

Did the Job Ladder Fail after the Great Recession?

Giuseppe Moscarini, *Yale University and National Bureau of Economic Research*

Fabien Postel-Vinay, *University College London and Sciences Po, Paris*

We study employment reallocation across employers through the lens of a dynamic job ladder model. Workers always agree on a ranking of employers at all points in time and search for better jobs both off and on the job. A parsimonious version of the model fits well the time series of gross worker flows by employer size from newly available US data from the Job Openings and Labor Turnover Survey. Focusing on the US experience in and around the Great Recession, our evidence indicates that the job ladder stopped working then and has not fully resumed yet.

I. Introduction

The persistence of high unemployment in the United States and in many other countries after the 2007–9 Great Recession is currently the central issue for macroeconomic policy around the world. In previous work (Moscarini and Postel-Vinay 2009, 2012, 2013), we document empirically and formulate a hypothesis to explain the pattern of employment decline and recovery during and after a typical recession. In a nutshell, in a

This is a contribution to the National Bureau of Economic Research's project titled "The Labor Market in the Aftermath of the Great Recession." We thank, for very useful comments and constructive suggestions, the editor, Alexander Mas; our discussant and referee, Rasmus Lentz; and participants of the project's conference at the NBER (Cambridge, Massachusetts, May 16–17, 2013) and numerous

[*Journal of Labor Economics*, 2016, vol. 34, no. 1, pt. 2]

© 2015 by The University of Chicago. All rights reserved. 0734-306X/2016/34S1-0009\$10.00

Submitted September 4, 2013; Accepted February 15, 2015; Electronically published December 16, 2015

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, rising 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.

We call this hypothesis a “dynamic job ladder.” The idea of a stationary job ladder, a uniform ranking of jobs by all workers, who climb it slowly via job-to-job quits while occasionally falling off it, is well established in the literature. Our previous work introduced a business cycle dimension to this hypothesis on worker turnover. In this paper, we confront this hypothesis with more demanding empirical tests. We still adopt employer size as an empirical measure of the job ladder “rung” based on the simple fact that employers higher up in a ladder tend to be larger, as they attract and retain more employment, and also based on the observed wage/size relationship. We go beyond the net worker flows by size that we studied in our previous work and here consider also the model’s implications for gross worker flows (hires, quits, and layoffs) and vacancy postings by employer size. These time series have been recently made available by the Bureau of Labor Statistics’ Job Openings and Labor Turnover Survey (JOLTS) program. Specifically, we calibrate the key turnover equations implied by a generic dynamic job ladder model to fit the monthly time series of net and gross employment flows by employer size. We extend our investigation to examine the Great Recession 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 rank of each employer in the ladder, does a remarkable job at fitting the dynamics of employment across size classes. The estimated hiring intensity by employer size resembles vacancies

seminars and conference presentations. Special thanks to Charlotte Oslund at the Bureau of Labor Statistics (BLS) for extensive assistance with JOLTS data. Moscarini also thanks the National Science Foundation for support to his research under grant SES 1123021. The usual disclaimer applies. Information concerning access to the data used in this article is available as supplementary material online. Contact the corresponding author, Giuseppe Moscarini, at giuseppe.moscarini@yale.edu.

by establishment size measured in JOLTS, but it 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 movements up the job ladder have slowed down considerably since the Great Recession. 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 to a typical recession, mostly through an increase in their layoffs, which were only partially compensated for by resilient vacancy posting and hiring.

Further support to our dynamic job ladder hypothesis has been recently offered by Kahn and McEntarfer (2014), who exploit the matched employer-employee microdata from the Longitudinal Employer-Household Dynamics at the US Census Bureau to isolate the firm component of wages and to track worker turnover over the period 1998–2011 at quarterly frequency. They find that high-paying firms grew faster during the aggregate expansion of the 2000s and shrank more quickly in the 2001 and 2008 busts. Low-paying firms were less sensitive to the aggregate unemployment rate. Furthermore, this pattern was due entirely to reduced separations to other jobs during recessions: while low-paying firms cut hiring more, their separations to other firms declined even more than at high-paying employers.

We now provide details on our contributions. From an aggregate labor market perspective, the Great Recession was no exception: job openings went down across the board, job finding rates plummeted, and layoff rates temporarily spiked, especially around the fall of 2008 when the financial crisis erupted. As a result, unemployment soared. As we argued and documented in our previous work, which covered the four previous recessions, this pattern created relatively favorable conditions for small, low-paying, less productive employers. High unemployment meant that there was plenty of cheap labor for them to hire. Vacancy yields soared as an army of the newly unemployed lined up for few available jobs. The collapse in aggregate job market tightness reduced not only the workers' exit rate from unemployment, as is well understood, but also the job-to-job quit rate.

Evidence on job openings and gross worker flows from JOLTS, the monthly Current Population Survey (CPS), and the Survey of Income and Program Participation (SIPP) largely corroborates this view. Job-to-job transitions indeed went down markedly during the Great Recession. The "poaching intensity" (share of new hires that originate from a job-to-job transition) declined sharply during and after the Great Recession, especially so for larger employers. Finally, while the share of small establishments in total job openings remained roughly stable throughout the Great Recession (if anything, it went up a little), the vacancy yield of small

employers sky-rocketed, in sharp contrast to the comparatively modest (and vanishing) increase in the vacancy yield of large establishments.

Yet—and this is where the Great Recession differs from previous recessions—small employers fared worse than large ones in terms of net employment growth. This unusually poor job creation performance 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.

These findings suggest the following interpretation of the Great Recession and of its aftermath. Small employers, especially existing ones, faced an unusual credit crunch that led to a wave of layoffs. To contrast this effect, the sharp increase in unemployment and relaxed hiring constraints kept small employers hiring at a relatively healthy pace. The collapse in hiring was concentrated among large employers and led to a deep freeze in job-to-job upgrading and attrition up the job ladder, taming the incentives of small employers to post vacancies and hire unemployed workers.

In Section II, we present descriptive evidence on labor market flows across employers of different sizes before, during, and after the Great Recession. In Section III, we present the turnover equations describing the business cycle dynamics of gross and net workers flows in a dynamic job ladder model. We also briefly discuss structural equilibrium foundations for this process and relate it to the descriptive evidence. In Section IV, we describe our methodology to calibrate turnover parameters and hiring intensity by firm size in the dynamic job ladder model so that it replicates the observed net and gross flows of employment by firm size. In Section V we discuss our empirical results. Section VI presents our conclusions.

II. The Dynamic Job Ladder: Descriptive Evidence

We examine the cyclical reallocation of employment among firms and establishments, especially around the Great Recession, through the lens of the job ladder, namely, the turnover process that occurs when all workers agree on a ranking of employers and face frictions in finding and retaining jobs. We begin with descriptive empirical evidence. In order to make the notion of a job ladder empirically operational, we need a measure of a ladder's "rung." As workers move up the ladder, employers high on the ladder tend to accumulate more employment, thus to be larger. We focus on an employer's size as the main empirical counterpart of its position on the job ladder because size is accurately and easily measurable in the data, unlike other natural candidates such as productivity or compensation

policy.¹ We present new evidence on the cyclicity of four relevant types of aggregate labor market statistics, all broken down by employer size: employment shares, net job creation, gross job flows (hires, quits, and layoffs), and vacancy postings.

Before we proceed, an important caveat: we emphasize that in our analysis we focus on continuing employers and abstract from entry and exit. The reason for this choice is threefold. First, and foremost, given our focus on cyclical employment variation, entry and exit play a relatively minor role. While they are extremely important to determine average job and worker flows,² their contribution to cyclical movements in aggregate employment is positive but modest: in the Business Employment Dynamics (BED),³ the standard deviation of the net job creation rate, passed through a Hodrick-Prescott filter with smoothing parameter 1,600, equals .48 for the whole economy and is just slightly lower, .435, for continuing establishments, which exclude openings and closings.⁴ Second, the prime novel data set that we employ in this paper, JOLTS by establishment size, is a survey of preexisting establishments, where exit is by and large offset by a monthly sample rotation/refreshment scheme, while entry does not contribute to the observations. Third, the equations describing workers' movement on a dynamic job ladder that we use for our calibration exercise are much simpler when ignoring entry and exit, although both of them could be accommodated in a limited sense.

To begin, we motivate our hypothesis that size is one relevant (albeit by no means the only possible) empirical counterpart of a job ladder rung. In an appendix, available online, we provide corroborating empirical evidence, drawing from the Quarterly Census of Employment and Wages (QCEW) for establishments and from Statistics of US Businesses (SUSB), an annual census of all employers, for firms.⁵ First, it has long been doc-

¹ In the United States, information on sales at the firm level, necessary to compute total factor productivity (TFP), is not available for a representative sample of firms from all industries.

² Haltiwanger, Jarmin, and Miranda (2013) document from an annual longitudinal census of US employers that in fact entrants create on average more jobs than the whole economy, as continuing establishments and exits on net destroy jobs.

³ The BED is a collection of quarterly gross job flows published by the Bureau of Labor Statistics, currently covering 1992Q3–2013Q1 and nearly the entire US private sector.

⁴ This way of measuring the contribution of entry is based on the extreme view that, one quarter after entry, new establishments are similar to incumbent ones of the same size. More generally, entrants may face a different growth process than incumbents early in their life cycle; in this case, a cyclical decline in entry may have long-run effects on aggregate job creation that are significantly larger than the small immediate impact that we document here.

⁵ The QCEW is the primary source of information on businesses from the Bureau of Labor Statistics. It publishes a quarterly count of employment and wages reported by establishment size, covering 98% of US jobs, both private and public sector, available at the county, metropolitan statistical area, state, and national levels by industry.

umented that employer size correlates positively with wage rates, after controlling for observable worker characteristics (Brown and Medoff 1989). We confirm that larger employers pay more. Second, as predicted by the dynamic job ladder hypothesis, the share of employment at large firms or establishments is procyclical: workers climb the job ladder faster in tight labor markets, when they can make contact with employers at a higher rate.

A. Worker Flows by Establishment Size

Our main focus in this paper is on business cycles and the resulting dynamic job ladder. In order to measure worker flows by employer size, we need at least a modicum of longitudinal links on firms/establishments and workers. 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 1, 10, 50, 250, 1,000, and 5,000 employees. This data set is central to our exercise.

In JOLTS, 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, independently of how its size changes while it is part of the sample.⁶ So, within each survey year, we know that the identity, hence the size quantiles of establishments in each size class, are fixed.⁷ 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 QCEW sample described above is 1,000 employees. As we get our shares of private sector establishment counts from QCEW, we will need to merge information from QCEW 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 covers few establishments), and the data pertaining to this category are somewhat noisy. The loss of

⁶ JOLTS (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 resampling.

⁷ Because this size classification follows an “initial employment” criterion, it is known to be subject to a mean reversion bias, creating the illusion of a negative size-growth relationship in the presence of a transitory component to firm size. This issue is likely to matter more in narrower size classes, at the bottom of the size distribution, where establishment size is more volatile. We will return to this issue when discussing size misclassification.

information implied by our aggregation of the largest two size classes into one is therefore arguably relatively minor.

Finally, and importantly, we should mention that 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 Great Recession, where the public sector played a disproportionate role in first buffering employment losses and then dragging on the employment recovery.

1. *Net Flows*

The cyclicity of employment shares of the various size classes of employers in our data, presented in the online appendix, provides limited information on the size of businesses that were most affected by the Great Recession. As we discussed in Moscarini and Postel-Vinay (2012), to avoid the so-called reclassification bias, we need to study business dynamics for at least two consecutive periods among classes to which employers are assigned based on 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 United States is strongly negatively correlated with unemployment in 1979–2010. Here we zoom in on the Great Recession using higher-frequency monthly data updated to cover the post–Great Recession recovery.⁸ Figure 1 repeats the exercise using JOLTS data by size of the establishment (this is an important distinction to which we will return later). The differential net job creation series in figure 1 follows a similar pattern as in previous recessions, but in the Great Recession it peaks later, in fact at the very end of the recession, than one would have expected based on the evidence reported in Moscarini and Postel-Vinay (2012) for previous recessions. It thus appears that small establishments were hit especially hard by the credit crunch.

2. *Gross Flows*

To examine in more detail the nature of these evolutions, we turn to 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 Great Recession. 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 2 plots hire rates (new accessions divided by employment) by establishment size.

⁸ All 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. 1.—Differential employment growth between establishment size classes (solid line) and unemployment rate (dashed line).

Hire rates began to decline before the Great Recession. Surprisingly, during the deepest phase of the financial crisis, following the Lehmann Brothers episode, hire rates collapsed at the larger establishments and not at the smaller ones; they even briefly spiked in the smaller class in late 2008




FIG. 2.—Hire rates by establishment size class




FIG. 3.—Layoff rates by establishment size class

to early 2009. Given that, in figure 1, smaller establishments fared worse 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 pace. We in fact see in figure 3 that layoff rates rose sharply and temporarily, especially at small establishments. Although not immediately evident from the figure, the increase in layoff rates was almost exactly proportional across all size classes. Because smaller establishments report higher layoff rates on average, the absolute increase in layoff rates during the Great Recession was more pronounced at the bottom of the size distribution.

The third gross worker flow available in JOLTS, the quit rate, is shown in figure 4. This flow conflates quits to nonemployment and quits to other employers. While quit rates fell markedly across the board both in 2001 and around 2008, the figure clearly suggests that the fall during the Great Recession was less sharp for small establishments than for large ones. This fact corroborates the hypothesis that the comparatively worse performance of small establishments during the Great Recession was entirely driven by a spike in layoff rates, as opposed to higher total quits or reduced hiring, which actually worked in the opposite direction.

JOLTS, as a survey of employers, provides a meaningful distinction between layoffs and quits, but not between quits to (or hires from) nonemployment, as opposed to (from) other jobs, a distinction that is crucial to the job ladder. We supplement JOLTS with information on gross worker




FIG. 4.—Quit rates by establishment size class

flows from the monthly CPS. Specifically, we use the hazard rates of transition between Employment (E), Unemployment (U), and Nonparticipation (N) estimated by Fallick and Fleischman (2004) from gross flows (using monthly matched files), starting in January 1994 and updated by the authors through May 2014. This series begins with the 1994 redesign of the CPS, which introduced a question on the change of employer that made it possible to measure the hazard and which greatly improved the reliability of employment status and thus reduced margin error. Figure 5 plots the transition rate, or hazard. While it is clearly procyclical and dropped significantly during the Great Recession, the most striking aspect is the declining trend. Off that trend, the decline during the Great Recession was not especially pronounced, and the recovery afterward was significant. But in absolute terms, that is, without detrending, the EE hazard remains at an all-time low almost 5 years into the post-Great Recession recovery. It is well known that EE transitions include involuntary reallocation and other events that reduce worker's earnings (our model explicitly accommodates this possibility through reallocation shocks—see Sec. III). Therefore, per se they provide only limited information on the extent to which workers climb the job ladder. It is, however, striking that the EE rate is the most lagging labor market indicator post-Great Recession.

The CPS contains no information on the size of a worker's employer. For this, we turn to the SIPP, starting with the 1996 panel. We exploit the




FIG. 5.—Total employment-to-employment hazard

availability of the start date and the end date of each job to construct EE transition rates by size of the hiring “workplace,” the phrasing in the SIPP questionnaire that we interpret to be an establishment. In figure 6, we show the share of all hires that originate directly from other employers, thus entail an EE transition, broken down by size of the hiring establishment. For readability, we aggregate size classes into two groups, separated by the 100-employee cutoff. The interpretation of the cyclical patterns is complicated by gaps in the time series, due both to genuine gaps between SIPP panels and to our decision to discard the first 4 months and the last 3 months in each panel, when no rotation group has a “seam” between interviews, leading to an underestimate of turnover rates. As predicted by the job ladder model, larger establishments appear to always hire more from other employers and less from nonemployment, especially late in expansions, when the market tightens and competition for workers grows stiffer. In the Great Recession, this “poaching” inflow-share collapsed for both size groups. Since total hires also declined sharply, this is the strongest evidence that the climbing of the job ladder came to a grinding halt.

To take stock, we showed that net job creation by small establishments was especially poor during the Great Recession, relative to larger establishments and to a typical US recession, and that this is due entirely to a spike in their layoffs, while hires and total quits declined much less at the

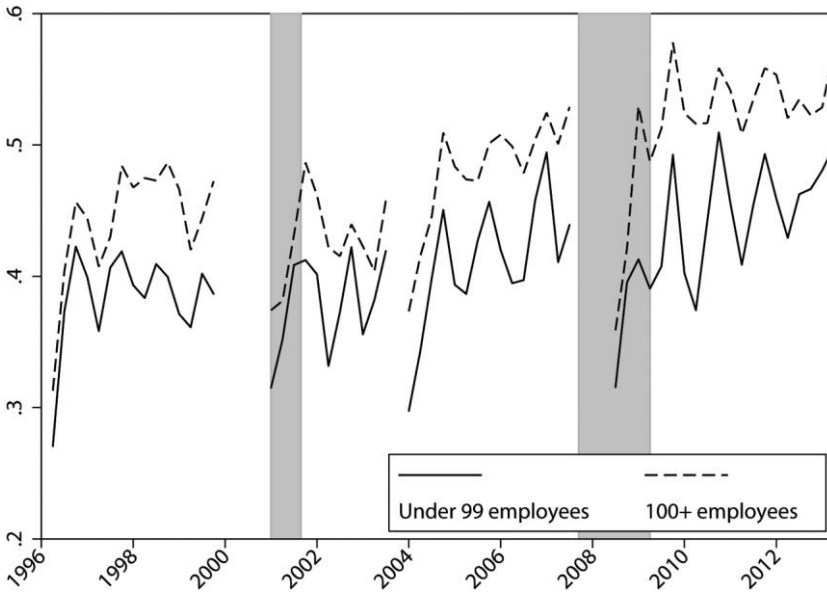


FIG. 6.—Share of hires from other employers, by employer size

bottom of the size distribution. Aggregate job-to-job transitions collapsed, even more so toward larger establishments, and never recovered.

B. Vacancies by Establishment Size

We now return to JOLTS to describe the behavior of measured job vacancies by size class. Vacancies are uniquely valuable as a direct measure of labor demand, or intensity of hiring effort, as opposed to outcomes. Figure 7 reports the time series of total job openings for each JOLTS size class. Figure 8 further shows vacancy shares by size class, that is, vacancies in each size class divided by total aggregate vacancies. If recorded job openings are an accurate measure of hiring effort,⁹ then the series plotted in figure 8 will represent the sampling probabilities of each size class. Next, figure 9 shows vacancy shares divided by the number of establishments in each class from QCEW, normalized at one in January 2001 to harmonize scales. We refer to those series as the *vacancy weights* by size class. These weights measure average hiring effort per establishment in a given size class, relative to aggregate hiring effort.

Figure 7 clearly shows that vacancies plummeted across the board during the Great Recession, with vacancy levels seemingly tracking each other across the various size classes. At first glance, figures 8 and 9 reinforce that

⁹ There are good reasons to believe that they are not, as we discuss below in Sec. IV.




FIG. 7.—Vacancies by establishment size class




FIG. 8.—Vacancy shares by establishment size class

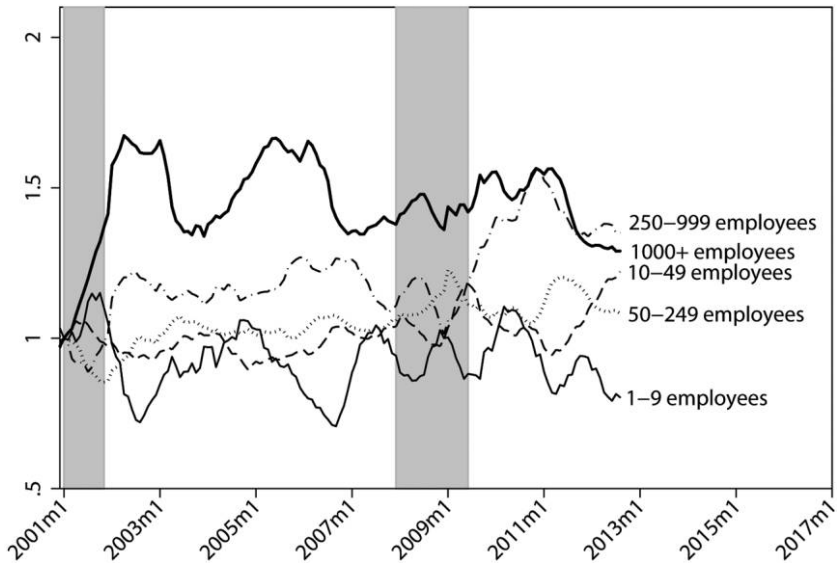


FIG. 9.—Vacancy weights by establishment size class

impression, as the movements in vacancy shares and weights appear small relative to the absolute decline seen in figure 7, which, to a first approximation, was uniform. On closer inspection, figures 8 and 9 further suggest that there is no evidence of a disproportionate impact of the financial crisis (post-September 2008) on the hiring effort of small establishments: the movements are relatively modest, and the 10–49 employee class shows the largest change, but upward. Overall, we conclude that hiring effort fell proportionally at establishments of all sizes.

Finally, figure 10 plots the vacancy yield, namely, the ratio between hires and vacancies reported a month before, by establishment size.¹⁰ Vacancy yields are countercyclical; specifically, during and after the Great Recession, the aggregate yield rose enormously with unemployment duration, and it became as easy for firms to fill vacancies as it was difficult for the unemployed to find work. Importantly, figure 10 shows that this

¹⁰ Note that the yield is greater than one for many dates and size classes (fig. 10), suggesting that the JOLTS measure of job openings misses something about true establishment hiring effort. This ties in with the results of Davis, Faberman, and Haltiwanger (2013), who report that around 40% of hires occur at establishments that do not report any job openings to JOLTS. We return to this issue below in Sec. V.




FIG. 10.—Vacancy yields by establishment size class

phenomenon was more pronounced the smaller the establishment. During the acute phase of the Great Recession, from the fall of 2008 onward, the vacancy yield literally took off at establishments employing 1–9 workers. At the largest establishments, however, the yield stopped rising. This surprising set of facts is consistent with the collapse in hires of employed workers, on which larger establishments rely more, but it can also be explained by tightening hiring standards by those large employers.

III. The Dynamic Job Ladder: Model

A. Flow Equations

In order to interpret the evidence laid out in Section II, we now propose a turnover accounting framework. This is a reduced-form model of employment dynamics, a set of equilibrium predictions shared by several models of the labor market with on-the-job search. Time $t = 0, 1, 2 \dots$ is discrete. The labor market is populated by a unit measure of workers, who can be either employed or unemployed, and by a unit measure of firms. Workers agree on a ranking of employers, which gives rise to a job ladder. Let $x \in [0, 1]$ be the rank of a firm in the job ladder: workers always prefer firms with higher x . The labor market is affected by search frictions in that unemployed workers can only sample job offers sequentially with probability $\lambda_t \in (0, 1)$ at time t . Employed workers draw each period with probability $s \in (0, 1]$ an independent and identically distributed (i.i.d.) op-

portunity to search on the job; thus they face a per-period sampling chance of job offers of $s\lambda_t$. Workers can only send one job application per period and can never receive more than one offer in any period. Conditional on a contact, workers draw offers from a sampling distribution with cumulative distribution function $F_t(\cdot)$, so $F_t(x)$ is the chance that the worker meets an employer of rank below x . An employed worker is exogenously separated from his employer and either, with probability $\delta_t(x)$, enters unemployment, or, with probability ρ_t , is immediately reallocated to another job, drawn randomly from the available ones according to $F_t(\cdot)$, without going through unemployment. The displacement shock $\delta_t(x)$ encompasses both layoffs and quits to nonemployment 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 rehiring by another employer. The objects that govern worker turnover, $F_t(\cdot)$, $\delta_t(\cdot)$, λ_t , and ρ_t are realizations of stochastic processes. We are particularly interested in their business cycle fluctuations.

Let $N_t(\cdot)$ denote the cumulative distribution function of employment over ranks at time t . So $N_0(x)$ is the date-0 measure of employment at firms of rank weakly below x , a given initial condition; $N_t(x)$ is the same measure at time t ; $N_t(1)$ is total employment; and $u_t = 1 - N_t(1)$ is the unemployment stock (or rate). Let

$$\bar{\delta}_t(x) = \frac{1}{N_{t-1}(x)} \int_0^x \delta_t(q) dN_{t-1}(q)$$

denote the average transition rate into unemployment by workers currently at employers of rank up to x . Applying a Law of Large Numbers to each firm rank and the definition of rank in a job ladder, we obtain equations for net and gross workers flows. We present the equations in terms of employment cumulated over ranks, N_t . Taking derivatives with respect to rank x would provide the equivalent equations at each x (for each employer) in the job ladder.

We start with gross flows, the inflow into (outflow from) unemployment from (into) employers of rank below x :¹¹

$$\begin{aligned} &\text{Employment to Nonemployment Flow:} \\ &EU_{t+1}(x) = \bar{\delta}_{t+1}(x)N_t(x). \end{aligned} \tag{1}$$

¹¹ In the notation just laid out, we use the letter U to imply nonemployment. The model is silent on any possible distinction between unemployment and nonemployment. We will return to this issue momentarily.

Nonemployment to Employment Flow:

$$UE_{t+1}(x) = \lambda_{t+1}F_{t+1}(x)[1 - N_t(1)]. \tag{2}$$

In (1), the chance of exogenous separation $\bar{\delta}_{t+1}(x)$ into unemployment multiplies the measure of employed workers. In (2), the chance of job contact times the chance that the contact is with a firm of rank below x multiplies the measure of unemployed job searchers.

The third gross flow comprises workers who leave employers of rank below x to join another employer of any rank. In turn, this flow includes forced reallocations with chance ρ_{t+1} and voluntary quits:

Employment to Employment Flow (Quits):

$$QE_{t+1}(x) = \rho_{t+1}N_t(x) + s\lambda_{t+1} \int_0^x \bar{F}_{t+1}(x')dN_t(x'). \tag{3}$$

To understand the integral term, note that a worker employed at rank $x' < x$ receives each period with chance $s\lambda_{t+1}$ an outside offer, which is above rank x' (so the worker accepts) with chance $\bar{F}_{t+1}(x') = 1 - F_{t+1}(x')$. A measure $dN_t(x')$ of workers were initially employed at rank $x' < x$. Here QE is a gross outflow; some of these workers join other employers whose rank is still below x , in some cases even below their current job's rank if the reallocation is forced.

The last gross flow is the inflow from other employers into firms of rank at most x . By an accounting identity, given the three gross flows above, this fourth one gives rise to net job creation by such firms. Since the net flow is easier to measure empirically, we focus on the latter, so the fourth gross flow is redundant. The net change in employment at firms of rank up to x evolves as follows:

$$N_{t+1}(x) - N_t(x) = -[\bar{\delta}_{t+1}(x) + \rho_{t+1} + s\lambda_{t+1}\bar{F}_{t+1}(x)]N_t(x) + \{\rho_{t+1}N_t(1) + \lambda_{t+1}[1 - N_t(1)]\}F_{t+1}(x). \tag{4}$$

The first line includes outflow from firms of rank below x due to either exogenous turnover, to unemployment $\bar{\delta}_{t+1}(x)$ and other employers ρ_{t+1} , or to outside offers received from firms of rank above x . The second line includes the inflow into firms of rank below x , which are sampled with probability $F_{t+1}(x)$ either by workers who are forced-reallocated or by the unemployed. Notice that the voluntary inflow from other employer is omitted from the second line, because it can only occur from below x , so it can at best reshuffle the mass of employment below x but not increase it.

To make equations (1)–(4) empirically operational, we need a measure of job ladder rank. We do not observe the workers' preferences that define the job ladder, so we rely on their revealed preferences. Because workers climb the job ladder, from lower-ranked to higher-ranked employers, while the contact rates $s\lambda_i$ and the forced reallocation rate ρ_i are rank-independent, this turnover process makes higher-ranked firms also larger in terms of employment measure. Thus, when given the opportunity, employed workers tend to move from smaller to larger employers. Exogenous forced reallocations to unemployment and to other employers interfere with this upgrading process and maintain a nondegenerate ergodic size distribution of employers. In order to guarantee that higher rank means larger size in the model, and thus to use firm size as an empirical proxy for rank, we further assume that the inflow rate into unemployment $\delta_i(x)$ is nonincreasing in rank x . This assumption encompasses as special cases exogenous separations at flat, rank-independent probability δ_i , as well as endogenous separations due to match-specific shocks, because workers must be more reluctant to endogenously give up higher-ranked jobs if they are more willing to accept them to begin with. We can then proceed to estimate turnover rates from equations (1)–(4) using data on employment stocks, net and gross worker flows, broken down by employer size. Before doing so, we briefly discuss structural foundations of the dynamic job ladder, namely, of the accounting equations (1)–(4), and how they relate to the descriptive evidence illustrated earlier.

B. Structural Foundations

Equations (1)–(4) describe the accounting of worker flows in a job ladder, namely, in an environment where all workers agree on the ranking of employers. This type of turnover process occurs in different frictional models of the labor markets. The prime, but by no means only, example is a wage-posting model. The canonical framework for the analysis of frictional wage dispersion with on-the-job search is Burdett and Mortensen (1998). This setup has strong implications also for worker turnover and for the distribution of firm size, where a firm is identified by a wage policy constrained to pay all workers the same. In particular, the unique steady state equilibrium of the Burdett-Mortensen model features a job ladder by employer size. In Moscarini and Postel-Vinay (2009, 2013), we introduce aggregate uncertainty in the Burdett-Mortensen model and accordingly identify a firm as a wage policy, which may now depend on the state of the aggregate economy, the size of the firm, and the distribution of wage offers by competitors.¹² In the ergodic steady state of the stochastic economy, the unique equilibrium is always rank-preserving. That is, a firm that

¹² This structural model does not, but can easily be extended to, include reallocation shocks with chance ρ_i .

is larger, and possibly permanently more productive, will always commit to a stream of payments of higher value to workers, who then move on a dynamic job ladder, from smaller, lower-paying to larger, high-paying firms, at all points of the business cycle. Because larger firms pay more and are ranked higher by workers, equilibrium preserves a stable ranking by size, although not necessarily a stable size distribution, for any history of aggregate shocks. In this model, firm-level productivity is a natural, although by no means the only, primitive that determines the rank on the ladder.

Coles and Mortensen (2011) introduce idiosyncratic shocks to firm productivity in a model that is very close to Moscarini and Postel-Vinay (2013)'s wage-posting framework and show the existence of a rank-preserving equilibrium (RPE). In other business cycle models of frictional labor markets with on-the-job search, workers agree in equilibrium on the ranking of jobs (matches) at each point in time. The allocation of jobs to employers is somewhat indeterminate, but it can be chosen to generate a dynamic job ladder and size distribution. Robin (2011) introduces aggregate uncertainty in Postel-Vinay and Robin (2002)'s sequential auction model of the labor market, where firms commit to wage offers but can respond to outside offers to their employees. These models feature random matching. Menzio and Shi (2011) obtain a job ladder by wage with aggregate shocks in a directed-search framework.

C. Revisiting the Descriptive Evidence

These structural models naturally dovetail with the stylized facts illustrated in the previous section. Wages are increasing in employer size, with causality running primarily from the former to the latter (paying workers more attracts and retains more of them) but also in the opposite direction. For example, in Moscarini and Postel-Vinay (2013), a larger firm, under the equal-treatment constraint, is willing to pay its new hires more than a smaller firm would, in order to pay more and retain its larger existing labor force. A procyclical job contact rate λ , and weakly countercyclical separation rate into unemployment $\delta_i(\cdot)$ then imply that workers climb the job ladder faster, and fall off the job/size ladder less often, in expansions, and vice versa in recessions. Hence, both the extra net job creation and the employment share of larger employers, those that are located higher on the ladder, are procyclical. Employer-to-employer transitions are directed up the size ladder. Job ladder models are mostly silent on separations into unemployment, which are assumed exogenous. The cyclicity of vacancy postings and hires by size are more difficult to discern qualitatively, and they require estimating the model, which is the objective of the next section.

An important role in our analysis is played by reallocation shocks, which move workers directly from employer to employer without any

measurable unemployment spell. These shocks are meant to capture in the data the sizable flows of workers who move in opposite directions among employers of different sizes, a phenomenon that is inconsistent with the idea of a job ladder in its most extreme form. One restriction imposed by equation (4) is that of a rank-independent chance ρ_i of reallocation shocks. Since employed workers voluntarily quit to accept an outside offer with probability that decreases in the rank of their current employer, they all move from job to job in the same direction (up, toward larger employers) on average, although not with probability one. This is a key prediction that we will test. Another restriction is the rank-independent relative efficiency of employed and unemployed job search, s . This can be interpreted as a time endowment available to all employed workers, no matter where currently employed, to search and interview for other jobs. An alternative interpretation, which would not be consistent with our assumptions, is that workers control their job search effort, in which case we should expect s to decline in rank x , as lower-ranked jobs, starting with unemployment at the bottom of the ladder, are less desirable and motivate more search effort. By assuming a constant s , we attribute all time-series variation in job contact rates from employment to that in job market tightness and all cross-sectional variation in turnover rates among workers to their different positions on the job ladder: all workers receive offers at the same rate, but they differ in their willingness to accept them. In the next section, we investigate whether the job ladder hypothesis can be rejected, or, conversely, there exists a calibration of model objects such that the resulting job ladder is consistent with gross worker flows by employer size each month over a long time period.

IV. The Dynamic Job Ladder: Calibration

We calibrate the job ladder model using 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 assumptions implying that employer size is a relevant rung of the job ladder, it is far from obvious that the job ladder dynamic equations (1)–(4) can replicate actual observations on gross and net flows, every month for 12 years, for several size classes. 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. In addition to the parameter s (the search intensity of employed relative to unemployed workers), equations (1)–(4) involve six time series— $\bar{\delta}_i(\cdot)$, λ_i , ρ_i , $F_i(\cdot)$, $N_i(\cdot)$, and the size of the labor force, which in the model we normalized to one but which is time-varying in the data, or, equivalently, the size of employment and unemployment, given the unemployment rate $1 - N_i(1)$. We now explain how we map monthly empirical

observations into our time series of interest. While some of them can be estimated directly, we need the model to back up ρ , $F_t(\cdot)$, and s .

A. Size Ranks

Assuming for the time being that employer size is correctly measured and that size does reflect rank in the job ladder (i.e., workers always prefer larger employers when they can choose), establishments in a given JOLTS size class $k = 1, 2, \dots, K$ will be representative of all establishments with ranks between two unobserved 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 resampling dates, 1/12 of the surveyed establishments are replaced with ex ante identical establishments, which had the same size and industry at the time of sampling; under the assumption, underlying this 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 resampling times.

The JOLTS sample thus provides observations at (almost) all dates of cumulated employment $N_t(X_k)$, layoffs, and total quits (and, potentially, sampling probabilities $F_t(X_k)$ —see below) for K job ladder rank quantiles $\{X_k\}_{k=1}^K$ corresponding to K size classes.¹³ In what follows, we should keep in mind that X_k is the cutoff 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 (1)–(3) with the JOLTS sample.

B. Separations into Nonemployment

As discussed earlier, a survey of employers like JOLTS reveals whether a separation is a quit or a layoff from the viewpoint of the surveyed establishment. As workers are neither interviewed nor tracked after a separation, measured quits are the sum of quits to unemployment and quits to other jobs, a distinction that is missing in the data but 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.

Focusing first on the aggregate separation rate (up to rank $x = 1$), we seek to construct $\bar{\delta}_{t+1}(1)$ based on equation (1) as the ratio between the total monthly flow from employment to nonemployment and the total

¹³ As discussed earlier, the raw JOLTS sample has six establishment size classes: 1–9, 10–49, 50–249, 250–999, 1,000–4,999, and over 5,000 employees. For reasons discussed earlier, we lump the largest two classes into one.

stock of employment. The flow consists of layoffs plus quits into nonemployment. We supplement the JOLTS data with the transition rates estimated from the CPS by Fallick and Fleischman (2004), updated by the authors through 2014. For every month t , we compute the share σ_t^{CPS} of total transitions that are employer-to-employer (EE), as opposed to transitions into nonemployment (say, EU):

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

All EE transitions are quits in the job ladder model; some are voluntary upgrades, and others are forced reallocations. Assuming that the CPS-based share σ_t^{CPS} applies to the workers employed by the JOLTS sample of establishments, we multiply total separations in JOLTS by $1 - \sigma_t^{CPS}$ to obtain an estimate of aggregate separations into nonemployment, $EU_i(1)$, that is consistent with the JOLTS data. The corresponding aggregate separation rate is then $\bar{\delta}_{t+1}(1) = EU_t(1)/N_t(1)$.¹⁴

This procedure further gives us the share of all 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 nonemployment in JOLTS. Subtracting the latter from total quits, we obtain a JOLTS-based measure of quits to other employers, or job-to-job outflow. We now introduce the ancillary—yet economically meaningful—parameter $\psi_t(x)$, defined as the share of total EU separations from employers of rank x that are quits to nonemployment, and

$$\bar{\psi}_t(x) = \frac{1}{\bar{\delta}_t(x)N_{t-1}(x)} \int_0^x \delta_t(x')\psi_t(x')dN_{t-1}(x'),$$

the same share from employers of rank up to x . In this notation, $\bar{\psi}_{t+1}(1)$ is the share of quits in aggregate separations into nonemployment, $EU_t(1)$, that we obtain from our procedure, the remaining share being layoffs. The aggregate layoff probability is then $\bar{\delta}_{t+1}(1)(1 - \bar{\psi}_{t+1}(1))$, and the probability of quitting into nonemployment is $\bar{\delta}_{t+1}(1)\bar{\psi}_{t+1}(1)$. Both of those, plus the total aggregate transition rate into nonemployment $\bar{\delta}_{t+1}(1)$, are plotted in figure 11. While most of this figure has the familiar feature of a largely

¹⁴ To 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 nonemployment, and quits to other employers. Worker surveys such as the CPS and the SIPP are notoriously plagued by noise in the layoff/quit distinction when the worker loses a job. Administrative data sets do not typically contain information about the reason for separation.




FIG. 11.—Average separation rate into nonemployment $\bar{\delta}_{t+1}(1)$ and its components

acyclical probability of transition into nonemployment, the Great Recession 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.

The Fallick and Fleischman (2004) series is only available at the aggregate level. Therefore, making our quit/layoff distinction operational at lower levels of size 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 nonemployment, $\psi_{t+1}(x)\delta_{t+1}(x)$, is independent of their employer’s rank x . That is to say, for all x , $\psi_{t+1}(x)\delta_{t+1}(x) \equiv \bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1)$. Since the total separation rate into nonemployment $\delta_{t+1}(x)$ is nonincreasing in (size) rank x , this assumption implies that both total separation rates and layoff rates are nonincreasing in x . Both implications hold in the JOLTS data.¹⁵ This additional identifying assumption enables us to construct total separations

¹⁵ Any 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 rank, 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 nonemployment is decreasing in employer 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.

into nonemployment from employers with rank up to X_k , namely, $\bar{\delta}_{t+1}(X_k)N_t(X_k)$, for all cutoff 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 nonemployment from those employers, equal to $\bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1)N_t(X_k)$ by assumption. Given observations on the cumulated employment distribution $N_t(X_k)$, this allows us to directly estimate the desired total probability of transition into nonemployment by size class, $\bar{\delta}_{t+1}(X_k)$.

C. Job Contact Probability

Equation (4) applied to 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 nonemployment. From this equation, we can back out the job finding rate from nonemployment, which is also the baseline job contact 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 nonemployed job seekers, U_t . Here again, we call on the Fallick and Fleischman (2004) CPS series, which offers a breakdown of the total nonemployment 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 to the former gives us an estimate of the relative job finding rate of inactive workers to the unemployed, say s_0 , so that the job finding probability of nonparticipants is $s_0\lambda_{t+1}$. We then construct the effective pool of nonemployed job seekers as

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

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 12. While it exhibits the familiar cyclicalty, including a vertiginous drop during the Great Recession, its level is fairly low because it includes transitions to employment from inactivity, which are a small fraction of the stock of inactive individuals.

D. Sampling Distribution and Employer-to-Employer Transitions

We now turn to the last, and arguably most salient, gross flow of workers predicted by the job ladder, namely, job-to-job quits $QE_t(x)$, given in equation (3). We show how this equation, combined with the net flow equation (4) and with the JOLTS data, allows identification of the




FIG. 12.—Job finding rate λ_{t+1}

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 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 et al. (2013) have recently forcefully argued that neither was true, especially at the low end of the establishment size distribution. Vacancies posted by different types of establishments may have different visibility or small establishments may rely more on informal hiring channels rather than vacancies.

Luckily, the law of motion of employment in rank-preserving equilibrium (see Sec. III.B) offers an alternative solution to estimate $F_{t+1}(\cdot)$. Equation (4) defines the sampling distribution at cutoff quantiles X_k and at all dates as

$$\bar{F}_{t+1}(X_k) = \frac{[N_{t+1}(1) - N_{t+1}(X_k)] - (1 - \rho_{t+1})[N_t(1) - N_t(X_k)] + \bar{\delta}_{t+1}(1)N_t(1) - \bar{\delta}_{t+1}(X_k)N_t(X_k)}{\rho_{t+1}N_t(1) + s\lambda_{t+1}N_t(X_k) + \lambda_{t+1}U_t}, \tag{5}$$

which we will use to estimate sampling probabilities $F_t(\cdot)$, employed search efficiency s , and reallocation shocks ρ_{t+1} , using the time series for separation and accession probabilities $\bar{\delta}_{t+1}(X_k)$ and λ_{t+1} , and the stock of nonemployment 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 (5), which is 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 (3), equal

$$QE_t(X_k) - QE_t(X_{k-1}) = \rho_{t+1}[N_t(X_k) - N_t(X_{k-1})] + s\lambda_{t+1} \int_{X_{k-1}}^{X_k} \bar{F}_{t+1}(x) dN_t(x). \tag{6}$$

The empirical counterpart is total quits in JOLTS size class k minus quits into nonemployment from employers in that size class, which were estimated in Subsection IV.B as $[\bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1)][N_t(X_k) - N_t(X_{k-1})]$. Fitting (6) to this JOLTS counterpart at each date t and size class k allows us, in principle, to identify 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 value of ρ_t for each month in the sample (140 in total) down to a single scalar, ρ . This restriction follows, for example, if ρ is the probability that the worker’s spouse is seeking a better job that would require the entire household to move, a job search that is successful with probability λ_t . Second, equation (6) is not exactly implementable, as the transformed net flow equation (5) 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 equation (6). We approximate the integral using a simple trapezoidal rule on the grid of points at which $F_t(\cdot)$ is known.

E. Misclassification

1. 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

rank of 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 resampling, 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, that is, to 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 a small establishment and also because the small-size classes (1–9 and 10–49 employees) are narrower than the larger ones, even in logs.¹⁶

The second reason to suspect that establishment size does not perfectly reflect the relevant rank in the ladder 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 rank for those establishments may be at the level of the parent firm, in which case the size measure that will best reflect rank is not the size of the establishment but that of the parent firm, which we do not observe in JOLTS. Indeed, in Moscarini and Postel-Vinay (2012), we document from the US Census’ Business Dynamics Statistics that the average size of an establishment first 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 workers per establishment at firms employing over 10,000 workers in total. So very large firms own hundreds or even thousands of separate, relatively small establishments (national banks and retailers come to mind), whose workers benefit from the productivity and compensation policy of the parent company.¹⁷

¹⁶ Mean-reverting innovations in establishment size are easily detected by the size/growth relationship. While growth in an establishment’s employment is strongly decreasing in its beginning-of-period size, it is nearly uncorrelated with the average size of the same establishment over the same period. Hence, Gibrat’s Law holds approximately, and the negative size/growth relationship originates from a classic regression to the mean fallacy, with the possible modest exception of very small establishments.

¹⁷ In his discussion of our paper, using administrative data *Integreret Database for Arbejdsmarkedsforskning* (Integrated database for labor market research, IDA) from Denmark, Rasmus Lentz reported that the variation of wages across the establishments of a typical firm, although not zero, is substantially lower than in the population of establishments as a whole. The variation of establishment size, on the other hand, is almost as large within a firm as in the wider population of establishments. We thank Lentz for pointing out this evidence, which speaks to the misclassification issue.

For both reasons, observed size classes in JOLTS and true rungs on the job ladder may not coincide. We propose to tackle those two issues and to reconcile size and rank classes by modeling misclassification explicitly. To avoid any 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., 50–249 employees), and *rank class* k , defined as the set of establishments whose unobserved rank on the job ladder falls within the quantile interval $[X_{k-1}, X_k]$.

2. Modeling Misclassification

Consider an establishment with job ladder rank x , whose “true” (or model-predicted) size at date t is $\ell_t(x)$. We assume that this establishment’s observed size is the true size $\ell_t(x')$ of an establishment with rank x' drawn at random from some conditional.

3. 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 $[\ell(X_{k-1}), \ell(X_k)]$. Observed employment in size class k is, therefore,

$$n_{kt}^o = \int_0^1 m_k(x) \ell_t(x) dx, \tag{7}$$

where $m_k(x) = M(X_k | x) - M(X_{k-1} | x)$ for all $x \in [0, 1]$ is the probability of an establishment of rank x being observed as belonging to size class k .

To gain some tractability and amenability to calibration, we further restrict misclassification weights $m_k(x)$ to be constant within rank classes, that is, we impose $m_k(x) \equiv m_{k|k'}$ for $x \in [X_{k'-1}, X_{k'}]$. With this approximation,¹⁸ (7) becomes

$$n_{kt}^o = \sum_{k'=1}^K m_{k|k'} n_{k't},$$

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

Collating all rank classes, our misclassification model implies

$$\mathbf{n}_t^o := \begin{pmatrix} n_{1t}^o \\ n_{2t}^o \\ \vdots \\ n_{Kt}^o \end{pmatrix} = \begin{pmatrix} m_{1|1} & m_{1|2} & m_{1|K} \\ m_{2|1} & m_{2|2} & m_{2|K} \\ \vdots & \vdots & \vdots \\ m_{K|1} & m_{K|2} & m_{K|K} \end{pmatrix} \begin{pmatrix} n_{1t} \\ n_{2t} \\ \vdots \\ n_{Kt} \end{pmatrix} = \mathbf{M} \mathbf{n}_t, \tag{8}$$

¹⁸ 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 resampling date.

which in turn implies that “true” employment in rank class k can be inferred from observed employment in size class k as $\mathbf{n}_t = \mathbf{M}^{-1} \mathbf{n}_t^o$. Misclassification weights $m_{k|k'}$ (the entries of the matrix \mathbf{M}) are unknown and are added to the set of parameters to calibrate.¹⁹

4. Measurement Equations with Misclassification

The transition rates $\lambda_t, \bar{\delta}_t(1)$ are estimated only off aggregate magnitudes and are not sensitive to size misclassification. With our assumption of a rank-independent probability of quitting into nonemployment, neither is 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, to nonemployment and to other jobs, from employees in rank class k are

$$Q_{kt}^o = \int_0^1 [\bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1) + \rho_{t+1} + s\lambda_{t+1}\bar{F}_{t+1}(x)]m_k(x)dN_t(x).$$

Under the assumption of constant misclassification weights in each rank class and over time, the expression for total observed quits from class k becomes

$$Q_{kt}^o = (\bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1) + \rho_{t+1})n_{kt}^o + s\lambda_{t+1} \sum_{k'=1}^K m_{k|k'} \int_{X_{k'-1}}^{X_{k'}} \bar{F}_{t+1}(x)dN_t(x).$$

This implies

$$s\lambda_{t+1} \begin{pmatrix} \int_{X_0=0}^{X_1} \bar{F}_{t+1}(x)dN_t(x) \\ \vdots \\ \int_{X_{K-1}}^{X_K} \bar{F}_{t+1}(x)dN_t(x) \end{pmatrix} = \mathbf{M}^{-1} \mathbf{Q}_t^*,$$

where \mathbf{M} is the conversion matrix defined in (8) and the vector \mathbf{Q}_t^* has K elements:

$$Q_{kt}^* = Q_{kt}^o - (\bar{\psi}_{t+1}(1)\bar{\delta}_{t+1}(1) + \rho_{t+1})n_{kt}^o.$$

Dividing by employment in rank class k and using $\mathbf{n}_t = \mathbf{M}^{-1} \mathbf{n}_t^o$, we thus obtain

$$s\lambda_{t+1} \int_{X_{k-1}}^{X_k} \bar{F}_{t+1}(x) \frac{dN_t(x)}{n_{kt}} = \frac{Q_{kt}^*}{n_{kt}}. \tag{9}$$

¹⁹ Note that, by construction, $\sum_{k=1}^K m_{k|k'} = 1$ for all k' .

This equation highlights the importance of introducing misclassification in our JOLTS data. The left-hand side of (9) is the conditional expectation of $\bar{F}_{t+1}(x)$ within rank class k ; the right-hand side 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 ranked higher 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 to decline in k . In the JOLTS data by establishment size, which is split into 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 during the sample period also between $k = 2$ and $k = 3$. We reconcile some of these observations with the job ladder by allowing some of the small establishments to be part of very large firms.

F. Implementation: Summary

For given reallocation shock arrival rate $\rho\lambda_s$, search efficiency s , and misclassification weights \mathbf{M} , using observations on employment stocks and total quits by size class, we can calculate Q_{kt}^* and the cumulated sampling probabilities at size cutoffs $\bar{F}_t(X_k)$ from equation (5) (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 (9) over the entire sample period.²⁰ Therefore, by construction, the only worker flow that our model can fail to replicate exactly are job-to-job quits by size of the current employer. This final stage of our calibration protocol thus uses $3K + 2$ parameters (the $3K$ independent entries of \mathbf{M} plus ρ and s) to match a number of moments which is equal to K times the number of months in our sample (the K moments in [9] in each month). With $K = 4$ size classes, this adds up to 14 parameters and 560 moments.

V. Results

We find that no sensible misclassification scheme can easily remedy the basic fact that the total quit rate, to nonemployment and to other establishments, originating 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 rank or desirability. The reason may be that small employers are largely

²⁰ In so doing, we add a penalty term to the criterion that we minimize (the norm of the difference between the two sides of [9]) 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 .

of a different nature than larger ones 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.²¹ 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, to calibrate the model we aggregate size of JOLTS establishments into $K =$ four classes: 1–49, 50–249, 250–999, and at least 1,000 employees. This partition, albeit coarser, still allows for significant heterogeneity and can be fitted 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 four very different size classes.

A. Calibration Results

Estimates of the various rates of separation into nonemployment and of the job finding rate were already shown in Subsections IV.B and IV.C, respectively. Here 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 conversion matrix \mathbf{M} , that is, the misclassification weights $m_{k|k'}$, $(k, k') \in \{1, \dots, K\}^2$. All those values are gathered in tables 1 and 2.

The misclassification weights in table 2 suggest that high-rank establishments (from class $K = 4$) have the largest (.65) probability of being misclassified, and they are almost always mistaken for establishments from size class 1 (1–49 employees). Apart from rank class 4, the estimated conversion matrix \mathbf{M} has most of its weight on the diagonal, suggesting that misclassification is less of an issue for low to intermediate rank levels (classes $k = 1$ –3). This finding is consistent with an interpretation of misclassification as arising primarily from the establishment/firm distinction, as some very productive—and large—firms are split into many small establishments, very often no larger than 50 employees. The calibrated matrix \mathbf{M} places some small weight on the subdiagonal, meaning that some establishments are actually seen as larger than their productivity would warrant under the job ladder assumption. We interpret this as a consequence of transitory noise or measurement error in establishment size:

²¹ As a manifestation of a similar phenomenon in the Danish matched employer-employee data set 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, Lentz, for pointing out this parallel.

Table 1
Estimated Parameters

Parameter	Estimate	Moment
ρ	.0145	Sample mean of $\rho\lambda_t = .0021$
s	.2034	Sample mean of $s\lambda_t = .0300$

Table 2
Estimated Misclassification Weights M

Establishment Size Range	Job Ladder Rank Class k			
	1	2	3	4
1-49	.977	.000	.000	.651
50-249	.023	.846	.000	.000
250-999	.000	.154	1.000	.000
1,000 plus	.000	.000	.000	.349

for example, an establishment whose long-run size is, say, 248 (and thus would normally belong to size class 2), can temporarily be seen reaching a size slightly above 250 and thus be misclassified into size class 3 (recall that JOLTS assigns establishments to size classes according to the largest size achieved over the 12 months prior to sampling).²²

The relative search intensity of employed workers is calibrated at $s = 0.203$, a value that is in the region of typical estimates based on worker microdata. This puts the sample mean monthly probability of receiving an outside offer to 0.03. Finally, the reallocation shock intensity is estimated to equal $\rho = 0.0145$. 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 forced reallocations (as opposed to voluntary transitions) is about a half (49.7% on average). This share is calculated as the sample mean of

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

The relatively large value of this share, given the relatively high odds of receiving an outside offer versus a reallocation shock ($s : \rho$ is about 14 : 1),

²² To adhere more strictly to the large firms/small establishments interpretation, we can also calibrate the model imposing that \mathbf{M} be upper-triangular. Imposing this constraint only affects the model fit very marginally, and it produces visually identical results (available upon request).

indicates that many offers are rejected by employed workers. This, in turn, is a consequence of the fact that the sampling distribution $F_t(\cdot)$ is skewed toward the lower end of its support. We now turn to the analysis of that distribution and the corresponding EE quit patterns.

B. Establishment Sampling Probabilities and Quit Patterns

Figure 13 plots the right-hand side of (9), namely, the estimated values of $s\lambda_{t+1}\bar{F}_{t+1}(X_k)$, for $k = 1, \dots, 4$ (solid lines), together with the left-hand side of (9), 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 14 further plots the estimated sampling cumulative distribution function $F_t(X_k)$ for $k = 1, \dots, 4$ (solid lines), together with $F_t^{JOLTS}(X_k)$ (dashed lines), the empirical cumulative distribution function 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 in figure 14 indicate JOLTS resampling dates.

We can see in figure 13 that our calibration ensures that the sampling distribution constructed by fitting the RPE dynamic equation (5) to net employment flow data from JOLTS is by and large consistent with the gross flow data on job-to-job quits by establishment size over the period covered by JOLTS. Although the data exhibit a slight downward trend in the job-to-job quit rates of the highest two rank classes (3 and 4), which




FIG. 13.—Rate of voluntary employer-to-employer quit. The Data series are corrected for misclassification.




FIG. 14.—The sampling distribution

the model fails to fully capture, we still conclude that the model, including its correction for the misclassification of employers into size classes, offers a remarkably good description of these data, especially considering its parsimony. In particular, EE transition rates, once corrected for misclassification, are indeed neatly ordered by rank 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 13 is that job-to-job exit rates from all but the highest rank class declined sharply during the Great Recession, especially at the lower end, and they remained low thereafter. Again, our simple job ladder model captures this pattern well, albeit with a slight lag for the lowest rank class, $k = 1$. This is one of our central findings: the Great Recession was a time when job-to-job quit rates declined sharply, not only in the aggregate, as 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 Great Recession and the recovery was almost absent.

Looking more closely at the calibrated sampling distribution (fig. 14), we first see that the empirical distribution of job openings, $F_i^{JOLTS}(\cdot)$, vastly underestimates our calibrated $F_i(\cdot)$ for all rank classes, but more severely so at the lower end of the job ladder. This is (qualitatively) consistent with the findings of Davis et al. (2013), who report that 41.6% of all hires occur at




FIG. 15.—Calibrated average class vacancy weights, normalized 01/2001 = 1

establishments with zero posted job opening in the microdata underlying JOLTS, with that proportion ranging from 76.9% for the small JOLTS size class down to roughly 7% for our largest size class. Second, there is a very slight upward time trend in the sampling distribution at all cutoff points X_k .²³ 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.

Finally, figure 15 shows the model counterpart of what we called average vacancy weights in our description of the data (Sec. II), that is, the sampling probabilities divided by the number of employers in each class,²⁴ normalized to one in January 2001 to harmonize scales. This is a measure of hiring effort by each employer per size class relative to the aggregate hiring effort. We can clearly see that, as the financial crisis unfolded, hiring effort by each employer rose in relative terms at the bottom of the size of the distribution and fell at the top. This is a symptom of a failing job ladder by employer size. Comparing figure 15 to its empirical counterpart

²³ A linear time trend is found to be positive and statistically significant for all k in both $F_t(X_k)$ and $F_t^{\text{JOLTS}}(X_k)$.

²⁴ Consistent with our procedure to correct for misclassification, we use the number of establishments in each size class in QCEW, corrected for misclassification using the conversion matrix \mathbf{M} , as our measure of the number of employers in each size class.

based on JOLTS vacancies (fig. 9), we see that our sampling weights are estimated to differ at the top of the size distribution from the JOLTS vacancy weights. In this sense, the model provides an important filter to the data.

C. Discussion

We now take stock of our results. Figure 15 indicates that, during both the 2001 recession and the first half of the 2008 recession, the vacancy weights and sampling probabilities of high-rank employers increased, while those of low-rank employers stayed flat or even declined. This fact in itself is striking in the light of Moscarini and Postel-Vinay (2012)'s finding that recessions are typically times when small (or low-rank) 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 those recessions, a hypothesis that finds some support in the raw data (fig. 10). 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, the sampling probability of the high-rank class collapsed, while that of the lowest-rank class soared, in relative terms. This, combined with a very low baseline job finding rate λ_t (fig. 12), suggests that at that point high-rank firms froze their demand for new labor and that whatever little hiring took place happened at the lower-rank end of the population of employers. This is indeed what we observe when examining JOLTS hire rates by employer size after reclassifications. Even more than in the raw data (fig. 2), hire rates rise sharply and temporarily at the lower end of the size distribution, while upgrading to better jobs slows down considerably, as evidenced by the durably low EE quit rates that ensued (fig. 13). In short, the job ladder failed, starting from the upper rungs.

Reclassification does not change, and if anything it reinforces, the qualitative time-series pattern of layoffs by establishment size that we found in the raw data (fig. 3). Layoffs significantly contributed to the increase in unemployment during the Great Recession, but the persistence of high unemployment in the 4 years after the end of the Great Recession is entirely accounted for by the failure of job finding rates to recover and the persistent increase in unemployment duration. After reclassifying establishments into rank classes so as to fit the job ladder model, the spike in layoff rates is much sharper among low-rank employers. The contemporaneous shift in sampling weights toward the bottom of the size distribution that we documented earlier suggests that the employers that were least affected by the Great Recession, especially after September 2008, took advantage of rising unemployment to hire. Because in the job ladder model each low-rank employer is more dependent on the reservoir of unemployed, it responded more, that is, cut its 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 they are unproductive. Others are temporarily small but are very productive and attractive because they are still growing. Indeed, Fort et al. (2013) draw a sharp distinction between the cyclical dynamics of net employment growth at young and old small firms in US Census data that break down net employment flows by age and size but lack 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 that it also hired more by taking advantage of high unemployment and the dynamism of young employers.

To summarize: during the Great Recession 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 hiring effort and 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 Great Recession relative to previous recessions was the credit crunch in late 2008 and early 2009. 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 renewing working capital to cover payroll at the end of each month, while attrition through quits to other employers and non-employment collapsed, so employers had to actively reduce their workforce through layoffs. 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.

VI. Conclusions

We study labor 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 employers that are ranked more highly by workers, for example, because they are higher paying, spend more hiring effort and, conditional on contacting another worker, are more likely to succeed in hiring. As a consequence, an employer's size is a relevant proxy for rank. We use newly

available monthly time series from JOLTS on employment net and gross flows by size of the establishment. We find that our parsimonious turnover model of a dynamic job ladder fits the facts well and that it implies “true” vacancy postings by size that are more in line with gross flows and intuition than JOLTS’ measures of vacancies, previously criticized by other authors. Our main finding is that the job ladder stopped working in the Great Recession and is yet to fully resume. Job-to-job quits, especially from the bottom of the size/rank distribution, collapsed, further reducing voluntary attrition and thus the incentives of small employers to post vacancies and to hire unemployed workers.

References

- Brown, Charlie, and James Medoff. 1989. The employer size-wage effect. *Journal of Political Economy* 97, no. 5:1027–59.
- Burdett, Kenneth, and Dale Mortensen. 1998. Wage differentials, employer size and unemployment. *International Economic Review* 39, no. 2:257–73.
- Coles, Melvyn, and Dale Mortensen. 2011. Equilibrium wage and employment dynamics in a model of wage posting without commitment. NBER Working Paper no. 17284, National Bureau of Economic Research, Cambridge, MA.
- Davis, Steven, Jason Faberman, and John Haltiwanger. 2013. The establishment-level behavior of vacancies and hiring. *Quarterly Journal of Economics* 128, no. 2:581–622.
- Fallick, Bruce, and Charles Fleischman. 2004. Employer-to-employer flows in the U.S. labor market: The complete picture of gross worker flows. Finance and Economics Discussion Series Working Paper no. 2004–34, Board of Governors of the Federal Reserve System, Washington, DC.
- Fort, Teresa, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2013. How firms respond to business cycles: The role of firm age and firm size. *IMF Economic Review* 61, no. 3:520–59.
- Haltiwanger, John, Ron Jarmin, and Javier Miranda. 2013. Who creates jobs? Small versus large versus young. *Review of Economics and Statistics* 95, no. 2:347–61.
- Kahn, Lisa, and Erika McEntarfer. 2014. Employment cyclical and firm quality. NBER Working Paper no. 20698, National Bureau of Economic Research, Cambridge, MA.
- Menzio, Guido, and Shouyong Shi. 2011. Efficient search on the job and the business cycle. *Journal of Political Economy* 119, no. 3:468–510.
- Moscarini, Giuseppe, and Fabien Postel-Vinay. 2009. The timing of labor market expansions: New facts and a new hypothesis. In *NBER macroeconomics annual 2008*, vol. 23, ed. Daron Acemoglu, Ken Rogoff, and Michael Woodford, 1–51. Chicago: University of Chicago Press.

- . 2012. The contribution of large and small employers to job creation in times of high and low unemployment. *American Economic Review* 102, no. 6:2509–39.
- . 2013. Stochastic search equilibrium. *Review of Economic Studies* 80, no. 4:1545–81.
- Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica* 70, no. 6: 2295–2350.
- Robin, Jean-Marc. 2011. On the dynamics of unemployment and wage distributions. *Econometrica* 79, no. 5:1327–55.