vanished; the job ladder slowed down at the bottom and almost stopped at the top.

We can briefly speculate on the reasons behind these events. One distinguishing feature of the GR, relative to previous recessions, was the credit crunch in late 2008 and early 2009, which may have impacted especially small employers. Our evidence is consistent with a credit crunch that affected more existing businesses, particularly the older and less productive ones, than new entrants and young growing but still small businesses. After the financial crisis exploded in Fall 2008, businesses, especially small ones, had a hard time finding and revolving working capital to pay their workers, so they had to actively reduce their workforce through layoffs, in part because attrition through quits to other employers and nonemployment collapsed. The contemporaneous reduction in vacancy postings that affected disproportionately large employers does not support more traditional theories of credit constraints, where firms, especially small ones, have a hard time securing new financing to invest and create new jobs.

6 Conclusions

We study employment reallocation, both through unemployment and directly from job to job, across employers of different productivities. We focus on the US economy around the Great Recession. In order to impose structure on our empirical investigation, we formulate a dynamic job-ladder model, where more productive employers spend more hiring effort and, conditional on contacting another worker, are more likely to succeed in hiring because they offer more. As a consequence, an employer's size is a relevant proxy for productivity. We use newly available monthly time series from JOLTS on employment net and gross flows by size of the establishment. We find that a parsimonious job ladder model fits the facts quite well, and implies 'true' vacancy postings by size that are more in line with gross flows and intuition than JOLTS' actual measures of vacancies, previously criticized by other authors. Our main finding is that the job ladder stopped working in the GR and never fully resumed. Job-to-job quits, especially from the bottom of the size/productivity distribution, collapsed, further reducing the incentives of small employers to post vacancies and hire unemployed workers to offset attrition. We speculate on the possible causes of this frozen reallocation, and conclude that the employment recovery in the US still has much mileage to cover, not only quantitatively but also qualitatively.

References

- [1] Burdett, K. and D. T. Mortensen (1998) "Wage Differentials, Employer Size and Unemployment", *International Economic Review*, 39, 257-73. [BM]
- [2] Bontemps, C., J.-M. Robin and G. J. Van den Berg (2000) "Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation", *International Economic Review*, 41(2), 305-358.
- [3] Brown, C. and J. Medoff (1989) "The Employer Size-Wage Effect", *Journal of Political Economy*, 97(5), 1027-1059.
- [4] Coles, M. and D. Mortensen (2011) "Equilibrium Wage and Employment Dynamics in a Model of Wage Posting without Commitment" NBER WP 17284.
- [5] Davis, S. J., R. J. Faberman and J. C. Haltiwanger (2010) "The Establishment-level Behavior of Vacancies and Hiring", NBER Working Paper 16265.

- [6] Fallick, B. and C. A. Fleischman (2004) “Employer-to-employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows,” Finance and Economics Discussion Series 2004-34, Board of Governors of the Federal Reserve System (U.S.)
- [7] Fort, Teresa, John Haltiwanger, Ron S. Jarmin, and Javier Miranda (2012), “How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size.” Unpublished manuscript.
- [8] Moscarini, G. and F. Postel-Vinay (2009) “The Timing of Labor Market Expansions: New Facts and a New Hypothesis”, *in* D. Acemoglu, K. Rogoff and M. Woodford, eds., *NBER Macroeconomics Annual*, 23, 1-51. [MPV09]
- [9] Moscarini, G. and F. Postel-Vinay (2012) “The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment”, *American Economic Review*, October, 102(6), 2509-2539. [MPV12]
- [10] Moscarini, G. and F. Postel-Vinay (2013) “Stochastic Search Equilibrium”. Forthcoming *Review of Economic Studies*. [MPV13]